91

Pulsations in the White Dwarf Primary of a Cataclysmic Variable Star

Liza M. van Zyl

Department of Astronomy
University of Cape Town
South Africa

A dissertation submitted to the University of Cape Town in partial fulfilment of the requirements for the degree of Master of Science

December 1998

Abstract

The field of Asteroseismology provides powerful techniques for the study of pulsating white dwarf stars. If these techniques could be applied to the white dwarf primaries of Cataclysmic Variable stars, the study of accretion in interacting binary systems would benefit enormously. Our understanding of normal modes in white dwarf pulsators would benefit likewise from the study of how they respond to accretion.

This dissertation presents the analysis of preliminary observations of the Dwarf Nova GW Librae, the first Cataclysmic Variable star found to have a DAV primary. A brief introduction to Cataclysmic Variables and the pulsating white dwarfs is given, and the significance of GW Librae to these fields of study is discussed.

In this dissertation I present the analysis of 7 weeks of high speed CCD observations obtained on GW Librae during 1997 and 1998. The power in the pulsation spectrum concentrates principally in regions near 376 s and 650 s. GW Librae has a pulsation spectrum that shows dramatic changes on a monthly basis, typical of the pulsating DA white dwarfs.

Contents

Acknowledgements

I am very grateful to my supervisor, Brian Warner, for teaching me everything I know about Cataclysmic Variables, and for his ability to work financial miracles; namely, for securing funding for my MSc year at UCT, and for finding money for the computer I needed for this analysis.

I would like to thank everybody who contributed observations to the May 1998 campaign; Denis Sullivan and John Pritchard (New Zealand), Stephane Vennes (Australia), Jonathan Kemp (USA), and Frank van der Hooft and Jorge Casares (Chile). I am very grateful to the South African Astronomical Observatory for the seven weeks of observing time over the last two years that I was able to dedicate to GW Librae.

This dissertation owes its existence largely to Darragh O'Donoghue, who built the UCT High Speed CCD Photometer, and who developed the time series analysis software which I made extensive use of in GW Librae's analysis. I am deeply grateful to Darragh for the invaluable help he gave me with software and all things computing, and for devoting so much of his time to my problems.

I have been extremely fortunate to receive a great deal of advice and assistance from various people: I am very grateful to Scot Kleinman for sharing his knowledge of the DAVs; to Atsuko Nitta for sending me her thesis on accreting DBVs; to Arno Lehrer for helping me install Linux on my computer; to Don Kurtz and François Vuille for useful discussions on Fourier analysis and resolvability; to Sonja Vrielmann for help with Latex; and to Jurek Krzesiński for invaluable help with Fortran and computer hardware.

List of Tables

List of Figures

Chapter 1

Introducing GW Librae

GW Librae is a faint (V$\sim$ 18.5) and unassuming Dwarf Nova (DN) which called attention to itself for the first time in 1983 when its brightness increased by $\sim$ 9 magnitudes (González 1983) - its only observed outburst. No spectroscopic observation at maximum is available (Duerbeck 1987), but the outburst amplitude led to the classification of Nova. However, when Duerbeck & Sitter (1987) obtained spectra of the star at minimum, Nova Lib 1983 revealed itself to be a very unenergetic system: narrow H in emission, but no HeII 4686 or triply-ionized lines that one would expect to see in the spectrum of a hot nova remnant. GW Librae is most probably a long outburst interval DN, or WZ Sagittae star[1].

Two important features of GW Librae's spectrum are the very broad H absorption troughs and the single-peaked hydrogen emission lines from the disc. A disc inclined to the line of sight produces double-peaked emission lines because the light coming from the approaching half of the disc is blueshifted, and the light from the receding half is redshifted. Single-peaked lines indicate that we are seeing the disc almost face-on. This means that the central white dwarf is not obscured, and the absorption troughs are pressure-broadened Balmer lines from the white dwarf photosphere. A large fraction of the light from this system is coming directly from the white dwarf.

Its extreme faintness precluded interest in the star for the next ten years, until

[1]The meanings of these terms will be made clear in the following three chapters.

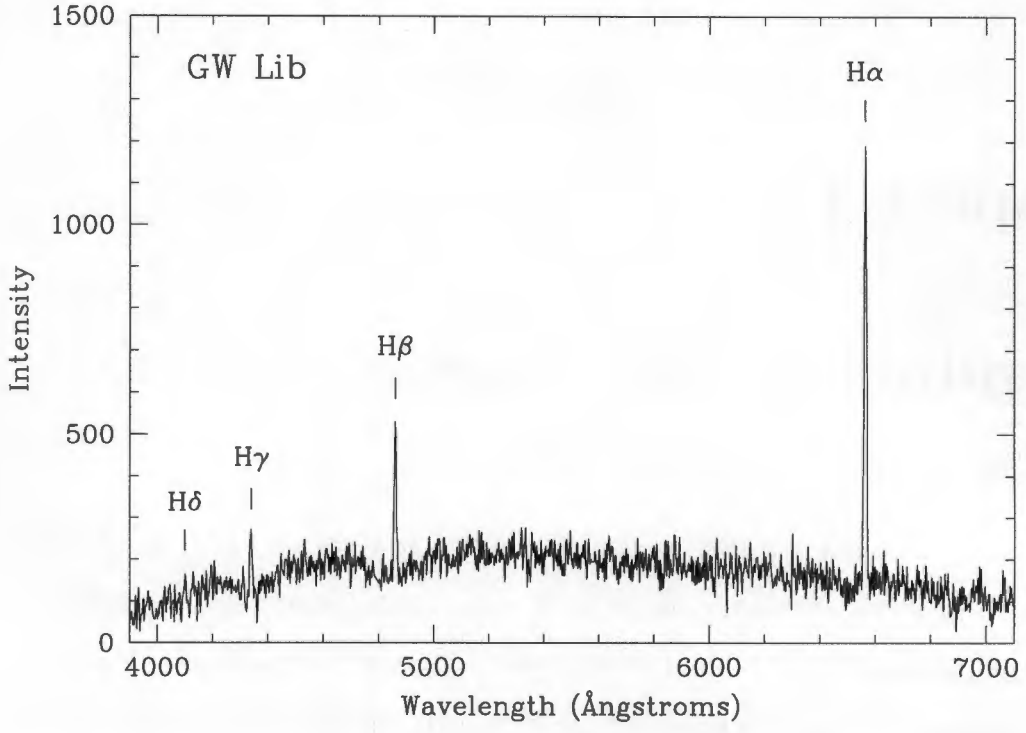

Figure 1.1: The spectrum of GW Librae at quiescence, obtained by P. Woudt with the SAAO 1.9-m telescope, 15/3/97.

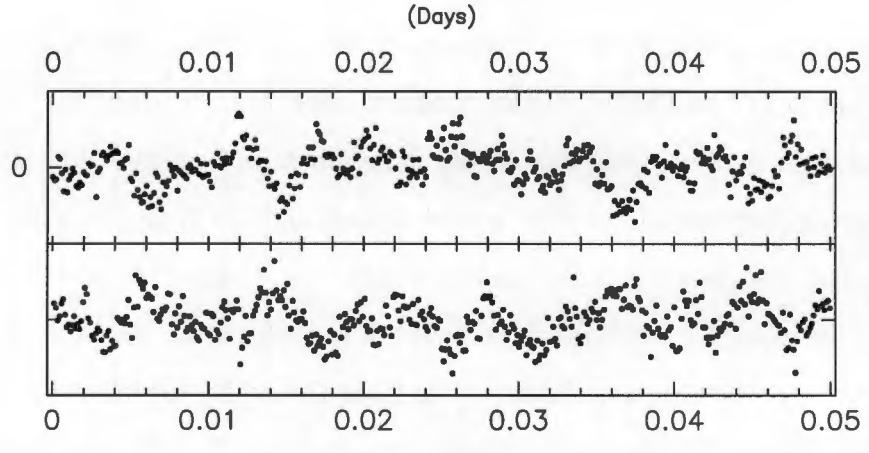

Figure 1.2: Part of the lightcurve of GW Librae, showing variations due to nonradial pulsations (obtained by the author with a CCD photometer on the SAAO 1.9m telescope, 20/5/98). The ordinate scale runs from 0.15 to -0.15 magnitudes.

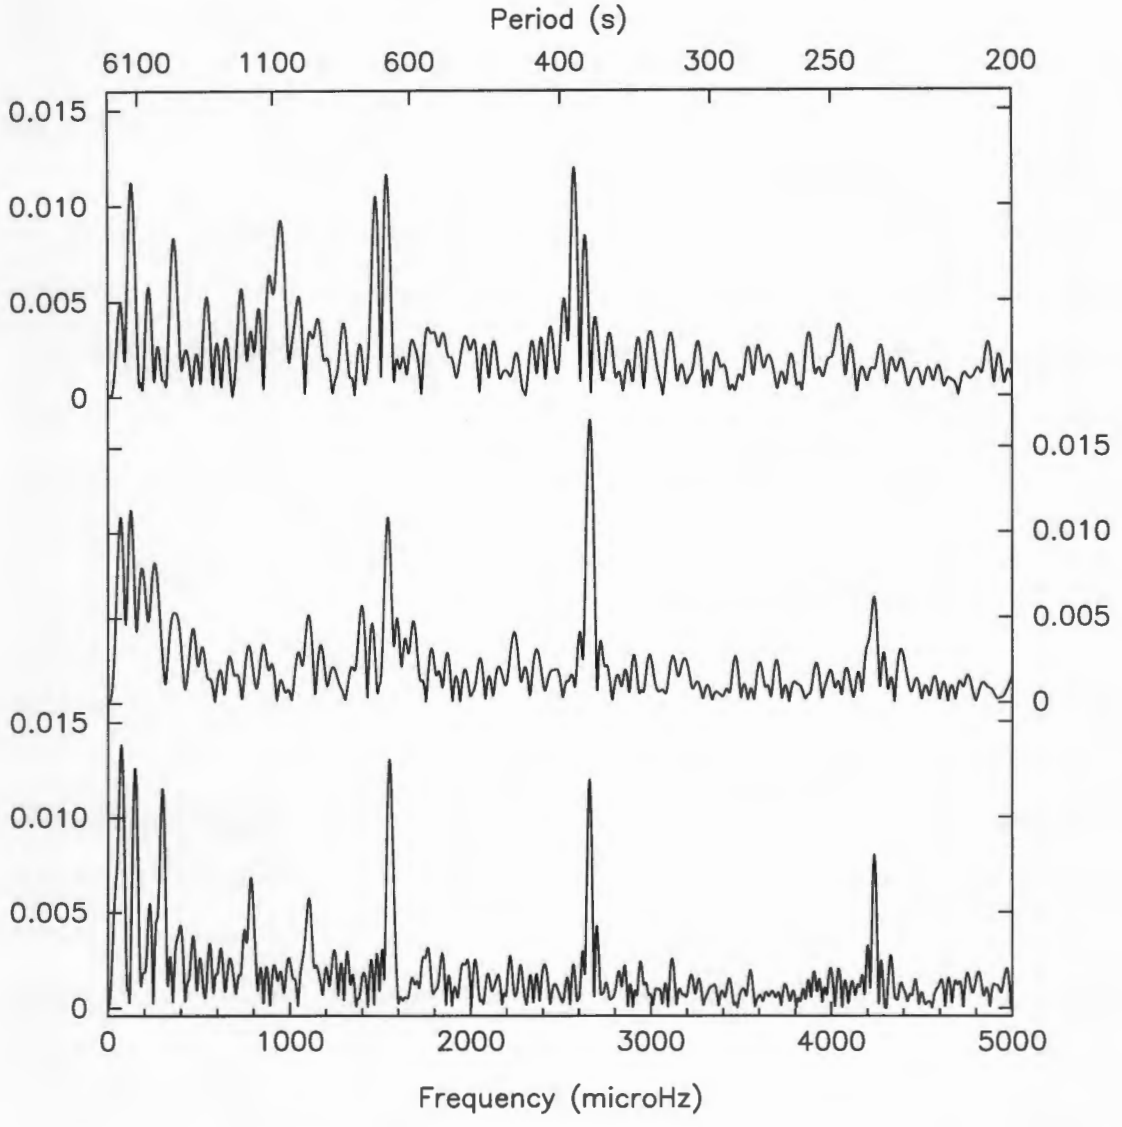

Figure 1.3: Fourier transforms of individual nights (2/4/97, 22/5/98, 25/5/98).

I chanced upon its unusual and completely unexpected properties in March 1997, when I made the first high speed photometric observations of it while investigating faint, understudied CVs with a sensitive high speed CCD photometer. The lightcurve showed what appeared to be regular periodicities, and the Fourier transform resembled that of a nonradially pulsating single white dwarf. This was the first time that nonradial pulsations had been observed in an accreting white dwarf. It had been thought that accretion would keep the primaries of Cataclysmic Variables (CVs) too hot to pulsate – only within a narrow temperature range (the *instability strip*) are white dwarfs unstable to pulsation.

GW Librae brings together two fields of research: the study of accreting binary stars, and asteroseismology, the study of pulsating stars. Because of the very different physical nature of these two classes of objects, each field employs very different research techniques and methodologies in its study. GW Librae opens up the possibility of combining research methods from each field, and perhaps extending the tools available for the study of both pulsating white dwarfs and CVs. It is likely to take many years of observation and modelling in order to understand GW Librae fully, but the potential rewards are many: asteroseismology will hopefully give us an "insider's view" of a CV, while accretion could test our understanding of white dwarf evolution.

The properties of the primaries of CVs are notoriously difficult to determine because obscuration by the disc and a complex system of emission lines due to the accreting hot gases make it very difficult to obtain meaningful radial velocity curves. If the primary is a non-radial pulsator, however, it should be possible to determine its mass, rotation period, hydrogen surface layer mass and magnetic field strength accurately from its pulsation amplitude spectrum, using analysis methods developed by the asteroseismologists.

As an Honours student at the University of Cape Town in 1997, I made a preliminary investigation of GW Librae (the results of which I presented as an Honours Project to the Department of Astronomy). As a Masters student, I conducted a much more extensive observing campaign in May and June of 1998. Sites in Chile,

New Zealand, Australia and Arizona contributed data, but due to *El Niño*, their contribution was unfortunately very small. In this dissertation I present the analysis of these data.

In the following three chapters, I present an overview of the research fields of Cataclysmic Variable stars and pulsating white dwarfs, and go on to discuss the case of GW Librae in particular. Chapter 5 deals with the observations and data reductions, and discusses the algorithms I used in the analysis. Chapter 6, the core of this dissertation, presents the analysis of the 1998 dataset. Chapter 7 examines the 1997 data, and Chapter 8 discusses possible future work.

This dissertation is a journal of the "voyage of discovery" that my MSc work has been. It charts my progress from novice to acolyte through a difficult apprenticeship in the mysterious, dark art of data analysis. In addition to learning the agonies of prising meaningful numbers from a Fourier transform, I have developed a strong appreciation for the keen sense of caution and judgement required to work with Fourier transforms that contain partially-resolved, multiperiodic signals.

This dissertation reflects my progression from a "state of innocence", in which I naively believed a Fourier transform to be a representation of stellar actuality on the same level that inscriptions on stone tablets carried down mountains by long-bearded men are a representation of divine will, to a "coming of age", in which the reality of any signal or structure found in a transform is treated to intense scepticism.

My "loss of innocence" came about with the realization, through experimentation with simulated lightcurves, that Fourier transforms of partially-resolved, closely-split signals are (at least to a very close approximation) not unique. Wildly different functions (ranging from pure harmonic series to single sinusoids modulated in amplitude, frequency or phase) can produce very similar Fourier transforms. It may not be possible to tell whether a cluster of partially-resolved, closely-split signals is caused by signals intrinsic to the star, or by a modulation of a single signal, or by artefacts in the data such as timing errors, or changes in amplitude caused by moving to telescopes of different apertures, to mention only a few possibilities.

Throughout my investigation of GW Librae, I have received invaluable advice

and assistance from astronomers who are *much* more expert than I, but the analyses and interpretations presented in this dissertation are entirely my own, and if future work proves me wrong, the fault is mine alone.

Chapter 2

CVology: the Cataclysmic Variable Stars

This chapter gives a brief intoduction to the Cataclysmic Variable stars, in order to aid later discussion on aspects on GW Librae's behaviour. Much more detailed discussions on all the issues touched upon in this chapter can be found in Warner (1995), my principal reference.

Cataclysmic Variables (CVs) are interacting binary stars, in which a white dwarf accretes material from a Roche lobe-filling secondary[1], which is usually a dwarf of spectral type K or M (but may be a red giant, for example T CrB, or a degenerate helium star as in the AM CVn stars). Mass flows from the inner Lagrangian point[2], and conservation of angular momentum leads to the formation of a disc around the primary.

The accretion stream impacts the rim of the disc in a region called the bright spot, which produces a characteristic orbital modulation in the light curve, the amplitude depending on the inclination of the system and on the relative luminosities of the

[1]The secondary is distorted by the gravitational field of the primary, and the *Roche lobe* is the largest closed equipotential surface that can contain the mass of the secondary.

[2]The inner Lagrangian point is the point on the secondary's Roche lobe that is closest to the primary. It is a saddle point of the binary's gravitational potential: here the gravitational fields of the primary and secondary are equal, and material from the secondary can spill over into the primary's potential well.

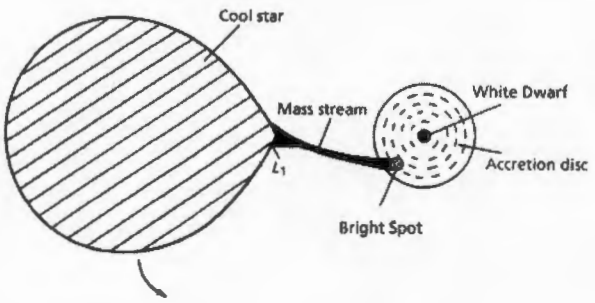

Figure 2.1: Schematic of a cataclysmic binary system. From Warner (1995).

bright spot and the disc. Blobs of gas in the accretion stream impacting the disc, or irregularities in the inner disc, can give rise to strong flickering in the light curve. Gravitational potential energy is radiated away due to viscous processes as the gas moves through the disc. Disc luminosities are typically $0.01\text{–}10 L_{\odot}$, often much higher than the combined luminosity of the component stars.

CVs are usually very blue objects (the exceptions are the long-orbital-period CVs, which have red giant secondaries and large, cool discs). The inner regions of the disc are typically $\sim 5 \cdot 10^4$K. The spectra of most CVs are dominated by chromospheric emission and by emission lines[3] from the bright spot and the boundary layer between the inner edge of the disc and the white dwarf, where infalling material moving with Keplerian velocities must match the much slower rotation velocities of the primary. Lines from the discs of intermediate to high inclination systems show characteristic double humps, due to doppler shifting of the light from the receding and approaching halves.

If the primary has a strong magnetic field ($> 10^7$G), the accretion stream is disrupted before a disc can form (Lamb 1989, Cropper 1990). The gas is threaded onto the field lines and accretes directly onto the white dwarf, forming accretion columns near the magnetic poles. Such systems, called polars, are copious X-ray and hard UV emitters because strong shocks form near the base of the accretion column as the material is accelerated to supersonic velocities by the white dwarf's

[3]Some CVs with optically thick, high viscosity discs show pure absorption lines, most notably the UX Ursae Majoris stars, a subclass of the Nova-Like Variables.

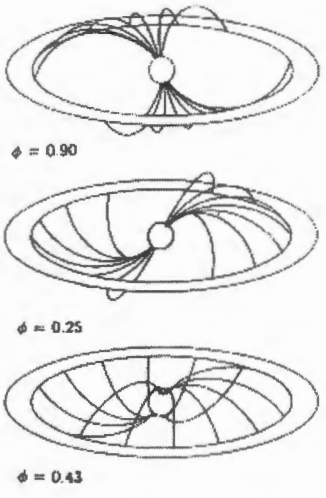

Figure 2.2: Magnetically-controlled accretion in an Intermediate Polar: schematic views of accretion from a disc into a magnetosphere, and the "accretion curtains" that form. From Hellier, Cropper & Mason (1991).

gravitational field. The magnetic field of the primary couples with the secondary, causing the primary to corotate at the orbital period. Polars show strong variable linear and circular polarization, modulated by the orbital period: the strongest linear polarization occurs when the field lines are perpendicular to the line of sight, and looking down the field lines gives maximum circular polarization.

If the primary has a weaker magnetic field, or if the mass transfer rate from the secondary, $\dot{M}(2)$, is very high, an intermediate polar – a system with a truncated disc – results. The primary does not rotate synchronously, so the lightcurves are multiperiodic, as light from the accretion columns is reprocessed by the secondary (Patterson & Price 1980, Warner, O'Donoghue & Fairall 1981).

Classification of CVs is based primarily on the morphology of their long-term light curves. The principal classes of CVs are the Novae, the Dwarf Novae, the Nova-like Variables and the above-mentioned magnetic systems.

Novae eruptions are thermonuclear runaways (Shara 1989, Livio 1993). The accumulation of accreted hydrogen onto the white dwarf eventually results in temperatures and pressures at the base of the hydrogen envelope that are high enough for nuclear burning to commence. The onset of nuclear burning occurs in a region of partial degeneracy, and since the equation of state for a degenerate gas has no

temperature dependence, the increase in temperature is not compensated for by a decrease in pressure, i.e. by the expansion of the envelope. Since the energy generation rate depends very sensitively on temperature, the onset of nuclear burning leads to a runaway increase in temperature and energy generation.

The temperature gets higher and higher until the Fermi temperature is reached and the degeneracy is lifted, at which point the equation of state becomes that of an ideal gas, $PV \propto T$. The material at the base of the hydrogen layer is then very seriously out of pressure equilibrium, and the envelope is violently blown off. The Eddington luminosity (about $10^4 L_\odot$ for white dwarfs) is exceeded, and the star brightens by 6 to 19 magnitudes.

Since their primaries are accreting hydrogen, all CVs[4] are expected to have nova eruptions. Time scales between eruptions are likely to be from 10^3 to 10^5y (although a small group of systems, the Recurrent Novae, undergo eruptions much more frequently).

Behaviour between eruptions is determined largely by the mass transfer rate from the secondary, $\dot{M}(2)$. The Nova-like Variables (NLs) have a high mass transfer rate: $\dot{M}(2) > 10^{-9} M_\odot y^{-1}$. Dwarf Novae have a lower mass transfer rate of $\dot{M}(2) \sim 10^{-11}$–$10^{-10} M_\odot y^{-1}$. The lifetime of a CV with an $0.5 M_\odot$ secondary is therefore around 10^7–$10^{10} y$. The $\dot{M}(2)$ is unlikely to be stable throughout the lifetime of the system, so a CV probably migrates from class to class during its lifetime as the mass transfer rate changes.

What drives the mass transfer? The separation of the components in binary systems in which $q = M_2/M_1 < 1$ (all CVs) should *increase*, in order to conserve angular momentum, if mass transfer occurs. The secondary would then detach from its Roche lobe and mass transfer would cease. It is likely therefore that magnetic braking by the secondary's magnetic field drives the mass transfer. Magnetic braking of single stars occurs because of the coupling of the stellar wind to the field lines which extend to infinity: outward-moving particles on the field lines are accelerated by the

[4]Other than the AM CVns, of course

star's rotation, resulting in a braking torque on and loss of angular momentum from the star (Weber & Davis 1967, Wilson 1966). In CVs, the secondary rotates synchronously with the orbital period, and so in order to conserve angular momentum in the system, magnetic braking reduces the orbital period and separation, which incidentally ensures that the secondary remains in contact with its Roche lobe (Verbunt & Zwaan 1981).

K and M main sequence stars are deeply convective and are expected to have strong, dynamo-generated magnetic fields. The magnetic fields are multipole and complex, and changes in the nature of the field due to magnetic cycles (sunspot cycles), and perhaps as a result of changes in the structure of the star due to the mass loss, will affect the braking rate, and will therefore result in long-term changes in $\dot{M}(2)$.

Gravitational radiation also removes energy and orbital angular momentum from a binary system (Kraft, Matthews & Greenstein 1962). Braking due to gravitational radiation is insignificant compared to the magnetic braking in longer orbital period CVs but may be the dominant braking force in the very short period systems.

Strong illumination of the secondary by a very hot primary or boundary layer heats the outer layers and increases $\dot{M}(2)$. After a nova eruption, the white dwarf is extremely hot (10^5–$10^6 K$) and cools over a few 10^2 years, during which time $\dot{M}(2)$ decreases by factors of $\sim 10^2$ (Kovetz *et al.* 1988; Duerbeck 1992). The Nova Remnants are a sub-class of the NLs, which include all the "non-eruptive" CVs: the UX Ursae Majoris stars, the RW Trianguli stars, the SW Sextantis stars, and the VY Sculptoris stars, and the magnetic systems. Polars and IPs exhibit both high and low states in their long-term lightcurves with gradual transitions between them, probably due to variations in $\dot{M}(2)$. The non-magnetic NLs all have a high $\dot{M}(2)$ but undergo long-term, (mostly) gradual changes in brightness. In particular, the VY Sculptoris stars spend most of their time in a "high state", varying little about a mean magnitude, but occasionally go into a "low state" for weeks to years, dropping in brightness by a magnitude or more. Again, changes in $\dot{M}(2)$ are the probable cause of these variations.

The low $\dot{M}(2)$ non-magnetic systems are the Dwarf Novae (DN), which spend most of their time in quiescence, but which have frequent outbursts recurring on timescales of weeks to years, typically with amplitudes of 2 to 5 magnitudes and lasting from 2 to 100 days. There are three subclasses of DN: the Z Camelopardalis (Z Cam) stars, the SU Ursae Majoris (SU UMa) stars, and the U Geminorum (U Gem) stars. Z Cam stars show occasional lengthy *standstills* about 0.7 mag below the maximum brightness. SU Ursae Majoris stars have occasional *superoutbursts* which are brighter and last longer than their normal outbursts. U Gem stars are all the DN that are not Z Cam or SU UMa stars. For DN in outburst the mass transfer through the disc is similar to that of the NLs: $\dot{M}(disc)$ is around $10^{-9}M_{\odot}y^{-1}$.

The difference in behaviour between the DN and the NLs can be explained by *Disc Instability* (Hoshi 1979, Cannizzo & Kaitchuck 1992). In the low $\dot{M}(2)$ systems, material accumulates in a low viscosity disc until a critical surface density is reached, usually at the outer or inner edge of the disc. A heating wave travels through the disc and it goes into a high viscosity "outburst" state. Mass flow through the disc is high, and therefore the luminosity increases by several magnitudes as the gas radiates half of its gravitational potential energy away before accreting onto the white dwarf. Mass drains out of the disc until the density drops below the critical density somewhere in the disc, at which point the outburst shuts off and the disc is back in its low viscosity, low $\dot{M}(disc)$ state, and the cycle begins again.

In the NLs, $\dot{M}(2) > \dot{M}_{crit}$, where $\dot{M}_{crit}$ is the mass transfer rate necessary to keep the disc at its critical surface density. NL discs are therefore luminous, steady state discs in which $\dot{M}(disc) = \dot{M}(2)$. In the Z Cam stars $\dot{M}(2)$ is equal or very close to $\dot{M}_{crit}$, and the standstills are periods in which the disc is in a steady state.

The SU UMa stars are very short orbital period, low q systems: all have $P_{orb} < 3h$ and $q < 0.3$. Tidal interactions with the secondary can therefore result in a precessing elliptical disc if the disc radius extends beyond the 3:1 resonance[5] (Osaki 1994). In

[5]i.e where a particle orbits the primary exactly three times in the time it takes for the secondary to orbit once. Resonance orbits are unstable because perturbations from the secondary are reinforced.

a normal outburst only around ten percent of the mass of the disc is accreted onto the primary, so the quiescent disc increases in radius after successive outbursts until an outburst extends the radius to the 3:1 resonance. This triggers a superoutburst: viscosity in a precessing elliptical disc is high because particle orbits intersect, and the disc is therefore kept in a quasi-equilibrium state in which $\dot{M}(disc) \gg \dot{M}(2)$ until enough mass has drained out to drop the radius below the 3:1 resonance.

GW Librae is a member of a subclass of the SU UMa stars, the WZ Sagittae (WZ Sge) stars. The WZ Sge stars are extremely low-mass-transfer systems. They exhibit no normal outbursts, only superoutbursts, and their superoutbursts are very rare, typically separated by tens of years. They have among the shortest known orbital periods for hydrogen-rich systems, and are therefore thought to be exceptionally old. The WZ Sge stars are discussed in detail in Chapter 4.

Chapter 3

Seismology: The White Dwarf Pulsators

3.1 Introduction

White dwarfs are the end points of stellar evolution for stars below about $8M_{\odot}$, which amount to 98–99% of all the stars in the Galaxy (Weidermann 1990; Kepler & Bradley 1995). A star on the main sequence burns hydrogen to helium, and when most of the hydrogen in the core is used up (about 10 per cent of the star by mass), the star evolves from the main sequence. Depending on the mass of the star, helium burning commences, and heavier elements are built up in the core, like carbon and oxygen (Böhm-Vitense 1989).

In the final stages of nuclear burning, a star sheds most of its mass and becomes a planetary nebula, leaving only its hot, degenerate core of spent fuel. Over billions of years, this core cools. The gravitational field within a white dwarf is strong enough to cause stratification of the elements: the heavier elements like carbon and oxygen settle to the core, and hydrogen floats to the surface where it lies on top of a helium mantle (Schatzman 1958).

White dwarfs hold many secrets in their interiors. They contain the products of nuclear burning, and an understanding of their composition would enhance our

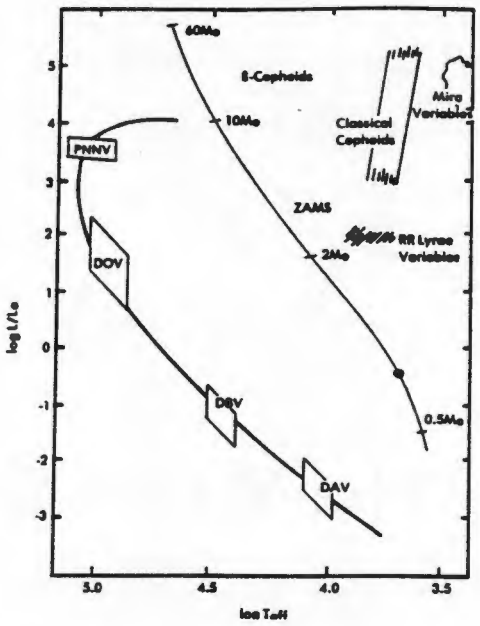

Figure 3.1: Pulsating stars in the Hertzsprung-Russell diagram. From Winget (1988).

understanding of stellar evolution. The oldest white dwarfs contain information about the first waves of star formation in the early universe. To understand the stars on the white dwarf cooling sequence, is to understand the evolution of the Galaxy (Winget *et al.* 1987; Wood 1992).

We are able to study the interiors of white dwarfs, because as they cool, they pass through four *instability strips*[1]; regions in which they become unstable to pulsations. The field of asteroseismology is the study of stellar pulsations as a means to explore stellar interiors, using pulsations as probes beneath their surfaces. White dwarfs pulsate when they have cooled enough to reach an instability strip; therefore, pulsating white dwarfs are otherwise normal stars, and the knowledge and understanding gained through the asteroseismological study of them applies to all white dwarfs[2].

The four classes of pulsating white dwarfs are the PNNVs (the planetary nebula nuclei variables, with temperatures near 100 000K), the DOVs (the hot DO pre-white

[1]There is evidence to suggest that individual white dwarfs may follow different evolutionary tracks; not all white dwarfs may pass through all four instability strips (Fontaine & Wesemael 1987; Nather, Robinson & Stover 1981).

[2]This is an oversimplification: there are some stars in the instability strips that do not pulsate (Kepler 1993, Dolez *et al.* 1991, Bradley & Winget 1994).

dwarf variables, also with temperatures near 100 000K), the DBVs (with helium-dominated spectra and temperatures near 30 000K), and the DAVs (with hydrogen spectra and temperatures near 13 000K)[3]. Figure 3.1 shows the location of these instability strips on the Hertzsprung-Russell diagram.

Oscillations are driven by the κ-mechanism: changes in ionization in partially ionized shells of the star. In white dwarfs, pulsations are driven by the ionization zones of hydrogen and helium, and carbon in the hottest stars (Dziembowski & Koester 1981, Dolez & Vauclair 1981, Winget 1981). The gas dams up radiation in the compression phase of the pulsation cycle (which heats it up and changes the ionization levels and therefore the opacity), and releases radiation in the expansion phase (which cools it, changing the ionization and opacity, and driving the next cycle). For a star to pulsate, the time scale for the partial ionization zone to store and release energy must be similar to the pulsation period of the star. This leads to a criterion for stellar pulsation that the partial ionization zone occurs at a particular mass shell – hence the instability strip (Hansen & Kawaler 1994)[4].

3.2 The Basics of Asteroseismology

All physical objects, whether they be stars, planets or Stradivari violins, have normal modes of oscillation. A bell, when struck, vibrates with a set of normal modes characteristic of its size, shape and the materials of which it is composed. Earthquakes excite modes of oscillation in the Earth, which carry information about its internal composition and structure.

Asteroseismology is the study of normal modes in stars. To determine the nature of a star's normal modes, we make use of the *local dispersion relation* for adiabatic

[3]The discovery of a pulsating white dwarf between the PNNVs and the DOVs suggests they may actually be a single class (Watson 1992, Vauclair *et al.* 1993).

[4]The cause of white dwarf oscillations is not quite as well-understood as I imply in this very brief review. It is very likely that convection drives the DA pulsations; detailed discussions are given by Brickhill (1983, 1990, 1991a,b) and Goldreich & Wu (1999).

oscillations (Unno *et al.* 1989),

$$k_r^2 = \frac{1}{\sigma^2 c_s^2}(\sigma^2 - N^2)(\sigma^2 - S_l^2)\,, \tag{3.1}$$

where σ is the angular frequency of vibration, k_r the wave number, c_s is the speed of sound in the propagating medium, N is the Brunt-Väisälä frequency (the oscillation frequency due to buoyancy), and S_l is the acoustic frequency for mode l.

$$S_l^2 = l(l+1)\frac{c_s^2}{r^2}\,, \tag{3.2}$$

$$N^2 = -g\left(\frac{dln\rho}{dr} - \frac{1}{\Gamma_1}\frac{dlnP}{dr}\right)\,, \tag{3.3}$$

and

$$\Gamma_1 = \left.\frac{dlnP}{dln\rho}\right|_{ad}\,. \tag{3.4}$$

The stellar medium is stable to pulsations for solutions of equation (3.1) that give $k_r^2 > 0$. The solution to (3.1) gives two frequency domains in which $k_r^2 > 0$:

$$\sigma^2 < N^2, \quad \sigma^2 < S_l^2 \qquad \text{the } g\text{−mode domain,} \tag{3.5}$$

and

$$\sigma^2 > N^2, \quad \sigma^2 > S_l^2 \qquad \text{the } p\text{−mode domain.} \tag{3.6}$$

There are two principal kinds of stellar pulsations: p-mode oscillations, in which pressure is the restoring force, and g-mode oscillations, in which gravity is the restoring force. The fundamental radial pulsation period for stars with the density of white dwarfs ($\bar{\rho} \sim 10^6 gcm^{-3}$) is around 10s, but the surface gravities are so high that radial oscillations are suppressed and only the non-radial g-modes can be excited. In the strong gravitational field of a white dwarf, the growth times of radial oscillations are much longer than for the nonradial g-modes, which move along the surface perpendicular to the gravitational field (Robinson, Kepler & Nather 1982). The short period p-modes are characterized by predominantly radial motions, while the long

period g-modes are predominantly horizontal motions. In white dwarfs, the periods of g-modes are typically on the order of several hundred seconds.

The intense gravitational fields of white dwarfs also ensure spherical symmetry to a very high order (except in rare cases where a star is rotating extremely rapidly), so we can model the pulsations as spherical harmonics on the stellar surfaces. The solution to the differential equations describing (low amplitude) stellar pulsations is also the one describing the eigenvalues of the hydrogen atom in atomic physics:

$$\zeta(r,\theta,\varphi,t) = \xi_k(r)\, Y_l^m(\theta,\varphi)\, e^{i\sigma t}\,. \tag{3.7}$$

The oscillations are described in terms of the indices (k,l,m). $\xi_k(r)$ is the radial eigenfunction with k giving the number of nodes between the centre and the surface of the star. $Y_l^m(\theta,\varphi)$ gives the spherical harmonics of indices l and m ($l = 0,1,2,...$; $m = 0,\pm1,...,\pm l$), and σ is the angular frequency of oscillation. l gives the number of nodal lines on the surface of the star, and m gives the number of nodal lines along a line of longitude. Each mode (k,l,m) is independent, or normal to, the others (at least in the linear approximation) (Kleinman 1995, Cox 1980).

Non-radial g-modes can be thought of as sloshing motions on the surface of the star. When material sloshes together, it compresses and heats, becoming brighter. Oscillations of g-modes have their maximum amplitudes in layers deep beneath the photosphere, and only a small contribution is seen on the surface of the star. The oscillations are visible because they produce patterns of compression and rarefaction across the photosphere, which cause changes in the temperature and optical depth. Provided that these brightness variations are not symmetrical across the visible hemisphere of the star (in which case they cancel out), we detect an overall increase and then decrease in the light from the star over each pulsation period. We only see the star as a point source, so we have a geometrically-averaged view of its pulsations. The recording of the brightness variations with time, or lightcurve, is then analysed by means of a Fourier transform to search for periodicities.

Identifying periods in the Fourier transform, and associating them with specific

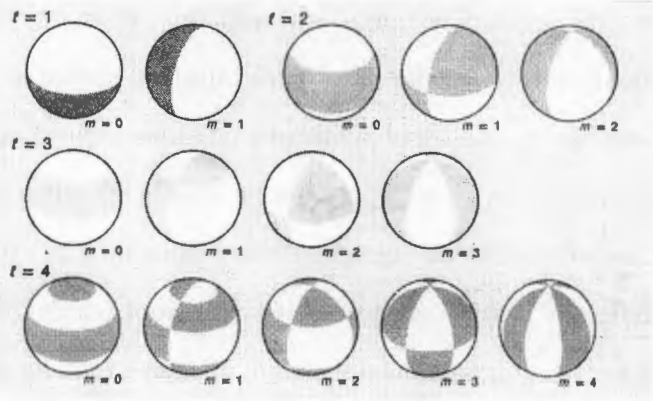

Figure 3.2: Some non-radial pulsation modes. From Winget & van Horn (1982).

values of (k, l, m) predicted by the model, yields direct measurements of various physical parameters of the star: the period spacing of high overtone g-modes of successive k is constant for a completely homogeneous star, and is related to its mass. The frequency spacings of the m-splitting of modes of the same k and l gives information about the rotation period and magnetic field strength.

In reality, the period spacing of successive k, same l modes is never constant, because white dwarfs are not chemically homogeneous. Even a pure carbon white dwarf would have a degeneracy boundary where the chemical potential changed. White dwarfs are heavily stratified: the intense gravitational field produces a relatively rapid vertical sorting of the elements, with hydrogen floating to the top and the denser elements settling inwards. It is easier for modes to oscillate if they have a node at a chemical boundary (the growth time of a mode is much shorter if it has a node at a boundary), and if a mode has a node near to such a boundary, it will tend to shift its frequency slightly to move into the boundary's "potential well". This effect is called "mode trapping", and results in shifts in the otherwise regular period spacings of successive k, same l modes (Winget *et al.* 1991, 1994; Bradley & Winget 1994). Oscillations are therefore seismic probes into the interior of the star and can be used to place bounds on the relative masses of the different compositional layers.

The mode identification is a very tricky task because modes of different (k, l, m) may have similar frequencies. In addition, the power at a particular frequency can vary from night to night, and this amplitude modulation produces sidebands in the

amplitude spectrum. The m-spacings of modes are typically very small, requiring observing runs on the order of several days to a couple of weeks to resolve. Unfortunately, any regular gap in the lightcurve is interpreted by the Fourier transform as a periodicity, and therefore the gaps in a week-long lightcurve resulting from the earth's rotation produce a pattern of aliases around any signal in the transform.

The frequencies of the aliases around a signal at frequency ν, as a result of single site observations, are given by

$$\nu_{alias} = \nu \pm \frac{n}{m} \frac{1}{\text{Sidereal Day}} \text{ Hz} \tag{3.8}$$

where m is the number of days between successive observing runs and $n = 1$ for the first alias from the highest spike, $n = 2$ for the second, etc. Aliasing represents an ambiguity in the number of cycles of a signal through a gap in the lightcurve: there may be n cycles, or $n \pm 1, \pm 2 \ldots$. The pattern of aliases produced by gaps in a lightcurve is called the "window pattern", which is simply the Fourier transform of the sampling function. In this dissertation, I calculate the window pattern by sampling a noise-free sinusoidal signal in exactly the same way as the data are sampled.

The resolution of a Fourier transform is proportional to the length of the observing run, so in order to be able to resolve closely-spaced signals in a star's pulsation spectrum, it is necessary to observe it continuously over a baseline of many days, which eliminates problems with aliasing and resolves the m-splitting of modes. The astrophysical returns of asteroseismology depend very sensitively on the number of modes identified in the transform, but the more the modes, the greater the confusion due to convolutions of interfering window patterns, and the harder it is to identify individual modes. For this reason multisite networks like the Whole Earth Telescope (Nather *et al.* 1990) are necessary for the study of nonradial oscillations in white dwarfs.

Two different physical effects can cause m-splitting of eigenmodes: rotation and magnetic fields. Both break the degeneracy in m that results from perfect spherical

symmetry. It is easier for oscillations to move parallel to field lines than perpendicular to them. Splitting due to slow rotation produces $2l+1$ m-modes with the splitting proportional to m, while splitting due to magnetic fields produces only $l+1$ m-modes with the splitting proportional to m^2 (Jones *et al.* 1989). The magnetic field increases the oscillation frequencies by "stiffening" the plasma.

The rotational m-splittings can be thought of as due to travelling waves moving around the star in either a prograde or retrograde direction. Under the slow rotation limit ($\Delta\nu/\nu \ll 1$), the m-splitting is given by

$$\Delta\nu = m(1 - C_{k,l})\Omega \tag{3.9}$$

where Ω is the stellar rotation frequency, or $\frac{1}{P_{rot}}$, and $m = 0, \pm 1, \ldots, \pm l$ (Kleinman 1995). $C_{k,l}$ is a function which depends on the star's density and composition, but in the asymptotic limit it can be approximated by $\frac{1}{l(l+1)}$ for g-modes. In practice, this limit is reached when $k = 4$, and the $C_{k,l}$ values do not deviate from this by more than 30 per cent for $k < 4$ (Bradley 1996).

For a uniformly rotating star, the m-splittings for modes of the same l will be constant. In the asymptotic limit (again this limit is reached very quickly), same l, successive k modes are equally spaced in period. The periods of such modes are given by

$$P = \frac{k\Delta\Pi}{\sqrt{l(l+1)}} + constant \tag{3.10}$$

(Tassoul 1980), where $\Delta\Pi$ is a function of the mass of the star, and is constant for a star of uniform composition. Compositionally stratified layers or any other radial discontinuities make the value of $\Delta\Pi$ different for each mode, but the mean is still related to the stellar mass. It is these deviations of $\Delta\Pi$ that give us information about the interior structure of a star: modes with nodes near the composition transition boundaries will shift their periods slightly in order that the nodes correspond as closely as possible to the composition discontinuities. This is the mode-trapping effect (Winget, van Horn, Hansen 1981).

In the asymptotic limit, equations 3.9 and 3.10 show that the ratio of the spacings for modes of different l, and the ratio of the m-splittings depend on l only. The period spacing ratio for $l = 1$ and $l = 2$ is $\sqrt{3}$ or 1.73, and the multiplet spacing ratio is 0.60. These simple relationships are very useful in the mode identification process.

Long term observations of eigenmodes in a star provide us with useful information about stellar evolution. From Winget, Hansen & van Horn (1983), the rate of period change with time, $\dot{P}$, for a mode with period P, is related to a star's temperature T and radius R in the following way:

$$\frac{\dot{P}}{P} \approx -\frac{1}{2}\frac{\dot{T}}{T} + \frac{\dot{R}}{R}. \tag{3.11}$$

As a star cools, the period of a mode increases. If a star is contracting, the period decreases (contraction dominates in the hottest and youngest of the white dwarf pulsators, but after that the radius remains constant). These tiny $\dot{P}$s are extremely important measurements because they are related to the ages of the stars, which can tell us something about the age of the local galactic disc or of the cluster to which the stars belong, and about their early star formation histories.

3.3 Case Studies

A further discussion of nonradial stellar pulsations would benefit from a discussion of the analyses of particular stars. In my analysis of GW Librae, I have drawn heavily on the work done by members of the WET on the analysis of non-radially oscillating white dwarfs. In the following sections I will briefly describe the results of the definitive papers that presented the first successful analyses of stars in the instability strips, and in doing so introduce concepts in asteroseismology that may have relevance to GW Librae.

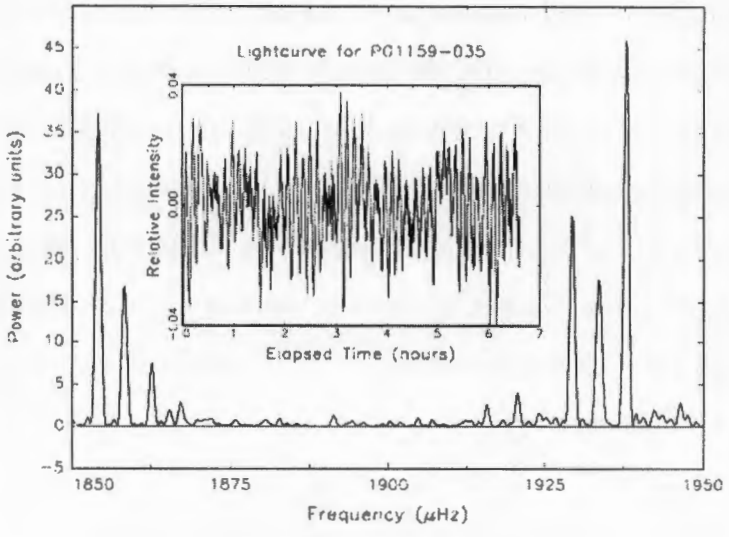

Figure 3.3: Part of the lightcurve and power spectrum of PG 1159-035. Two m-split eigenmodes are shown. From Hansen and Kawaler (1994).

3.3.1 The DOVs: PG 1159-035

PG 1159-035 was the first star successfully analysed by the WET and was a triumph for asteroseismology[5]. In 1989 the WET obtained 264.1 hours of nearly continuous data on this star, the power spectrum of which was completely resolved into 125 individual frequencies. 101 of these frequencies were identified with specific quantized g-modes. The power spectrum contained series of $l = 1$ triplets and $l = 2$ quintuplets. The m-splittings, the m-component ratios and the $(l = 1)/(l = 2)$ ratios were almost exactly what theory predicted.

The mode identification permitted the precise determination of several physical parameters of the star. The mass was determined to be $0.59 \pm 0.01 M_{\odot}$, and the rotation was found to be uniform with a period of 1.38 ± 0.01 days. The deviations from uniform period spacing for modes of the same l permitted an investigation of the compositionally stratified interior. Model fits gave a mass of $\sim 0.004 M_{\star}$ for the helium-rich surface layer. From the mode identification, an absolute luminosity could be derived; this gave a distance to PG 1159 of 440 ± 40 parsec, much more precise than the previously measured parallax of 800 ± 400 parsec.

[5]Principal references for this section are Winget *et al.* (1991) and Kleinman (1995).

The m-splittings were found to be due entirely to rotation and not to magnetic fields. The m-splittings of individual $l = 1$ modes were uniform to better than 5% (a value constrained by the resolution and therefore by the run length). This gives an upper limit to the magnetic field strength of < 6000 Gauss. In addition, the m-splittings were found to be the same within $\sim 1\%$ for both the triplets and the quintuplets; this implies uniform rotation to within observational limits. Modes of the same l but different k sample the interior of the star at different depths, therefore a nonuniformly rotating star would show different m-splittings for modes of the same l but different k (Winget *et al.* 1994).

Using the assumption that the m-split components of a single mode have intrinsically the same (long-term) amplitude, and that an average of the m components of all the modes is symmetric about $m = 0$, the investigators determined the inclination of PG 1159 to be $\sim 60°$ to the line of sight. The rationale for this is easily seen in Figure 3.2: the observed amplitude for the $m = 0$ component of the $l = 1$ mode would be zero if the star were observed equator on, and at full amplitude if the star were observed pole on (the $(k, 1, 0)$ mode is an alternating brightening and darkening of each hemisphere). The observed amplitudes for the $m = \pm 1$ components of the $l = 1$ mode would be zero if the star were observed pole on, and at full amplitude if observed equator on. Therefore, provided that the intrinsic means of the m-components of the $l = 1$ modes are equal (as yet, theory cannot tell us if they should be or not), the inclination of the star can be measured.

The analysis of PG 1159 was a great success for asteroseismology because of the close match between the observations and the theory. For example, the ratio of the m-splittings of the $l = 1$ and $l = 2$ modes was found to be 0.61, close to the ratio of 0.60 predicted by theory, and the ratio of the mean k period spacings of the $l = 1$ and $l = 2$ modes was found to be 1.72, close to the predicted value of 1.73. However, the analysis still left many questions unanswered; for example, why were the average m-components of the $l = 2$ modes so asymmetric about $m = 0$? For very similar pulsation time scales and eigenfunctions, we would expect very similar growth rates, and the amplitude of a mode is related monotonically to its growth

rate (Dziembowski 1988). However, the $l = 2$ modes of PG 1159 show that this is apparently not the case; there is a lot that we do not understand about growth rates and the nonlinear mechanisms that determine the amplitudes of individual modes, or about the mode selection mechanisms due to adiabatic processes.

The relationship between the amplitude and the growth rate of a mode has led to speculation in the literature that modes with the largest amplitudes probably have the fastest growth rates. The modes with the fastest growth rates should be those which are trapped in the outermost layers (Winget, van Horn, Hansen 1981; Dolez & Vauclair 1981; Winget & Fontaine 1982). However, in PG 1159 there appears to be no relation between trapped modes and modes of the highest amplitude: the expected correlation was not found.

3.3.2 The DBVs: GD 358

In 1990 the WET obtained 154 hours of nearly continuous high-speed photometry on the DBV GD 358[6], in the first successful asteroseismological analysis of a DBV. More than 180 signals were identified in the power spectrum. The power spectrum is dominated by $l = 1$ triplets; although there is evidence for some $l = 2$ modes, they are extremely low in amplitude and barely above the noise.

By constructing pulsation models and adjusting their parameters until their eigenfrequencies matched those of the star as closely as possible, Bradley and Winget (1994) were able to identify each triplet uniquely with a specific value of k. From the mean k period spacings they determined a mass of $0.61 \pm 0.03 M_{\odot}$, and from the deviations of each period spacing ΔP from the mean, they found a mass for the helium layer of $2.0 \pm 1.0 \cdot 10^{-6} M_{\star}$, a value around 10 000 times smaller than the value predicted by stellar evolution theory (Iben 1989; D'Antona & Mazzitelli 1991).

They were also able to find the absolute luminosity to be $0.050 \pm 0.012 L_{\odot}$ and a distance to the star of 42 ± 3 parsec, very close to the parallax measurement of 36 ± 4 parsec (Harrington *et al.* 1985). This derived luminosity provides a check

[6]The principal references for this section are Winget *et al.* (1994), and Kleinman (1995).

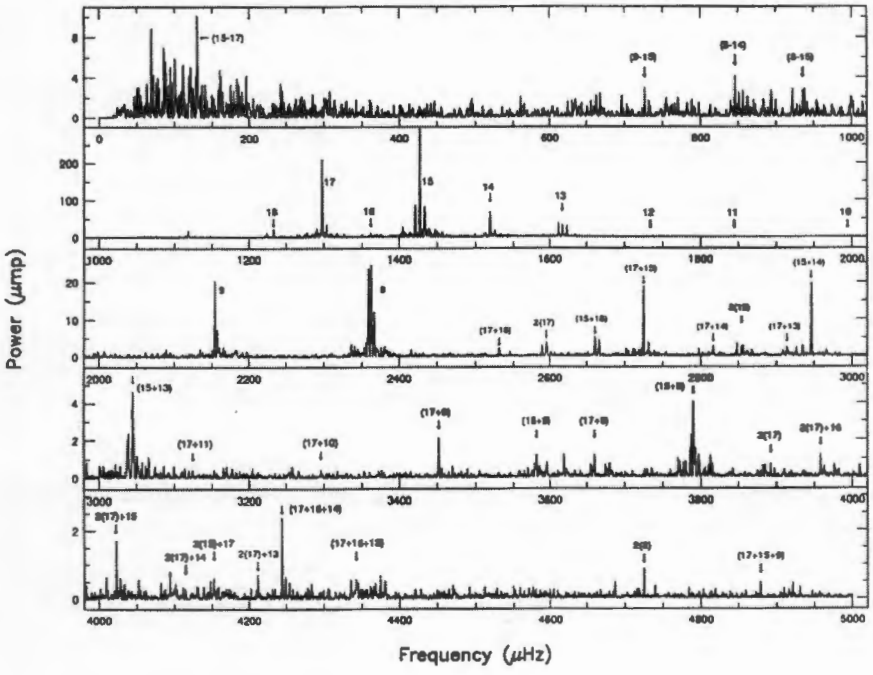

Figure 3.4: The power spectrum of GD 358, showing a succession of $l = 1$ multiplets and their linear combination modes. The index numbers correspond to the values of the radial overtones k, or their combinations. From Winget *et al.* (1994)

on the mode identification: if all the triplets were actually $l = 2$ or higher modes with only three components in each excited, then the derived luminosity would give a distance of ~ 75 parsec or greater.

A very interesting finding was that the m-splittings of the $l = 1$ modes increased with k. Modes of the highest k (16,17) had m-splittings ~ 1.8 times those of the lowest k (8,9). Since modes of higher k preferentially sample shells closer to the surface, it would appear that GD 358 is rotating differentially: the outer envelope is rotating 1.8 times faster than the part of the outer core sampled by the low k modes[7]. It is possible that as GD 358 cooled on an evolutionary time scale, its outer layers contracted much faster than the degenerate core, conserving angular momentum.

In addition, the splittings for the $m = +1$ (prograde) modes are systematically

[7]This interpretation is now considered to be too simplistic (Kleinman 1998a, Vuille 1998a). We need to consider the integrated effect of differential rotation on the mode throughout the star, not just on its outermost nodes. In this case, the difference in m-splittings is far too large to be due to differential rotation alone, but as yet, there are no alternative explanations.

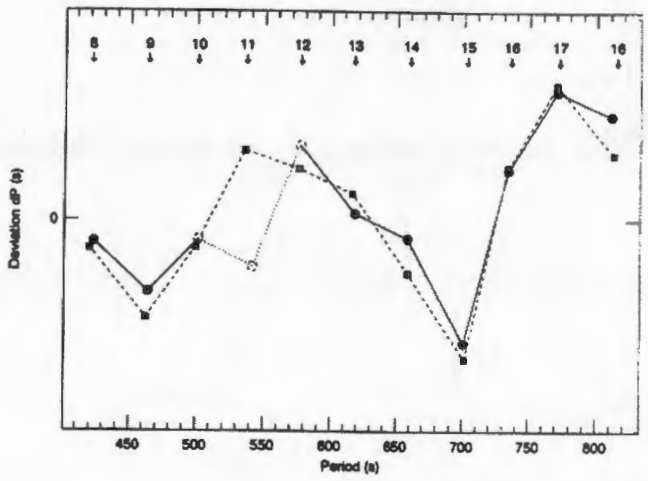

Figure 3.5: The interior of GD 358: Observed deviations of individual period spacings from the mean (circles; open circles for multiplets in which the $m = 0$ frequency identification is uncertain), with the theoretical model that best fits the data (squares). These deviations give information about the interior structure of the star. From Winget *et al.* (1994).

larger than those for the $m = -1$ (retrograde) modes. This asymmetry indicates the presence of a weak magnetic field – one which perturbs but does not dominate the motions of the stellar plasma, and which modifies the rotationally split modal frequencies by an amount proportional to m^2, hence the asymmetry with respect to the central frequency (Jones *et al.* 1989). The field strength in GD 358 was found to be 1300 ± 300 Gauss, a value two orders of magnitude smaller than can be measured from Zeeman splitting.

Follow-up investigations of the star reveal that some modes changed their amplitudes but not their frequencies, while others changed slightly in frequency but not in amplitude. These effects suggest the presence of a changing magnetic field associated with the outer layers. The observed magnetic field may possibly be internally generated by pulsation-driven plasma motions, convection or shearing due to differential rotation (Markiel, Thomas & van Horn 1994).

A very noticeable feature of the power spectrum is the multitude of linear combination modes. If the medium in which two oscillations propagate can respond in a linear manner, the frequencies retain all of their power. If it cannot, they lose some of their power to other frequencies than are the sums and differences of the

original frequencies. This is a non-linear effect, and since the luminosity changes in the lightcurve result from changing temperature (Robinson, Kepler & Nather 1982), it is unsurprising to find linear combinations in the transform, as $L \propto T^4$ is a nonlinear transformation. In addition to harmonic distortion (the inability of the propagating medium to respond linearly to the full amplitude of the oscillation), resonant mode coupling could be operating here (existing eigenmodes near the combination frequencies of active modes could be excited by resonance (Dziembowski 1982)). However, the plethora of linear combination modes in GD 358 suggests they are caused primarily by harmonic distortion.

The exact mechanism for these nonlinear effects is not fully understood. There are several possibilities; for example, the nonlinear transformation from temperature to luminosity mentioned above. The temperature variations themselves may be nonlinear: if the pressure variations in the driving region are sinusoidal (linear), then the temperature variations are not (Brickhill 1992).

As with PG 1159, the amplitudes vary with m within and between multiplets. The reasons for this are still a mystery.

One aspect of DBV research which may be of relevance to GW Librae involves the evolutionary processes that produce DB white dwarfs. It is thought that the DBs can form in two possible ways; either from single stars (Fontaine & Wesemael 1987, Shipman 1989), or as the end product of interacting binary white dwarf (IBWD) evolution (Nather, Robinson & Stover 1981, Provencal 1993). An IBWD, or AM CVn star, is a CV consisting of a pair of white dwarfs, which ultimately evolves into a DB. Work by Nitta (1996, 1998; Nitta & Winget 1997) indicates that a history of accretion should have an observable effect on the eigenmodes of a DBV. I will discuss this further in Chapter 4.

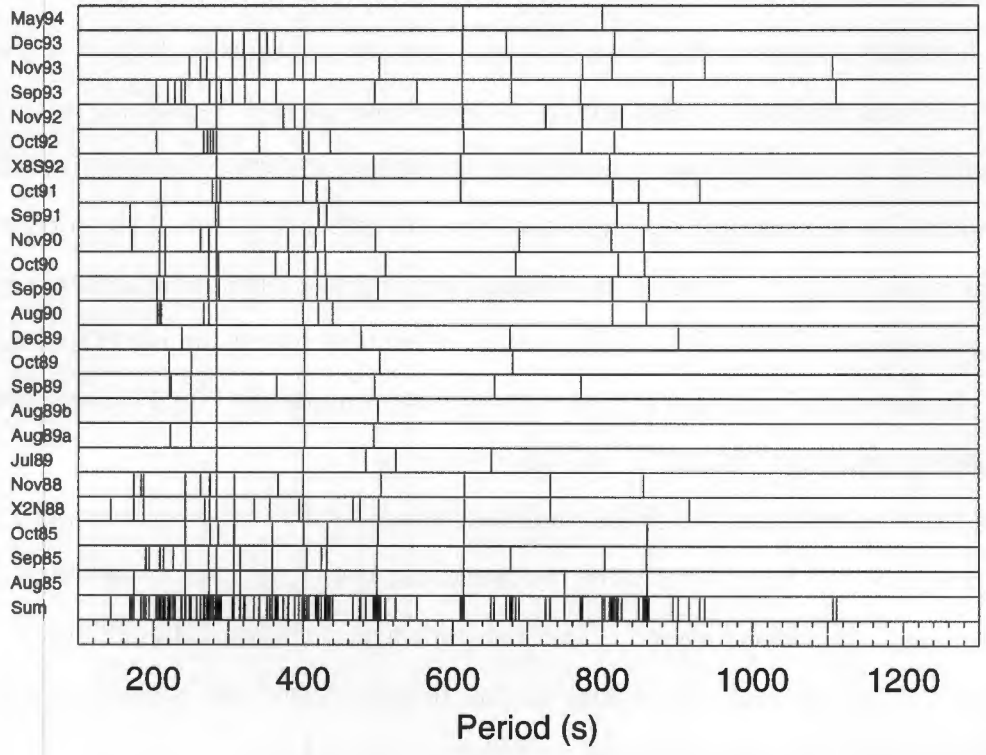

Figure 3.6: Schematic diagram of the CDAV G29-38's periodicities for the entire dataset. It is very noticeable how unstable the pulsation spectrum is from year to year.

3.3.3 The DAVs: G29-38

The DAVs (also called the ZZ Ceti stars; this is the class of which GW Librae is a member), are the coolest and oldest of the white dwarf pulsators[8]. These stars have spent most of their existence on the white dwarf cooling track. Their ages provide important constraints on the ages of the galactic disc and bulge, and of globular and open clusters. The age of a DAV can, in principle, be measured directly from its cooling rate, which is a function of its mass and internal structure. This is where asteroseismology comes in. Unfortunately, the DAVs are also the most complex and difficult to understand of the white dwarf pulsators. Despite the DAV HL Tau 76 being the first pulsating white dwarf discovered (Landolt 1968), the DAVs are still far from being fully understood.

[8]Principal references for this section are Clemens (1993, 1994), Kleinman (1995, 1998a,b), Kleinman *et al.* (1998a), and (Bradley 1996).

The DAVs are often divided into two camps according to temperature. The hot DAVs (HDAVs) tend to have only a few low amplitude, short period modes, while the cool DAVs (CDAVs) have many modes, with longer periods and higher amplitudes. The modes of the HDAVs tend to be stable from season to season, while the modes of the CDAVs are typically very unstable: they tend to come and go, and wander about. The amplitude spectrum of a single CDAV can change its appearance drastically from season to season. The CDAVs typically show linear combination modes (sums and differences of other modes) in their amplitude spectra, thought to be signs of non-linearity and which are not very useful for traditional linear asteroseismological theories. The linear combination modes almost certainly contain a lot of information about pulsation driving mechanisms and the response of the stellar material in the driving cavity, but as yet we do not have access to it.

The success of an asteroseismological study of a star is strongly dependent on the number of modes identifiable. The HDAVs each have too few modes to be solved individually, but Clemens, using a technique called *ensemble asteroseismology*, was able to solve the HDAVs as a class by looking at the ensemble of individual pulsators. He discovered properties global to all the HDAVs, and was then able to determine properties of individual stars in relation to this. He found the modes in the HDAVs to be mostly $l = 1$ eigenmodes, but only a few are present in any given star. Adding the modes of all the stars together gave an understandable pattern of successive k modes.

The averaged modes of all the HDAVs can be used to determine group properties and mean parameters of the stars, and the deviation of each individual star's modes from the superset can be used to measure individual parameters: we are able to perform asteroseismology on an HDAV despite never seeing its complete set of modes.

The implication of Clemens's success is that the HDAVs are all very similar in structure. They have masses of $\sim 0.6 M_{\odot}$ and H envelopes of $\sim 10^{-4} M_{\star}$. The strong similarities between the HDAVs support the asteroseismology credo that the pulsators are "otherwise normal stars". Clemens's work highlights our lack of understanding of the mode selection process. Why do stars show some modes but not

others? Why do we not always see multiplets, and what determines the amplitudes of the modes, and of the multiplet components?

Clemens was unable to fit the CDAVs, with their complex and variable amplitude spectra, into the picture. Kleinman, however, was able to show that a single CDAV, G29-38, can be analysed in terms of a series of successive k, $l = 1$ g-modes, and therefore, by implication, that the cool DAVs are the same kind of animal as the other white dwarf pulsators.

The amplitude spectrum of G29-38 is not stable: it changes dramatically from one observing season to another, with smaller changes within a season. Kleinman was able to come to a partial understanding of the star by combining many seasons of intensive observations (spanning ten years), and by being able to identify the true eigenmodes from the plethora of linear combination modes in the combined spectrum.

The combined amplitude spectrum, minus the linear combination modes, was found to be consistent with a succession of 19 $l = 1$ modes (not all of the eigenmodes were able to be identified with certainty or unambiguity). Some of these modes seem to wander around a bit in frequency from season to season. Only a few modes showed multiplet splitting, and the splitting size, $\Delta\nu$, could vary by $\sim 10\%$ from season to season. The reason for this is not known, but indicates that something fundamental about our understanding of normal mode pulsations in white dwarfs is amiss.

Kleinman emphasizes the need for a "different kind of dataset" in order to be able to understand these stars:

> A single set of observations over a single observing season, no matter how well resolved it is, will not suffice. There must be well-sampled observations over many observing seasons.
>
> *Kleinman et al. 1998a*

Kleinman *et al.* (1998b) have attempted ensemble asteroseismology on CDAVs for which there are many seasons of data. They find concise groupings of modes every 50s from 350 s to 650 s, a pattern which suggests we may be seeing $l = 1$,

successive k modes. Combining these results with Clemens's HDAV ensemble of modes reveals an almost exact overlap between the longest period HDAV modes and the shortest period CDAV modes. These results indicate that all the DAVs have similar masses and structures: as a DAV cools (progresses from being a HDAV to a CDAV), the pulsation driving zone moves deeper into the star, providing more energy (and therefore higher amplitudes) for the pulsations, and an increased thermal timescale (and therefore longer periods) for the driven modes.

Chapter 4

CV-seismology: The Linking of Two Fields of Research

4.1 GW Librae and the WZ Sagittae Stars

GW Librae is a member of a small subclass of the Dwarf Novae, the WZ Sagittae stars[1]. The WZ Sagittae stars are characterized by exceptionally long outburst intervals, T_s, and almost all of them exhibit only superoutbursts. They have among the shortest known orbital periods for Hydrogen-rich CVs, as well as extremely low mass secondaries, indicating that probably most of these systems are exceptionally old. They have extremely low mass transfer rates, $\dot{M}(2)$, on the order of $1\text{–}3{\cdot}10^{-11} M_{\odot} y^{-1}$.

The shortest possible orbital period, P_{min}, for a hydrogen-rich CV is $\sim$ 75 min (Paczynski & Sienkiewicz 1981, Rappaport, Joss & Webbink 1982). When the mass of the secondary drops below $0.08 M_{\odot}$, it becomes fully degenerate and stops burning hydrogen, becoming a very low mass hydrogen white dwarf. White dwarfs in thermal equilibrium obey an inverse mass-radius relationship, and therefore a decrease in mass leads to an *expansion* of the secondary. This process, combined with the fact

[1]The WZ Sagittae stars are not really a distinct subclass anymore, as they have been assimilated into the SU Ursae Majoris stars (O'Donoghue *et al.* 1991), but for the purposes of this discussion it is convenient to continue to treat them as such. Various definitions have been used to classify stars as members of the WZ Sge subclass; here I use a very broad definition, in which I consider members of the WZ Sge subclass to be those CVs which are most likely to have DAV primaries.

that mass transfer leads to an increasing orbital period, results in a period minimum. (Due to orbital angular momentum loss, the secondary makes contact with its Roche lobe at smaller separation if it shrank after transferring mass, or at larger separation if it expanded.)

Since the WZ Sagittae stars have periods close to $\sim$ 75 min, they are either close to or have passed through the period minimum, indicating that (unless they formed with P_{Orb}s very close to P_{min}) they are probably very old ($\sim 10^9$ years or more).

In Table 4.1 I list the known WZ Sge stars. Not all these stars are pure superoutbursters; some show normal outbursts too (Warner 1998), but all are characterized by exceptionally long outburst intervals, indicating very low mass transfer rates.

Table 4.1: The WZ Sagittae stars

Star	P_{orb}(h)	T_s(d)	Reference
WX Cet	1.25:	~1000	1
WZ Sge	1.36	12000	1
SW UMa	1.36	954	1
HV Vir	1.39	~3500	1
GW Lib	?	>5000	
UZ Boo	3?	~360	2,3
RZ Leo	1.70	?	3
VY Aqr	1.52	?	2,3
EG Cnc	1.38	?	3
IO Del	?	?	2
DV Dra	?	?	2
Hya 1	?	?	2
V522 Sgr	?	?	2
YY Tel	?	?	2

References: 1. Warner 1995; 2. Downes, Webbink and Shara 1997; 3. Ritter and Kolb 1997

There are various theories proposed for the extremely long T_s (Smak 1993, Osaki

1994). As mentioned before, DN discs in quiescence have a very low viscosity, so the accreted material tends to pile up on the outer rim of the disc until the critical surface density is reached and an outside-in outburst ("inburst"?) is triggered. However, in the case of the WZ Sagittae systems, it seems likely that because of the extremely low $\dot{M}(2)$, the incoming material has time to distribute itself more evenly through the disc despite the low viscosity. In this way the discs can delay the onset of outburst for a considerable length of time, and in addition they can accumulate a great deal of mass by the time they eventually do go into outburst.

This would help to explain other anomalous features of the WZ Sagittae stars, namely their very large outburst amplitudes ($\sim$ 9 mag), the long duration of the outbursts (a lot more mass needs to be drained away), and the fact that almost all outbursts trigger superoutbursts (the disc is so massive that an outburst inevitably extends the radius beyond the 3:1 resonance[2]).

Although the above argument gives long T_s values it cannot fully account for the length of the observed intervals, so there are other factors at work. The critical surface density required to trigger an outburst is roughly proportional to the radius of the disc, so if mass is distributing itself evenly through the disc, an inside-out outburst ("outburst" as opposed to "inburst") will result. If the primary of a WZ Sagittae-type system has a weak magnetic field, the innermost parts of the discs would be removed, thereby requiring a higher surface density, and thus a longer accumulation time, before an inside-out outburst is triggered.

Another possible way to get very large values for T_s is if the viscosity of the discs is far lower than that of other DN discs in quiescence, causing material to pile up slowly on the outer rim (Smak 1993, Osaki 1994). Since the mechanisms for generating viscosity in CV discs are not fully understood, it is not known why the discs of WZ Sagittae systems should have viscosities much lower than the discs of other DN, but the suggestion should be kept in mind as a possibility until more is known. There are,

[2]When the disc's viscosity increases, the radius also increases, because conservation of angular momentum requires that some material gains angular momentum and moves outwards to counteract the loss of angular momentum from the material moving inwards.

however, predictions that can be made: if a magnetic field removing the innermost part of the disc is responsible for the large outburst intervals, then we would expect inside-out outbursts, but if ultra-low viscosities are the culprit, then the outbursts would be outside-in[3].

An important consequence of the low mass transfer rate is that the primary in a WZ Sagittae-like system is not excessively heated by accretion. Typical temperatures of the outer layers of CV primaries (T_{eff}) range from 10–50·10^{3}K because accretion heats the outer layers. These temperatures are maintained even in Dwarf Novae in quiescence when $\dot{M}(disc)$ is minimal because (T_{eff}) depends on the *time averaged* accretion and therefore on $\dot{M}(2)$ rather than on $\dot{M}(disc)$. The luminosity of the primary depends on the rate at which mass is falling into its potential well plus whatever its intrinsic internal luminosity L_{int} may be:

$$L_{WD} = 4\pi R^2 \sigma T_{eff}^4 \propto \frac{GM\dot{M}(2)}{R} + L_{int}\,. \tag{4.1}$$

In practice, L_{int} is insignificant, because these systems are usually very old, and so the interior of the white dwarf will have cooled. (The isothermal core would not have cooled below the surface temperature, but the flow of heat to the surface when the gradient becomes low will be negligible compared to the accretion luminosity.) Because the outer layers conduct heat relatively poorly, very little accreted heat will reach the deep interior of the star.

With their low values of $\dot{M}(2)$, the WZ Sagittae stars have cooler primaries, with accretion keeping T_{eff} at around 10–15·10^{3}K, very close to the instability strip for the DA white dwarfs. Except in the case of extremely low $\dot{M}(2)$ systems like the WZ Sagittae stars, accretion keeps the primaries too hot to pulsate.

In Table 4.2, I list the long-term average mass transfer rates, $\dot{M}$, and the temperatures of the primaries, $T_{eff}(1)$, for a selection of CVs, in order to show the

[3]This can be determined if the rise of an outburst is observed simultaneously in the visible and in the UV (Smak 1984a,b, 1987): in inbursts the UV rise lags the visible by a few days (temperature in the disc is inversely related to radius because material at smaller radii is deeper in the primary's potential well).

relationship between the mass transfer rate and the primary's surface temperature. The values for $\dot{M}$ are notoriously difficult to determine, so the numbers I quote here are approximate only, but the trend is clear: higher mass transfer rates lead to hotter primary photospheres. Figure 4.1 shows this trend clearly.

Table 4.2: Mass transfer rates and white dwarf temperatures.

	P_{orb}	$\dot{M}$	$T_{eff}(1)$
Star	(h)	$M_{\odot}y^{-1}$	10^3K
WZ Sge	1.36	$1 \cdot 10^{-11}$	12.5
SW UMa	1.36	$\sim 10^{-11}$	10-15
Z Cha	1.79	$3 \cdot 10^{-11}$	15.6
OY Car	1.52	$3 \cdot 10^{-11}$	16.5
U Gem	4.25	$5 \cdot 10^{-10}$	30
MV Lyr	3.21	$\sim 10^{-9}$	50
TT Ari	3.30	$\sim 10^{-9}$	50

From Warner (1995) p82, 114, 186, 479.

For a given mass transfer rate, the temperature of the primary's photosphere depends on the structure of its boundary layer (a large, diffuse boundary layer will radiate accretion-dumped energy more efficiently than a dense boundary layer that forms a thin shell above the primary's photosphere; a boundary layer with a larger "surface area" radiates more effectively), and on its mass. Treating the white dwarf as an ideal gas of α-particles supported by a massless Fermi-Dirac degenerate electron gas, one can from Statistical Mechanics derive a (first order) mass-radius relationship for white dwarfs; see for example the derivation given in Cleymans (1997):

$$R_{WD} \propto M_{WD}^{-\frac{1}{3}}. \tag{4.2}$$

Equations 4.1 and 4.2 then give a relationship between the temperature, mass and

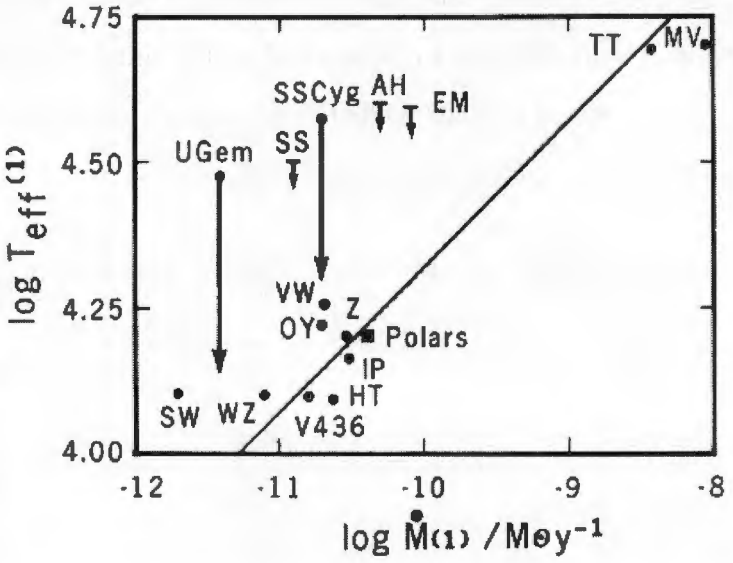

Figure 4.1: Effective temperatures of CV primaries as a function of mass transfer rate. The arrows show the effect of allowing for the larger masses of the primaries of U Gem and SS Cyg. From Warner (1995).

mass transfer rate for CV primaries:

$$T_{eff}^4 \propto M_{WD}^2 \dot{M}(2)\,. \tag{4.3}$$

Therefore, for a given mass transfer rate, a more massive white dwarf is heated more, first because the accreted material releases more energy, and secondly because it has a smaller surface area from which to radiate the accreted energy.

If the primary is a DAV, its T_{eff} lies in a very restricted range, and we get the following approximate relationship between the primary's mass and the observed mass-transfer rate:

$$M_{\rm WD} \propto [\dot{M}(2)]^{-\frac{1}{2}}\,. \tag{4.4}$$

Therefore, since the outburst interval is proportional to $\dot{M}(2)$, a DAV CV with a shorter outburst interval than GW Librae should have a lower mass primary. An independent determination of the mass can be made from the pulsation eigenfrequencies, and will be consistent with the value given by equation (4.4) unless our understanding of the physics of CVs or DAVs is at fault. This is one example of

how the study of GW Librae and any other stars like it will bring together two independent fields of research, and provide a testing ground for the currently accepted models of CVs and pulsating white dwarfs.

Now that we know that DAV behaviour in CVs is possible, high-speed photometry of the other WZ Sagittae stars should be obtained. In Table 4.3 I list some understudied CVs which have undergone infrequent, large amplitude outbursts and which need further study.

Table 4.3: Possible ultra-low $\dot{M}(2)$ CVs.

V592 Her	V630 Cas	GR Ori	V522 Sgr
V632 Her	SS LMi	V344 Ori	V511 Sgr
WY Sge	BK Lyn	AS Psc	YY Tel
KY Ara	V358 Lyr	Psc 3	BC UMa
CU Vel	PU Per	RW UMi	BZ UMa

References: Downes, Webbink and Shara 1997; Ritter and Kolb 1997

Some polars also have low $\dot{M}(2)$ and hence low T_{eff} because magnetic braking in these systems is weak and therefore the mass transfer rate is low (Wickramasinghe & Wu 1993). This is because the fields of the primary and secondary will couple, leaving fewer field lines which extend to infinity and so the braking rate is lower. However, it is possible that strong magnetic fields suppress nonradial oscillations, because material will be unable to move across the field lines, although according to Dziembowski (1998), there is no reason why a white dwarf with a strong magnetic field should not pulsate.

In addition, the primaries of dead or hibernating CVs, which appear as single white dwarfs because the mass transfer has shut off, will pass through the instability strip as they cool. A CV may become dormant if changes in the structure or in the magnetic field multipolarity of the secondary cause it to detach from its Roche lobe. A CV dies when the Jupiter-like secondary eventually detaches from its Roche lobe

after the system has passed through the period minimum. How many known DAVs are actually dead or dormant CVs?

In GW Librae we are lucky because the low $\dot{M}(2)$, together with its low to intermediate inclination and optically thin disc, mean that a substantial contribution to the light of the system is made by the white dwarf itself. Nonradial pulsations should be looked for in CVs with spectra that clearly show the pressure-broadened Balmer absorption lines, which indicate that the white dwarf (and therefore its pulsations if it has any) is visible.

4.2 Accreting White Dwarf Pulsators

The study of accreting DAVs suffers at present from a lack of theoretical work and modelling. Before the discovery of GW Librae's DAV properties, there was no incentive to model accreting DAVs, since it was assumed that accretion would keep the white dwarf too hot to pulsate.

What are the effects, if any, of accretion on the normal modes of a DAV? Can we treat GW Librae as an isolated DAV, at least to a first approximation, or will its history of accretion make it fundamentally different from the other DAVs?

4.2.1 The DBVs and the AM Canum Venaticorum Stars

The only work to date on the effects of accretion on the normal modes of white dwarfs was a study by Nitta and Winget (Nitta 1996, 1998; Nitta & Winget 1997), investigating the possibility of DBVs evolving from interacting binary helium white dwarfs.

Nather, Robinson and Stover (1981) point out that the double-degenerate helium interacting binary systems, the AM Canum Venaticorum stars (AM CVns), must inevitably produce DBs as the end products of their evolution. AM CVn itself has a temperature of 25 000 K, which places it in the DB instability strip. If AM CVn remains at this temperature at the end of its mass transfer phase, will it pulsate? Indeed, what stops it from pulsating now? If it did pulsate, how would its normal

modes differ from that of a single DAV, or one that had evolved entirely as a single star?

These questions lead Nitta and Winget to make a preliminary investigation of the effects of a history of accretion on the DBVs (none of which is known to be in a binary system). They find that even after accretion has terminated, a DB produced by an AM CVn system retains a "memory" of accretion for a very long time, in the form of an altered thermal structure, which has an observable effect on its normal modes.

Nitta and Winget created "hybrid" stars by combining the cores of hot model DBs with the envelopes of cool model DBs, in this way simulating the thermal structure of DBs evolving from AM CVns. The mass replaced by the hot envelopes is related to the duration of the accretion phase τ_{acc}: the longer the accretion phase, the greater the mass fraction of the star heated up. The accretion phase of an AM CVn star is thought to last $10^6 - 10^9$ years (Faulkner, Flannery & Warner 1972; Nather, Robinson & Stover 1981; Wood *et al.* 1987; Provencal 1994). Nitta and Winget find that the accretion time scale necessary to heat the star throughout (so that it has the same thermal structure as a hot, single star) is $\sim 10^8$ years or more, which is roughly equal to the thermal, or Kelvin-Helmholtz, time scale of a hot, single DBV.

To obtain the normal mode periods of the models, they used a pulsation code which solves the adiabatic nonradial oscillation equation using the Runge-Kutta-Fehlberg scheme (Kawaler, Hansen & Winget 1985). It was found that all the hybrid models are unstable to g-mode pulsations. The results are shown in Figure 4.2, which plots the periods of the $l = 1$, k modes for different accretion time scales. It is clear that there are significant differences in the period spacings of the consecutive $l = 1$, k modes, even though all the hybrids are spectroscopically identical.

The thermal inversion that results from accretion will increase the periods and the period spacings of the modes. The lower k modes especially will have significantly longer periods than what one would expect from the surface temperature.

Comparing these results with the best-studied DBV, GD 358, indicates that GD 358 either evolved as a single star, or went through an extremely long accretion phase

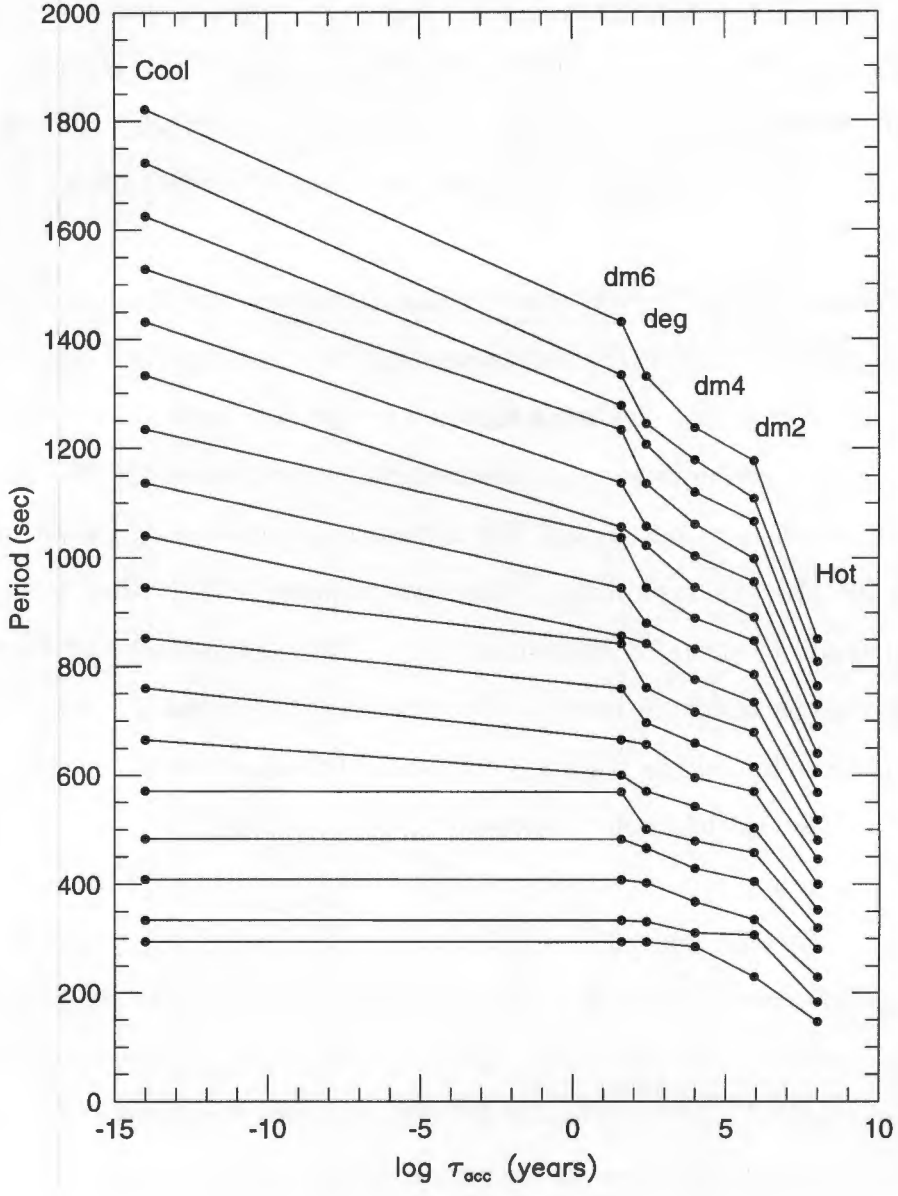

Figure 4.2: The periods of the $l = 1$, k modes for model DBVs with different accretion histories. From Nitta & Winget (1997).

(pushing the theoretical limit).

Are these results relevant to the DAVs, and to GW Librae in particular? The models described above are only simple approximations, but with no self-consistent interacting binary evolution code available, they are currently the best models for accreting white dwarf pulsators. A model for an accreting DAV would need to include surface (longitudinal) differential rotation and surface temperature and density gradients (a disc accreting onto the equator would heat an equatorial band and spin it up). Accretion torque may spin up the star to angular velocities several orders of magnitude greater than those seen in single DAVs (Bildsten 1998) – the m-splittings given by equation (3.9) apply to slow rotation only.

Nitta and Winget's models have relied on the fact that the AM CVns evolve quite quickly, with $\tau_{acc} \sim 10^6$ to 10^9 years, so that their primaries receive more heat from accretion than they can radiate away. This may not be the case for the hydrogen-rich CVs, however, for which $\tau_{acc} \sim 10^8$ to 10^{11} years. In stars like GW Librae, we may not find long-lasting temperature inversions.

4.2.2 Speculations

Let us consider the behaviour of a DAV in an accreting environment. As mentioned above, accretion disrupts the spherical symmetry of the white dwarf's photosphere. If the star has a negligible magnetic field, the disc accretes onto the equator (if the star's and disc's rotation axes are parallel), spinning up the equatorial layers. If the white dwarf has a magnetic field, the inner disc will be channelled by the magnetosphere down accretion curtains and funneled into accretion columns that form above the magnetic poles. Accreted mass and heat is deposited on the star at the "footpoints" where the accretion columns meet the photosphere.

It is not known how quickly the accreted heat, angular momentum and mass is distributed from the accreting regions to the rest of the stellar surface and interior (Pringle (1998) gives a discussion of lateral heat transfer). Sparks *et al.* (1993) find that they can only fit WZ Sge's primary's post-outburst cooling curve if the

hot accreted material is confined to a broad equatorial band for 10 years. Therefore we need to assume that the outer envelope of an accreting white dwarf will have a non-spherically symmetric rotation, temperature, density and chemical composition structure. The analysis of normal modes of oscillation in white dwarfs has thus far been heavily dependent on assumptions of spherical symmetry.

The effects of these non-symmetric structures on the normal modes of a star depend on the depth to which they extend. The pulsation driving regions are buried deep in the envelope. If accretion affects only the upper layers of the photosphere, perturbations on the normal modes may be small, but if the effects of accretion extend deeply into the nondegenerate envelope, the normal modes may be fundamentally changed. Perhaps the hot DAVs, with their thin, shallow driving zones, may be more strongly affected by accretion than the cool DAVs, which have deep, broad driving zones. The effect of accretion on the pulsations may also differ greatly from star to star; depending on the specific accretion geometry, boundary layer structure, inner disc structure and various other stellar and system parameters.

Normal mode analysis has also assumed slow rotation. Single white dwarfs have rotation periods on the order of ~ 1 day (Kepler 1990). For a CV primary accreting angular momentum from a disc for a few $\times 10^9$ years, the time taken to spin it up to break-up velocity is relatively short, but its actual P_{rot} depends on a number of competing and uncertain factors (Warner 1995). The shortest rotation period for a white dwarf is given by Warner (1995) as

$$P_{rot,min} = 15.1 M_{WD}\ \mathrm{s} \qquad (0.4 M_\odot \leq M_{WD} \leq 0.7 M_\odot)\ , \tag{4.5}$$

(any faster than this and they break up and fly apart). The general rotation periods for the outer envelopes of non-magnetic CV primaries can be found from the maximum periods of Dwarf Nova Oscillations (Warner 1995), giving $P_{rot} \sim 30$s! Direct measurements of $v\ sin\ i$ for CV primaries give $P_{rot} \geq 40$s for TT Ari (Shafter *et al.* 1985), $P_{rot} \sim 110$s for U Gem (Sion *et al.* 1994), and $P_{rot} \sim 29$s for the matriarch of GW Lib's family, WZ Sge (Patterson 1980). Accreted angular momentum may

not necessarily be transferred to the deep interior of the star (King, Regev & Wynn 1991), in which case there may be fearsome radial differential rotation gradients through these stars.

GW Librae may be spinning $\sim$ 3000 times faster than a single DAV. The slow rotation model for normal modes may not apply. In addition, rapid rotation distorts the primary: the equatorial radius of a white dwarf spinning at $P_{rot,min}$ is 1.255 times its average radius (Tassoul 1978); yet another departure from spherical symmetry.

Because of the low mass transfer rate requirement, a nonmagnetic DAV CV will be a Dwarf Nova. This is potentially the most exciting feature of accreting DAVs: it may be possible for us to watch these stars move through the instability strip in a matter of weeks or years, far faster than the $(5-10) \times 10^8$ years required for single DAVs (Wood 1990; Kleinman 1998a).

The dumping of hot material into the primary's photosphere during a DN outburst temporarily increases its T_{eff}. For example, Wood *et al.* (1993) find that Z Cha's primary has a T_{eff} equal to 17 400 K immediately after a normal outburst, cooling to its quiescent temperature of 15 600 K in $\sim$ 16 days; and Long *et al.* (1994) find that for U Gem's outbursts, its primary has a T_{eff} equal to 39 400 K after 13 days, and 32 100 K after 70 days. Of greater relevance, WZ Sge takes $\sim$ 3000 days ($\sim$ 8 years) to cool after its superoutburst (Sion & Szkody 1990).

These cooling time scales can give the depth of heating of the white dwarf outer envelopes; Sion & Szkody (1990) find the heated masses to be $(5-20) \times 10^{-10} M_{\odot}$. Are these masses sufficient to give the temperature inversion effects on normal modes predicted by Nitta & Winget (previous section)?. These masses are for spherically symmetric heating, however, whereas accretion heat is more likely to be deposited in equatorial bands. An equatorial band cutting deeply into the driving zone may have a much more dramatic effect on a star's g-modes than a spherically symmetric temperature inversion in the top $(5-20) \times 10^{-10} M_{\odot}$ layers. With no theoretical work or modelling, it is not possible to say just how outbursts will affect the g-modes, but the possibilities are extremely exciting.

Could accretion be a stochastic or chaotic pulsation driving force? While ac-

creting material may not penetrate the envelope deeply enough to affect the driving region directly, it may drive r-mode oscillations on the surface, which in turn could be a perturbing force on the g-modes. r-modes (*Rossby* waves) are surface oscillations which are driven by "winds" of a less dense fluid blowing over it (ocean waves are an example of this). Accretion may be seen as a high velocity, low density wind blowing across the surface of the much denser white dwarf photosphere, and it could generate r-mode oscillations in the photosphere. For primaries in the DA instability strip, these r-modes could excite the deeper, larger amplitude, self-sustaining g-modes, or perturb them in a stochastic or chaotic manner, resulting in continuous phase, frequency or amplitude modulation of each normal mode.

Chapter 5

Observations, Reductions and Analysis Techniques

Most of the observations of GW Librae presented in this dissertation were acquired by the author at the Sutherland observing station of the South African Astronomical Observatory, using the University of Cape Town High Speed Wright CCD photometer on the 0.75-m, 1.0-m and 1.9-m telescopes. This chapter outlines the operation of the UCT CCD photometer, presents the observations, gives an introduction to Fourier analysis and *temporal spectroscopy*, and discusses some of the more important algorithms used in the analysis of multiperiodic signals.

5.1 High Speed CCD Photometry: The UCT CCD Photometer

5.1.1 The Instrument

High speed photometry, or the study of rapid light variations in stars, has traditionally been carried out with the use of photomultiplier tubes (PMTs). In a PMT, light strikes a photosensitive anode, releasing electrons, which are accelerated through a high voltage to strike a series of cathodes, releasing more electrons in each cascade. In this way, a photon is read out as a pulse of electrons.

A disadvantage of the PMT is that it is a "single pixel" device: there is no spatial information in the data recorded, so a simultaneous measurement of a program star, a comparison star and the background sky requires three independent PMTs operated simultaneously.

Another disadvantage is their low quantum efficiency (QE)[1]. Peak QEs for PMTs are around 25%. Charge Coupled Devices (CCDs), on the other hand, can have peak QEs of more than 80%. In a CCD camera, the sky is imaged onto a two-dimensional array of pixels, and therefore sky subtraction is much easier and more accurate. These two factors make it possible to obtain high speed photometry on stars much too faint to study with PMTs.

Further advantages of observing with CCDs rather than PMTs, is the reduction of shot noise from the sky, and the fact that one can vary the apertures during the reductions (i.e. *after* the data have been obtained).

CCDs are widely used in astronomy because of their high quantum efficiency, linearity and imaging capabilities. A CCD is a metal oxide semi-conductor solid state integrated circuit composed of a two-dimensional array of discrete electrodes (pixels). Each pixel is in effect a potential well which can accumulate negative charge. When a photon strikes the surface of the CCD array and is absorbed by the silicon, an electron is released from the valence band into the conduction band where it is attracted by and stored in the nearest potential well. During an exposure, electrons accumulate in each potential well. The total charge in a potential well is proportional to the number of photons incident on that pixel.

Each row of pixels feeds into the output register, a single perpendicular row of electrodes, which leads to an output amplifier. At the end of an exposure the two-dimensional charge distribution is read out. At each step of the readout the voltage applied to each electrode is changed to that the potential wells with their accumulated charge move towards the output register; this is called "clocking" the device. The output register is clocked repeatedly to move the potential wells to

[1]The *quantum efficiency* of a detector is its sensitivity: a detector with a QE of 25% in a certain wavelength band detects a quarter of the photons in that wavelength band striking it.

the output amplifier. The final signal is given in ADU (analogue-to-digital units). The UCT CCD is operated with a gain of 10 electrons per ADU. The UCT camera sends its output signals to an acquisition PC which passes the image frames to a Sun workstation or Pentium II PC where the reductions are performed.

Silicon is more transparent to light of longer wavelengths: red photons penetrate deeper into the chip, while blue photons are more likely to be absorbed by the electrodes covering the surface of the chip, resulting in fewer electrons from shorter-wavelength photons landing in the potential wells. A CCD chip is therefore less sensitive in the blue. To improve its blue response the chip can be thinned and used upside down so that starlight falls on the back of the chip (*back illumination*). Since the UCT CCD photometer is designed to observe very blue stars like CVs, the UCT Wright CCD is a thinned, back-illuminated device. Its peak quantum efficiency is 68% at 5500Å.

When charge is transferred across a CCD, not all the electrons in each pixel will advance when the potential well is moved, leaving a residual charge which appears in the following pixel. The UCT CCD has a 1% charge loss for a transfer from the diagonally opposed corner to the output amplifier.

When the charge stored in a pixel approaches the trapping potential of the well, the response of the pixel becomes nonlinear, and when it exceeds the trapping potential, the pixel becomes saturated and charge can spread to the surrounding pixels. The UCT CCD saturates at 32 000 ADU, and nonlinear effects set in above 25 000 ADU. GW Librae's brightness distribution never exceeded more than a few hundred ADU per 20-s to 30-s exposure.

The intrinsic noise in the signal is the square root of the number of electrons in a pixel (*Poisson noise*). Poisson noise cannot be eliminated, but the higher the photon count, the better the signal-to-noise. Another source of noise is the *dark current* which is given by thermally excited electrons. Cooling the chip decreases the dark current. The UCT CCD is cooled by a 4 stage Peltier device which gives an operating temperature of $\sim$ 200K and a dark current of $\sim$ 0.03 electrons per pixel per second. Although the dark current is fairly uniform over the chip there

are dark current spikes at some pixels in which the dark current is unusually high. It is important to note the positions of these "bright" pixels to avoid imaging stars there. The output amplifier produces a *readout noise* which increases as the readout rate is increased; and the more pixels covered by a stellar image, the larger the readout noise. The UCT CCD has a readout noise of less than 10 electrons RMS. The analogue-to-digital conversion introduces a *digitization noise*, which depends on the gain of the CCD: a gain of 10 electrons per ADU yields an output signal of 1 ADU for 10 electrons as well as for 19 electrons, so in the latter case, 9 are lost. *Random noise* is introduced by cosmic rays and other ionizing particles. The reduction software must be able to distinguish between cosmic ray hits and stars.

The UCT CCD is a frame transfer device: half the chip is masked and used as an image storage buffer while the other half is exposed to starlight. When an exposure is finished the image is transferred to the storage half in milliseconds and then read out and transferred to the workstation in 8–12 seconds while the next exposure is in progress (during 1998, the reduction software was recompiled to run on newly-available Pentium II PCs, which not only have much faster CPUs than the workstations, but also much faster hard discs, so the transfer time has been reduced by about a factor of 3, enabling much shorter exposure times). This means there is almost no dead-time between frames, and the integration time can be as short as the sum of the readout and transfer times. Increasing the speed of the readout will increase the readout noise.

Shorter integration times can be obtained by pre-binning (in which a group of pixels becomes a single pixel) which significantly reduces the readout and transfer times because the number of pixels is reduced. In 2x2 binning (used with the 1.0-m) the number of bytes in an image is reduced by a factor of four; 3x3 binning (used with the 1.9-m) gives a 9x reduction. Prebinning is required on the 1.0-m and 1.9-m telescopes, because the 1.0-m's arcsecond to pixel ratio is smaller than for the 0.75-m; and for the 1.9-m, smaller still. This means that a stellar image[2] observed with

[2]The true image of a star is an unresolvable point, but atmospheric scintillation and seeing broaden the image. Under excellent seeing conditions, a stellar image can have a diameter of less

the 1.9-m covers a much larger area of the chip than when observed with the 0.75-m. Prebinning therefore reduces the readout noise, and should also be used when the seeing is very poor.

5.1.2 Data Reduction

Before the reductions proper can begin, CCD frames need some preparing and cleaning. Firstly, the images need to be *debiased.* Because of the random nature of Poisson and readout noise, negative values can occur in the signal which would cause problems during analogue-to-digital conversion. To avoid this an electrical offset or *bias* of constant charge (the exact value of which varies from frame to frame) is applied to each pixel before readout. The bias for each frame is calculated by measuring the bias in the overscan region of the chip (which is an "imaginary" region of the chip, subject to the same instrumental effects as the chip itself, created by clocking for a few extra cycles after the real charge distribution has been read out). A mean ADU for each row of the overscan is calculated and subtracted from the image.

Frames need to be *flatfielded* because there is a non-uniform response across the chip caused by sensitivity variations due to manufacturing imperfections or dirt on the chip. If flatfielding is not done the brightness of a star depends on its position on the chip, which is disastrous because the stars drift around slightly on the chip during the night. Flatfields are images of uniform light sources, like the twilight or predawn sky, or a uniformly illuminated screen in the dome. Each object frame is then divided by a correction image which is prepared from an average of several flats.

Once the frames have been cleaned, the brightness of the stars is determined in two ways: aperture photometry and profile fitting. In the first, the total ADU inside a software-defined aperture is summed. Background subtraction is done by subtracting the average ADU of an annulus around the star, chosen in such a way as to avoid the wings of the stellar image. In the latter method, a profile is fitted to the point-spread-function (PSF) of the star. The software that does this is DoPHOT, a

than 1 arcsec. The worst seeing conditions experienced while observing GW Librae gave stellar images ~8 arcsec across.

fast and highly automated program originally designed to reduce observations of star clusters (for a complete description see Schechter, Mateo and Saha 1993). For very faint stars (GW Librae, for example), I find that the PSF-derived magnitudes give lightcurves with a much higher signal-to-noise than the lightcurves obtained from the aperture magnitudes.

Using an elliptical Gaussian with power law wings, DoPHOT extracts a magnitude from the PSF of each star in regions of the image specified by the user. The results of the fit are listed in a *sum* file, and the process is repeated for each of the $\sim$ 3000 frames acquired in a typical high speed photometry run. The magnitudes in the *sum* files are then extracted, giving lightcurves of the selected stars. The lightcurves with the best signal-to-noise can be subtracted from the lightcurve of the target star, removing first order extinction variations and any variations due for example to thin cloud. In this way good quality lightcurves can be obtained even in non-photometric conditions, which in my experience as an observer, nearly doubles the amount of usable data obtainable in an average observing week.

5.2 The Observations

The best time of year to observe GW Librae, $\alpha = 15\ 19\ 46$, $\delta = -25\ 00\ 25$ (2000), is in May, but with the discovery of its unusual properties only in March 1997, it was too late to organize a multisite campaign on the star for May 1997. In 1997 I obtained single site data in March (4 short runs), April (4 runs) and September (7 short runs). The 1997 runs revealed "bands of power", or the gross structure in its amplitude spectrum, but were too short to resolve individual signals, or to give any clue as to the nature of the fine structure of features in the transform. It was clear that, like the other DAVs, GW Librae would need a multisite campaign (and possibly many) in order for any progress to be made in unravelling its amplitude spectrum.

A multisite campaign was organized for the following May. The aim of a multisite campaign is to provide continuous coverage of a star, in order to eliminate the aliasing

Table 5.1: 1997 Observing Log

Run name	*Date*	*start HJD* (-2450000)	t_{int} (s)	t_{run} ($h\,m\,s$)	*Tele-scope*	*Ob-server*
			March 1997			
GWLIB01	13/3/97	521.58710	6	1 59 53	SAAO 1.0-m	LvZ
GWLIB02	14/3/97	522.49980	12	3 49 46	SAAO 1.0-m	LvZ
GWLIB03	15/3/97	523.56415	12	2 28 55	SAAO 1.0-m	LvZ
GWLIB04	16/3/97	524.55071	12	2 25 07	SAAO 1.0-m	LvZ
			April 1997			
GWLIB05	1/4/97	540.50686	20	1 44 12	SAAO 1.0-m	LvZ
GWLIB06	2/4/97	541.39409	20	6 30 32	SAAO 1.0-m	LvZ
GWLIB07	4/4/97	543.40310	20	4 33 27	SAAO 1.0-m	LvZ
GWLIB08	7/4/97	546.39368	20	5 58 09	SAAO 1.0-m	LvZ
			September 1997			
GWLIB09	31/8/97	692.27492	20	2 11 16	SAAO 1.0-m	LvZ
GWLIB10	1/9/97	693.22053	20	2 59 21	SAAO 1.0-m	LvZ
GWLIB11	4/9/97	696.22408	10	3 27 58	SAAO 1.9-m	LvZ
GWLIB12	5/9/97	697.21810	10	3 31 48	SAAO 1.9-m	LvZ
GWLIB13	6/9/97	698.22028	10	3 14 46	SAAO 1.9-m	LvZ
GWLIB14	7/9/97	699.22226	10	2 49 33	SAAO 1.9-m	LvZ
GWLIB15	8/9/97	700.21981	10	2 39 31	SAAO 1.9-m	LvZ

Table 5.2: 1998 Observing Log

	Run name	*Date*	*start HJD* (-2450000)	t_{int} (s)	t_{run} $(h\,m\,s)$	*Telescope*	*Observer*
			May				
*	JK15	15/5/98	949.77526	30	3 02 42	ARIZ. 1.3-m	JK
*	JK18	18/5/98	952.73836	30	3 05 34	ARIZ. 1.3-m	JK
	GWLIB16	19/5/98	953.28861	10	5 51 32	SAAO 1.9-m	LvZ
	GWLIB17	20/5/98	954.30589	10	7 26 37	SAAO 1.9-m	LvZ
	GWLIB18	21/5/98	955.28927	10	9 07 51	SAAO 1.9-m	LvZ
*	JK21	21/5/98	955.82384	30	1 31 15	ARIZ. 1.3-m	JK
	GWLIB19	22/5/98	956.34921	10	7 15 35	SAAO 1.9-m	LvZ
	MTJ23	23/5/98	956.92086	40	4 00 48	MTJ 1.0-m	DS&JP
	GWLIB20	23/5/98	957.24252	10	10 12 38	SAAO 1.9-m	LvZ
	MTJ24	24/5/98	957.93488	40	5 22 41	MTJ 1.0-m	DS&JP
	GWLIB21	24/5/98	958.23836	10	10 08 27	SAAO 1.9-m	LvZ
	GWLIB22	25/5/98	959.23145	20	10 18 56	SAAO 1.9-m	LvZ
*	GWLIB23	26/5/98	960.31582	30	6 23 59	SAAO 0.75-m	LvZ
	GWLIB24	28/5/98	962.35375	30	5 53 25	SAAO 0.75-m	LvZ
	GWLIB25	29/5/98	963.24397	30	8 30 09	SAAO 0.75-m	LvZ
*	SV30	30/5/98	963.91411	60	7 18 00	MSSSO 1.0-m	SV
	GWLIB26	30/5/98	964.20302	30	9 32 27	SAAO 0.75-m	LvZ
	GWLIB27	31/5/98	965.20068	30	9 11 25	SAAO 0.75-m	LvZ
*	GWLIB28	1/6/98	966.23016	30	3 17 46	SAAO 0.75-m	LvZ
			June				
	GWLIB29	16/6/98	981.22138	30	8 31 14	SAAO 0.75-m	LvZ
	GWLIB30	17/6/98	982.20179	30	7 49 05	SAAO 0.75-m	LvZ
	GWLIB31	17/6/98	982.53037	30	0 52 04	SAAO 0.75-m	LvZ
	GWLIB32	18/6/98	983.24030	30	5 34 10	SAAO 0.75-m	LvZ
	GWLIB33	18/6/98	983.47518	30	2 35 08	SAAO 0.75-m	LvZ

DS = Dennis Sullivan, JP = John Pritchard, JK = Johnathan Kemp, SV = Stephane Vennes, LvZ = Liza van Zyl. MTJ = Mt John Observatory (New Zealand), MSSSO = Mt Stromlo & Siding Spring Observatory (Australia), SAAO = South African Astronomical Observatory.

* indicate runs not included in the final analysis due to poor data quality (usually because of cloud).

inherent to single-site amplitude spectra, and to obtain unbroken coverage of the beat cycles between closely-spaced signals. A deconvolution of closely-spaced signals (in the limit of high signal-to-noise) is only successful if at least 90% of their beat cycle is sampled.

With a magnitude of ~18.5, GW Librae is too faint to be observed with the WET, which operates with PMTs. We needed a CCD-equipped multisite campaign. We encountered several problems: very few CCD cameras are designed for high speed photometry of blue objects. With the exception of the UCT CCD, the chips used in the campaign had very poor QEs in the blue (they were not thinned or back-illuminated devices), and did not have frame transfer capabilities, necessitating long dead-times between exposures. In addition, many of the telescopes best equipped for CCD photometry were unavailable to us, as May is the height of the "Bulge Season," and these telescopes were dedicated to the MACHO and PLANET microlensing projects.

Contributing sites other than SAAO in the May 1998 campaign were Arizona, USA (observer: Jonathan Kemp), Mt. John Observatory, New Zealand (Dennis Sullivan, John Pritchard), and Mt. Stromlo & Siding Spring Observatory, Australia (Stephane Vennes). Frank van der Hooft and Jorge Casares, observing from Cerro Tololo Inter-American Observatory, Chile, attempted some observations, but obtained less than an hour of data before being thwarted by weather. Unfortunately, all the sites other than SAAO were plagued by extremely poor weather throughout the campaign. When observing was possible, the poor blue-responses of the chips resulted in very poor data quality.

Bad weather and chip limitations experienced by sites other than SAAO in the May 1998 GW Librae multisite campaign resulted in almost all the data being contributed by SAAO alone. While this is by far the best dataset acquired to date, it suffers from severe aliasing problems – it is essentially a single-site dataset.

The observing log for all GW Librae observations to date is presented in Table 5.1 and 5.2. Run names preceded with an asterisk denote runs in which the data quality is very poor, and which I have excluded from the analysis.

The May dataset spans 13 days. The resolution of an amplitude spectrum[3] is inversely proportional to the length of the dataset. Investigation of the amplitude spectrum of the May data revealed fine frequency splittings on the order of $\frac{1}{10}$ to $\frac{1}{14}$ days^{-1}, or $\sim 1\mu$Hz (discussed in detail in the next chapter), so a resolution greater than the inverse of 13 days was highly desirable.

We obtained three nights of data in June, two weeks after the May campaign. A transform of the June data shows the gross features of the pulsation spectrum to be unchanged from May to June, so I have added the June data to the May dataset, in order to get improved resolution. All the results from the 1998 observations presented in this dissertation come from the combined May and June dataset.

Figure 5.1 presents the lightcurves of all GW Librae observations to date. Before being combined into the total dataset, ready for analysis, each lightcurve is prepared by subtracting the lightcurves of the brightest comparison stars, to remove first order extinction effects and atmospheric effects like cloud or transparency variations. Next, I subtract from each lightcurve its mean: if the mean is not zero, the Fourier transform interprets the runs and the gaps between them as step functions, which generates forests of high amplitude harmonics which drown out the intrinsic, low amplitude pulsations.

Most CVs are very blue stars, so the comparison stars used for differential photometry are typically much redder, and therefore not as badly affected by extinction at large zenith angles. This results in a large differential extinction effect over the length of the run (differential photometry only takes care of first order extinction effects[4]). To remove these second order extinction effects, I fit and subtract a second order polynomial from each lightcurve. This reduces the power at the low-frequency end of the amplitude spectrum.

[3]The definition of ***resolution*** used throughout this dissertation is the following: if signals in a transform are seen as separate "spikes," they are resolved; if they are merged into a single spike, they are unresolved.

[4]In principle, all extinction effects can be removed if accurate UBV photometry is done on all stars in the image, and the lightcurves of the comparison stars are individually corrected for extinction. In practice, this would waste valuable observing time since we are usually only interested in variations on much shorter time scales. In addition, accurate extinction corrections for CVs are tricky to achieve because CV colours are very non-blackbody and vary over the orbital period.

In Figure 5.1, I have plotted each lightcurve on an x-axis of 3 hours (0.125 days), with a y-axis running from 0.15 to -0.15 mmag (millimagnitude).

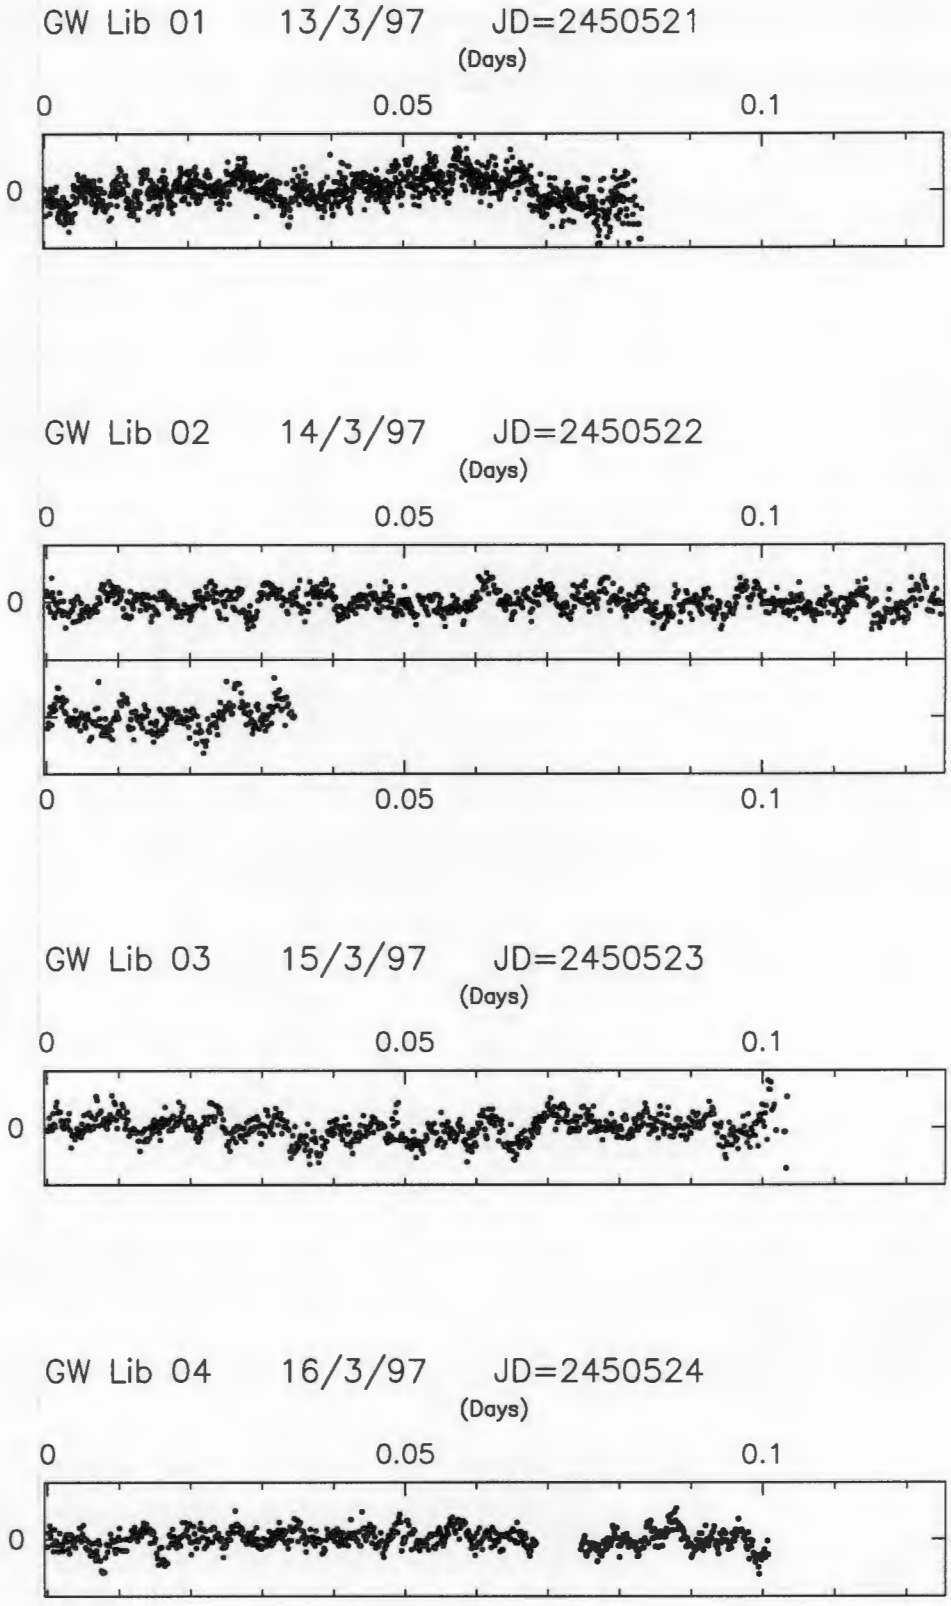

Figure 5.1: The lightcurves of all the GW Librae observations used in this dissertation. Each lightcurve is plotted on an x-axis of 3 hours (0.125 days), with a y-axis running from 0.15 to -0.15 mmag.

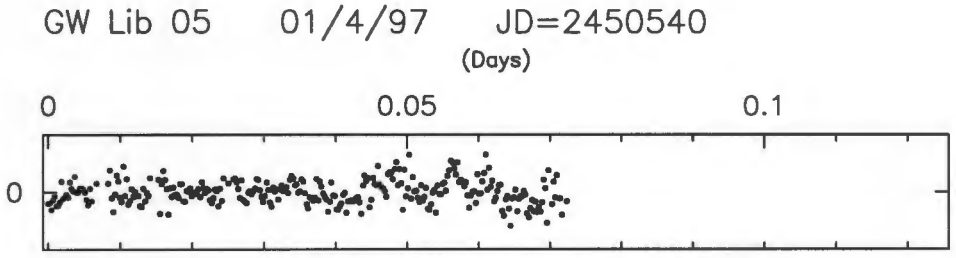
GW Lib 05 01/4/97 JD=2450540
(Days)
0
0.05
0.1
0

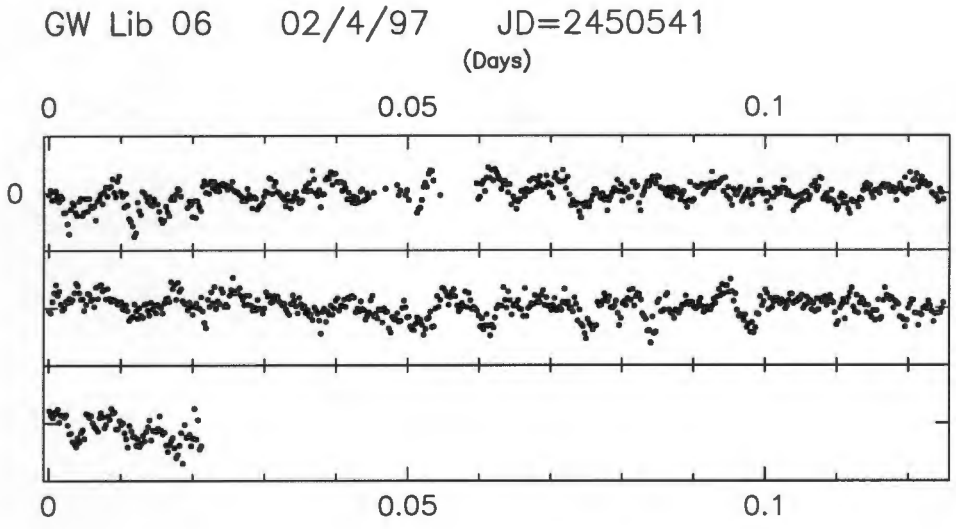
GW Lib 06 02/4/97 JD=2450541
(Days)
0
0.05
0.1
0
0
0.05
0.1

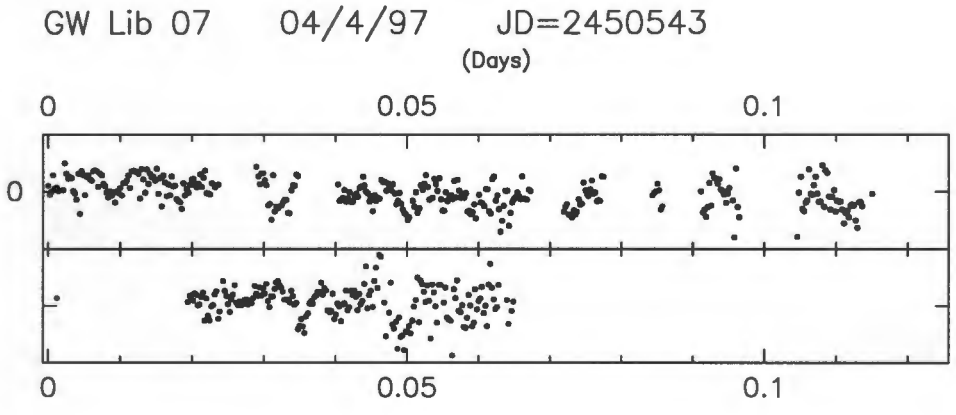
GW Lib 07 04/4/97 JD=2450543
(Days)
0
0.05
0.1
0
0
0.05
0.1

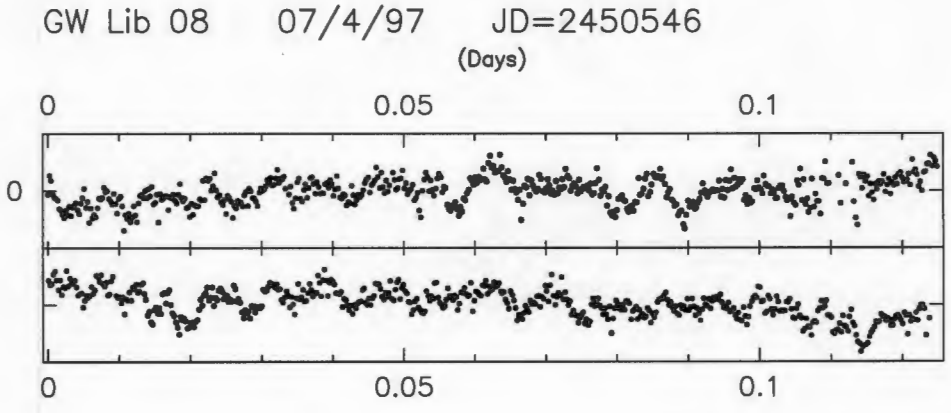
GW Lib 08 07/4/97 JD=2450546
(Days)
0
0.05
0.1
0
0
0.05
0.1

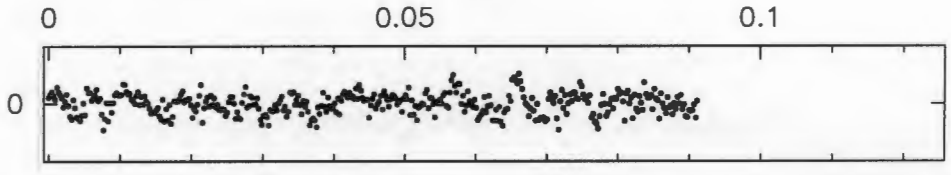
GW Lib 09 31/8/97 JD=2450692
(Days)
0
0.05
0.1
0

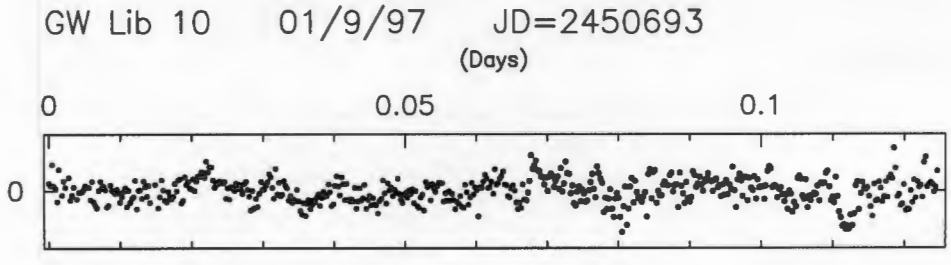
GW Lib 10 01/9/97 JD=2450693
(Days)
0
0.05
0.1
0

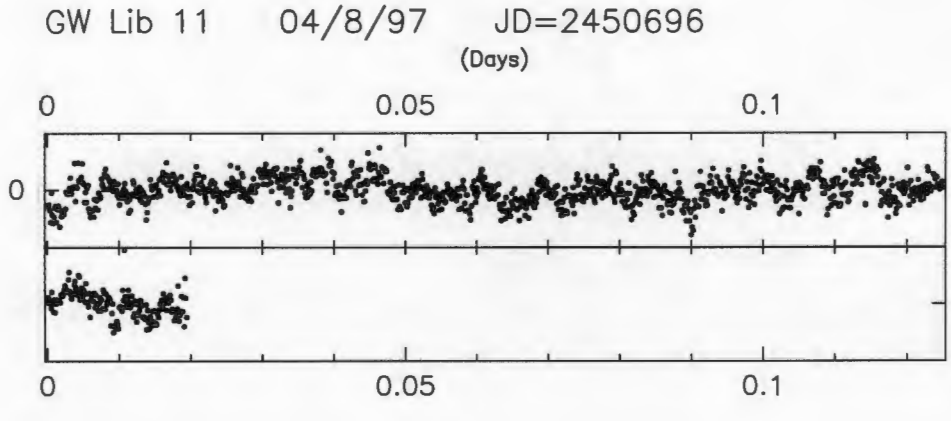
GW Lib 11 04/8/97 JD=2450696
(Days)
0
0.05
0.1
0
0
0.05
0.1

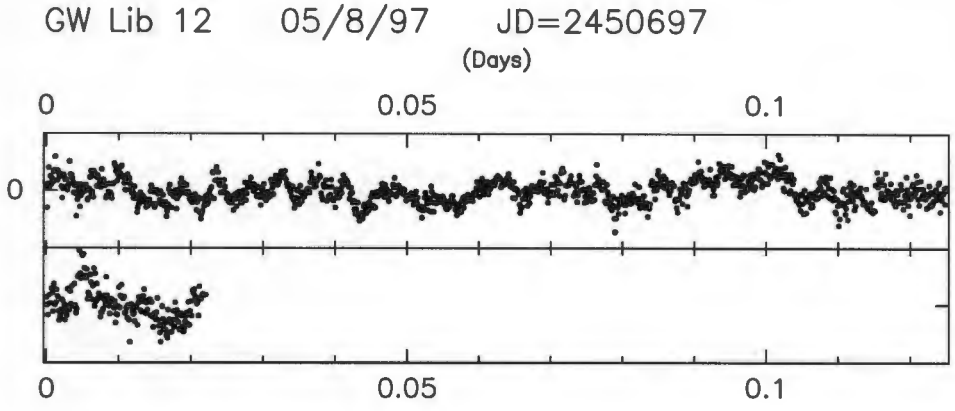
GW Lib 12 05/8/97 JD=2450697
(Days)
0
0.05
0.1
0
0
0.05
0.1

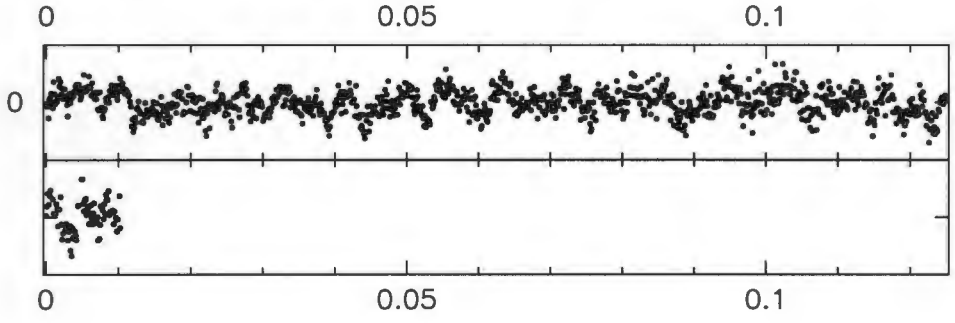
GW Lib 13 06/8/97 JD=2450698
(Days)
0
0.05
0.1
0
0
0.05
0.1

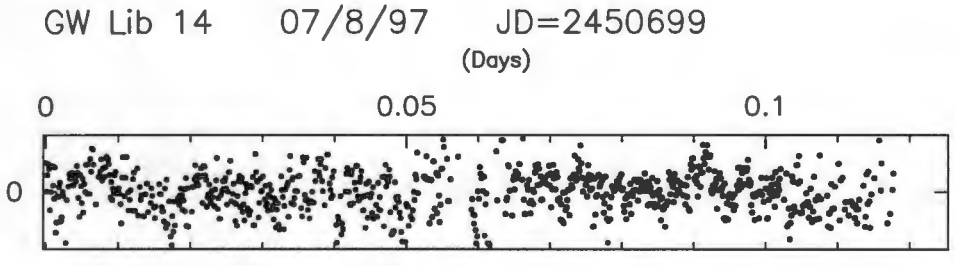
GW Lib 14 07/8/97 JD=2450699
(Days)
0
0.05
0.1
0

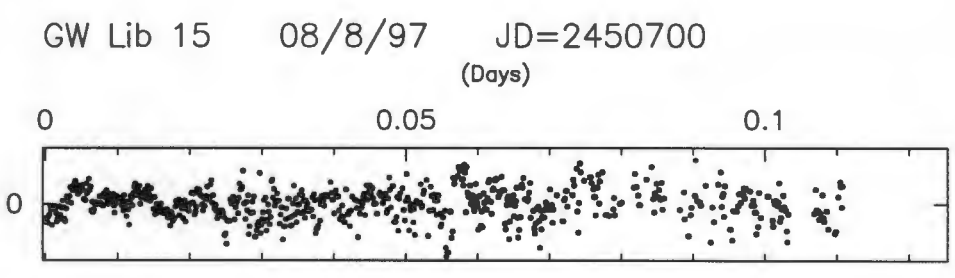
GW Lib 15 08/8/97 JD=2450700
(Days)
0
0.05
0.1
0

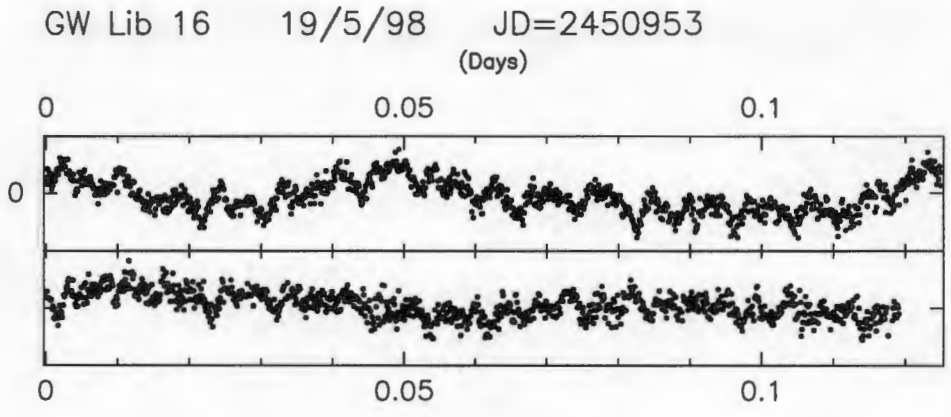
GW Lib 16 19/5/98 JD=2450953
(Days)
0
0.05
0.1
0
0
0.05
0.1

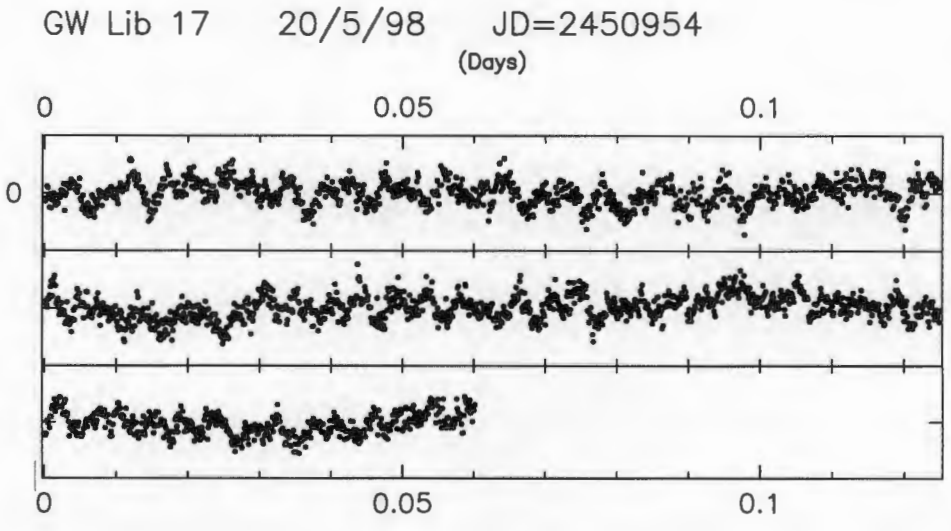
GW Lib 17 20/5/98 JD=2450954
(Days)
0
0.05
0.1
0
0
0.05
0.1

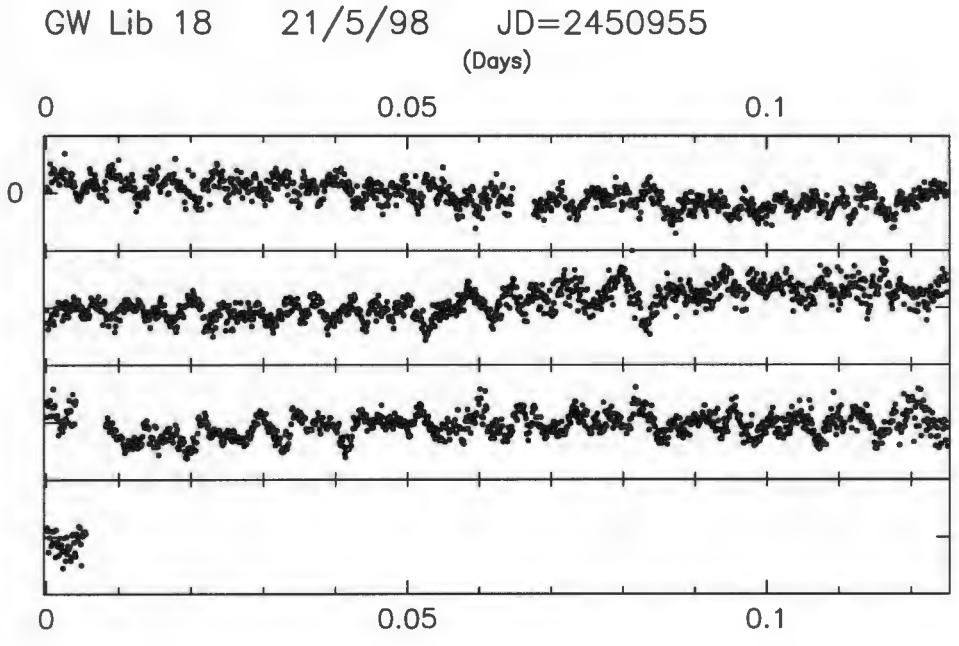
GW Lib 18 21/5/98 JD=2450955
(Days)
0
0.05
0.1
0
0
0.05
0.1

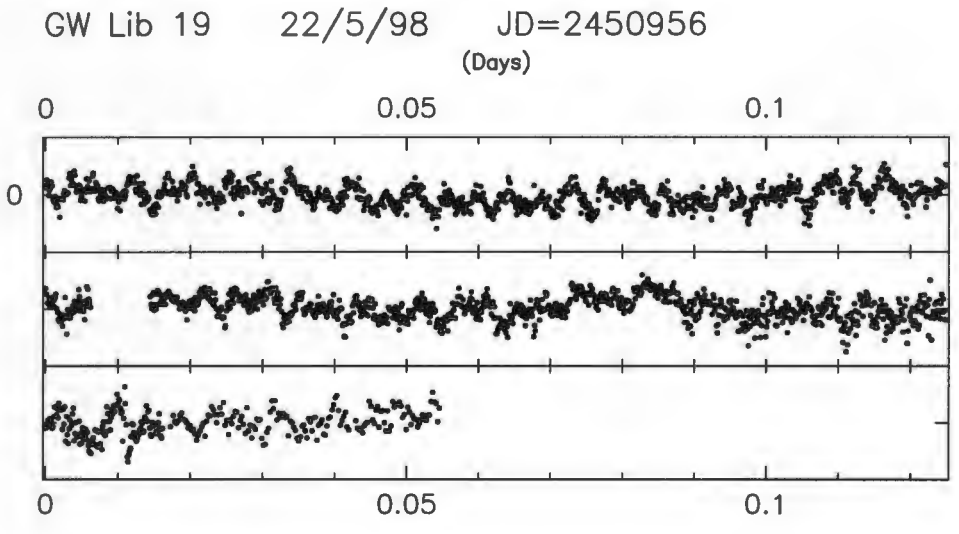
GW Lib 19 22/5/98 JD=2450956
(Days)
0
0.05
0.1
0
0
0.05
0.1

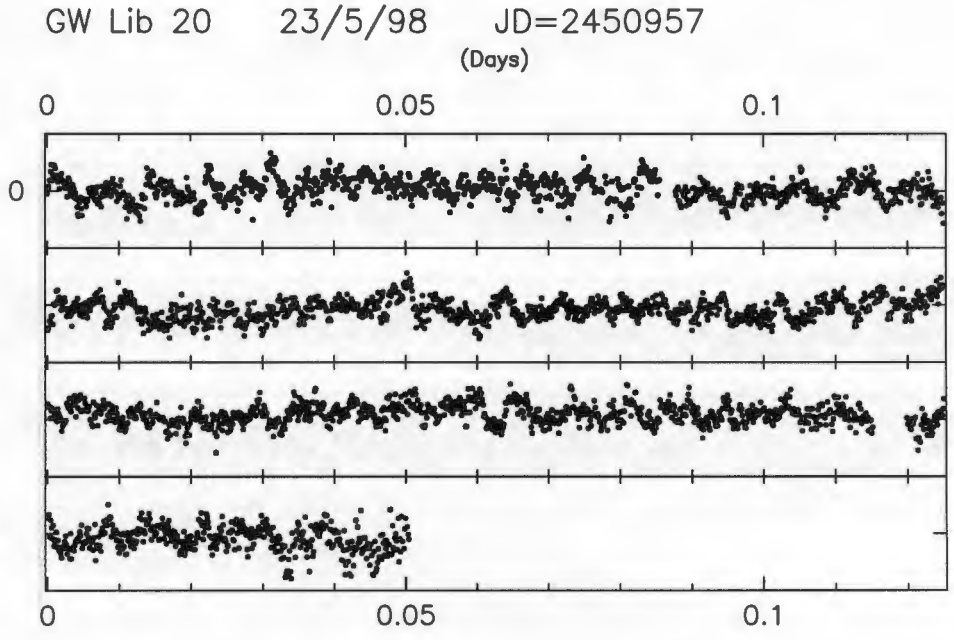
GW Lib 20 23/5/98 JD=2450957
(Days)
0
0.05
0.1
0
0
0.05
0.1

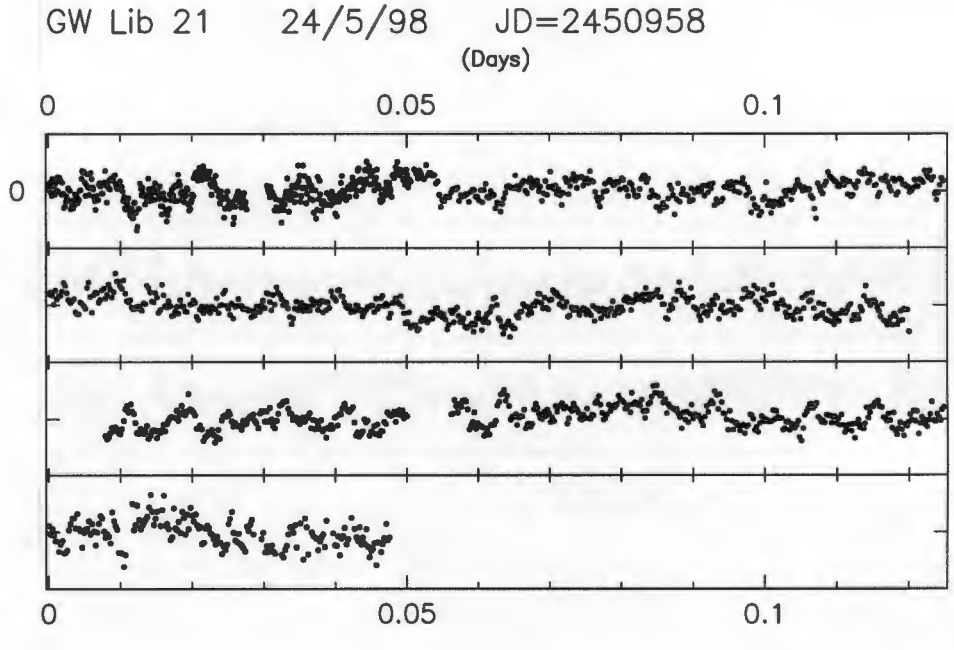
GW Lib 21 24/5/98 JD=2450958
(Days)
0
0.05
0.1
0
0
0.05
0.1

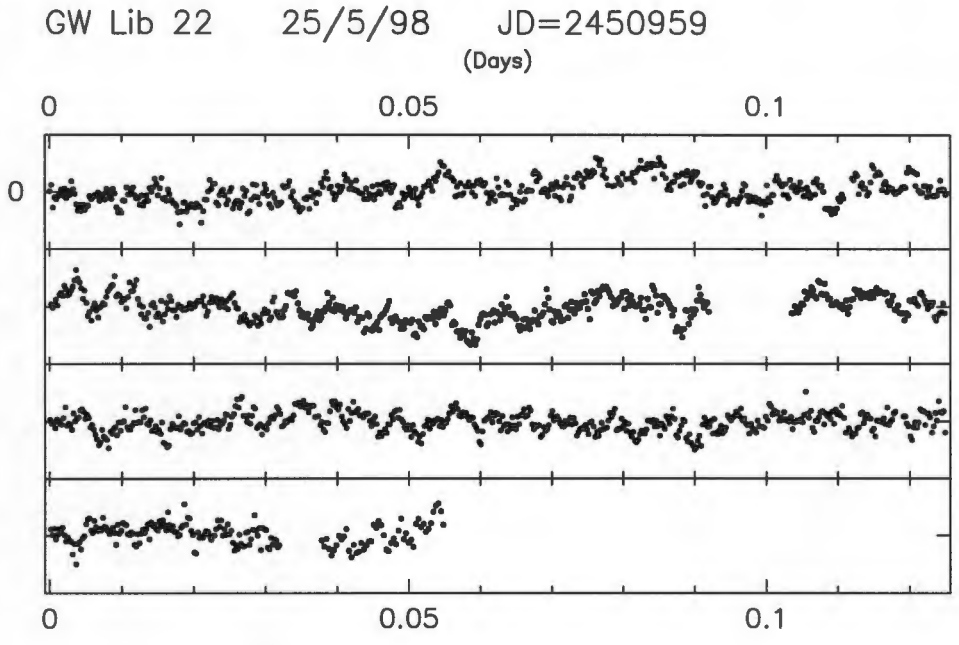
GW Lib 22 25/5/98 JD=2450959
(Days)
0
0.05
0.1
0
0
0.05
0.1

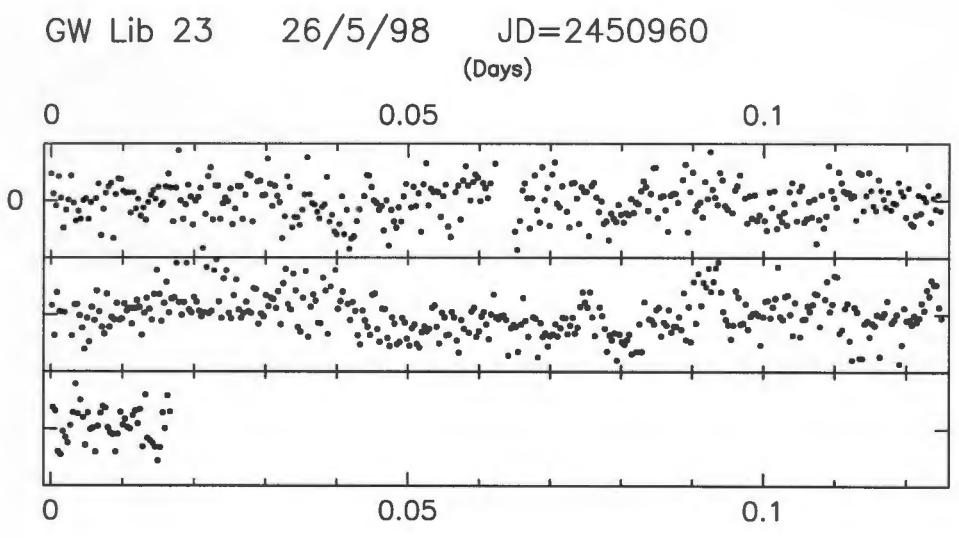
GW Lib 23 26/5/98 JD=2450960
(Days)
0
0.05
0.1
0
0
0.05
0.1

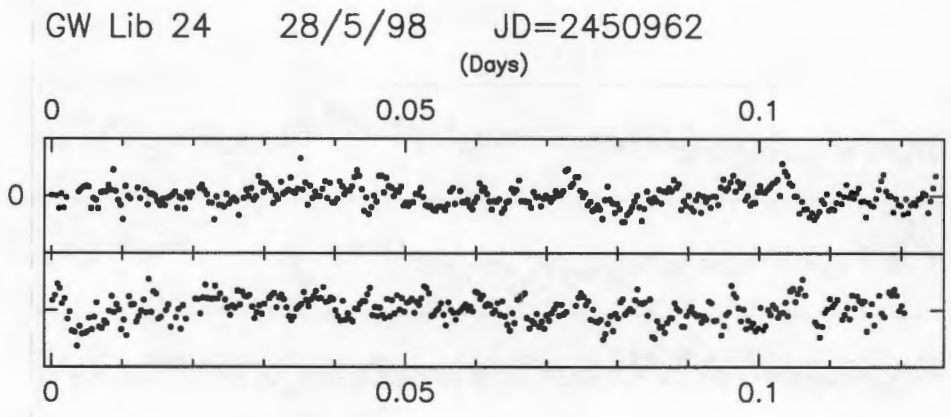
GW Lib 24 28/5/98 JD=2450962
(Days)
0
0.05
0.1
0
0
0.05
0.1

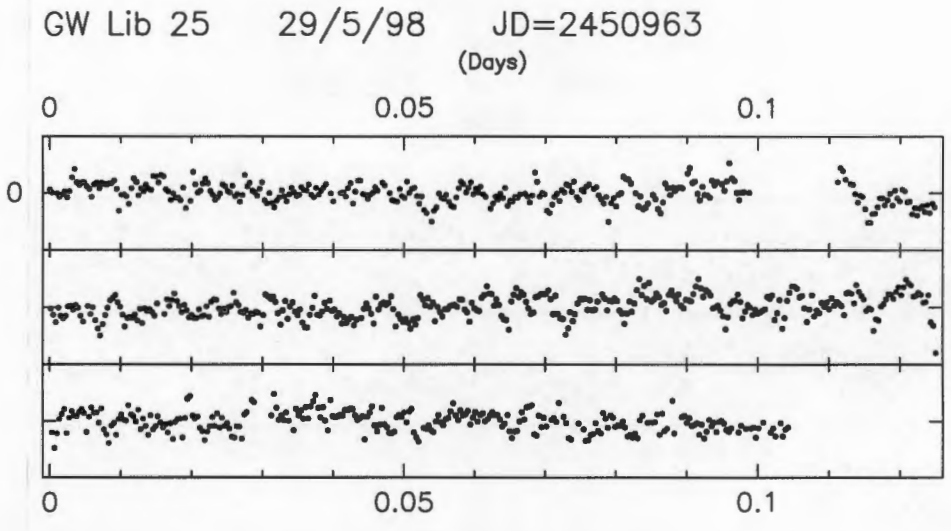
GW Lib 25 29/5/98 JD=2450963
(Days)
0
0.05
0.1
0
0
0.05
0.1

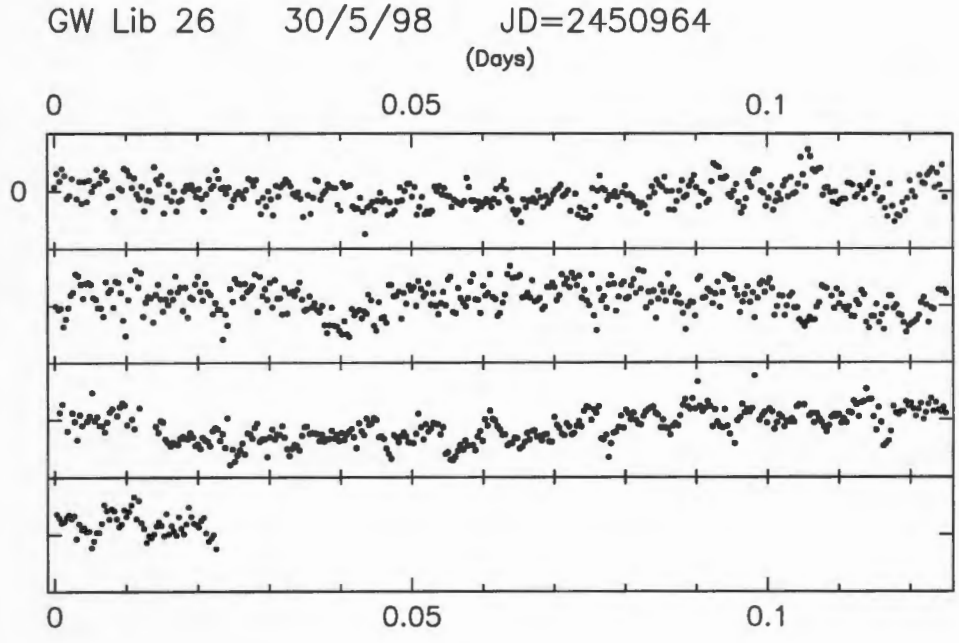
GW Lib 26 30/5/98 JD=2450964
(Days)
0
0.05
0.1
0
0
0.05
0.1

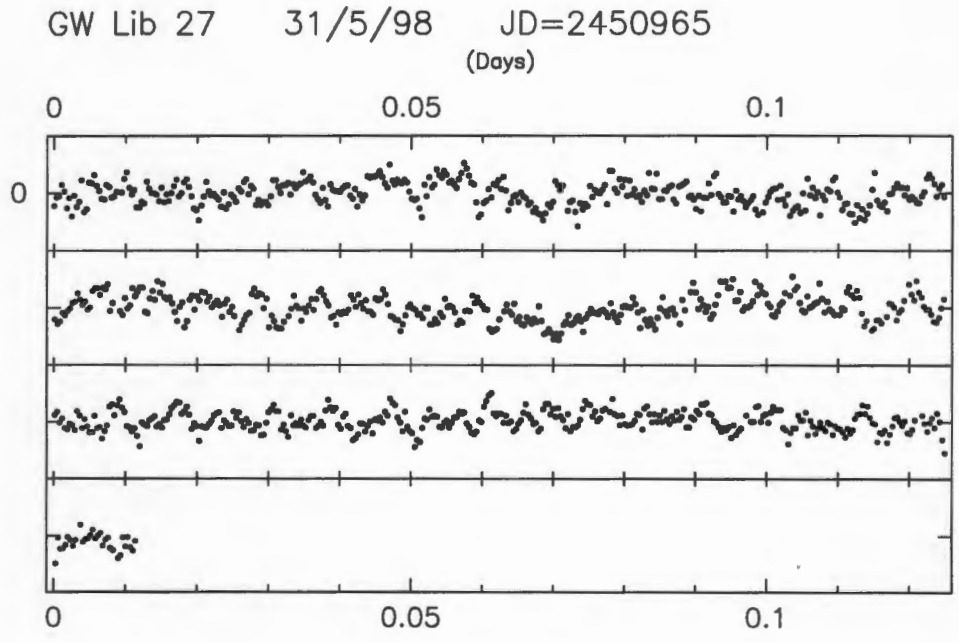
GW Lib 27 31/5/98 JD=2450965
(Days)
0
0.05
0.1
0
0
0.05
0.1

GW Lib 28 01/6/98 JD=2450966

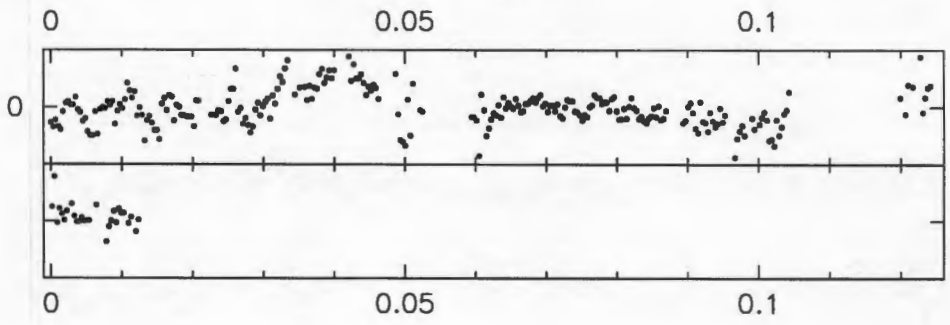

GW Lib 29 16/6/98 JD=2450981

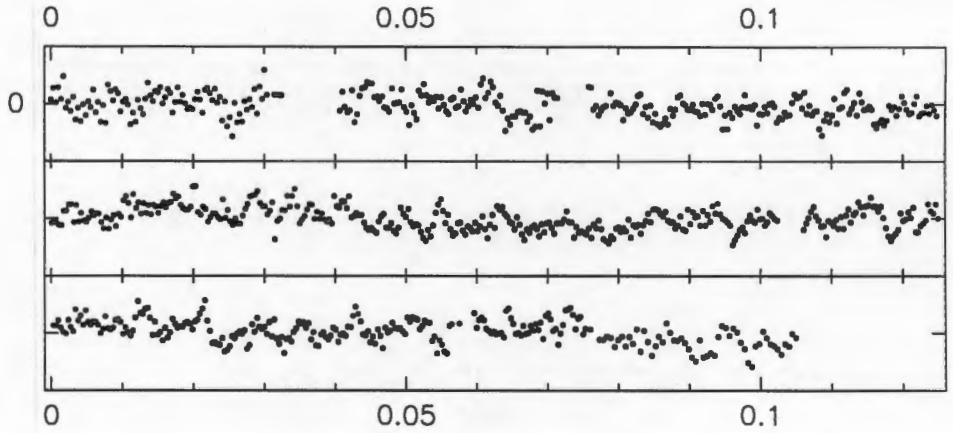

GW Lib 30+31 17/6/98 JD=2450982

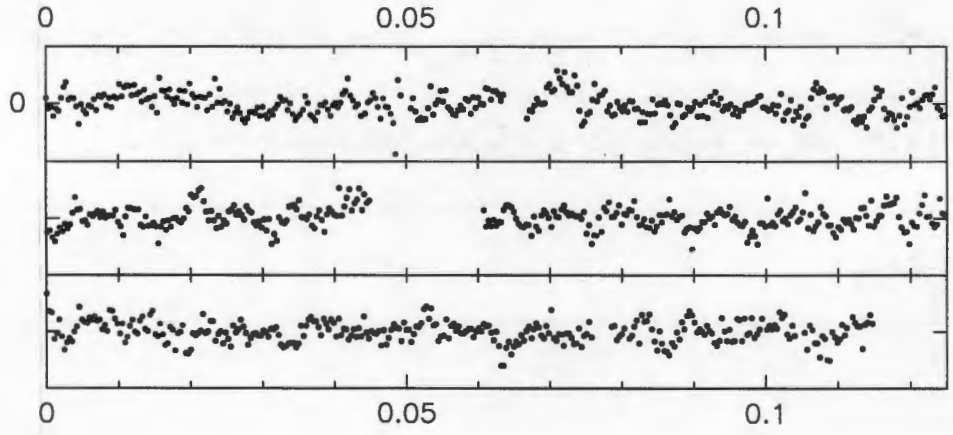

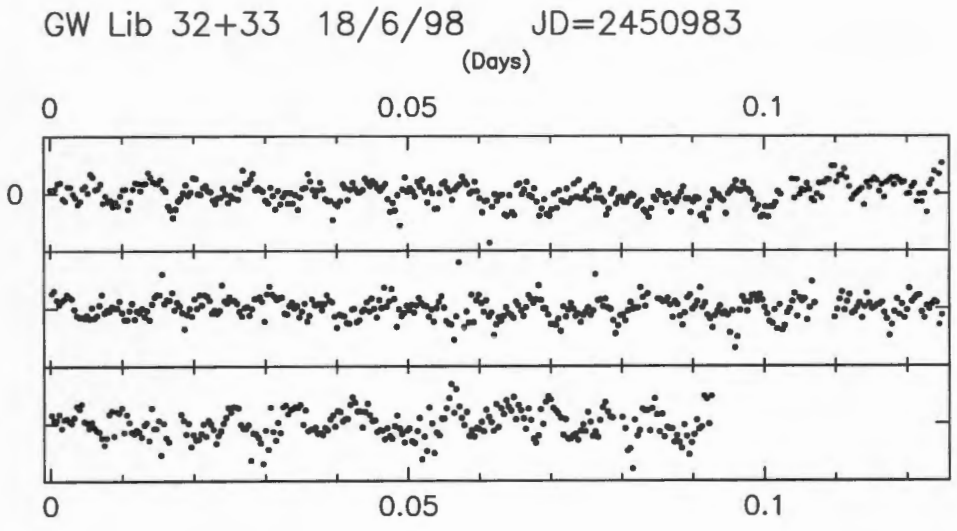
GW Lib 32+33 18/6/98 JD=2450983
(Days)
0
0.05
0.1
0
0
0.05
0.1

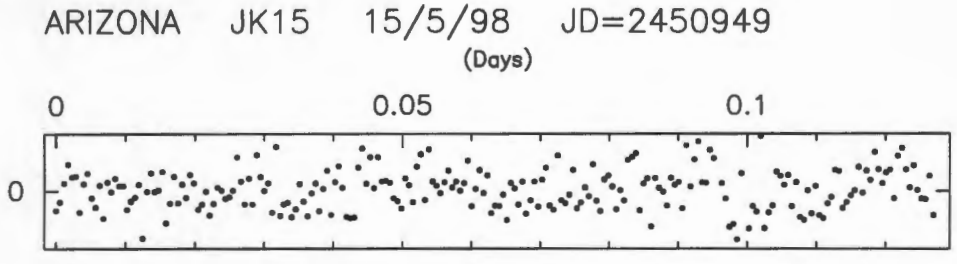
ARIZONA JK15 15/5/98 JD=2450949
(Days)
0
0.05
0.1
0

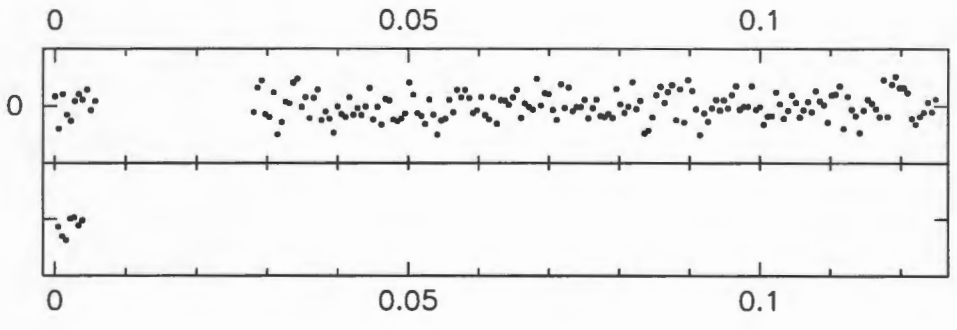
ARIZONA JK18 18/5/98 JD=2450952
(Days)
0
0.05
0.1
0
0
0.05
0.1

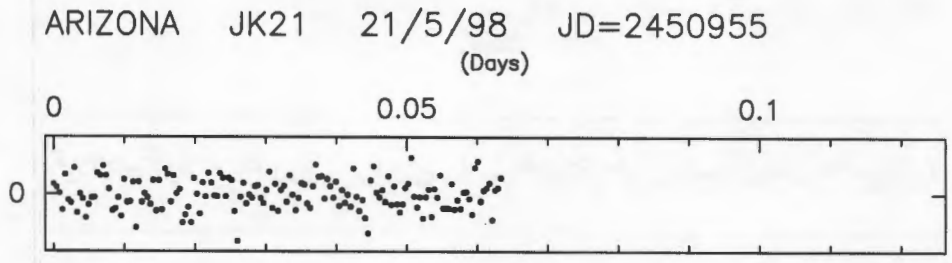
ARIZONA JK21 21/5/98 JD=2450955
(Days)
0
0.05
0.1
0

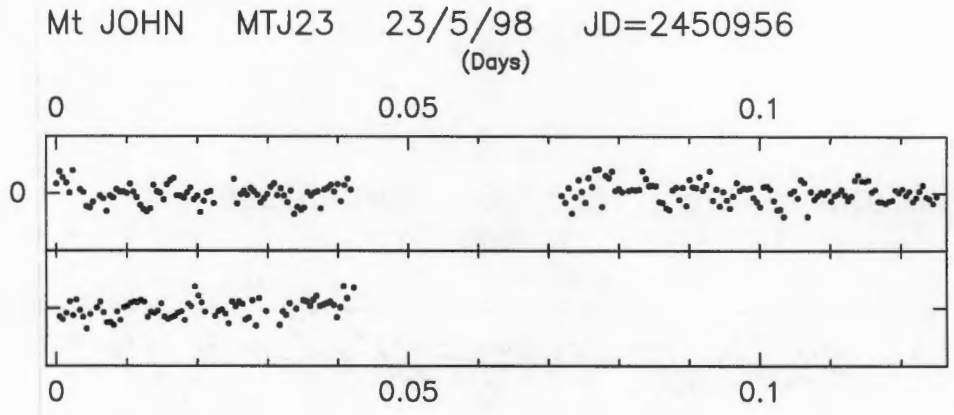
Mt JOHN MTJ23 23/5/98 JD=2450956
(Days)
0
0.05
0.1
0
0
0.05
0.1

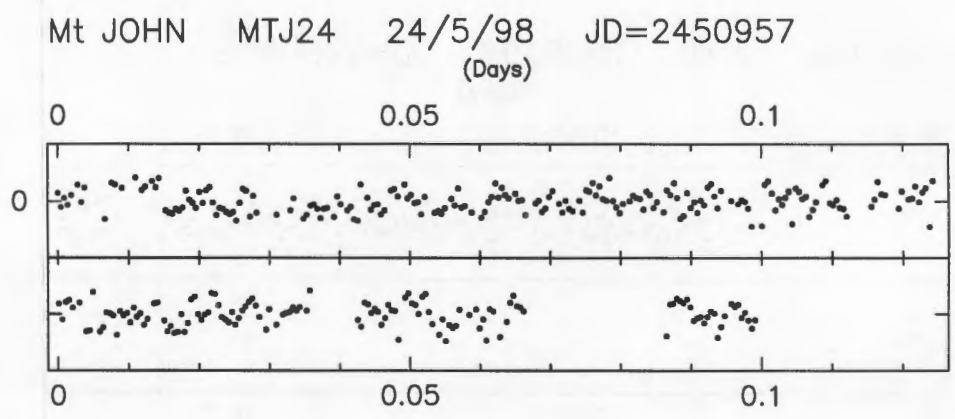
Mt JOHN MTJ24 24/5/98 JD=2450957
(Days)
0
0.05
0.1
0
0
0.05
0.1

5.3 Temporal Spectroscopy: The Investigation of Multiperiodic Signals

5.3.1 Fourier Analysis

The brightness variations of stars with time are studied by means of Fourier transforms. The *finite Fourier transform* $F_T(\nu)$ of a time-varying function $f(t)$ is defined (see for example Deeming 1975) as

$$F_T(\nu) = \int_{-T/2}^{+T/2} f(t) e^{i2\pi\nu t} dt \,. \tag{5.1}$$

Since $f(t)$ is a discrete dataset we define the *discrete Fourier transform* as

$$F_N(\nu) = \sum_{k=1}^{N} f(t_k) e^{i2\pi\nu t_k} \,. \tag{5.2}$$

This is simply a transformation of the data to a new representation – in frequency instead of time. If the data set $f(t_k)$ is a multiperiodic function with frequencies $\nu_1, \nu_2, \nu_3, \ldots$, $F_N(\nu)$ will be nonzero in the vicinity of $\nu = \pm\nu_1, \pm\nu_2, \pm\nu_3, \ldots$

Due to the finite spacing and finite length of the data there will be confusion in identifying frequencies in the observed $F_N(\nu)$ because of *aliasing* and the *spectral window*, which makes exact identification of frequencies difficult. The spectral window function for an infinitely long, uninterrupted data set is a Dirac delta-function; for finite data sets it broadens into an Airy function. Aliasing is the result of an uncertainty as to whether a gap in the data bridges n cycles (highest probability), or $n \pm 1$ cycles (slightly smaller probability), or $n \pm 2$ (smaller probability) ... , so that each frequency identified in the Fourier transform of two consecutive nights of observations, for example, is a small thicket of aliases. The observed Fourier transform $F_N(\nu)$ is the *convolution* of the true Fourier transform $F(\nu)$ with the spectral window $\delta_N(\nu)$:

$$F_N(\nu) = F(\nu)^* \delta_N(\nu) \,, \tag{5.3}$$

where the spectral window $\delta_N(\nu)$ is a function only of ν and the times of observation

t_k:

$$\delta_N(\nu) = \sum_{k=1}^{N} e^{i2\pi\nu t_k} . \tag{5.4}$$

To Fourier analyse my data, I wrote a computer program in the programming language *TruBASIC*, using the algorithm developed by Deeming (1975) for Fourier analysis with unequally-spaced data. When I was satisfied that my program was working properly (and therefore that I had grasped the principles involved), I switched to using Darragh O'Donoghue's *EAGLE* program, a much more powerful time series analysis tool, written in Fortran and running under Linux, also based on Deeming's algorithm. All amplitude spectra presented in this report were obtained with this program. In figure 5.2 I present the Fourier transforms of the total datasets.

Before calculating the Fourier transform I subtract the mean from each contributing lightcurve of the full dataset. If this is not done, the data would contain large amplitude "boxcar" and triangular functions, the transforms of which are very complex and multiharmonic, and would drown any contribution from the low amplitude variations in the star. I also remove any spikes from the data, which the Fourier transform would interpret as delta functions and try to fit with a multitude of harmonics. Spikes arise if the profile fitting routine in the CCD reductions is unable to give a good fit to a stellar image, which can happen frequently in cloudy conditions or if the seeing is bad.

5.3.2 Some *EAGLE* Algorithms

To determine initial values for the frequencies, I used the technique of *prewhitening*; the Fourier transform of a dataset is calculated to determine starting values of the frequency, amplitude and phase of the highest amplitude signal. I then use a least-squares routine to fit a sinusoid with these parameters to the data set, which is then subtracted from the data, point by point. Removing a frequency removes its entire window pattern. The Fourier transform of the residuals is then calculated, and the next highest frequency found and subtracted.

When working with partially-resolved, closely-spaced signals, it soon becomes

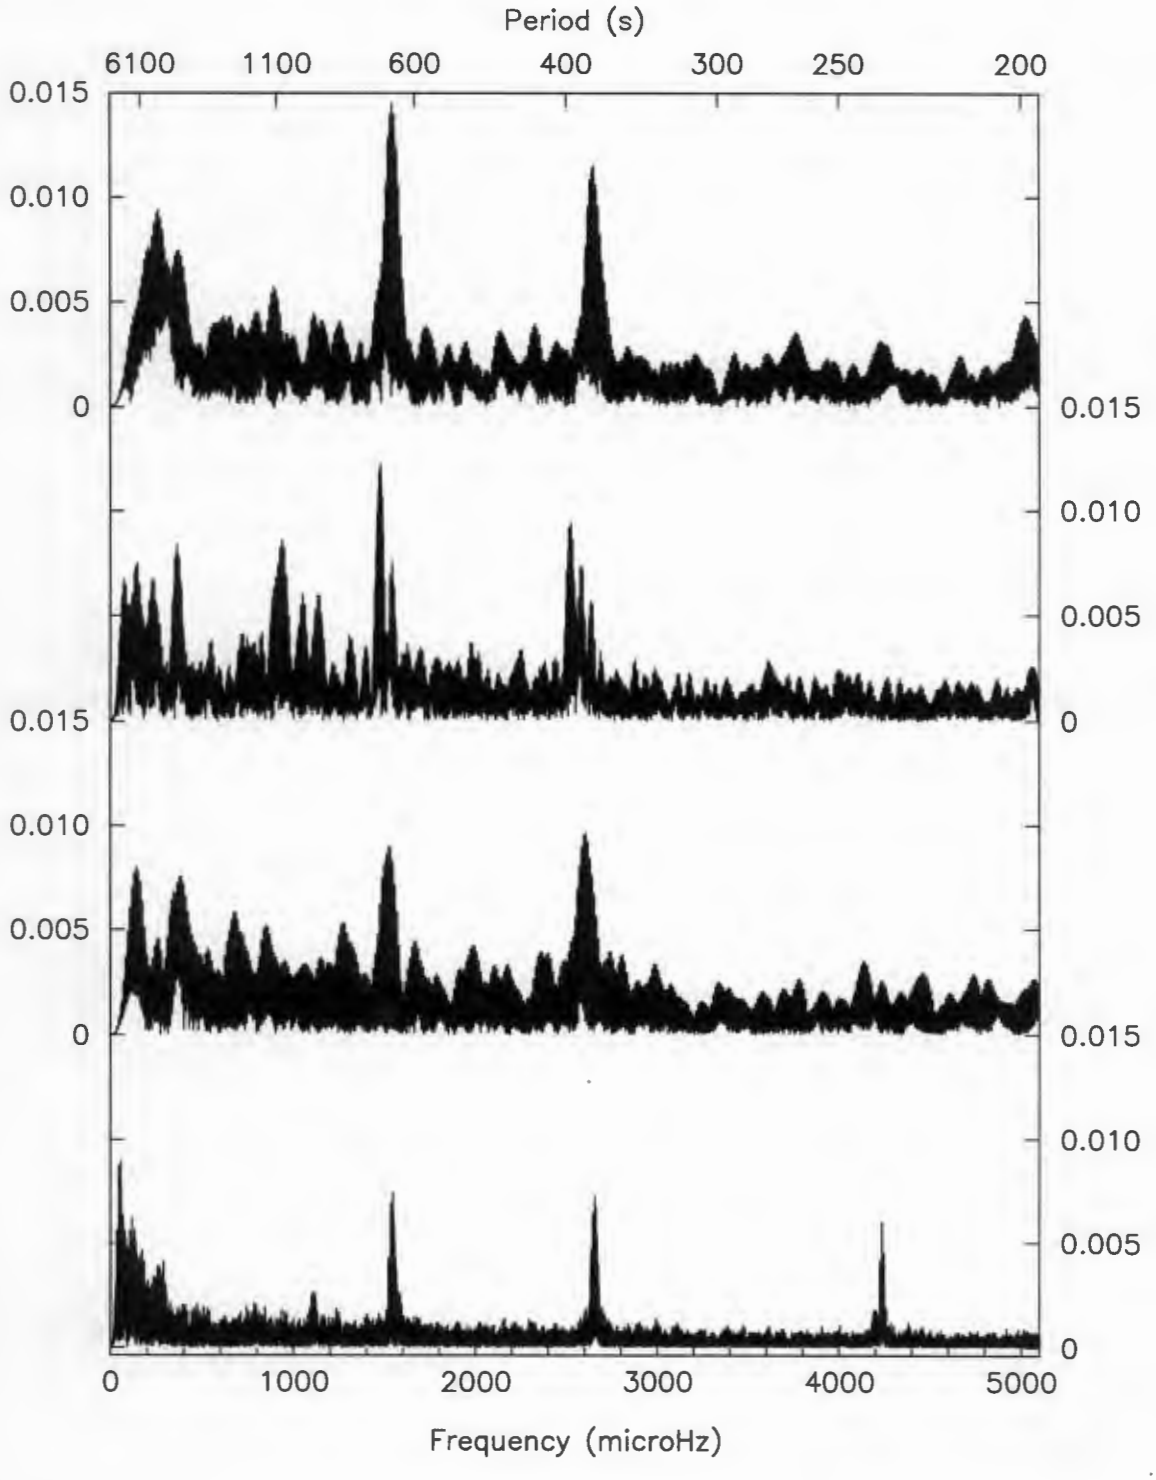

Figure 5.2: The amplitude spectrum of GW Librae in (*a*) March '97, (*b*) April '97, (*c*) September '97, and (*d*) May/June '98.

clear that least squares prewhitening of signals one by one ceases to give meaningful results after the second iteration (often sooner). This is because the algorithm cannot completely deconvolve one signal from the others, due to the phenomenon of *spectral leakage*: window patterns of frequencies interfere with each other. If a dataset covers 90% or more of their beat period[5], the spectral leakage between two signals will be very small, but for datasets that only partially sample beat periods, separating one frequency from another becomes very difficult or impossible (Kurtz 1998, O'Donoghue 1998, Kleinman 1998a, Vuille 1998).

The only way to hope to analyse a transform of partially-resolved, closely-spaced frequencies is to fit the frequencies simultaneously. One starts by finding the first two signals using the method described above, and then fitting them simultaneously with a nonlinear fitting routine that looks for the best possible fit, and then prewhitening the original (unprewhitened) data with these refined parameters. A third signal is then located in the residuals, and all three are fitted simultaneously to the data, first with a linear least squares routine to find better values for their parameters, which are then refined by fitting with the nonlinear algorithm. Only when all the signals present in the data can be extracted, parameterized and simultaneously fitted to the data successfully, can a scientific analysis of them begin.

The *EAGLE* routines that perform these fits are *lstsq* (for least squares fitting) and *nonlin* (for nonlinear fitting). I will give a brief outline of the mathematics behind these routines, from Press *et al.* (1992).

A linear least squares routine finds the best-fitting values of the parameters a_k of a set of functions

$$y_i(x_i) = \sum_{k=1}^{M} a_k X_k(x_i) \tag{5.5}$$

fitted to a set of data points. The *basis functions*, $X_1(x), \ldots, X_M(x)$, are arbitrary fixed functions of x. In our case, the functions are sines and cosines, the y_is are the magnitudes, the x_is are the times, and the a_k are the amplitudes and phases (in the

[5]The beat period P of two closely-spaced signals with periods P_1 and P_2 is $\frac{1}{P} = \frac{1}{P_2} - \frac{1}{P_1}$. The more closely spaced two frequencies are in period, the longer their beat period is.

lstsq routine, the frequencies are fixed).

We define the *merit function*

$$\chi^2 = \sum_{i=1}^{N} \left[y_i - \sum_{k=1}^{M} a_k X_k(x_i) \right]^2 . \tag{5.6}$$

The best values for the parameters a_k are those that minimize χ^2. To minimize (5.6), we construct a matrix $\mathbf{A}$ whose $N \times M$ components are the M basis functions evaluated at the N abscissas x_i, or

$$A_{ij} = X_j(x_i) , \tag{5.7}$$

called the *design matrix* of the fitting problem.

The minimum of (5.6) occurs where the derivative of χ^2 with respect to all M parameters a_k vanishes. This condition yields the M equations

$$0 = \sum_{i=1}^{N} \left[y_i - \sum_{j=1}^{M} a_j X_j(x_i) \right] X_k(x_i) \qquad\qquad k = 1, \ldots, M , \tag{5.8}$$

which, by interchanging the order of summations, we can write as

$$\sum_{j=1}^{M} \alpha_{kj} a_j = \beta_k \qquad \text{or equivalently} \qquad [\alpha] \cdot \mathbf{a} = [\beta] \tag{5.9}$$

where

$$\alpha_{kj} = \sum_{i=1}^{N} X_j(x_i) X_k(x_i) \qquad \text{or equivalently} \qquad [\alpha] = \mathbf{A}^T \cdot \mathbf{A} \tag{5.10}$$

an $M \times M$ matrix, and

$$\beta_k = \sum_{i=1}^{N} y_i X_k(x_i) \qquad \text{or equivalently} \qquad [\beta] = \mathbf{A}^T \cdot \mathbf{b} \tag{5.11}$$

a vector of length M. We have defined a vector $\mathbf{b}$ of length N by $b_i = y_i$, and denoted the M vector whose components are the parameters to be fitted, $a_1, \ldots, a_M$, by $\mathbf{a}$.

Equations (5.8) or (5.9) are called the *normal equations* of the least squares

problem. They can be solved for the vector **a** by any of a selection of standard methods well-suited to computers, see for example Chapter 2 of Press *et al.* (1992).

In the *nonlin* routine, the frequencies as well as the amplitudes and phases are allowed to vary as the routine hunts for the best fit. I will now outline the mathematical techniques underlying nonlinear fitting. The approach is the same as outlined above – defining a χ^2 merit function and determining the best-fit parameters by its minimization – except that now the model depends *nonlinearly* on the parameters $a_1, \ldots, a_M$. With nonlinear deependencies, the minimization must proceed iteratively. Given a set of starting values for a_k (I obtain this by first invoking the *lstsq* command to obtain a simultaneous linear fit), one needs a procedure that improves the initial solution. The procedure is repeated until χ^2 stops (or effectively stops) decreasing.

The procedure amounts to taking a step down the gradient, towards a "minimum energy solution", in other words,

$$\mathbf{a}_{next} = \mathbf{a}_{current} - const \times \nabla\chi^2(\mathbf{a}_{current}) \tag{5.12}$$

where the constant is small enough not to exhaust the downhill direction. Derivations and detailed discussion can be found in chapters 10 and 15 of Press *et al.* (1992).

Translating this equation into a computer-usable form, and fitting a large set of sinusoids to the vast datasets generated by multisite campaigns, generates very large matrices, which are very demanding in terms of computer memory and processing power. Prior to 1998, we ran these routines on an alpha machine, but this year Darragh O'Donoghue helped me recompile the code to enable them to be run on a Pentium II PC, whose greater power allowed a refinement of *nonlin*, which has increased its ability to find solutions by a factor of 5.

Chapter 6

GW Librae in 1998

The resolution of the 1998 amplitude spectra is much higher than those of the previous year, so I shall first try to interpret GW Librae's behaviour in 1998 before re-examining the 1997 data. The results of the 1998 analysis may make it easier to interpret the unresolved eigenmodes in the previous year's data.

When GW Lib was observed in May and June of 1998, the power in the amplitude spectrum concentrated in three principal regions (Figure 5.3): at 236 s (4237 μHz), around 376 s (2660 μHz), and around 650 s (1545 μHz). I will examine each of these regions in turn, and then move on to the rest of the amplitude spectrum to look at the linear combination modes.

6.1 The Mode At 236 s

In 1998, the power in GW Librae's pulsation spectrum concentrated in roughly the same regions of the Fourier transform that were occupied in 1997, with one exception: the 236-s region showed no hint of power in 1997, but in 1998 a strong, stable signal was present here throughout the May and June observing runs - in fact, it is the most stable and constant feature of the 1998 observations.

The 236-s structure (Figure 6.1) is by far the simplest feature in the amplitude spectrum. It appears to be composed of only one dominant signal, and can therefore be used as a diagnostic of the health of the rest of the transform. A noticeable feature

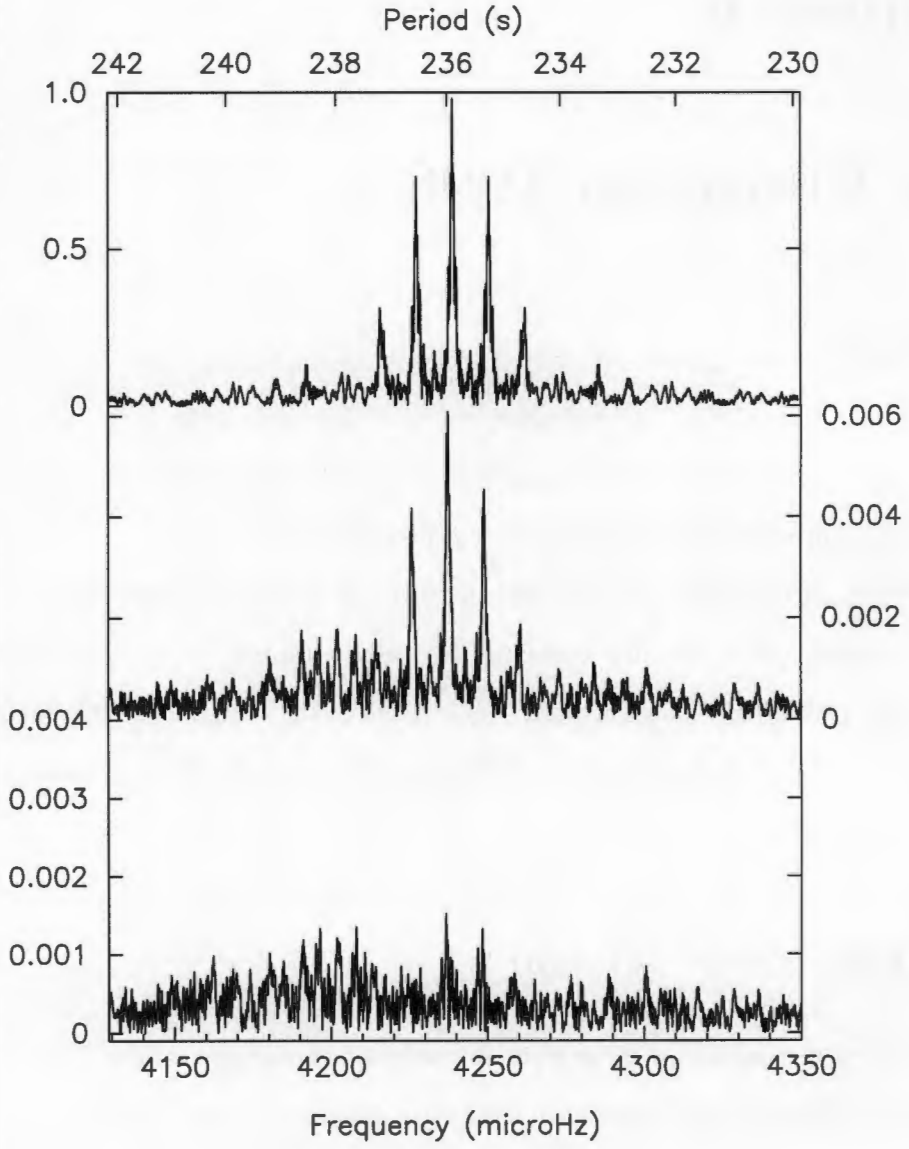

Figure 6.1: The 236-s mode. The first panel shows the spectral window, the second the mode at 236 s, and the third shows the results of prewhitening with the 236-s signal. (The bump left of centre is probably a linear combination mode.)

of many of the spikes in the transform is the fine splitting they show at their tips. The 236-s signal shows this effect clearly (Figure 6.2). Such splittings could indicate the presence of two barely resolved signals, or of a modulation in the amplitude or phase of a single signal, or of a timing error in the data.

The possibility of a timing error is a very serious concern; no analysis results can be trusted until this possibility is ruled out. The time of each observation is fed into the *.fits* file of every CCD frame directly from the Observatory's time service, which is calibrated on the Global Positioning Satellite system's signals. Although I always check the time service at the start of each run, it has occasionally been wrong (it may not be correctly reset after a power failure, for example), and my check is accurate only to about a tenth of a second.

The best way to check for timing errors is to make use of an O-C ("observed minus calculated") diagram. I chopped each night of data into four chunks, and used the linear least squares routine described in Chapter 5 to determine the phase of the 4237.5 μHz (236-s) signal in each chunk. O-C would then be the phase of a chunk minus the phase for the entire dataset. Any timing errors would be revealed as clumps of phases offset from the main O-C curve (which, for a noise-free, stable oscillation, would be flat, with O-C equal to zero).

Figure 6.3 shows the resulting O-C diagram. There are no sudden jumps that would indicate a timing error, but the curve is not flat; there is a definite phase variation over the length of the dataset. It is a relief to be able to confine any timing error to be smaller than the scatter in the diagram[1], but the O-C diagram presents another mystery: what is the reason for the phase drift?

A slight variation in the frequency of the eigenmode would cause the splitting we see in the peak, and would result in a phase drift in the O-C diagram. If the June data have a slightly different frequency from the May data, then the phases of the 4237.5 μHz signal (which is calculated from the combined dataset) in the May data would show a systematic drift. This is exactly what is observed in the O-C diagram.

[1]Unless the timing error is exactly in phase with the 4237.5 μHz signal...

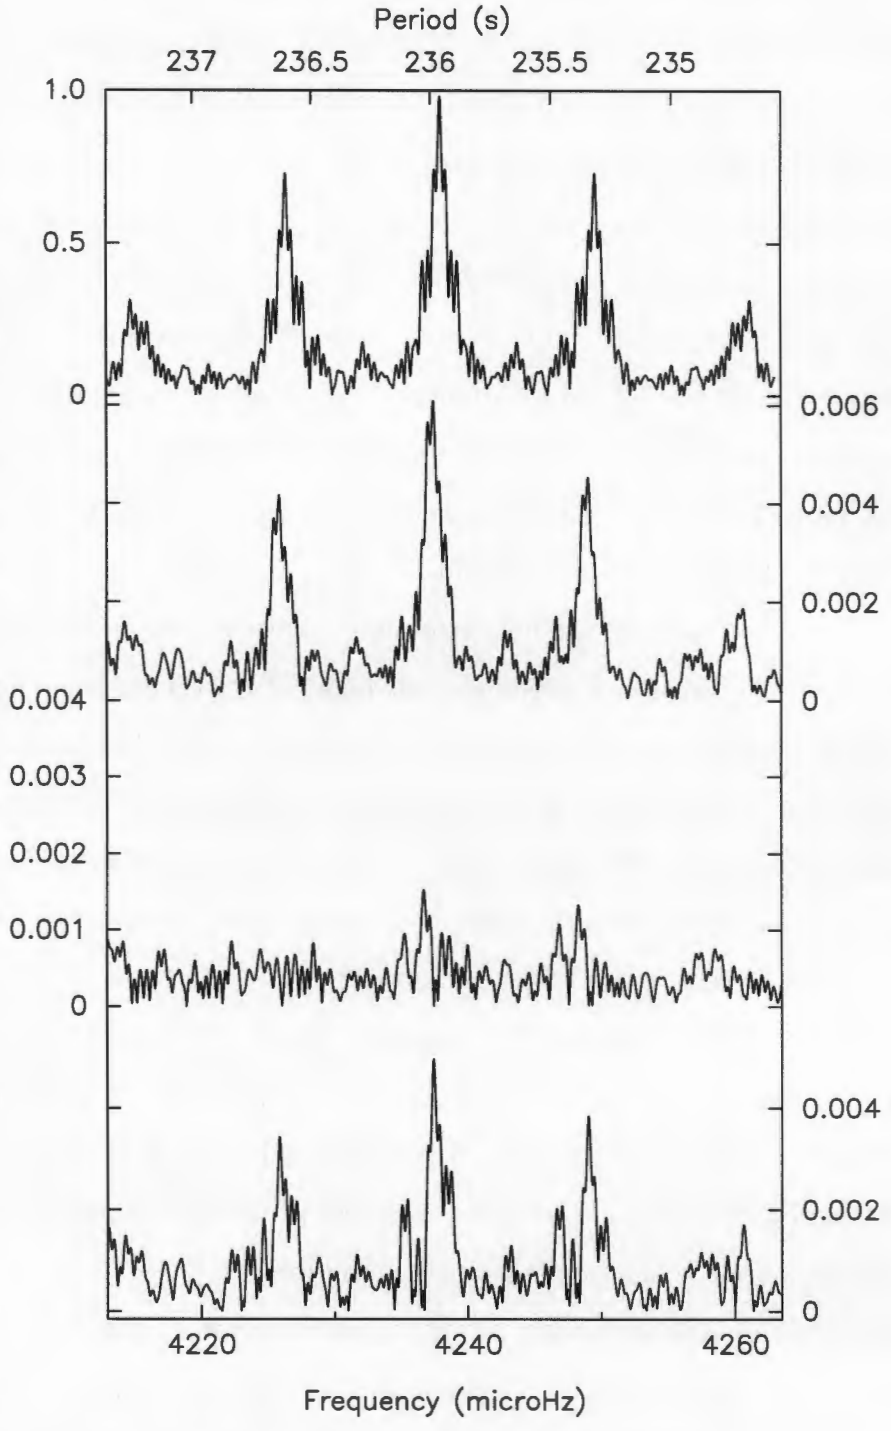

Figure 6.2: A magnification of the 236-s region. The top panel is the spectral window, the second panel shows the asymmetry in the 236-s mode. The third panel shows the low-amplitude 4236.7-μHz signal revealed after prewhitening with the dominant 4237.5-μHz signal. The fourth panel shows how the dominant signal becomes much more symmetric after the (partial) removal of the 4236.7-μHz signal.

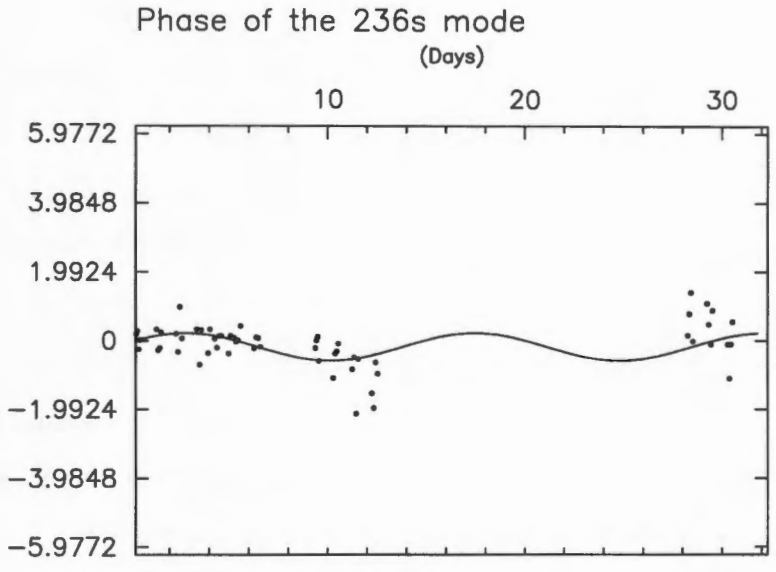

Figure 6.3: The O-C diagram for the phase of the 236-s mode over the observing period. A sinusoid with a frequency of 0.79 μHz is fitted to the phases (see discussion in the text).

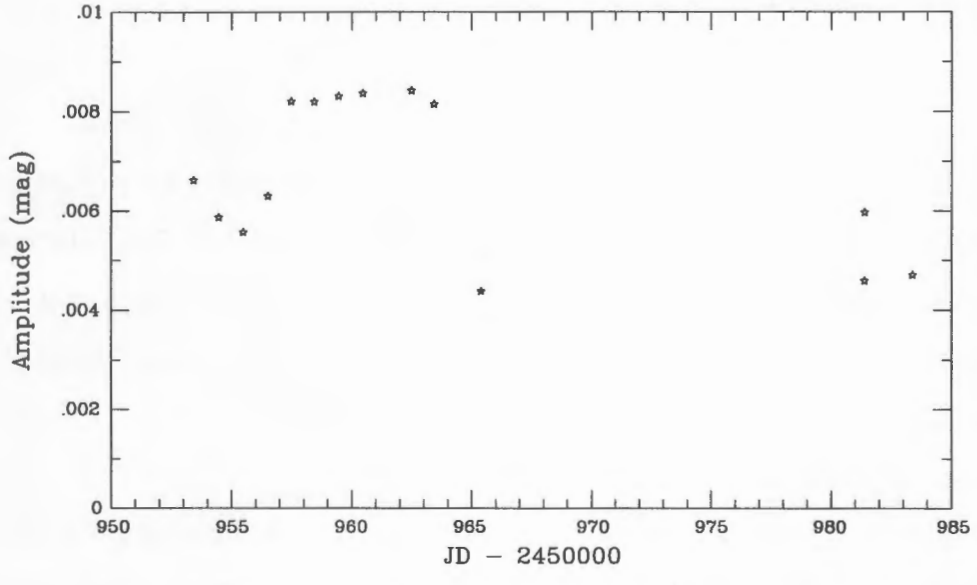

Figure 6.4: Amplitudes of the 236-s mode on individual nights.

The frequency from the combined dataset is 4237.48 μHz, and 4237.35 μHz for the May data alone. This corresponds to a 1 in 300 000 change in frequency[2].

Since GW Librae is an accreting system, there may be perturbations on the temperature of the surface layers. Are we seeing the effects of accretion heating on the eigenmodes? If the modes are unstable in frequency, they will present a formidable challenge to analysis, because frequency modulation will produce a thick forest of sidebands, making identification of closely spaced real modes extremely difficult. However, in the case of the behaviour of the 236-s region, there may be a simpler explanation. Figures 6.1 and 6.2 reveal a possible low amplitude signal in the residuals when the data are prewhitened with the 4237.5 μHz signal. This signal, at 4236.7 μHz, is only 0.79 μHz away from the dominant signal. A simultaneous nonlinear fitting of both frequencies gives the following results:

Table 6.1: The 236-s mode.

Frequency μHz	*Amplitude* mMag	*Spacing* μHz
4237.49	6.03	
4236.70	1.91	0.79

The presence of this second signal neatly explains the asymmetric appearance of the dominant mode. I prewhitened the data with the second signal alone (the fourth transform in Figure 6.2) in the hope that the dominant mode would then match the window exactly. Unfortunately (and this will become a recurring theme throughout this analysis), the least squares and non-linear fitting routines do not cope well with such closely-spaced frequencies.

A least squares fit of the second signal alone gives it an amplitude that is 1.6 times too large: the two signals are too close to be properly deconvolved. For this

[2]It must, however, be noted that we have no way of knowing whether the phase shift was ϕ, or $\phi + 2\pi$. The latter possibility implies a much more rapid variation over the observing gap.

reason, a nonlinear fit to the low amplitude signal alone does not converge; only a simultaneous nonlinear fit of the two frequencies will fit the weak signal properly. Because of this, it is not possible to fit and extract the weak signal alone. The fourth transform in Figure 6.2 shows an attempt to do this; while the result is untidy, the dominant mode does show much greater symmetry and a closer resemblance to the window pattern.

These two signals will beat against each other with a beat period equal to the inverse of their separation in frequency: 14.7 days. Therefore we expect to see not only a variation in phase but also a sinusoidal variation in amplitude of the combined signal over the run length. Figure 6.3 is a plot of the amplitude of the 236-s signal for each night of data. While the scatter is large, a variation is indeed seen. I have plotted a best fit sinusoid of 0.79 μHz through the phase points in Figure 6.3. The large scatter in the phase and amplitude plots is due to the much higher noise level in the transforms of the short data segments relative to the noise level in the transform of the whole dataset.

Kepler *et al.* (1995) give a formula for the phase variation in an unresolved doublet

$$tan\ \phi_{max} = \frac{A_1/A_0}{[1-(A_1/A_0)^2]^{\frac{1}{2}}} \tag{6.1}$$

where A_0 and A_1 are the amplitudes of the largest and smallest peaks. Using the amplitudes given in the results of the simultaneous nonlinear fit, $\phi_{max} = 0.33$ radians. Fitting a sinusoid with the beat frequency to the phase points gives an amplitude of 0.40 radians. This number is not likely to be very accurate because of the sparsity of points and their large scatter, but it does show that, within errors, the result is not inconsistent with the expected variation. The discrepancy between the calculated and observed variations may also be due to an incorrect deconvolution of the two amplitudes by the simultaneous nonlinear fit, or perhaps to the presence of a third low-amplitude, closely spaced component.

Prewhitening both frequencies reveals an extremely low amplitude possibility for a third component. A simultaneous nonlinear fit of all three components gives the

following results:

Frequency	*Amplitude*	*Spacing*
μHz	mMag	μHz
4237.56	4.83	
4236.92	2.77	0.64
4234.79	0.43	2.13

My gut feeling is that the third signal is not real: with such a low amplitude it is statistically insignificant - the amplitude of the noise is higher. If the 236-s signal is a rotationally-split $l = 1$ mode, then either there is a third component hidden in the noise, with an amplitude of 0.5 mMag or less, or the third component is not excited.

There is one final point of concern to be addressed in this region: at first glance, aside from its asymmetry, the splitting of the 236-s mode looks very similar to the splitting in the window pattern (Figure 6.2). The two identified components have a separation of 0.79 μHz, or 14.7 days, and the dataset consists of two weeks of data, a two week gap, and three additional nights of data. One should always regard with great suspicion any signals or separations having periods close to the run length or to significant patterns in the dataset, as they may be artefacts of the window pattern. Is it possible that the the 236-s signal is a single frequency with a pure window pattern that has been distorted by noise, misleading me into thinking that it consists of two closely spaced frequencies?

Aside from the evidence against this in the phase and amplitude variations (which are noisy and therefore open to challenge), a close examination of the "hair" on the principal spike in the window pattern shows that the splittings are separated by 0.41, 0.46 and 0.43 μHz, i.e., different from the splitting of the 236-s signal. The full width at half maximum of the principal spike of the window pattern, equal to the inverse of twice the run length, is a rough guide to the resolution of a Fourier transform. The run length is 31 days, so the FWHM is 0.19 μHz. In practice the resolution is less than this because of the messiness of the window pattern (most of the data are in

the first two weeks, so the dominant structure in the window is a 14 day pattern, and only the top third of the aliases has the 31 day resolution), but the 236-s structure is simple enough to be easily resolved.

6.2 The Power Near 376 s

The 376-s region shows much greater complexity than the 236-s region. It evidently contains many more signals than the 236-s region. Extracting the signals in this region has proved to be a nightmare: there seem to be many very closely spaced, barely resolved frequencies here. Least squares and nonlinear fitting and prewhitening fail because it is not possible to deconvolve one frequency completely from all its friends: prewhitening one frequency also removes parts of the other frequencies, leaving in the end an uninterpretable tangle of partially resolved frequencies and the ghosts of signals not properly prewhitened.

In Utopian Fourier Space, each signal is independent and orthogonal, but since we are not working with a noise-free transform, and since there are nonlinear effects inherent to the pulsations themselves, the signals are not independent and orthogonal (Kleinman 1995). Therefore, the only way to find their true values (*if* the dataset is good enough, i.e. it adequately samples the beat cycles of closely-spaced signals, and the signal-to noise is high) is to fit the whole lot simultaneously using a nonlinear fitting routine. This is not trivial to do, because for nonlinear fitting routines to work, one has to supply them with input parameters very close to the actual values; a catch-22 situation. I attempted to solve this problem by imposing an interpretation and using it as a guide to finding starting frequencies for the nonlinear fit.

6.2.1 A first hypothesis

The 376-s region (Figure 6.5) appears to contain two dominant structures: a triplet of resolved signals of more or less the same amplitude and with separations of $\sim 1\mu$Hz, and a convolution of unresolved signals that resembles the stump of a tree trunk after a forest fire. Our first hypothesis was that we were seeing two rotationally-split

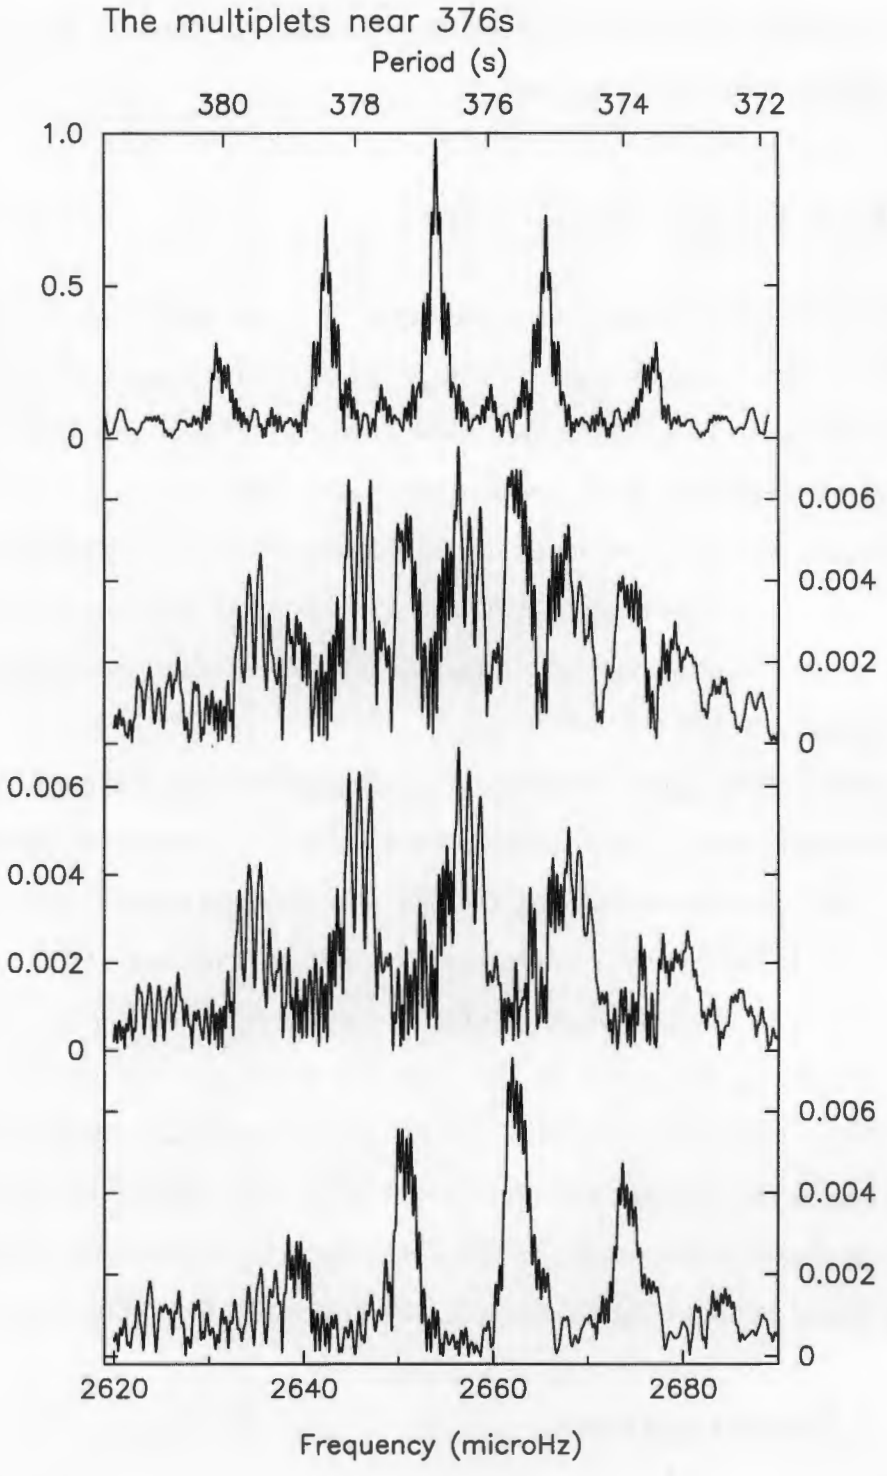

Figure 6.5: The power at 376 s. The first panel shows the spectral window; the second shows the structures around 376 s; in the third and forth panels I isolate the "quintuplet" and "triplet" respectively.

modes; one whose components were resolved, the other whose components were not, possibly because they had a smaller separation.

From equation (3.9), the rotationally-split components of a mode of l equal to 2 would have a separation in frequency 1.67 times the separation of the components of an $l = 1$ mode. Therefore, we reasoned, if the resolved triplet is an $l = 2$ multiplet with $\sim 1\mu$Hz m-splitting, which is on the order of the resolution of the transform, then the $l = 1$ multiplets would have m-splittings of $\sim 0.6\mu$Hz, and would be unresolved. It would be very unlikely to find two multiplets of the same l in this region; in the DAVs, successive $l = 1$ modes are ~ 50s apart (Clemens 1993, 1994; Kleinman 1998b), and from equation (3.10), the $l = 2$ multiplets would have a spacing $\frac{1}{\sqrt{3}}$ of that, or ~ 29s. The two structures in the 376-s region are separated by 2s of less, depending on which are the dominant aliases.

If this region consists of a rotationally-split $l = 1$ triplet and $l = 2$ quintuplet, then all the power in this region must be contained in 8 signals or less (not all possible m-components may be excited to observable amplitude), with m-spacings consistent with equation (3.9). I will now set out to prove or disprove this hypothesis. I will refer to the tree stump structure as the "triplet" and the three resolved signals as the "quintuplet," but I stress that I use the terms for convenience only; I am still a long way from being able to assign eigenvalues to modes.

There is ambiguity as to which set of alias is the true "quintuplet". Without longer runlengths, there is no way to be sure which alias is the true one, but for the purpose of constructing this hypothesis I will choose the left hand alias as the true one, because this gives the quintuplet a more comfortable 2s separation from the $l = 1$ mode. Bradley's models show that the smallest period spacings between modes of different l is on the order of a few seconds (Bradley 1995).

Isolating the "triplet" by removing the "quintuplet" and prewhitening it, shows that three frequencies seem to be sufficient to account for all the power in the triplet. Figure 6.6 shows the three signals: in the second transform I have removed the central frequency; in the third, the central and left hand frequencies; and in the final transform I have removed all three frequencies. The signals look very messy and not

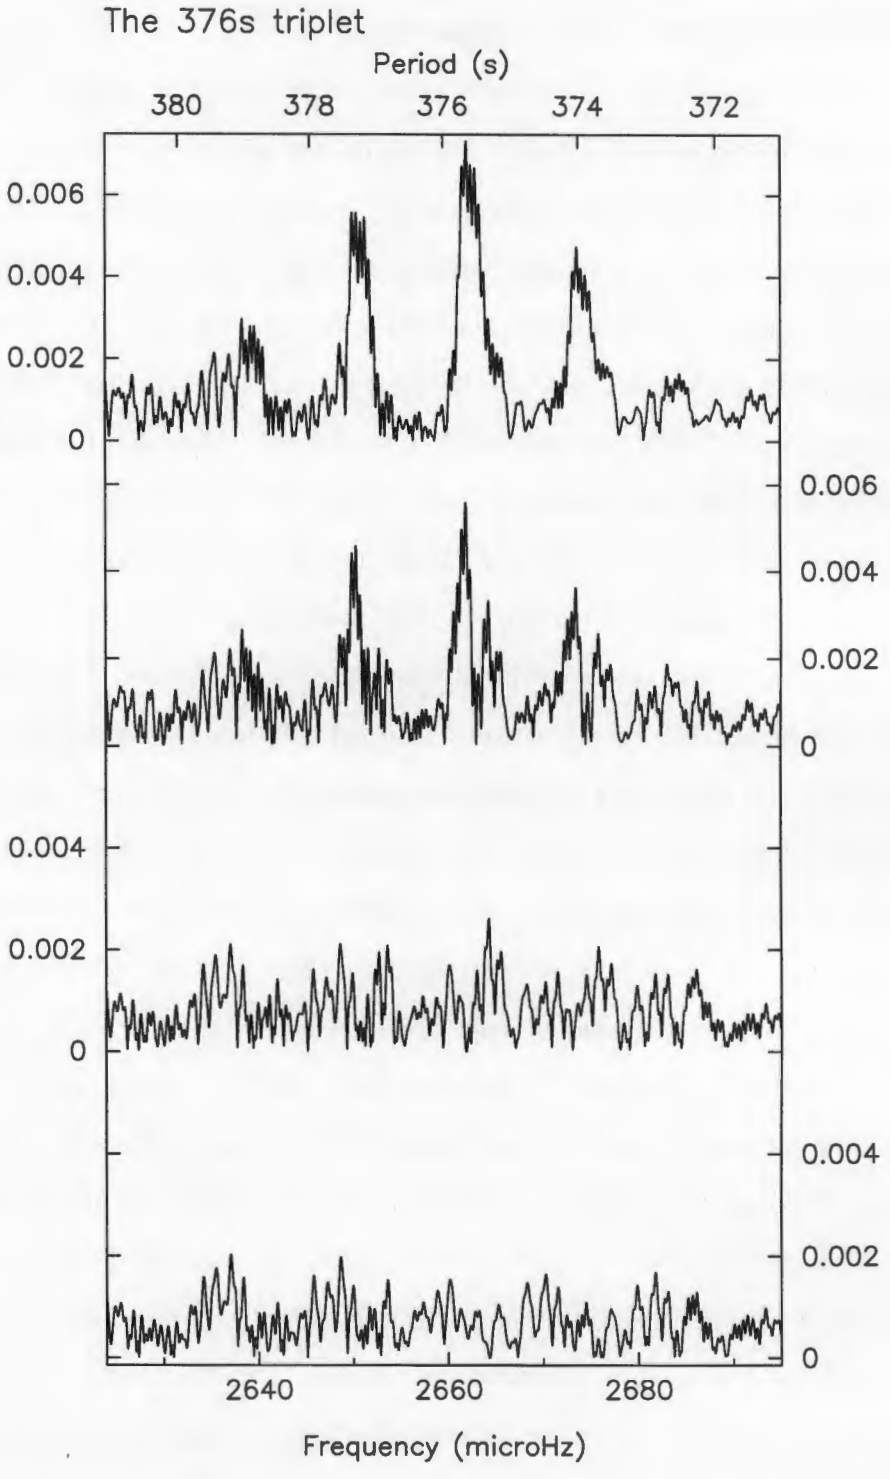

Figure 6.6: The components of the "triplet". In the second transform the central frequency has been removed. In the third transform the central and left hand components have been removed, and in the final transform all three components have been removed.

at all like the window pattern, but removing only three frequencies seems to account for everything in this structure. So far so good.

The real test of the hypothesis is to attempt a simultaneous nonlinear fit of the three "quintuplet" signals and the above three "triplet" components. Doing this was a tedious and time-consuming process, because the *nonlin* routine needs starting values as close as possible to the final values. It does not cope very well with so many closely spaced frequencies, so one needs to give it all the help one can. I started by fitting the highest amplitude signals first (stronger signals are less affected by the noise and are therefore more well determined), prewhitening the residuals, adding the highest amplitude frequencies to the set, attempting a simultaneous nonlinear fit ..., gradually building up a set of signals whose values come closer and closer to the overall best fit solution.

In fact this needs to be done for all signals simultaneously over the entire dataset, because of the problem of "spectral leakage": signals can affect each other even if they are far apart in frequency. This procedure is enormously time consuming: even with a fast Pentium it requires weeks of work. The final results are tabulated below.

Table 6.2: The signals near 376 s.

Quintuplet			*Triplet*		
Frequency	*Amplitude*	*Spacing*	*Frequency*	*Amplitude*	*Spacing*
μHz	mMag	μHz	μHz	mMag	μHz
2643.11	3.06	1.72			
2644.83	4.61	1.07	2661.91	4.71	1.03
2645.90	5.39		2662.94	5.51	
2646.96	4.35	1.06	2663.98	5.51	1.04
2648.47	2.23	1.51			

In Figure 6.7 I have isolated the "quintuplet" and "triplet" again, this time using

the best fit frequencies.

The spacings of these frequencies show that our hypothesis is in trouble. To examine equation (3.9) in greater depth, the spacings of the m splittings produced by a rotation frequency Ω (we assume solid body rotation) are given by

$$\Delta\nu = (1 - C_{kl})\Omega m \tag{6.2}$$

(equation 3.9 reproduced), where

$$m = 0, \pm 1, \pm 2, \ldots \pm l, \qquad \text{and} \qquad C_{kl} = \frac{1}{l(l+1)} \tag{6.3}$$

in the asymptotic limit, which in practice means for k higher than ~ 4 (Kleinman 1998a). This gives

$$\Delta\nu = \frac{1}{2}\Omega m \quad \text{when} \quad l = 1 \tag{6.4}$$

and

$$\Delta\nu = \frac{5}{6}\Omega m \quad \text{when} \quad l = 2\,. \tag{6.5}$$

Therefore, in the asymptotic limit, the ratio of the m-spacings of $l = 1$ triplets to the m-spacings of $l = 2$ quintuplets is 0.60. The spacings in Table 6.2 show that the splitting of the three best determined frequencies of the "quintuplet" is the same as that of the "triplet" components. Evidently we are not seeing two modes of different l split purely by solid body rotation.

If the values for the signals in the above table are accurate (that is, if the analysis routines have successfully deconvolved the signals from each other), and if our hypothesis is correct, then any explanation for the behaviour in the 376-s region will have to account for the following features:

1. The equal splitting of the triplet and the central components of the "quintuplet."

2. The wider spacings of the outermost components of the "quintuplet".

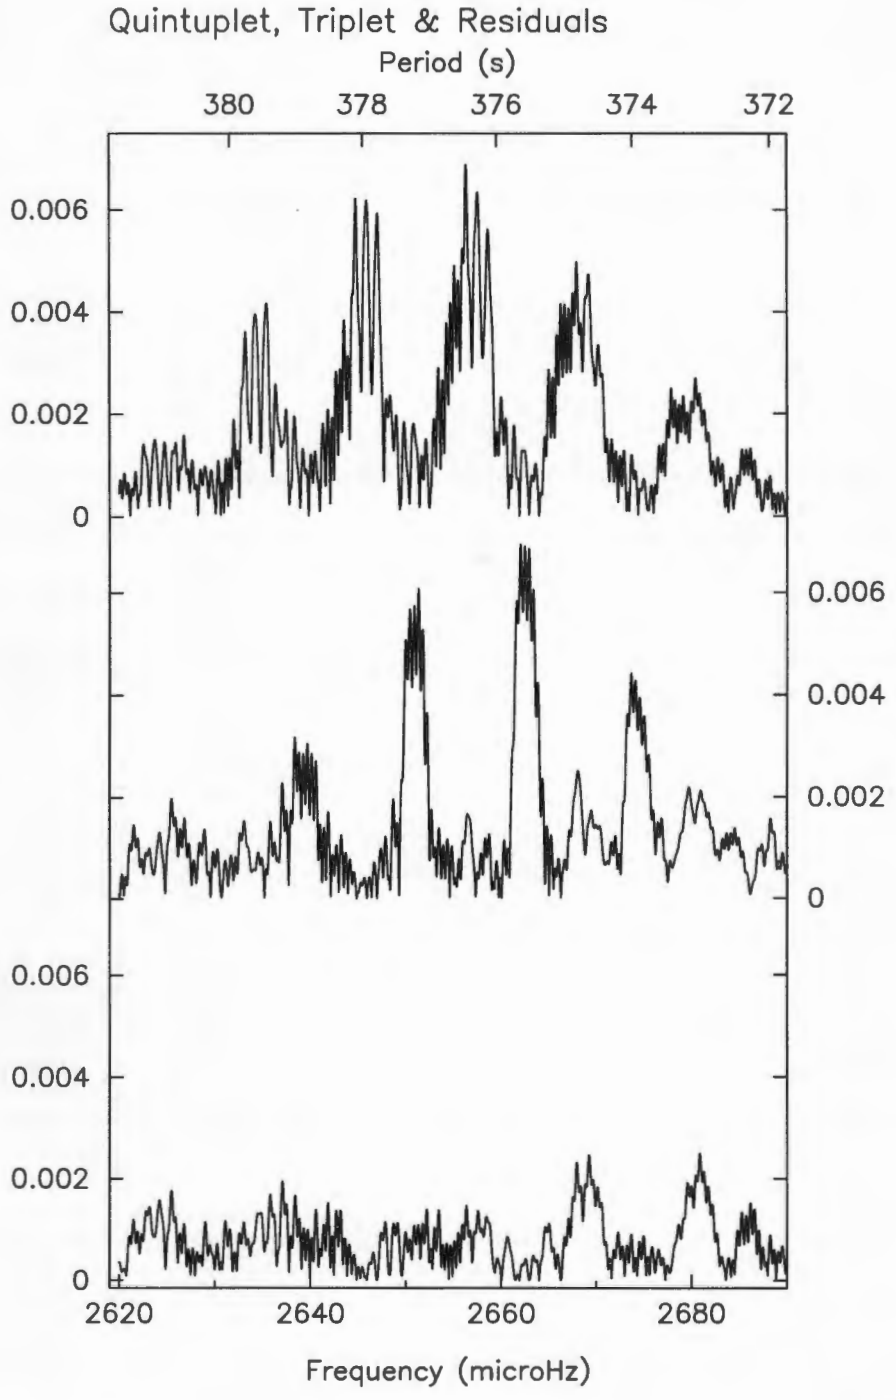

Figure 6.7: The 376-s region. *(a)* I have isolated the quintuplet by removing the 3 best-fit frequencies of the triplet. *(b)* I have isolated the triplet by removing the quintuplet. *(c)* The residuals after all 8 frequencies are removed. The structure at 2680 μHz (along with its alias at 2668 μHz) is a linear combination mode.

3. The 0.79 μHz splitting of the doublet at 236 s (along with the spacings of any other multiplets found in other regions of the transform).

4. The small separation in period between the "quintuplet" and the "triplet."

5. The very different appearance of the two modes, even though their components appear to have the same spacings and similar amplitudes.

With the sparsity of modes in this dataset I have very little to work with, since at any given time, a cool DAV puts power in only a few of its many eigenmodes, and one only sees a tiny part of the big picture.

In constructing an hypothesis, I have to try to explain GW Librae's pulsation spectrum in terms of what has been seen in other white dwarf pulsators, trying not to resort to any new physics. Any claim for seeing something never seen before has to be backed up with observational evidence much more extensive than what I currently have.

In addition, a central tenet of the asteroseismological study of white dwarfs has been that the pulsators are "otherwise normal stars", so that the insights provided by asteroseismology are applicable to all white dwarfs. GW Librae is the only accreting DAV known, but in order to study it using the asteroseismological tools made available by the study of other white dwarfs, I have got to assume that it is an "otherwise normal pulsator."

Upon finding that a star has a set of pulsational frequencies that does not resemble anything observed before, it is a good idea to take a close look at the observations and the analysis techniques to examine them for possible sources of error. In section 6.1 I searched the data for timing errors. I now examined the data, point by point, to eliminate bad points and scruffy regions in which the signal to noise was poor (due to poor observing conditions, large airmasses or bad guiding).

A useful way to test whether signals in a dataset are real, is to divide the dataset in two and then to see if the same signals are present in both halves. Unfortunately, I need the entire length of the dataset in order to resolve these closely split signals.

An alternative is to look at alternate days, or, to avoid horrible aliasing, to split the data into first and second halves of each night.

I split each individual run in half and then arranged the portions into two data sets, each with half the number of points but with the same total baseline as the full dataset. Figure 6.8 shows the spectral windows of the two datasets. With half the number of points, the Fourier transforms of each are a lot noisier. I analysed the half-datasets in the same way as the full dataset. With the higher noise in each, I found fewer signals.

Since the mode at 236-s (Section 6.1) is the simplest structure in the transform, the values for its frequencies in each half-set should be nearly identical to the numbers found in the total dataset, if anything else found in it is to be believed. I found that the values for the dominant component of the 236-s mode in each half-set differed by 0.01 and 0.06 μHz from the frequency in the full dataset, and 0.04 and 0.22 μHz for the low amplitude component. Except for the last number, these values are an order of magnitude smaller than the resolution, so we can consider the values of the 236-s mode yielded by each half to be identical to the signal in the full dataset.

The 376-s area is not so simple. I have listed the frequencies and amplitudes of all the best-fitting signals in the 376-s region from each dataset in Table 6.3 below. Where a matching frequency is not the highest alias, I have put it in italics. A glance at the numbers gives one great relief, because similar signals appear in all three transforms. The three dominant components of the quintuplet structure give the best matches, most likely because they seem to be the only signals that are adequately resolved. Except for one case, they match to within 0.1 μHz in all three datasets.

As for the reality of the two outer, low amplitude components of the "quintuplet," only one candidate was found in a half-dataset. This may simply be due to the higher noise levels of the half-datasets and that they just are simply not resolved, but I think it wise to be cautious of believing their existence for now.

The interpretation of the tree-stump structure as an unresolved triplet is supported by the closely matching frequencies that come out of the first half-dataset,

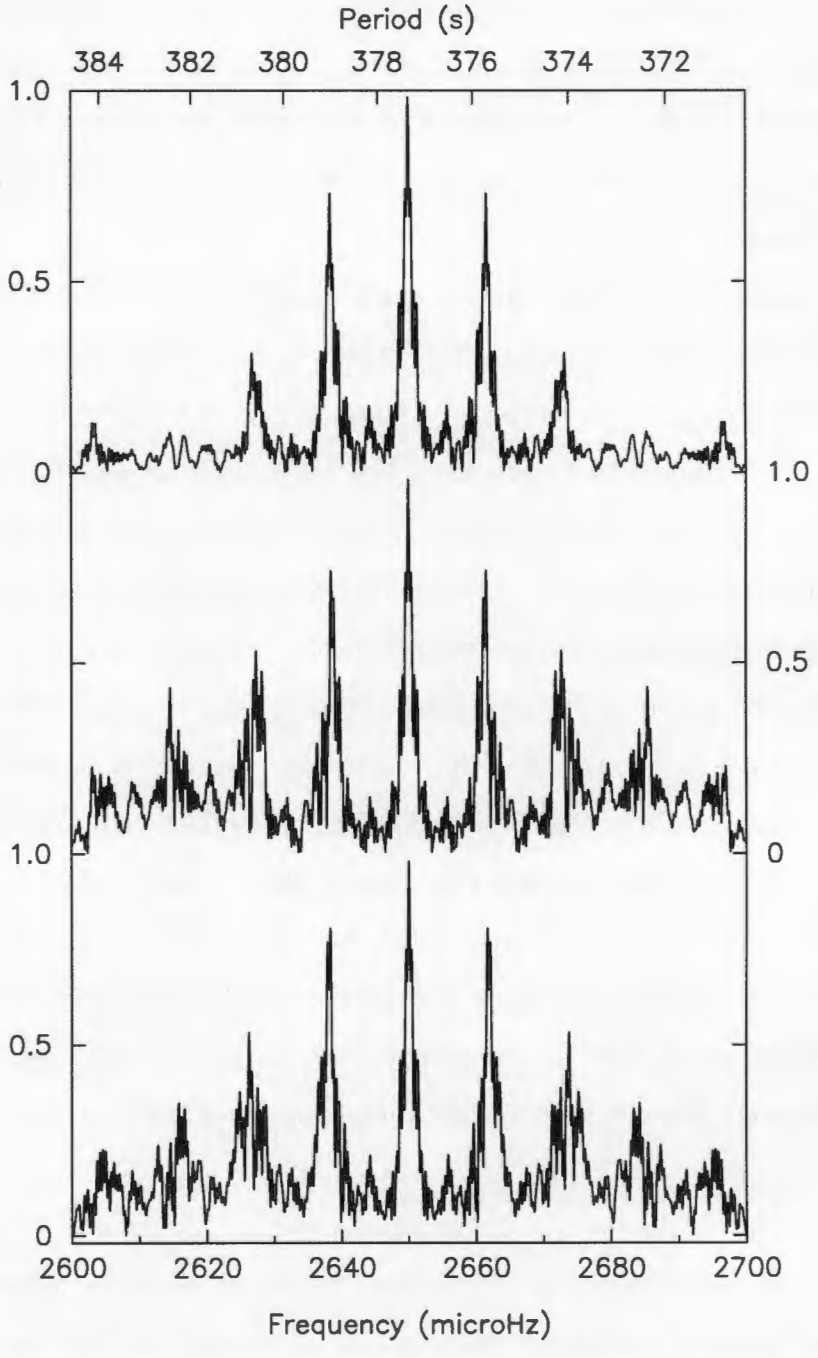

Figure 6.8: The Spectral Windows of the full dataset and the half datasets.

but in the second half, the central frequency does not appear and only one component matches that of the full dataset closely. Unless we find a better way to analyse this unresolved structure, we will have to reserve judgement on the accuracy of the frequencies for the "triplet."

It is interesting to note how different the values for the amplitudes are in each dataset. I have not included the relative phases for each signal, but their differences are even greater. The phase of a signal is determined from the start of the dataset and is therefore meaningless, since one can start observing a star at any time. However, the phase *spacings* are important[3].

Table 6.3: The 376-s signals in the full and half datasets.

Total		*First Half*		*Second Half*	
Frequency	*Amplitude*	*Frequency*	*Amplitude*	*Frequency*	*Amplitude*
μHz	mMag	μHz	mMag	μHz	mMag
2643.11	3.06	*2642.68*	2.66		
2644.83	4.61	2644.86	6.54	2644.76	6.72
2645.90	5.39	2645.94	5.02	2645.44	3.95
2646.96	4.35	2647.06	3.71	2647.03	7.47
2648.47	2.23				
2661.91	4.71	2661.80	5.21	*2661.53*	7.87
2662.94	5.51	2662.92	6.43		
2663.98	3.18	*2664.33*	2.66	2663.91	4.39

The disparity in the amplitudes and phase spacings between the full dataset and its halves suggests that the analysis techniques are unable to deconvolve closely

[3]The phase spacings of three phases ϕ_1, ϕ_2, ϕ_3 are $\phi_1 - \phi_2$ and $\phi_2 - \phi_3$, independent of the actual values of ϕ_1,ϕ_2, ϕ_3, which are determined by the start of the dataset.

spaced signals completely in a dataset that only partially covers the duty cycle. I hope that I have now sufficiently examined the data for errors and artefacts that could be generating false signals; it is now time to investigate the analysis techniques.

First of all, I checked the precision in the code: information can be distorted or lost if the routines don't make sufficient provision for significant figures. However, all the arrays for times and time series ordinates are assigned double precision in *EAGLE*, and provision is made to prevent accumulating rounding errors.

Next, I analysed my data using a completely independent and well-tested Fourier analysis routine (that of Don Kurtz). The structures in the resulting transform were the same as those produced by *EAGLE*.

Having (hopefully) eliminated all sources of error or false signals in the data and in the analysis code, the next step was to try to understand what the analysis routines were actually doing. It is well known from research on solar oscillations that it is not always possible to deconvolve closely spaced signals if there are large gaps in the observations (Kurtz 1998). Could this be what is happening here?

The way to test this is to see what output the analysis programs give for a known input; i.e., by analysing Fourier transforms of simulated lightcurves.

6.2.2 Analysing transforms of synthetic data

I simulated a dataset by writing a Fortran program that calculates, for a given set of input sinusoids, the magnitude for each point in the time series array of the true dataset. In this way I can make synthetic lightcurves with the exact spectral window of the true data. In effect, I am generating multiple signal spectral windows. This is an extremely useful technique in the investigation of crowded amplitude spectra (Kleinman 1995, 1998a).

I created a synthetic dataset containing the best fit frequencies, amplitudes and phases of the three resolved components of the "quintuplet" and the values given for the tree-stump "triplet." The dataset contained only these six sinusoids and was noise free. I then calculated the Fourier transform of this synthetic dataset. Figure

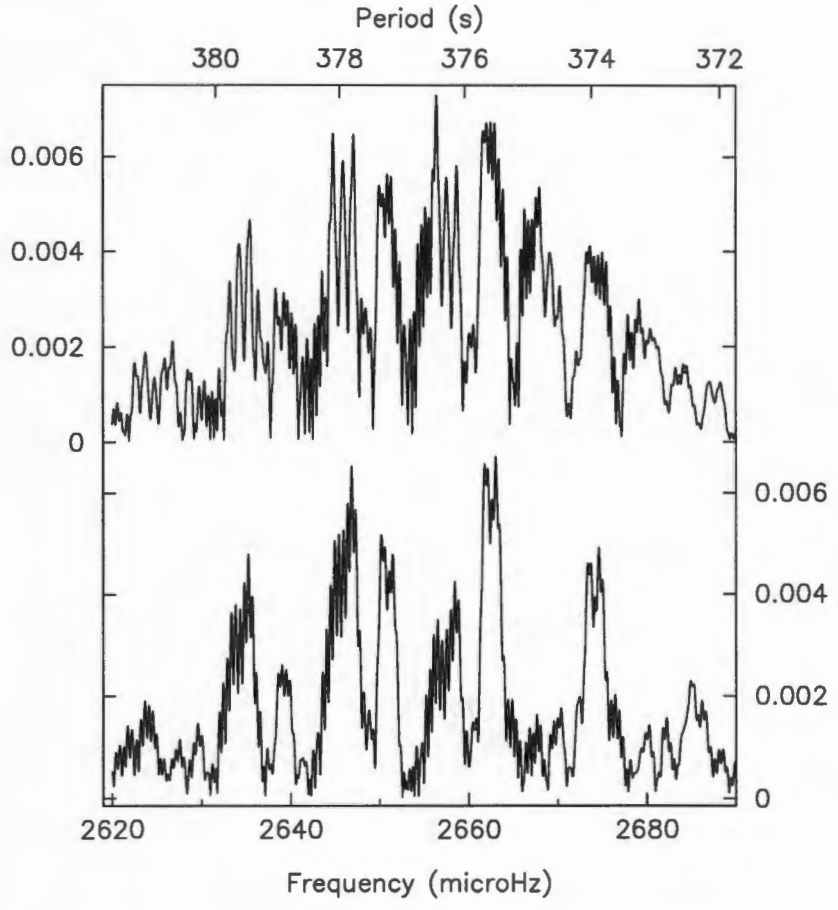

Figure 6.9: Recreating the light curve. The first panel is the Fourier transform of the real data in the 376-s region. The second panel is the transform of a synthetic noise-free light curve made up of the six highest amplitude, best fitting signals in the first transform.

6.9 is the result. The top panel is the transform of the true data, the bottom one, the synthetic data.

It is clear that the transform of a lightcurve made up of the best fit signals does not resemble the real transform. The resolved "quintuplet" is now completely unresolved, and the unresolved tree-stump now appears to be partially resolved! If their beat cycles are inadequately sampled, closely spaced signals cannot be completely deconvolved. From the results obtained by splitting the dataset in two, I guessed that this might be happening here, and this simulation confirms it. What is needed now is to try to determine the extent of the problem, in order to be able to try to obtain an "index of believability" for each signal in the transform.

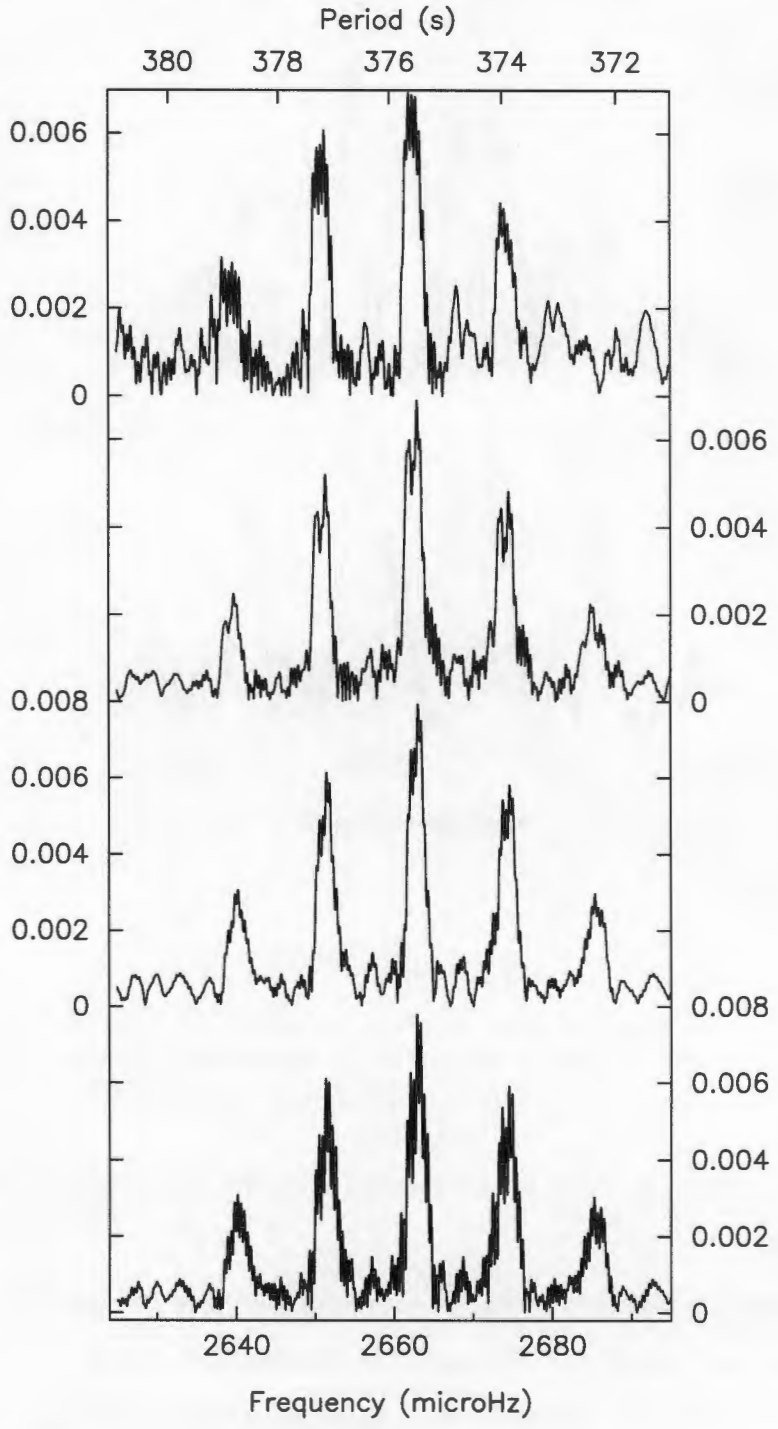

Figure 6.10: The Unresolved Triplet. The first panel is the transform of the real data with the "quintuplet" removed. The second panel is a transform of a synthetic light curve composed of the three best fitting signals from the first panel. The third and fourth panels are transforms of synthetic lightcurves containing three sinusoids of arbitrary amplitude, spaced 0.64 and 0.79 μHz apart respectively, and with relative phases chosen to ensure that the triplets will be unresolved.

If we assume that the 236-s pair with the 0.79 μHz spacing to be an $l = 1$ mode, and the three resolved components with average spacing 1.065 μHz to be an $l = 2$ mode, then if the tree stump is an $l = 1$ mode, its components should have a spacing somewhere between 0.64 μHz (from equations 6.2 to 6.5 and the $l = 2$ spacing) and 0.79 μHz.

Figure 6.10 shows the transform of a synthetic dataset containing the three best-fitting frequencies and amplitudes extracted from the tree stump. I also show two transforms of synthetic datasets consisting of three frequencies with spacings of 0.64 μHz and 0.79 μHz. While these simulations prove nothing, the appearance of their transforms is consistent with the appearance of the transform of the real data. While such simulations cannot prove anything, they can *disprove* hypotheses.

Table 6.4: Analysing simulated data.

Input		*Output*	
Frequency	*Amplitude*	*Frequency*	*Amplitude*
μHz	mMag	μHz	mMag
2662.30	4.71	2662.58	5.97
2662.94	5.51	2663.09	8.42
2663.58	3.18	2660.52	0.70

To test how well the analysis routines are able to deconvolve the components of an unresolved triplet, I analysed one of the synthetic unresolved triplets with the simultaneous nonlinear fitting routine, in the same way that I analysed the real dataset. The results are shown in Table 6.4.

The routine fails to find and fit signals accurately in a noise-free dataset whose contents are known exactly. It is now quite apparent that it will not be possible to extract closely spaced unresolved signals in the GW Librae – their window patterns cannot be deconvolved. All that can be concluded about the tree stump structure

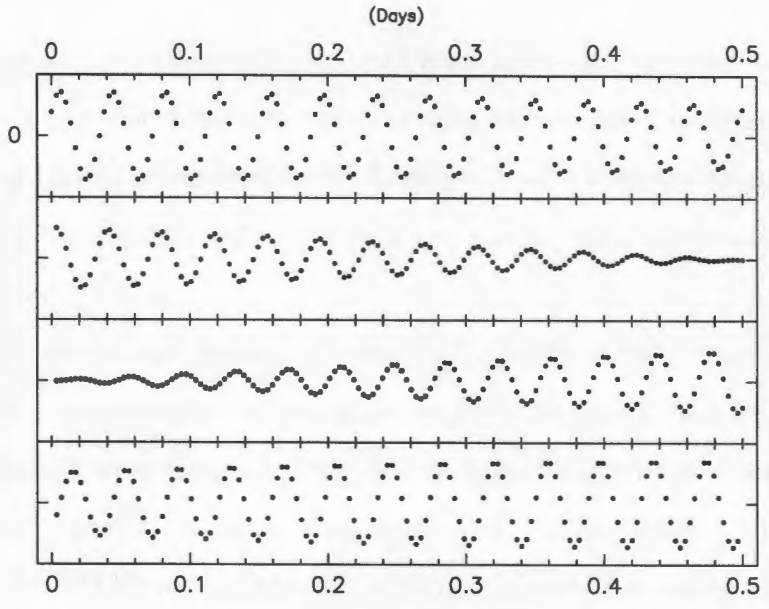

Figure 6.11: A beat cycle of two sinusoids closely spaced in frequency.

is that its appearance is consistent with that of an unresolved triplet with an m-splitting of $0.7 \pm 0.1 \mu$Hz.

I discovered while experimenting with synthetic data that the most important factor determining whether three signals with spacings on the order of the resolution will be resolved, is their relative phases in the sampled sections of their beat cycles. The relative phases determine how the window patterns of three closely-spaced signals will convolute. It may be difficult to picture this; it is unfortunate that our perceptions of life in Fourier Space are very much limited by the fact that we see only the frequencies and amplitudes in a transform, but no information about phase.

One can picture the effect of phase in the following way: suppose you have two frequencies that beat. If you sample the part of the beat cycle in which they are in-phase or out-of-phase, the amplitude will be constant. If you sample that part of the beat cycle in which they are going from in-phase to out-of-phase, detection of more than one component will be that much easier. Therefore, when one has a window pattern in which only parts of the beat cycles are sampled, the relative phases do become important.

Figure 6.11 shows a simulated lightcurve containing a single beat cycle of the

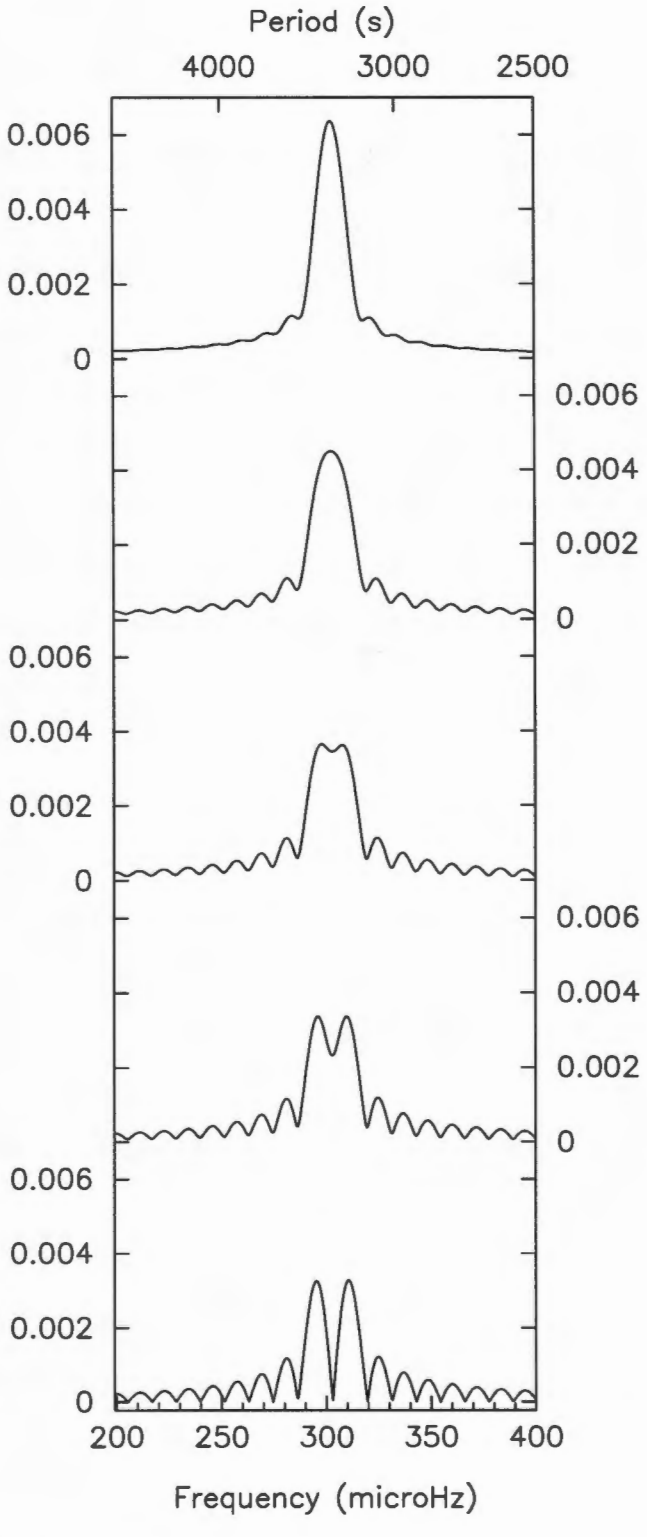

Figure 6.12: Transforms of half a beat cycle of two closely spaced signals for phase spacings $\phi_2 - \phi_1$ equal to zero, $\frac{\pi}{6}$, $\frac{\pi}{4}$, $\frac{\pi}{3}$, and $\frac{\pi}{2}$.

function

$$f(x) = A_1 \sin(2\pi\nu_1 t + \phi_1) + A_2 \sin(2\pi\nu_2 t + \phi_2) \; . \tag{6.6}$$

I have used arbitrary values of $A_1 = A_2 = 0.005$, and chosen a frequency spacing $\nu_2 - \nu_1$ of $\frac{1}{2 \text{ days}}$, or 5.787 μHz.

Figure 6.12 shows the transforms of *half* the beat cycle for different phase spacings $\phi_2 - \phi_1$. Changing the phase spacings is equivalent to sampling different parts of the beat cycle, for example, $\phi_2 - \phi_1 = 0$ samples the first two panels of Figure 6.11, and $\phi_2 - \phi_1 = \frac{\pi}{2}$ samples the middle two panels. Figure 6.12 shows clearly that the "resolvability" of two closely spaced signals depends on where in the beat cycle they are sampled. (It is interesting to note that prewhitening the first transform ($\phi_2 - \phi_1 = 0$) gives two signals whose separation is smaller than the true value, and prewhitening the last transform ($\phi_2 - \phi_1 = \frac{\pi}{2}$) gives a separation more than two times too large).

The same effect occurs if I sample half the beat cycle in small sections (simulating a window of successive observing nights). The "resolvability" of a lightcurve containing three closely spaced signals is much more complicated, but again the phase spacings play a crucial role. The beat period for a pair of signals separated by $\sim 1\mu$Hz is $\sim$ 10 days, and therefore the GW Librae data, with a duty cycle of 38.2% over 14 days, samples only $\sim$ half this beat period. It would seem that the components of the "quintuplet" have phase spacings that allowed them to be resolved, while the phase spacings of the "triplet" did not.

Figure 6.13 shows the transform of a lightcurve containing the best-fit signals for the three highest-amplitude components of the "quintuplet". I found that by manipulating the relative phases (keeping the frequencies and amplitudes fixed), I could produce a transform whose appearance was very similar to that of the transform of the real data (compare the second panel with Figure 6.7). Notice how there appear to be five components, even though there are only three signals in the data. I cannot therefore be sure that the fourth and fifth components I found for the quintuplet really exist - they may be artefacts. If there are more than three components to this

structure, they are either merged with these artefacts or lost in the noise.

I prewhitened this synthetic lightcurve and analysed it using the simultaneous nonlinear fitting technique, and the best-fit values agreed with the input signals to better than 1 in 100 000. Therefore, we can be confident that the analysis techniques are giving believable results for closely-spaced resolved signals. However, I only got matching results when I simultaneously fitted *all* the signals. If one was left out, the others would not give accurate solutions. Is it possible to obtain best-fit solutions for a resolved triplet if there is an unresolved structure nearby?

I analysed a synthetic lightcurve containing a resolved triplet with a spacing of 1.065 μHz and an unresolved triplet with a spacing of 0.64 μHz, with a transform closely resembling the real one. The input signals and the best-fit results for the resolved triplet agreed to within 0.01 μHz. One would never achieve such good results in reality, however; my synthetic data are noise-free.

Having critically examined the signals in the 376-s region, my final conclusion is that this region contains two principal structures: the first being three components of a possible quintuplet with an average spacing of 1.065 μHz, located at 377.94s (but with an aliasing ambiguity; it could be located at 376.30s). The second structure, at 375.53s, is unresolved, but is consistent with a triplet with three components spaced $0.7\pm0.1\mu$Hz apart. There are no other signals in this region whose reality can be reliably established.

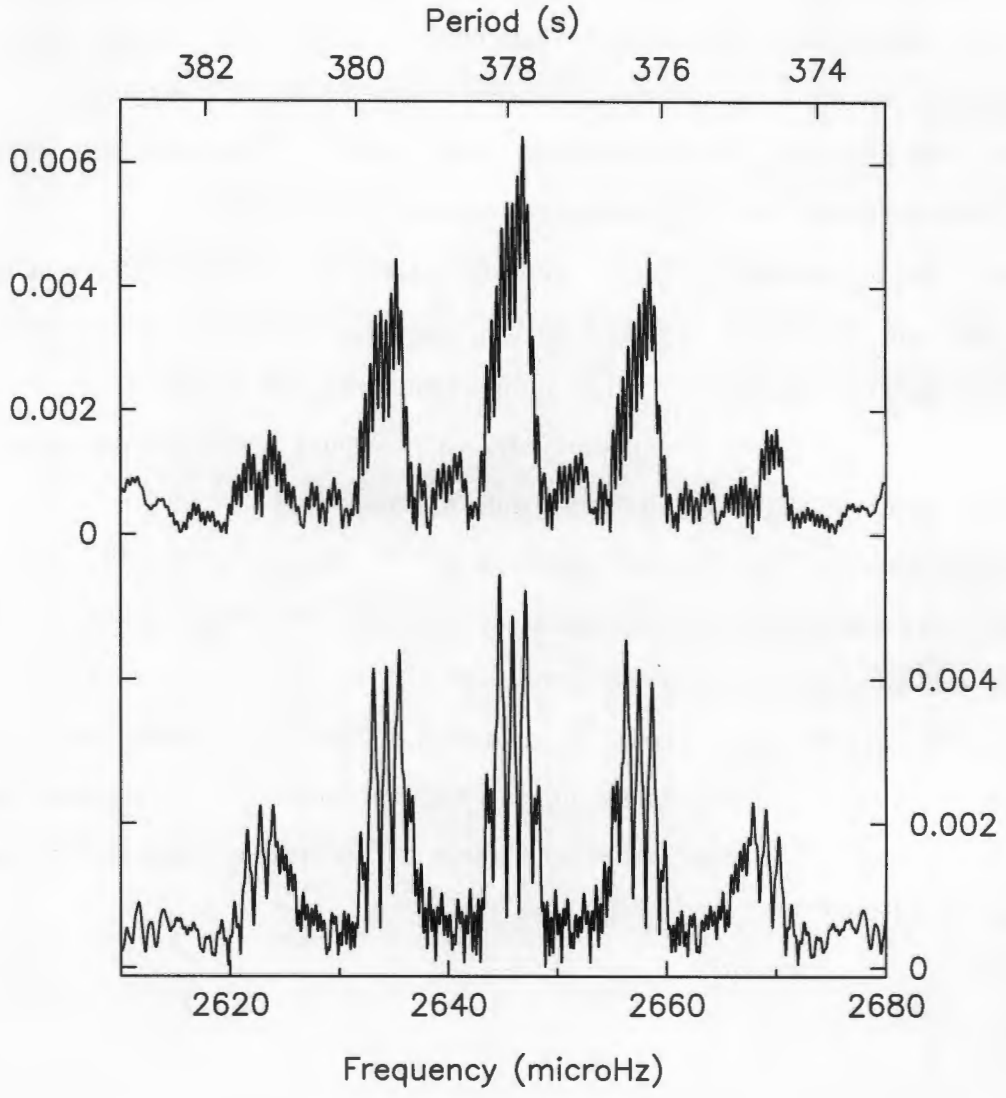

Figure 6.13: The Quintuplet. The first panel is the transform of a synthetic lightcurve containing the best-fit frequencies, amplitudes and phases for the three highest amplitude components of the structure I have labeled a quintuplet. The second panel is a transform of the same signals except that the relative phases have been altered so that it matches the real transform.

6.3 The Power Near 650 s

As we move upwards in period we seem to enter increasingly complex and chaotic territory. The 650-s region of the amplitude spectrum is a real "confusiogram" of frequencies without any discernible pattern (Figure 6.14). Prewhitening and simultaneous nonlinear fitting of the signals in this region gave (after considerable effort) the following values:

Table 6.5: The signals near 650 s.

Frequency μHz	*Amplitude* mMag	*Spacing* μHz
1536.07	3.04	5.99
1542.06	5.29	0.58
1542.64	7.71	2.06
1544.70	5.50	
1546.90	7.28	2.20
1549.42	7.84	2.52
1551.28	2.90	1.86

In Figures 6.15 and 6.16, I have isolated each of these seven signals. Many of them show the strong asymmetry that led me to believe that the signal at 236-s is a doublet. In this case, I am reluctant to believe that these signals are doublets, for three reasons. First, this region is obviously not very well-resolved. If frequencies are not properly fitted and removed, there will be ghosts in the residuals which will affect the symmetry of the appearance of a signal. The 236-s signal was in an isolated region with no interference from other signals.

Secondly, the existence of a forest of doublets in a small region of frequency space is physically highly improbable. Doublets would have to be part of $l = 1$ or 2 modes,

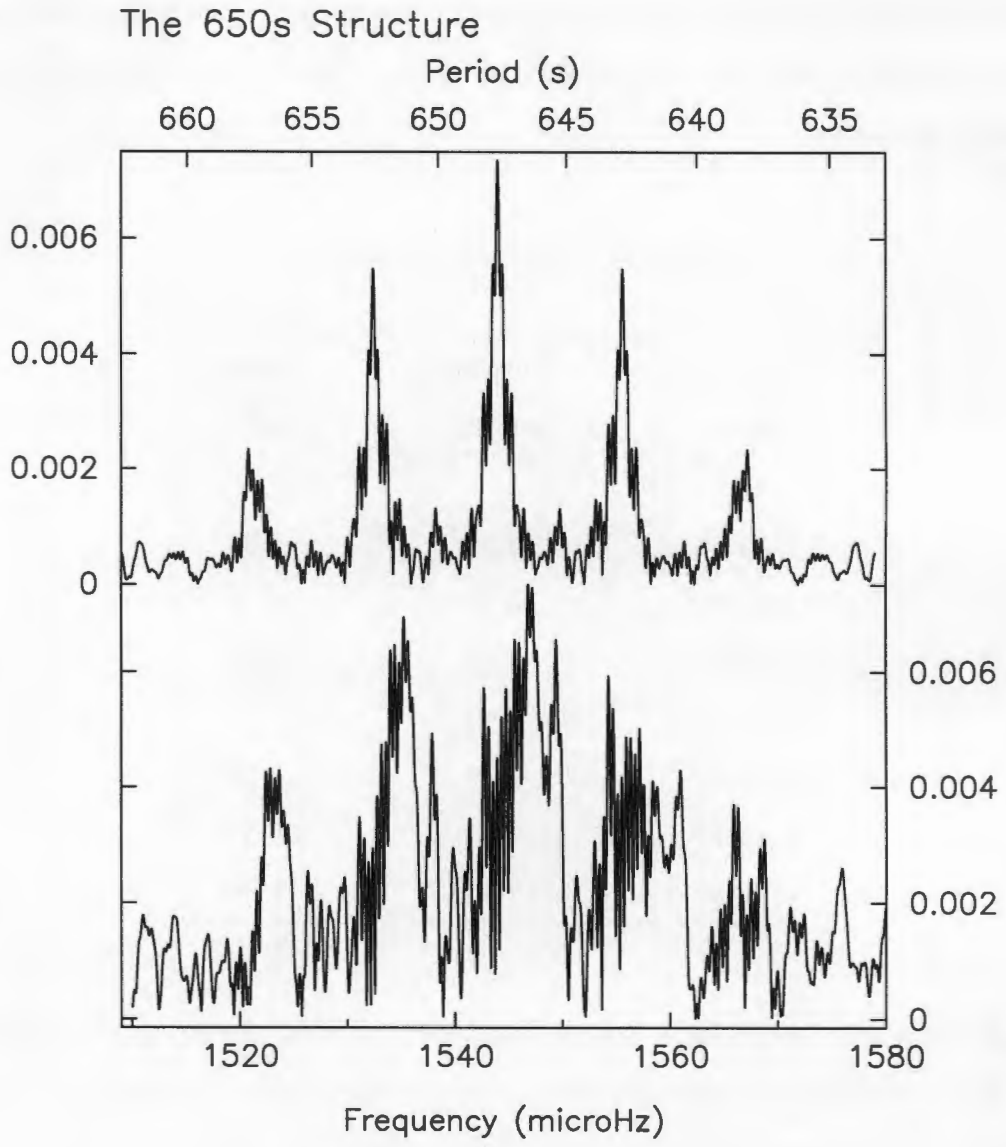

Figure 6.14: The power near 650 s. The first transform is the window. There is a linear combination mode at 1587 μHz.

which are roughly equally spaced in period. Mode-trapping could pull modes closer together, but there would be a problem explaining the coexistence of more than two modes in the same region of frequency space.

Thirdly, when analysing simulated data with identical windows to the real data, I have found that it is very common to get close doublets to appear even if the data contain very few frequencies. It seems to be easy to generate close doublets with the particularly awful window pattern of the GW Librae data.

Table 6.6: The 650-s signals in the full and half datasets.

Total		*First Half*		*Second Half*	
Frequency	*Amplitude*	*Frequency*	*Amplitude*	*Frequency*	*Amplitude*
μHz	mMag	μHz	mMag	μHz	mMag
1536.07	3.04				
1542.06	5.29	*1541.57*	4.17	1542.69	2.13
1542.64	7.71	1543.09	6.35	*1543.27*	5.60
1544.70	5.50	1544.81	6.81	*1545.34*	3.51
1546.90	7.28	1546.80	8.34	1546.80	6.80
1549.42	7.84	1549.38	7.77	1549.45	5.17
1551.28	2.90			*1552.18*	3.11

First of all, I need to establish the reality of these signals by performing the half-dataset test, described in section 6.2.1. The results are given in Table 6.6. Where a frequency is not its highest alias, I have put it in italics. The frequencies found in the first half-dataset are a better match to the full data than those found in the second half-dataset. Even so, there are only two frequencies that match well in all three datasets, but their amplitudes differ substantially.

Interestingly, while analysing the half-datasets, I found several frequencies which differed from 0.1 to 0.02 μHz from other frequencies. Such frequencies are statistically

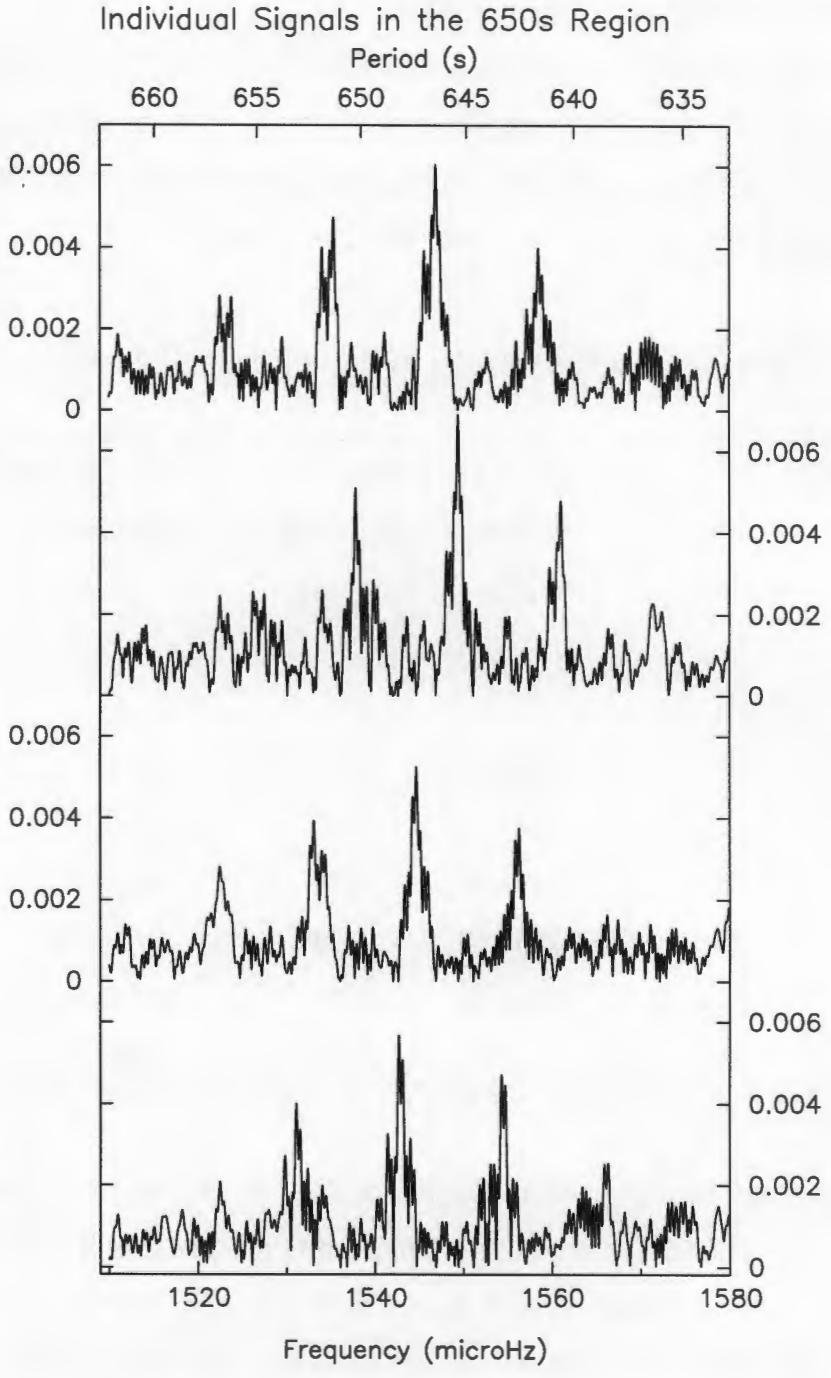

Figure 6.15: The four highest amplitude signals in the 650-s region.

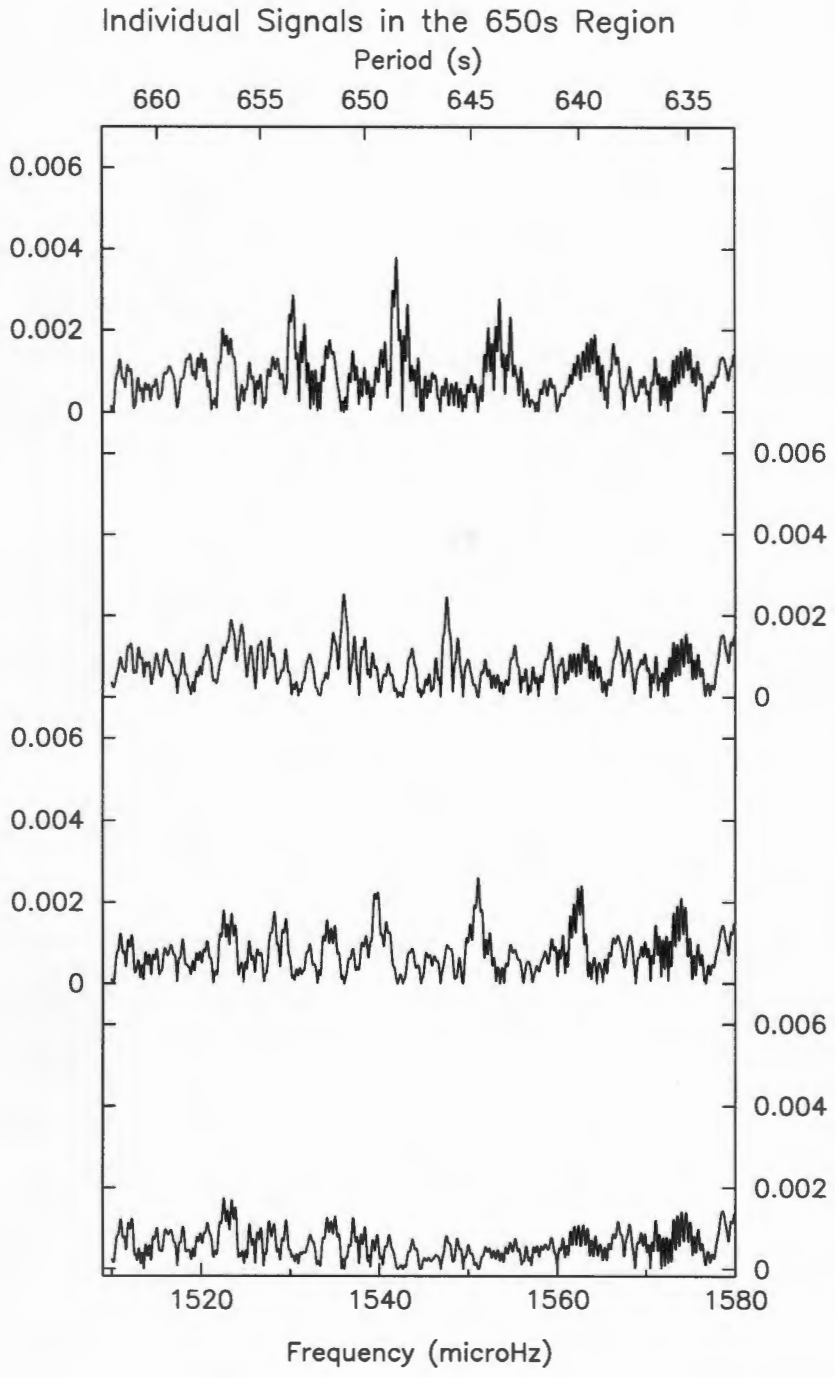

Figure 6.16: Three additional low amplitude signals, and the residuals after removing 7 signals.

identical, since the separations are very much smaller than the data resolution. I interpret this phenomenon as a symptom of the analysis technique's inability to deconvolve partially resolved, closely-spaced signals properly in this transform. By analysing noise-free synthetic data with the same window pattern, I have found that the analysis routines generate doublets or bristling thickets of artefacts with separations sometimes two orders of magnitude *smaller* than the resolution, even when the data consist of only a few sinusoids.

Due to the differences in the results for the half-datasets, I would be cautious about believing the values for the spacings given in Table 6.5. There is no obvious pattern in this region that could give a clue as to the physical nature of the pulsations generating this power.

Perhaps the 650-s region is a mode-trapping region. The strength of a trap is proportional to the power in a mode (Kleinman 1995, Kleinman 1998a). If a mode near 650-s is growing (or diminishing) in amplitude, its frequency would change as it is pulled into (or pulls away from) the trap. A single mode (with its m-split components) changing in frequency and amplitude could produce the tangled forest we are seeing in the 650-s region: sidebands produced by frequency and amplitude modulation could create a confusion of false signals.

Successful Fourier Analysis depends on signals in a lightcurve having fixed frequencies, amplitudes and phases. Variations in these parameters produce forests of sidebands. There are better techniques than Fourier analysis for analysing modulated signals or nonlinear effects (i.e. nonsinusoidal signals), such as wavelet analysis, but in the end there is no substitute for better data. I do not believe than any further progress can be made understanding the 650-s region until data with much better coverage of the duty cycle are obtained. Even then, if the eigenmodes change their power or frequency on a time scale shorter than that needed to resolve their m components (which seems to be at least a month), it is going to be extremely difficult to identify them and analyse their behaviour.

The other DAVs have shown that the transforms of their lightcurves can radically change appearance on time scales of weeks. The 1997 data show this to be the

case in GW Librae too. Its faintness will make it much harder to observe with the coverage needed; it is going to present a very big challenge to observational asteroseismology. As the only accreting DAV known, it presents a big challenge to theoretical asteroseismology too.

6.4 Linear Combinations and Harmonics

A striking feature of the amplitude spectra of the DAVs, especially the cooler ones, is the profusion of harmonics and linear combination modes present. These modes are often more numerous than the true modes, and it can sometimes be difficult to identify the parent modes.

Harmonics and linear combinations are a result of nonlinearities. The Fourier transform interprets a non-sinusoidal pulse shape as a series of sinusoidal harmonics. Non-sinusoidal pulse shapes may be a result of a non-linear response of the stellar medium to the perturbation travelling through it (Brickhill 1992), or of the outgoing Eddington flux to the surface temperature variations (Brassard *et al.* 1995).

Resonant mode coupling, often invoked to explain the exchange of energy between eigenmodes, could produce linear combination modes, but it fails to account for the very rapid changes in amplitude seen in the modes of G29-38 (Viulle 1998a,b) - some unknown nonlinear effect is at work in these stars. Combination modes probably carry a lot of information about the mechanisms and energetics of the processes that drive nonlinear oscillations, but until we learn what causes combination modes, this information will be hidden from us (Kleinman 1995). For example, we see in G29-38 that some modes always produce combinations when they are excited, while others never do. An $m = +1$ mode of a triplet may combine with many other modes, while the $m = -1$ mode never does. Why this should be, we don't yet know.

At present, the only use we can find for the linear combination modes and harmonics is as a diagnostic tool to examine signals whose reality is in doubt. For example, if we suspect that a mode is a doublet with a very low amplitude component, then if this low amplitude component is also seen in the mode's first harmonic,

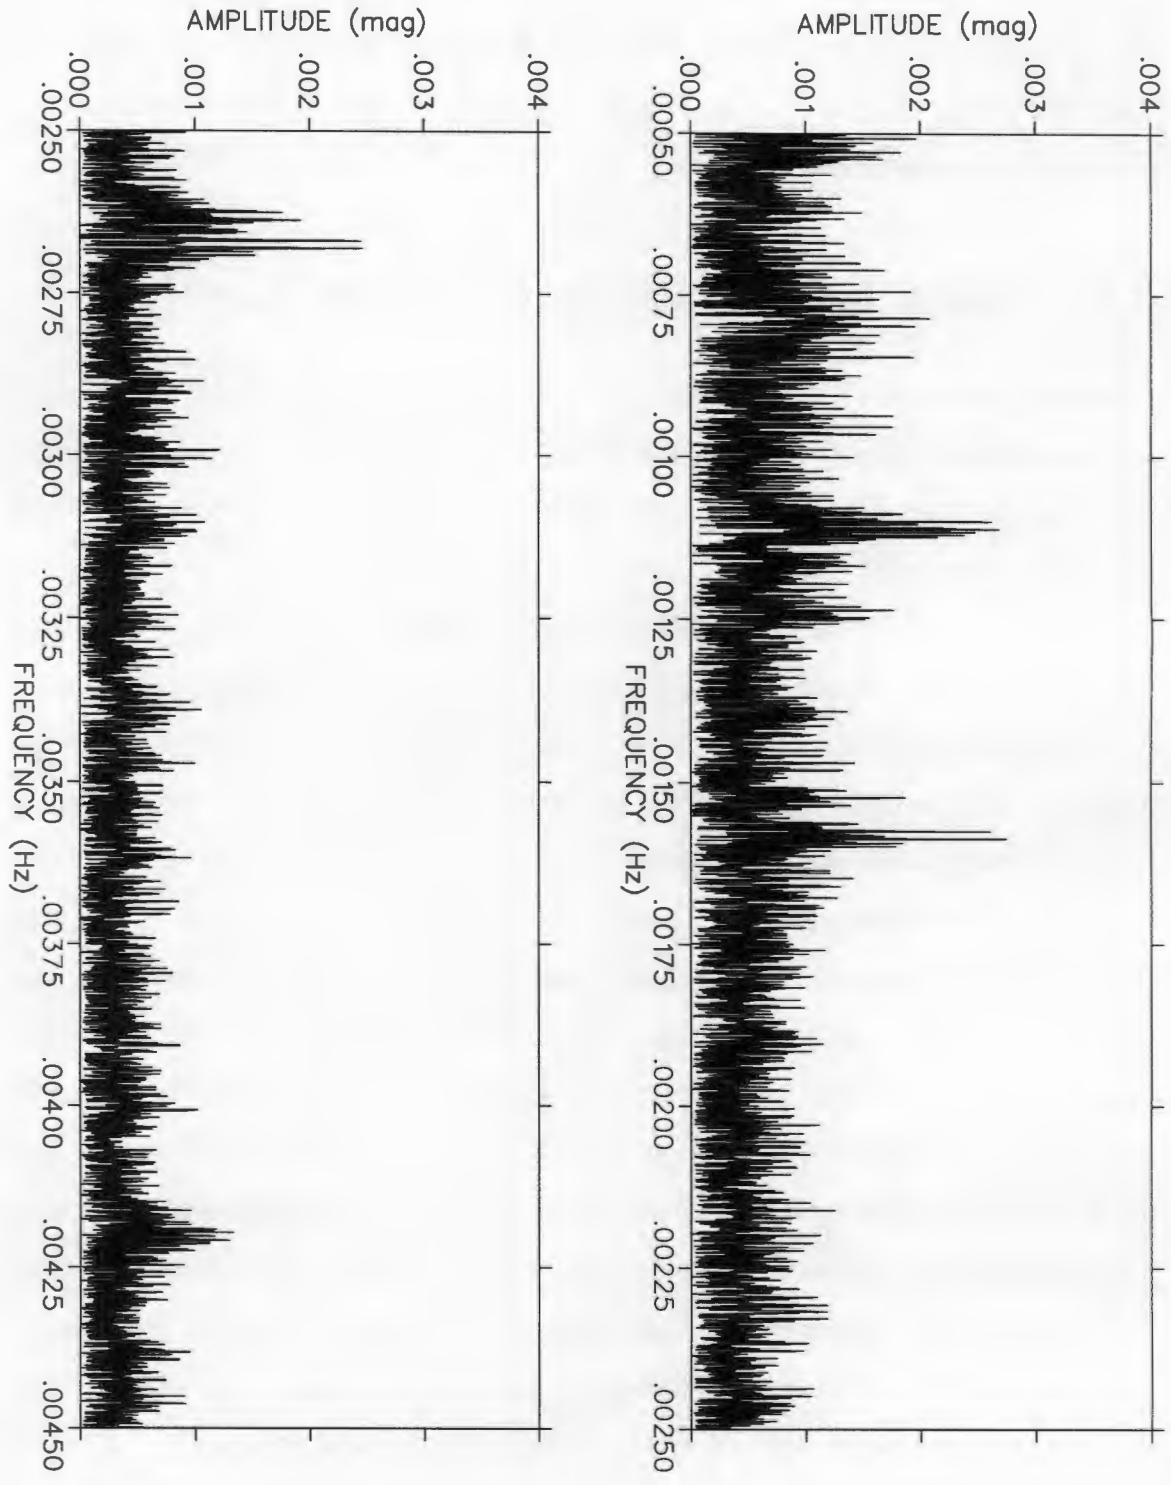

Figure 6.17: The residuals of the 1998 amplitude spectrum, showing four linear combination modes near .0011Hz, .0015Hz, .0027Hz and .0042Hz.

it is all that more likely to be real. This can only be done, of course, if the signal-to-noise of the harmonics and linear combination modes is high enough: they often have low amplitudes.

Figure 6.17 shows the residuals of GW Librae after the removal of the signals discussed in the preceding sections. There are no harmonics evident, but there is power at four linear combination frequencies. These modes have a low amplitude, and the signal-to-noise is poor. I analysed these regions of power with the simultaneous nonlinear fitting routine described above.

When comparing the signals found in these regions with the eigenmodes, I initially considered only combinations containing the four highest amplitude components of the 650-s mode(s) (labeled a_1 to a_4), the three resolved components of the 376-s quintuplet (b_1 to b_3), the centre of the unresolved 376-s triplet (b_4), and the highest amplitude component of the 236-s doublet. The other signals, I assumed, would be too low in amplitude or too uncertain for there to be believable matches containing them in the low signal-to-noise regions of the combination modes. The matches (and possible matches) I found are listed in Table 6.7. Where a signal is not its highest alias, it is in italics.

Although I expected not to find combinations containing the low amplitude signals of the 650-s region, I did in fact find signals in the $c - a_n$ region that matched them very closely. In the $c - b_n$ region I found three signals which could be combinations with the "quintuplet" components, with separations of 0.77 and 1.52 μHz, and three signals which could be combinations with the unresolved "triplet" components, with spacings of 1.08 and 1.02 μHz.

The components of the combination modes are likely to have different relative phases to their parent modes (Kleinman 1998a, Vuille 1998b), so they could in principle provide us with a "different view" of the parent modes: components of a parent mode may have relative phases which cause them to be unresolved, but the relative phases of the components in combination with another mode may allow them to be resolved; i.e., a parent mode may be unresolved while its combination is resolved.

If the above numbers are to be believed, then perhaps the low amplitude signals in

Table 6.7: Linear combinations.

Combination	*Calculated* (μHz)	*Found* (μHz)	$\Delta\nu$ (μHz)
$a_2 + b_3$	4196.38	4196.47	0.09
$a_3 + b_2$	4190.60	4191.06	0.46
$a_4 + b_3$	4189.60	*4189.67*	0.07
$a_1 + b_4$	4209.83	4210.64	0.81
$a_2 + b_4$	4212.36	*4212.77*	0.41
$a_1 + b_3$	4193.86	4193.60	0.26
$a_4 + b_1$	4187.48	*4185.46*	2.02
$a_2 + b_1$	4214.26	*4216.86*	2.60
$a_3 + b_4$	4207.63	*4208.83*	1.20
$b_3 - a_1$	1100.06	*1100.02*	0.04
$b_3 - a_4$	1104.32	*1104.26*	0.06
$b_1 - a_1$	1097.94	*1098.35*	0.41
$b_2 - a_1$	1099.01	*1098.35*	0.66
$b_1 - a_2$	1095.41	1095.00	0.41
$b_2 - a_3$	1101.21	*1101.74*	0.53
$b_4 - a_4$	1120.29	1120.77	0.48
$b_4 - a_2$	1113.51	1113.64	0.13
$b_4 - a_3$	1118.24	1117.62	0.62
$b_3 - a_2$	1097.54	1098.75	1.21
$c - a_3$	2692.80	*2692.54*	0.26
$c - a_1$	2690.60	*2691.03*	0.43
$c - a_7$	2686.37	2686.11	0.26
$c - a_4$	2694.85	*2693.24*	1.61
$c - a_2$	2688.07	*2688.21*	0.14
$c - b_4$	1574.56	*1575.56*	1.00
$c - b_4$	1574.56	1574.48	0.08
$c - b_2$	1591.59	1591.49	0.10
$c - b_4$	1574.56	1573.46	1.10
$c - b_3$	1590.53	*1589.97*	0.56
$c - b_1$	1592.66	1592.26	0.40

a_n = the signals near 650s
b_n = the signals near 376s
c = the dominant 236s signal

the 650-s region are real. The spacings between the components of the combinations of the two modes near 376-s could indicate that the modes have the same l. If true, these modes would not resemble anything known to asteroseismology.

6.5 Devil's Advocate

Under what circumstances could signals in a transform pass the half-dataset test, combine with each other to produce combination modes, and yet not be true pulsations in the star?

Up to now, I have largely been assuming that I have been dealing with stationary signals; that is, signals that maintain a constant frequency, amplitude and phase throughout the observing campaign. This assumption may well turn out to be too simplistic: the DAVs (indeed, the DBVs too) are well-known for the non-stationary behaviour of their pulsations (Kleinman 1995, 1998a; O'Donoghue 1998; Vuille 1998a; Nitta *et al.* 1998b). Modes can display dramatic changes in amplitude over only a few days, or show baffling phase drifts, and most peculiar of all, show steady changes in frequency of up to $\sim 5\mu$Hz over several days.

What are the effects of simple modulations in amplitude and frequency on the transform of a dataset that has a fractional coverage, or duty cycle, significantly less than 100%? To investigate this question, I constructed a synthetic lightcurve that contained a single sinusoid modulated in amplitude and phase, using the following function:

$$\Delta mag(t) = A(t)\ sin\ (2\pi\, t\, \nu(t) + \phi) \tag{6.7}$$

where $A(t)$ and $\nu(t)$ are simple first-order modulations to the amplitude and frequency,

$$A(t) = A_\circ + \frac{t}{T}\,\delta A \tag{6.8}$$

and

$$\nu(t) = \nu_\circ + \frac{t}{T}\,\delta\nu\ . \tag{6.9}$$

I chose arbitrary values for $A_\circ$ and $\nu_\circ$, and chose values for δA and $\delta\nu$ that are typical for the amplitude and frequency shifts seen in the DAVs:

$$\begin{aligned} A_\circ &= 5 \text{ mmag} \\ \delta A &= 2 \text{ mmag} \\ \nu_\circ &= 300\ \mu\text{Hz} \\ \delta\nu &= 4\ \mu\text{Hz} \\ \phi &= 0 \\ T &= 20 \text{ days} \end{aligned}$$

ie, over a runlength of 20 days, the amplitude increases by 40% and the frequency increases by 4 μHz.

I removed regularly-spaced chunks from the resulting lightcurve, varying their size in order to vary the fractional coverage or duty cycle, and then calculated the Fourier transforms of these synthetic datasets. Figure 6.18 shows some of the transforms. The results are surprising, to put it mildly. Even a small modulation in data with a poor duty cycle can generate a dreadful mess in the amplitude spectrum.

The duty cycle of the principal two weeks of GW Librae observations is 38.2%. In Figure 6.19 I show a transform of a triplet of signals separated by 1.2 μHz with arbitrarily chosen amplitudes and phases, sampled with the spectral window of the real data. I modulated the frequencies and amplitudes of the three signals in the same way as described by equations (6.7) to (6.9), until I obtained transforms that resembled the real transform in the 650-s region. It does not take much modulation to produce synthetic lightcurves with transforms that look very similar to the real one.

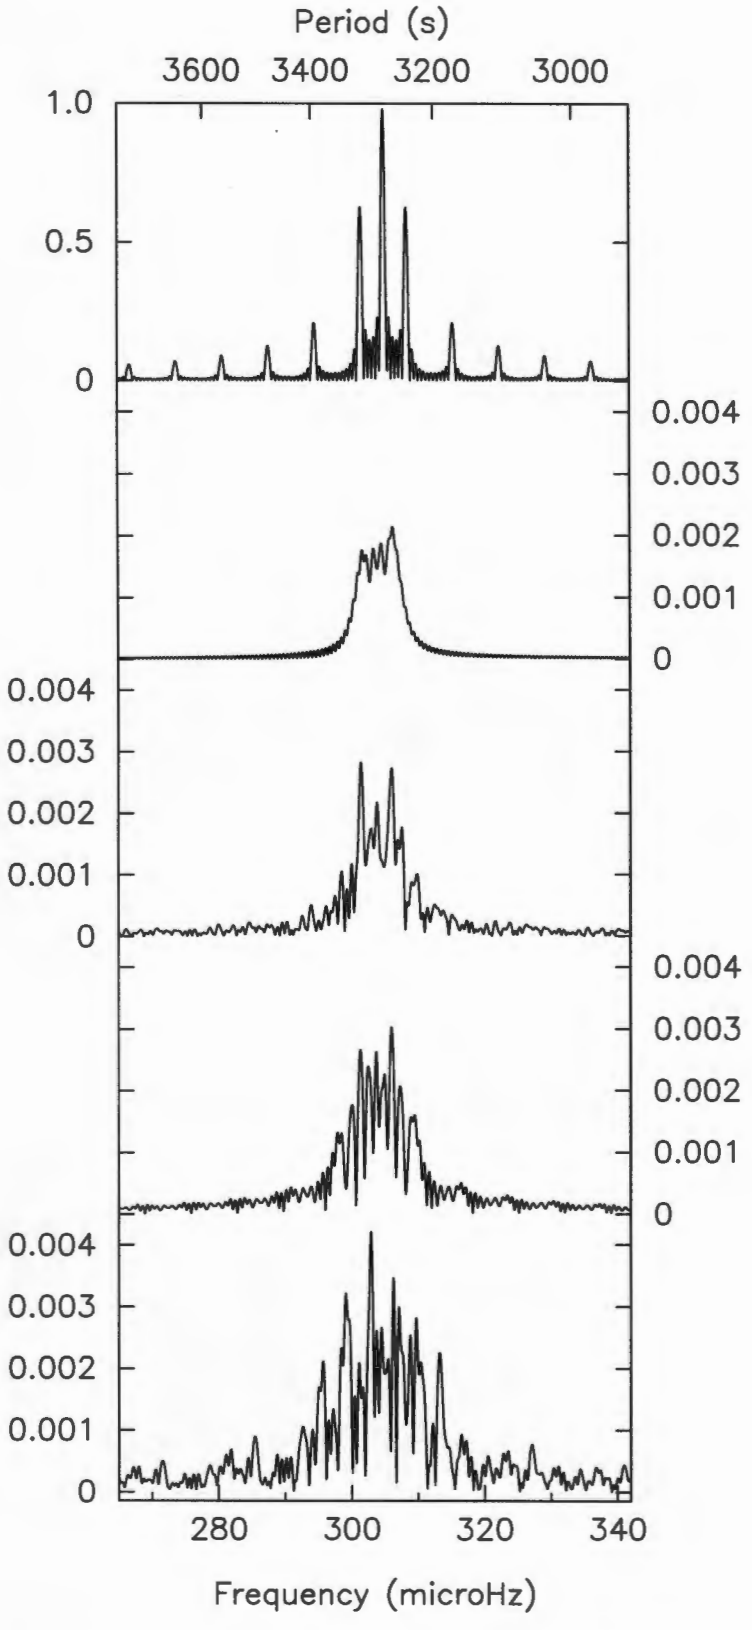

Figure 6.18: Transforms of synthetic lightcurves containing a single sinusoid modulated in amplitude and frequency, with fractional coverages (duty cycles) of 100% (second panel), 75%, 50%, and 25% (bottom panel). The top panel is the spectral window for the 50% duty cycle data.

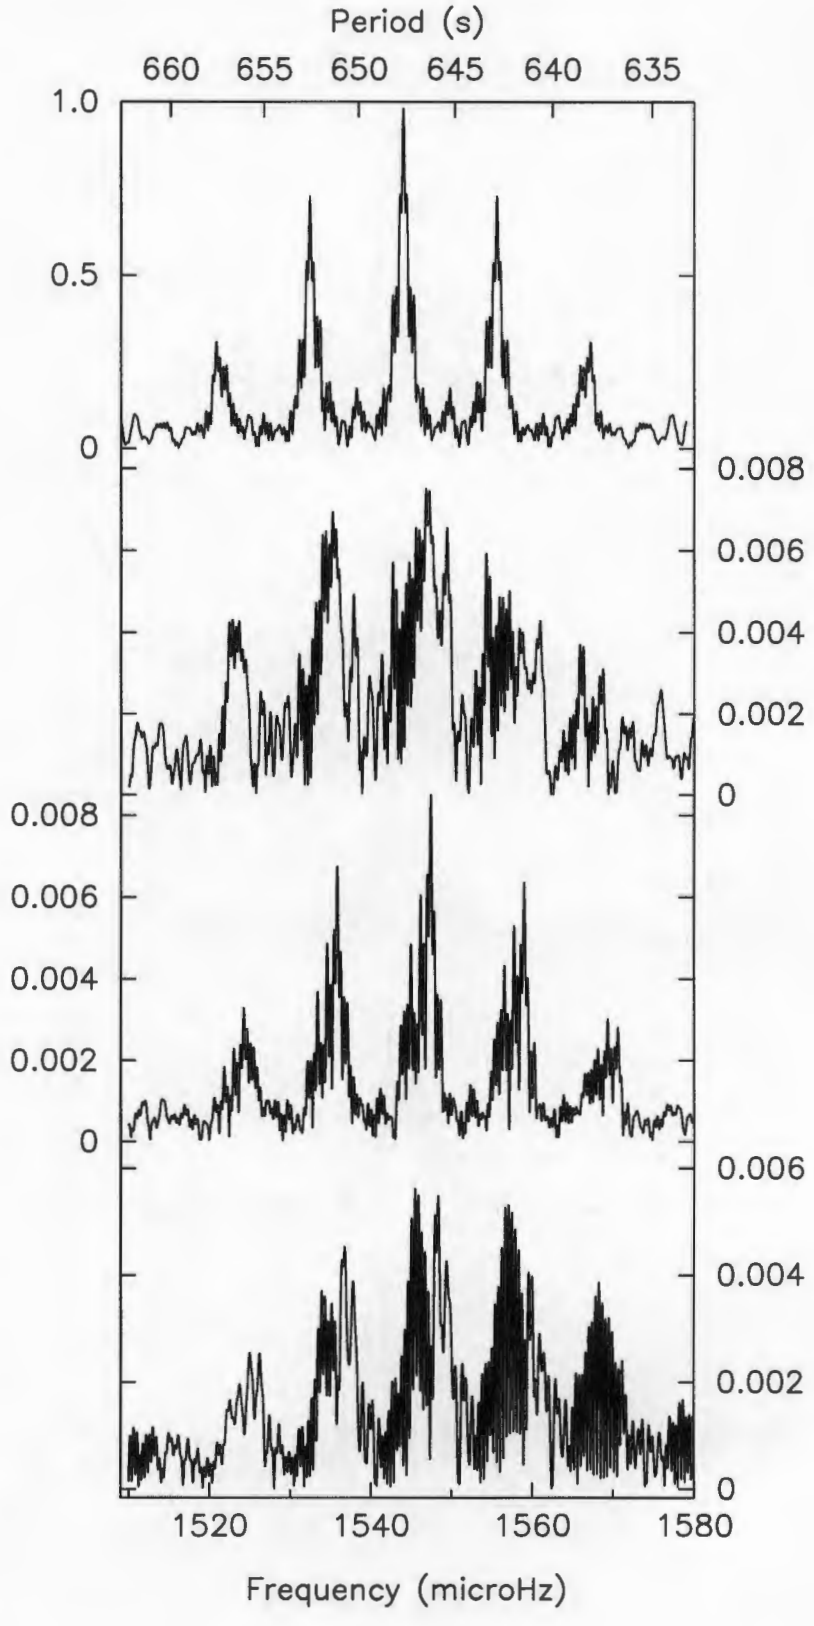

Figure 6.19: Transforms of modulated signals with the same spectral window as the real data. First panel is the spectral window; second panel is the transform of the real data in the 650-s region; third panel is a triplet of arbitrary pure sinusoids; fourth panel is the same triplet modulated in frequency and amplitude.

Chapter 7

A Bigger Picture

In this chapter I discuss the results of previous observations of GW Librae, and I attempt to find a "bigger picture" that could aid in the interpretation of the pulsation spectrum.

7.1 The 1997 Amplitude Spectra

The May/June campaign in 1998 produced by far the best observations of GW Librae to date, but I also observed GW Librae in March, April and September of 1997. The runs were too short to resolve the fine structure in the transform, but they serve to show the gross "bands" of power, and reveal the presence of structures not present in 1998. In Figures 7.1 to 7.4, I compare the transforms of all the observing runs to date.

The most noticeable difference between the amplitude spectra in 1997 and 1998 is the complete absence of the 236-s signal in 1997. It is also very obvious how the star is changing the power in its modes on time scales of less than a month. This is a classic signature of the cool DAVs (Kleinman 1995, 1998a,b; Kleinman *et al.* 1998a,b).

In Table 7.1 I list the results of prewhitening and simultaneous nonlinear fitting performed on each dataset. I have added the central or highest amplitude component of the 1998 modes for comparison. The significant figures quoted are probably far

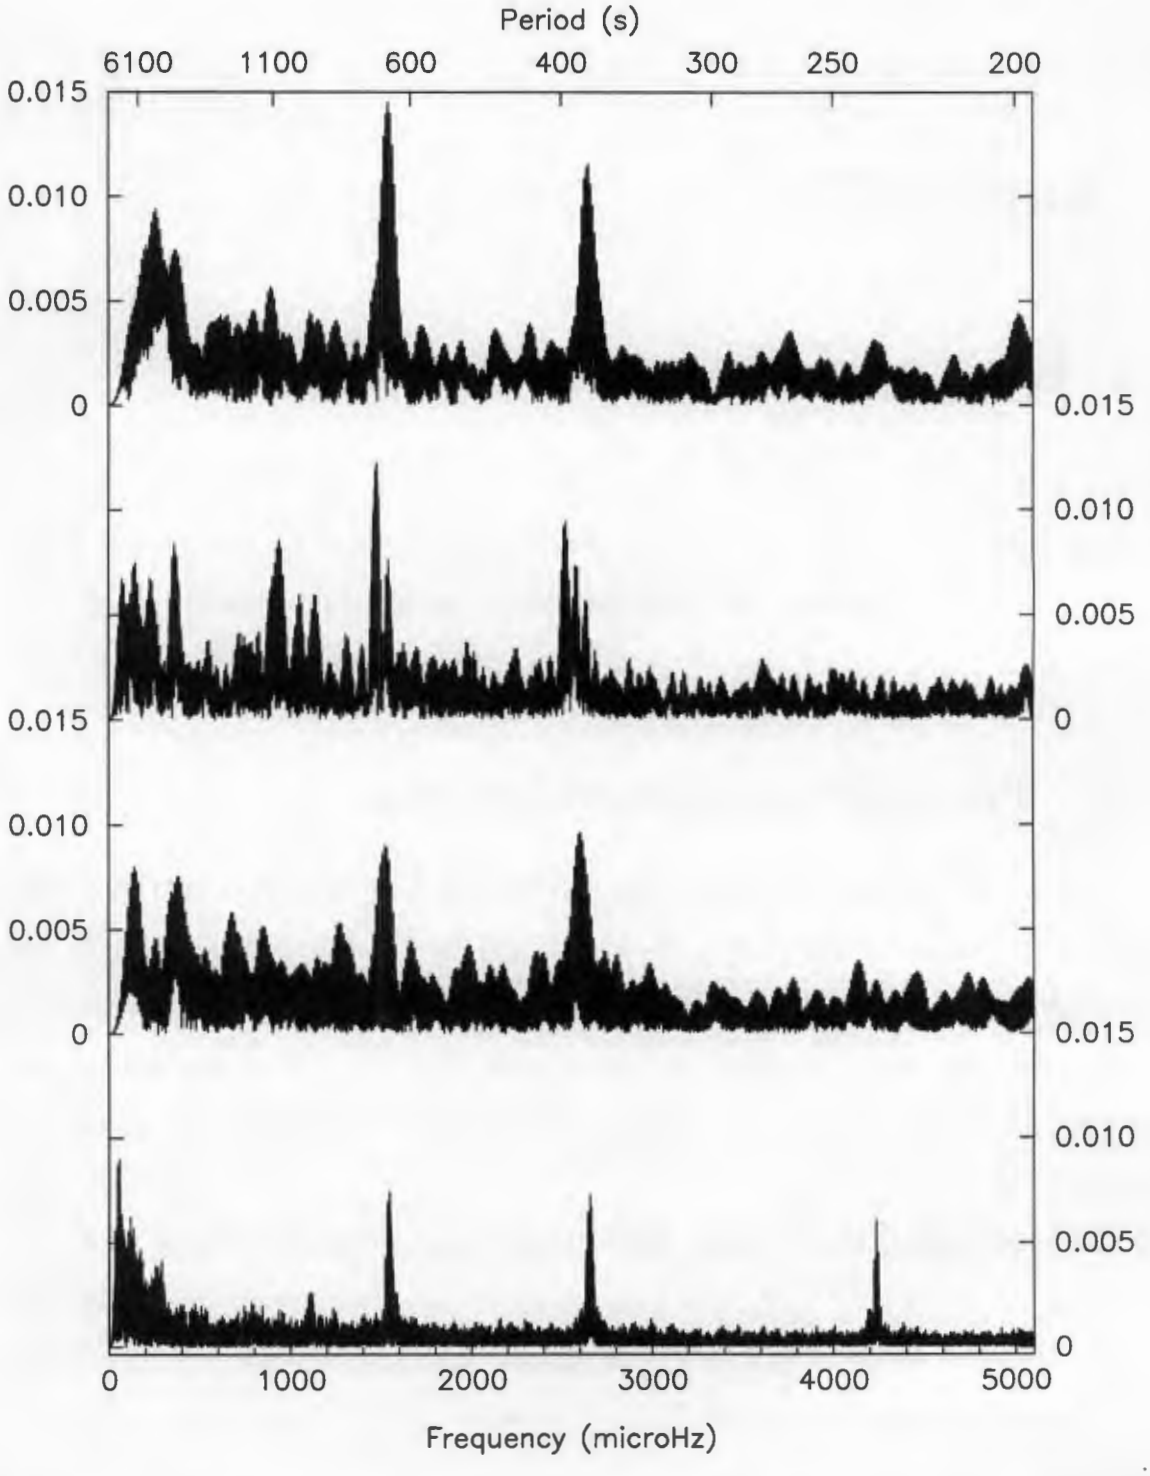

Figure 7.1: The amplitude spectrum of GW Librae in (*a*) March '97, (*b*) April '97, (*c*) September '97, and (*d*) May/June '98.

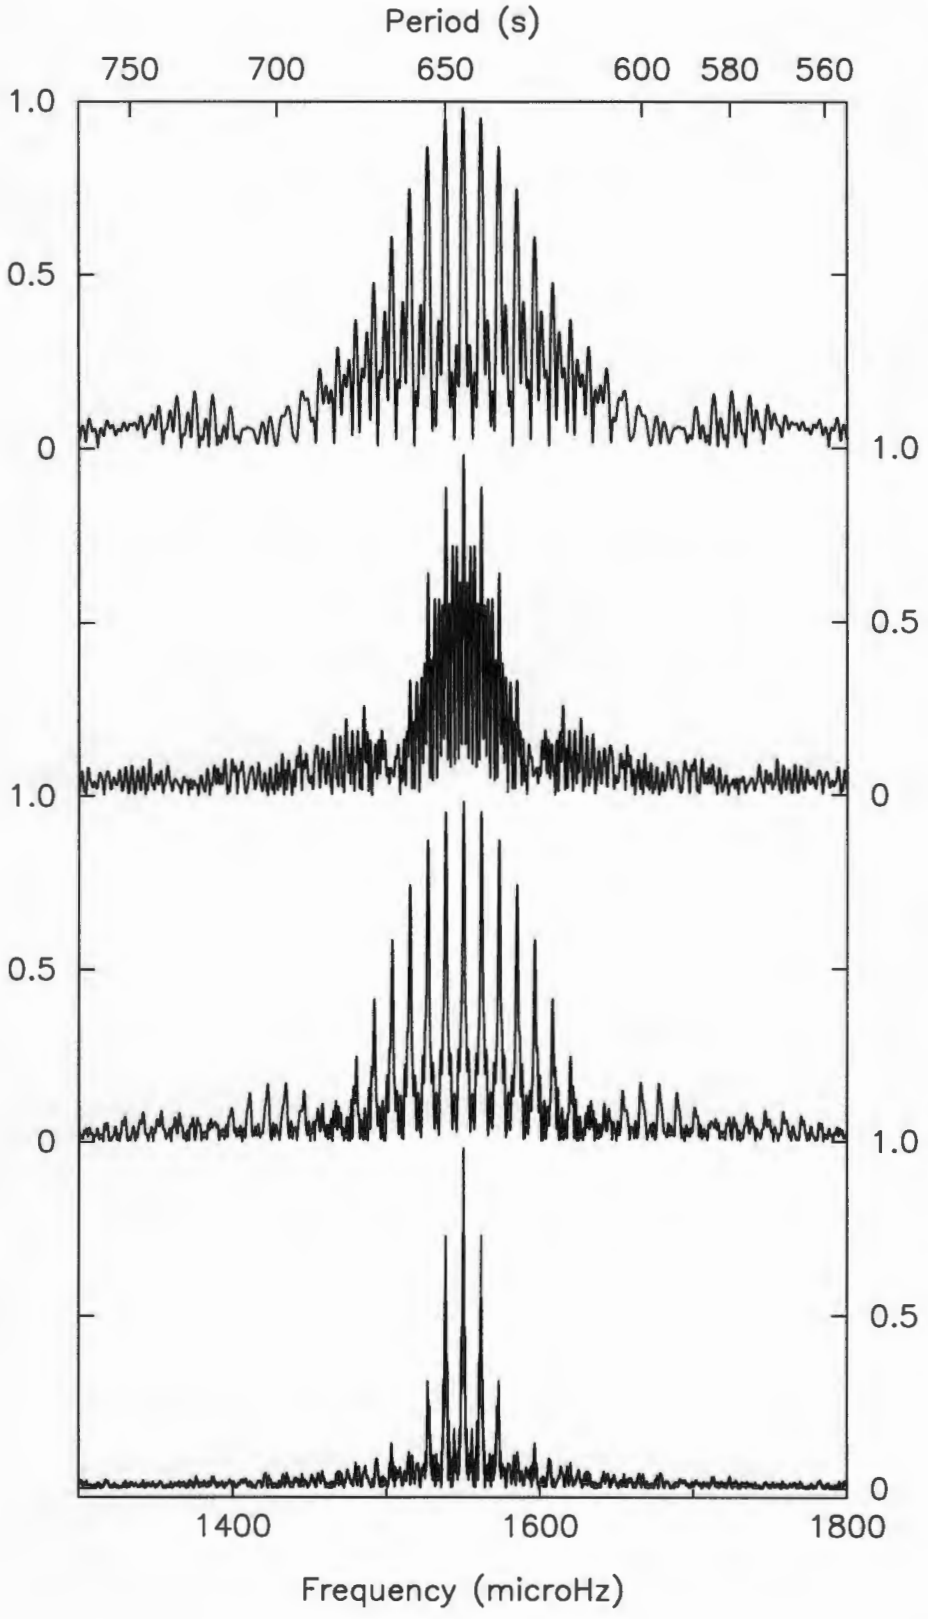

Figure 7.2: The spectral windows of the data in (*a*) March '97, (*b*) April '97, (*c*) September '97, and (*d*) May/June '98.

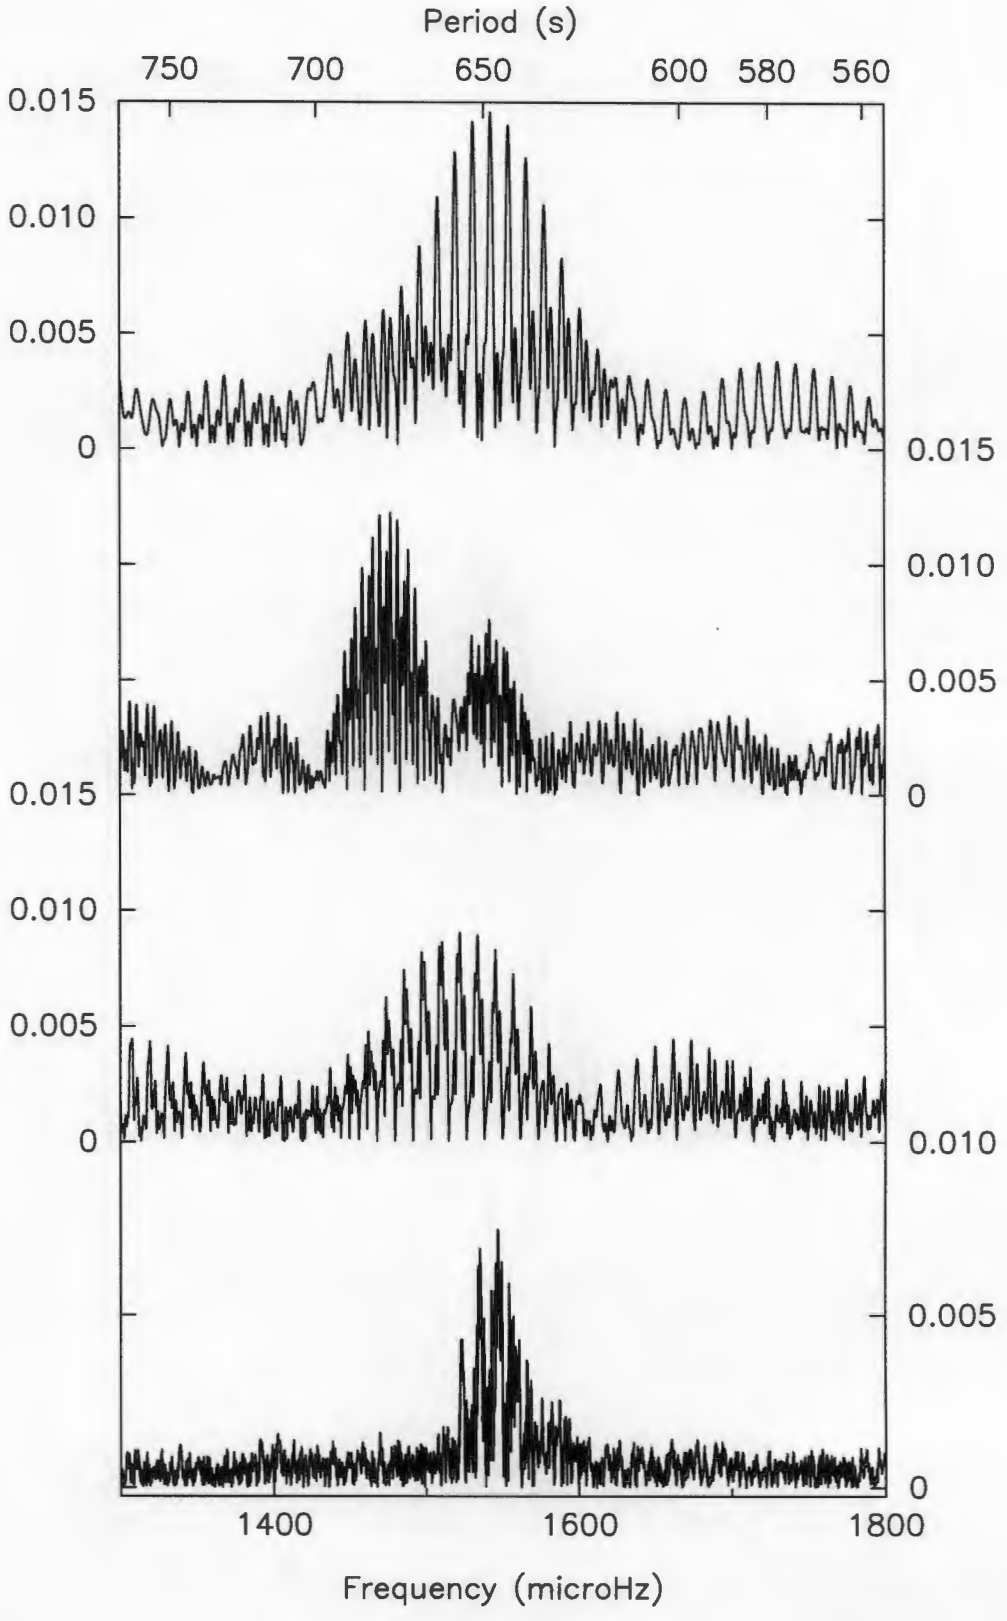

Figure 7.3: The 650-s region in (*a*) March '97, (*b*) April '97, (*c*) September '97, and (*d*) May/June '98.

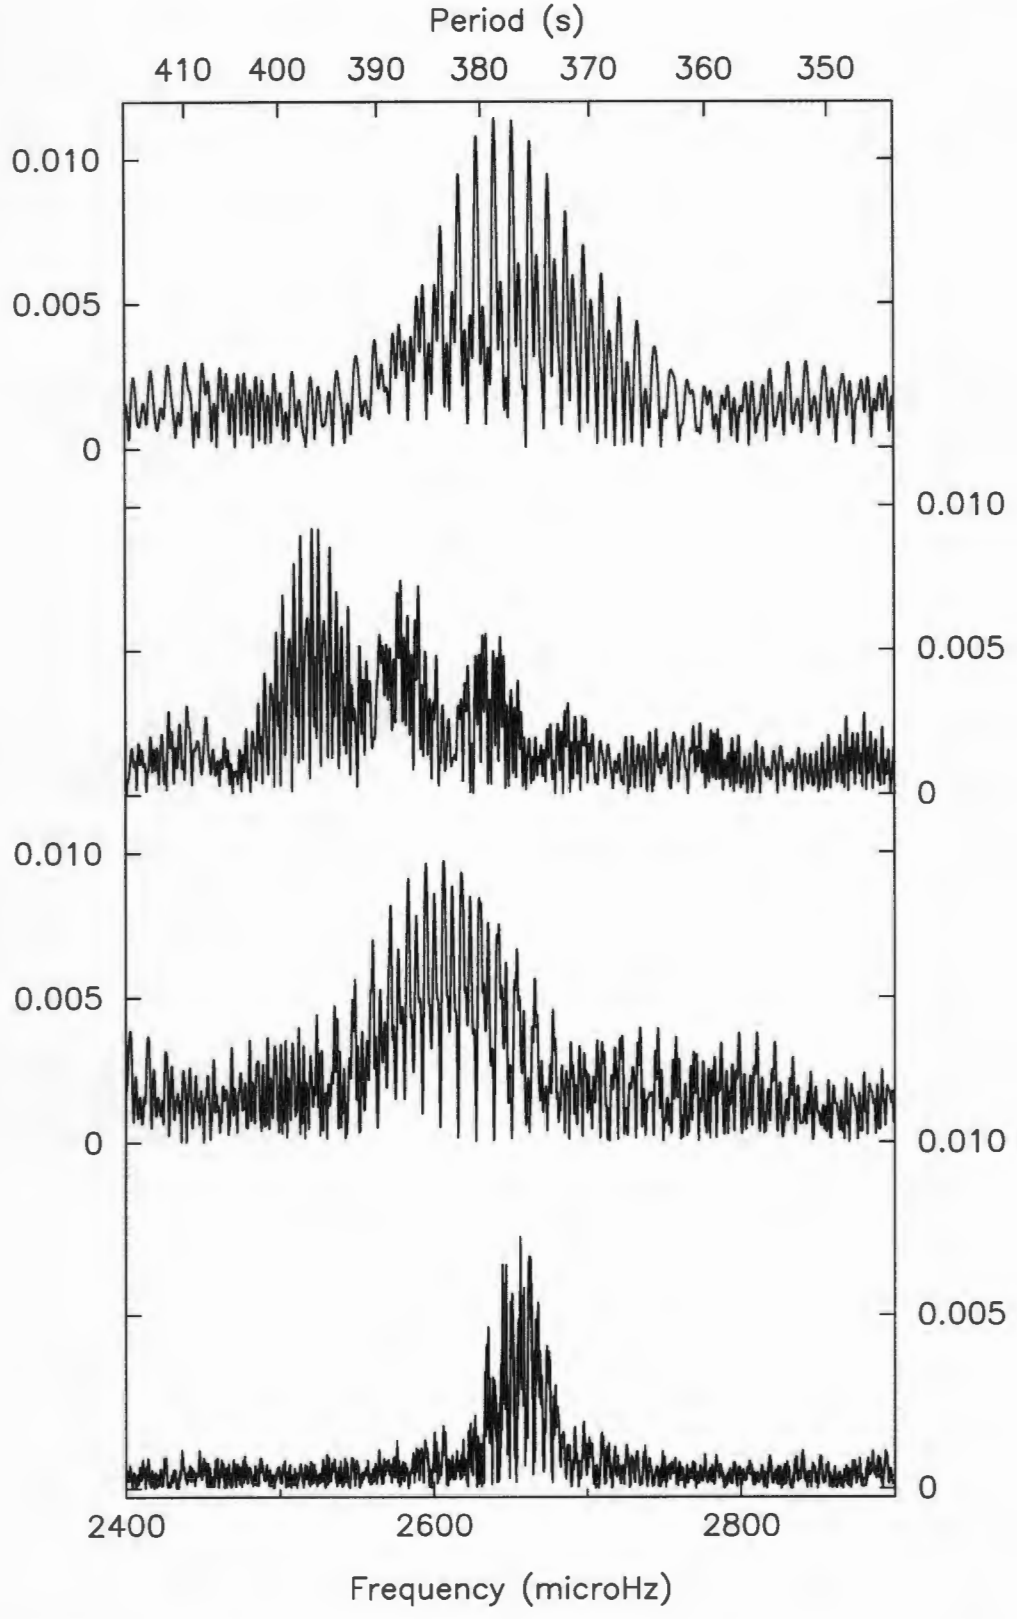

Figure 7.4: The 376-s region in (*a*) March '97, (*b*) April '97, (*c*) September '97, and (*d*) May/June '98.

in excess of the true accuracy of the numbers. The $\frac{1}{2T}$ rule-of-thumb for judging resolution only gives a guide to the errors on isolated, single peaks. The signals in these transforms are in all likelihood made up of several components, and therefore the frequency found for an unresolved blend of components gives only a rough guide to the true signals in that region.

Table 7.1: Signals in the 1997 data (with the 1998 power for comparison).

	Freq	*Amp*	*P*		*Freq*	*Amp*	*P*
	μHz	mMag	s		μHz	mMag	s
			March 1997				
⋆	1543.12	14.45	648.04		2640.53	11.39	378.71
			April 1997				
	1477.20	**9.97**	**676.96**		**2520.95**	**8.68**	**396.68**
	1474.58	9.66	678.16		**2578.78**	**6.12**	**387.78**
⋆	**1542.14**	**6.02**	**648.45**		2570.22	5.13	389.07
○	1529.78	4.83	653.69	•	**2638.99**	**4.47**	**378.93**
					2529.61	4.83	395.32
			September 1997				
○	1522.14	12.23	656.97		2606.74	9.34	383.62
	1511.96	7.18	661.39		2612.12	8.39	382.83
	1519.01	5.59	658.32				
			May/June 1998				
⋆	1544.70	5.50	647.38	•	2645.90	5.39	377.94
	4237.49	6.03	235.99		2662.94	5.51	375.53

Signals in boldface are the dominant signals in "multiplet components," and are spaced roughly 60 μHz apart.

⋆ Signals in the 650-s region that appear to correspond.

• Signals in the 376-s region that appear to correspond.

○ Possible correspondence?

Most interesting is the amplitude spectrum for April 1997. There is a "doublet" in the 650-s region and a "triplet" in the 376-s – 400-s region. The components are separated by $\sim 60\mu$Hz. The signals in boldface in table 7.1 are the highest amplitude signals in each component of the "triplet" or "doublet". The other signals in the April dataset are located "underneath" the boldface ones. So we appear to have multiplets in which each component is itself made up of several components. Unfortunately the resolution is too poor to allow any interpretation of the signals that make up the individual components.

Could the "multiplet components" actually be individual k modes? They have separations in period of $\sim$ 28.51s (650-s region), and $\sim$ 8.90s and $\sim$ 8.85s (376-s region), far smaller than the apparent k-spacings of $\sim$ 50s in the DAV ensemble (Clemens 1993, 1994; Kleinman 1995, 1998a,b; Kleinman *et al.* 1998b).

The signals at 650 s and 376 s in the March 1997 and the May/June 1998 data seem to correspond to the right hand, lowest amplitude April "multiplet components" (indicated with "$\star$" and "$\bullet$" in Table 7.1. The September 1997 signals, however, do not correspond convincingly to any structures seen in any other dataset (I have indicated a possible correspondence with "$\circ$" in Table 7.1).

I have failed to fit the April results to the September data: the signals in the September data appear to be independent, and I am unable to explain them in terms of unresolved mixtures of signals from previous runs. The simulated lightcurve technique I used in the previous chapter is extremely useful for "establishing consistency" between signals in different datasets of the same star (Kleinman 1998a): I modelled the April signals with the September window, using a wide range of different phase spacings between the signals, but obtained nothing that resembled the structures in the September transform. As mentioned already in Chapter 6, simulated lightcurves prove nothing, but are very useful for eliminating possibilities.

It is disappointing not to be able to match signals convincingly in the September data with signals in the other datasets. It would be useful to be able to find common signals in all the datasets. The dominant September signals lie $\sim$ midway between the April "multiplet components," which will present quite a challenge to attempts

to interpret the pulsation spectra in terms of normal mode oscillations.

7.2 Ideas and Speculations

Ignoring the hypothesis of the previous chapter, what interpretations can we give to the April structures? If we consider each "clump" of signals to be individual k modes, they are far too closely spaced in period. The possibility that we are seeing two modes of different l only really works for two modes.

Perhaps the "doublet" near 650 s is actually two modes of different l, or two successive k, $l = 1$ modes (separated $\sim$ 28.5s in period), and perhaps the outermost members of the "triplet" near 376–400 s are two successive k, $l = 2$ modes (separated $\sim$ 17.8s in period), with the middle one being an $l = 1$ mode. Maybe not entirely impossible, but certainly somewhat contrived. Such an interpretation would have to account for the misfit September signals, and explain why the lowest amplitude member of the April "triplet" appears to consist of at least 6 signals (according to the 1998 data).

What are the smallest possible period spacings for the DAVs? According to Bradley (1996), a DAV with a $10^{-10}M_{\odot}$ hydrogen envelope will have an average period spacing for $l = 1$ modes of 65s if its mass is $0.4M_{\odot}$, and 43s if its mass is $0.95M_{\odot}$. If it has a hydrogen envelope of $1.5 \times 10^{-4}M_{\odot}$ (likely for CV primaries), then its $l = 1$ mode average period spacing is 58s for a mass of $0.4M_{\odot}$, and 34s for a mass of $0.95M_{\odot}$.

For stars more massive than $\sim 1M_{\odot}$, crystallization becomes important, with period spacings increasing as the degree of crystallization increases (Kleinman 1998a). The most massive DAV known, the crystallized star BPM37093, which has a mass of $1.1M_{\odot}$ and which appears to pulsate with $l = 2$ modes only (Nitta *et al.* 1998a), has an $l = 2$ mode spacing of $\sim$19s, which would give it an $l = 1$ mode spacing of $\sim$32s (Section 3.2). So if GW Librae has a mass of $\sim 1M_{\odot}$, the modal arrangement suggested above may not be impossible. However, it would be pushing the limits of theory; let us first see if it is possible to find GW Librae a place among the normal

DAVs before assigning it to the asylum.

The ensemble of observed modes in the cool DAVs show concise groupings every 50 s from 300 s to 650 s (Clemens 1993, 1994; Kleinman 1995, 1998a,b; Kleinman *et al.* 1998b). If we consider the triplet in the April transform to be an m-split, single k mode centred at $\sim$ 388 s (the period of the middle component), and the doublet likewise to be an m-split, single k mode centred at $\sim$ 650 s or $\sim$ 675 s (there is no way to know which is the central component), then we have a pattern of modes that is consistent with what is seen in the other cool DAVs.

Rotational splitting of $\sim 60\mu$Hz, or $\sim$ 2.3 hours (or $\sim$ 3.9 hours if the modes are $l = 2$), is *much* more comfortable than the rotational splitting of $\sim 1\mu$Hz, or $\sim$ 5.7 days, considered in Chapter 6. As discussed in Section 4.2.2, single white dwarfs have rotation periods on the order of a day, but CV primaries, accreting angular momentum for $\sim 10^9 - 10^{10}$ years, are likely to be spinning much more rapidly. We would expect GW Librae to have a rotation period anywhere from 28s (WZ Sge's rotation period) to a $\sim$few hours (if the rotation period was locked to the orbital period, i.e., if GW Librae is a polar).

There are two major problems with an interpretation of the April "doublet" and "triplet" as rotationally split multiplets:

Problem 1: The September data are not consistent with a picture of m-splittings of 60 μHz: the dominant signals fall between the April modes. Is it possible for eigenmodes to shift $\sim 30\mu$Hz in frequency? If yes, this would be $\sim$ 5 times greater than any frequency shift observed in any other star.

Problem 2: The "doublet" and "triplet" are themselves composed of several signals. If the 376 and 650-s structures in the 1998 data are single rotationally-split m modes, we need to find a convincing explanation for why they are each composed of at least six signals. Any explanation will also have to account for the simplicity of the 236-s signal in the 1998 data.

Under what circumstances would it be possible for an individual (k, l, m) mode to exhibit finer splitting?

One possibility is an oblique rotator: while not observed up to now in a white dwarf, it is possible for a star to have a pulsation axis which is not aligned with the rotation axis. This is observed, for example, in the rapidly oscillating Ap stars (Kurtz 1982). In a roAp star, the pulsations are symmetrical to the magnetic axis, which is not aligned with the rotation axis, with the result that as the star rotates, a pulsation mode is seen from varying aspect, inducing an amplitude and phase modulation. In these stars, a single pulsation mode, $(k, 1, 0)$ for example, is made up of many components, because the magnetic field distorts it, so that it cannot be described by a single normal mode (i.e. by a single spherical harmonic (Kurtz 1999).

While no pulsating white dwarf observed up to now has a magnetic field strong enough to to align the pulsation axis with the magnetic axis, there is no reason why such stars shouldn't exist (Dziembowski 1998). The polars and intermediate polars (IPs) show us that it is not uncommon for white dwarfs to have strong magnetic fields whose axes are not aligned with the rotation axes. Perhaps GW Librae is the first oblique rotator among the DAVs. At present there is no way to tell if it is or not[1].

How else can we get the sort of variable pulsation geometry that could produce fine structure splitting (or sidebands) of single normal modes?

If GW Librae's rotation axis is not aligned to the binary orbital axis, is it possible to generate components due to rotational splitting, each one of which is split because of an orbital modulation (or vice versa)?

It would be possible to modulate the pulsations and generate sets of sidebands for each pulsation in a similar way as in the roAp stars if the white dwarf's rotation axis is precessing, but it is difficult to imagine how this could happen unless there is a third star in the system, very rare among CVs (see footnote).

The observed amplitude of a mode (and its chances of detection) depends on the pulsation geometry, that is, on the distribution of bright and dark patches over

[1]But if GW Librae does turn out to be the first DAV oblique rotator as well at the first DAV CV, I suspect it would be violating the Law of Conservation of New Phenomena Discovered in a Single Star.

the surface of the star, and to what extent they cancel out. Throw in an accretion disc that produces a variable obscuration of the stellar surface (either through the geometry of the binary system, or because of changes in the disc, like variable optical thickness), and perhaps you get the kind of modulation that produces the amplitude spectra we are seeing. A partially obscuring, semi-transparent accretion disc with a stochastically varying optical depth will produce modulations of the normal modes which will be very difficult to analyse.

IPs have multiperiodic amplitude spectra, because reprocessing of radiation from the accreting regions off surfaces like the secondary produces sidebands (Warner 1995). However, in order to see signals from reprocessing in the amplitude spectrum, there has to be some kind of beaming source (or at least a nonsymmetrically radiating source, otherwise one gets geometric cancellation), and the radiation has to be extremely hot (hard UV or X-rays). As a DAV, GW Librae is too cool to do this, and its spectrum shows no sign of the hot accreting regions seen in IPs and polars.

All this assumes that the normal modes themselves are stationary. If they are wandering around in frequency, amplitude and phase, perhaps an extreme form of the wanderings seen in the other DAVs, or due to stochastic modulation or driving by accretion, then it would be quite possible to generate the kind of amplitude spectra we're seeing, but they would most likely be almost impossible to analyse. If each mode has a stochastically modulated frequency, amplitude and phase, there would be too many variables to ever find a unique fit to the lightcurve.

7.3 An Orbital Modulation?

The lightcurves of most low $\dot{M}(2)$ CVs show characteristic "orbital humps" (unless they're seen face-on); a brightness modulation resulting from the Bright Spot on the rim of the accretion disc rotating in and out of view (Chapter 2). Finding GW Librae's orbital period would help us place bounds on the physical parameters of the

system, but its orbital modulation has proved elusive[2].

We would expect GW Librae to have an orbital period P_{orb} somewhere between $\sim$ 1.25 to $\sim$ 1.5h, typical of the WZ Sge stars (Table 4.1). However, while individual lightcurves show very definite modulations (see for example run GWLIB16, JD=2450953), it has been difficult to find consistency from run to run. Figure 7.5 shows the low frequency end of the transforms. There is no signal that stands out as being common to all datasets.

The $\frac{1}{2T}$ rule-of-thumb is only a rough guide to the "believability" of a frequency; there is in addition a strong dependence on where in the amplitude spectrum the frequency is located. For example, if a 90 min and 5 min signal are found in the transform of a 2h run, the 5 min period is more likely to be a real periodicity, because its power has been reinforced by many cycles, whereas the run contains only one cycle of the 90 min period. If only a few cycles of a period are present in a run, there is no way to tell if the supposed period is a result of stochastic processes or if it is a true periodicity of the star. For example, during the 2h run there may be two large flares spaced 90 min apart; a purely stochastic process that the Fourier transform will nevertheless try to fit, producing a large amplitude peak at 90 min.

CVs typically give very noisy amplitude spectra because of their flickering and flaring, caused by blobs of gas in the accretion stream smashing into the disc, and by the violence in the boundary layer where material in the inner disc is decelerated from Keplerian velocities to match the slower rotation of the white dwarf photosphere. For this reason, run lengths as long as possible are needed so that the Fourier transform can "iron out" the stochastic variations and only show the variations that are truly periodic. Flickering and flaring occurs on time scales of minutes to hours, which drives up the low frequency noise in CV amplitude spectra, making it hard to find true signals like the orbital modulations of low inclination systems.

In addition, amplitude spectra of stellar observations contain *red noise* due to sky transparency variations. (In red noise, the height of the grass has an inverse de-

[2]Subsequent to the submission of this dissertation, Szkody *et al.* (1999) found a spectroscopic orbital period for GW Librae of 78.8 minutes.

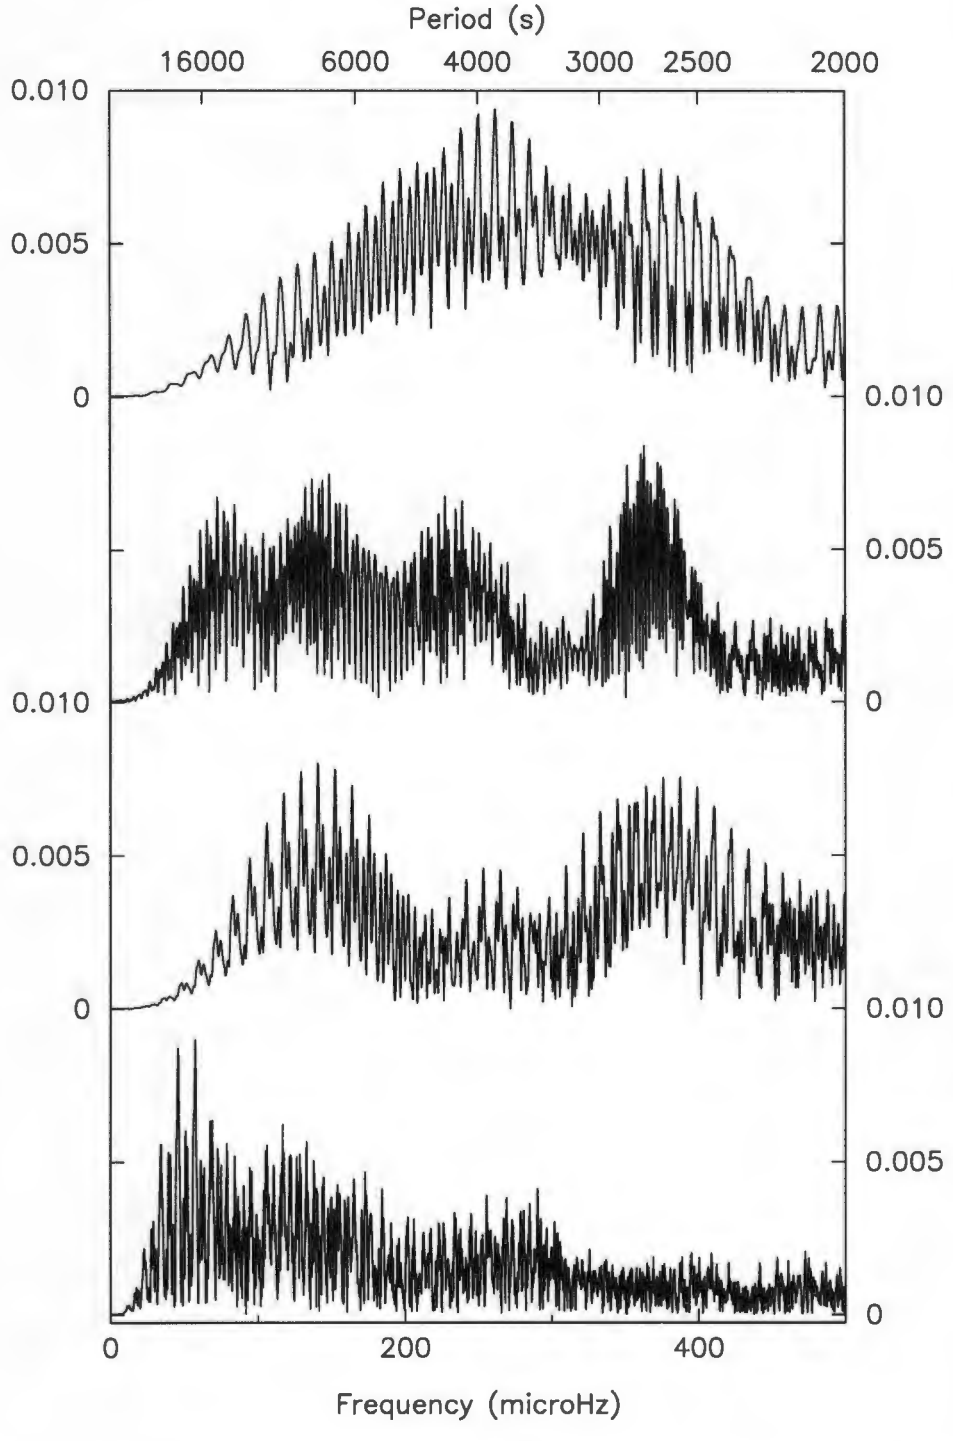

Figure 7.5: The low frequency end of the transforms: (*a*) March '97, (*b*) April '97, (*c*) September '97, and (*d*) May/June '98.

pendence on frequency, while *white noise* is by definition independent of frequency.) This means it can be tricky to distinguish if spikes occurring at frequencies smaller than $\sim 10^{-3}$Hz (corresponding to periods of $\sim$ a few hours) are due to effects originating at the distance of the star or are due to purely local atmospheric effects. The longer the baseline of the data set, the more likely that transparency variations will be averaged out.

The lightcurves in Figure 5.1 show no evidence for a consistent orbital modulation. On some nights (for example, the lightcurve of 19/5/98), there appears to be a low amplitude modulation of around 2 hours, but on other nights, there is no sign of it. This may be due to changes in the optical thickness of the disc. If the disc becomes optically thin it is possible to see the bright spot through the disc throughout the orbital cycle, and therefore there will be no orbital modulation. Discs of extremely low $\dot{M}$ CVs, like those of the WZ Sagittae stars, are expected to be optically thin, at least in their inner regions.

The 1998 dataset, with its much longer run lengths, is the most likely to obtain the orbital modulation if there is one. I have removed second order trends from each run, to remove differential extinction effects (see discussion in 5.2). Prewhitening the low frequency end of the amplitude spectrum (fourth transform in Figure 7.5) gives almost 20 signals above 3 mmag in amplitude, most of which must be due to transparency variations, incorrect extinction subtraction, or flickering and flaring in the star itself. At present, there is no way to tell which of these signals is the orbital period.

The highest amplitude signal has a period of 4.85 hours (57.3 μHz). This is much longer than the orbital period we would expect GW Librae to have. If GW Librae is an almost face-on system, the orbital modulation would have a very low amplitude, possibly lower than the red noise amplitude.

The second highest amplitude signal has a period of 2.38 hours (117 μHz), close to the white dwarf rotation period we'd get if the April "triplet" and "doublet" are rotationally-split modes. If this is the true orbital period, then the white dwarf rotation period is equal to the orbital period, making GW Librae a polar. The very

low white dwarf temperature is consistent with GW Librae being a polar, but its >9 magnitude 1983 outburst is not.

Warner (1996) proposes a subclass of polars, which he calls the "polaroids," in which the white dwarfs rotate synchronously with the orbital periods, but which nevertheless have accretion discs. i.e., they are IPs in which the white dwarf's rotation is locked to the orbital period. IPs can have DN outbursts (Warner 1995), but could such a system produce outbursts of 9 magnitudes or greater?

This is, however, only one of many possible scenarios that could account for the splitting we see. Anything that raises the degeneracy of the m-modes would split them apart. Perhaps we are seeing the effects of the tidal pull of the secondary?

With $\sim$ 20 signals above 3 mmag in amplitude in the red noise, a chance match with a possible rotation period is not unlikely. Indeed, the fifth-highest signal has a period of 3.8 hours (72.6 μHz), which matches the rotation period of 3.9 hours which we get if the April structures are $l = 2$ multiplets.

The only way to nail down the orbital period properly is to obtain radial velocity curves for the system. Unfortunately, being so faint, GW Librae would require a very large telescope in order to obtain the necessary spectroscopic precision.

Chapter 8

Future Work

This study of GW Librae has been limited by insufficient data, and further studies will be severely hampered by its faintness. GW Librae is not suitable for observation by the WET, but this may change in the not too distant future.

8.1 A CCD-Equipped Whole Earth Telescope

It is clear that the study of GW Librae will require lengthy observing runs with duty cycles of $\sim$ 100%, but at mag =$\sim$18.5, it is too faint for most of the telescopes operated by the WET. However, there are plans to introduce high speed CCD photometers into the WET arsenal (Kleinman 1998a; O'Donoghue 1998), which will greatly increase the magnitude limit for smaller telescopes. A campaign on GW Librae consisting of a combination of small (1m-class) CCD-equipped telescopes and larger (4m-class) PMT-equipped telescopes may be feasible in the not too distant future.

How long would such a multisite campaign need to be? If the $\sim 1\mu$Hz splittings are due to stellar rotation, then campaigns with a very high fractional coverage of at least $\sim$ 12 days are mandatory. Single site extensions of a week on either end of the main would be very useful, to extend baselines and to check the frequency, amplitude and phase stability of the pulsations. For this reason, single site observations for a week or two the month before and after the main campaign will probably also be a

good idea.

One unfortunate aspect is that the best time of year to observe GW Librae is in May, which is near the height of the "Bulge season" (the optimum time for observing the Galactic Bulge), so that most of the 1m-class CCD-equipped telescopes most ideal for long observing campaigns are dedicated to the MACHO and PLANET gravitational lensing projects.

In addition, judging from what we've learned from the other DAVs, it will take many more than just one intensive multisite campaign to "crack" GW Librae. Now that we know that DAVs can be found in accreting systems, it would be worth looking for other GW Lib stars among the polars and the WZ Sge stars. Hopefully we will find others that are brighter and more suited to WET campaigns.

8.2 The WET and the HST

In order to be able to extract useful information about the interior of a DAV, it is crucial to be able to assign correct k, l and m values to its pulsation modes. As discussed in detail elsewhere in this dissertation, the principal method of identifying the (k, l, m)-values of the modes is to measure accurately the period spacings between consecutive radial overtone modes, and the fine structure splitting caused by stellar rotation. This *period distribution method* can only be applied to stars for which a great many modes have been observed. For a cool DAV, decades of high resolution, multisite data may be required before enough modes are observed to make the period distribution method successful.

For stars like the DAVs, which do not show many modes at any given time, the period distribution method is not efficient. Robinson *et al.* (1995) and Kepler *et al.* (1995, 1997) applied another mode identification method, the *limb-darkening method*, to the DAVs G117-B15A and G226-29. Limb-darkening has a strong wavelength dependence, and therefore (due to geometric cancellation), modes of different l have different amplitudes in the UV relative to those in the optical.

Simultaneous observations in the UV with the Faint Object Spectrograph (FOS)

on board the Hubble Space Telescope (HST), and in the optical with the WET, can give UV/optical ratios for individual modes, sufficient to determine their l-values. Clemens *et al.* (1999) used the same theory, obtaining mode identifications of G29-38 without the use of the HST, but rather by obtaining high precision optical measurements with the Keck telescopes.

Could this method be applied to GW Librae? It could potentially aid greatly in deciphering the mysterious 650-s and 376-s regions of the amplitude spectrum. However, the limb-darkening method could be tricky to apply in the case of GW Librae, because unlike the other pulsating white dwarfs, it has an accretion disc. The inner parts of an accretion disc can be much hotter than the primary's photosphere (Warner 1995), dominating in the UV.

Fortunately, observed temperature profiles of Dwarf Nova discs in quiescence give temperatures of 3000 to 4000 K for the inner disc (Wood *et al.* 1986, Warner 1995 p80). Applying the limb-darkening method to GW Librae could be very rewarding, but to be successful, we would need spectroscopy in both the UV and optical of sufficiently high resolution to separate the disc's contribution from the white dwarf's light, in order to obtain accurate pulsation amplitudes. The disc's narrow emission line spectrum should not be too difficult to separate from the broad absorption line spectrum of the white dwarf photosphere.

8.3 A Dearth of Modelling

This study of an accreting DAV has been rather limited by the fact that we do not know what we are looking for in its transform. We do not know how accretion is going to affect the normal modes of the star; indeed, we have no idea whether accretion effects will dominate the pulsation spectrum, or whether the pulsations of an accreting DAV will be indistinguishable from those of a single one.

The infant field of accreting DAVs is in desperate need of the support of theoreticians and modellers. What are the normal modes of a DAV with an outer envelope that has a strongly non-spherically symmetric temperature, rotation, density and

chemical composition structure?

It may be useful to model normal modes of DAVs which are very different from the DAVs currently known. How do the normal modes of a DAV respond to extremely rapid rotation? (For example, Bildsten *et al.* (1996) discuss how neutron star g-modes are affected by rapid rotation.) What are the normal modes of a DAV with a magnetic field strength similar to those observed in polars? What is the pulsation spectrum of a DAV close to the Chandrasekhar mass? Is accretion a perturbing or stochastic driving force on g-modes?

I sincerely hope that the discovery of the DAV nature of GW Librae will stimulate investigation into these questions. GW Librae is a tough star to study, but the potential rewards to the fields of Asteroseismology and Cataclysmic Variables make it well worth the effort.

References

Bildsten, L., 1998, private communication.

Bildsten, L., Ushomirsky, G., Cutler, C., 1996, *Ap.J.*, **460**, 827.

Bohm-Vitense, E., 1989, *Introduction to Stellar Astrophysics, Vol. 3*, Cambridge University Press.

Bradley, P.A., 1996, *Ap.J.*, **468**, 350.

Bradley, P.A., Winget, D.E., 1994, *Ap.J.*, **421**, 236.

Bradley, P.A., Winget, D.E., 1994, *Ap.J.*, **430**, 850.

Brickhill, A.J., 1983, *MNRAS*, **204**, 537.

Brickhill, A.J., 1990, *MNRAS*, **246**, 510.

Brickhill, A.J., 1991a, *MNRAS*, **251**, 673.

Brickhill, A.J., 1991b, *MNRAS*, **252**, 334.

Brickhill, A.J., 1992, *MNRAS*, **259**, 529.

Cannizzo, J.K., Kaitchuck, R.H. 1992, Sci. Amer., **266**, **1**, 42.

Chen A.L., *A Survey of Cataclysmic Variables from the Edinburgh-Cape Blue Object Survey*, PhD Thesis, UCT, 1994.

Clemens, J.C., 1993, *Baltic Astronomy*, **2**, 407.

Clemens, J.C., 1994, PhD Dissertation, University of Texas at Austin.

Clemens, J.C., *et al.*, 1999, under referee.

Cleymans J., *Statistical Mechanics*, Honours course notes, Dept. of Physics, UCT, 1997.

Cox, J.P., *Theory of Stellar Pulsation*, Princeton University Press, 1980.

Cropper M.S., *Sp.Sci.Rev.*, **54**, 195, 1990.

D'Antonia, F., & Mazzitelli, I., 1991 in *IAU Symp. 145, Evolution of Stars: The Photospheric Abundance Connection*, ed. G. Michaud & A. Tutukov (Dordrecht:Kluwer), 399.

Deeming T.J., *Astr.&Sp.Sci* **36**, 1975.

Dolez, N., Vauclair, G., 1981, *A&A*, **102**, 375.

Dolez, N., Vauclair, G., & Koester, D., 1991, in *White Dwarfs*, ed., Vauclair, G. & Sion, E., Dordrecht:Kluwer, 361.

Downes R.A., Webbink R.F., Shara M.M., *PASP* **109** 734, 1997.

Duerbeck, H.W., 1987, *Sp. Sci. Rev*, **45**, 68.

Duerbeck, H.W., 1992, *MNRAS*, **258**, 629.

Duerbeck, H.W., Seitter, W.C., 1987, *ApSS*, **131**, 467.

Dziembowski, W., 1982, *Acta Astron.*, **32**, 147.

Dziembowski, W., 1998, private communication.

Dziembowski, W., Koester, D., 1981, *A&A*, **97**, 16.

Dziembowski, W., 1988, in *Workshop Proceedings Multimode Stellar pulsations*, ed. G. Kovacs, L. Szabados & B. Szeidl (Budapest: Konkoly Observatory), 127.

Faulkner, J., Flannery, B.P., Warner, B., 1972, *Ap.J.*, **175**, L79.

Fontaine, G., Wesemael, F. 1987, in *IAU Colloq. 95, 2nd Conf. on Faint Blue Stars*, ed A.G.D. Philip, D.S. Hayes, J. Liebert, (Schenectady: Davis), 319.

Ghosh P., Lamb F., *Ap.J.*, **223**, L83, 1978.

Goldreich, P., Wu, Y., 1999, *Ap.J.*, **511**, 904.

González., L.E., 1983, *IAU Circ 3854.*

Hansen, C.J., Kawaler, S.D., *Stellar Interiors*, Springer-Verlag, 1994.

Harrington, R.S., *et al.*, 1985, *AJ*, **90**, 123.

Harrop-Allin M., 1996, *Superhumps in AM Canum Venaticorum Stars*, MSc Thesis, University of Cape Town.

Hellier C., Cropper M., Mason K.O., 1991, *MNRAS*, **248**, 233.

Hoshi, R., 1979, *Prog. Theor. Phys.*, *61*, 1307.

Iben, I.,Jr., 1989, in *IAU Colloq.106, Evolution of Peculiar Red Giant Stars*, ed. H.R.Johnson & B. Zuckerman (Cambridge: Cambridge Univ. Press), 205.

Jones, P.W. *et al.*, 1989, *Ap.J.*, **336**, 403.

Kawaler, S.D., Hansen, C.J., Winget, D.E., 1985, *Ap.J.*, **295**, 547.

Kepler, S.O., 1990, *Rev. Mex. Astr. Astrophys.*, **21**, 335.

Kepler, S.O., 1993, *Baltic Astronomy*, **2**, 444.

Kepler, S.O., *et al.*, 1995, *Baltic Astronomy*, **4**, 302.

Kepler, S.O., *et al.*, 1997, *Proceedings of "A Half Century of Stellar Pulsation Interpretation,"* ed P.A. Bradley and J.A. Guzik, 140.

Kleinman, S.J., 1995, *G29-38 and the Advent of Cool DAV Asteroseismology*, PhD Dissertation, University of Texas at Austin.

Kleinman, S.J., 1998a, private communication.

Kleinman, S.J., 1998b, Proceedings of the Norway White Dwarf Conference, in press.

Kleinman, S.J., *et al.*, 1998a, *Ap.J.*, **495**, 424.

Kleinman, S.J., *et al.*, 1998b, Submitted to *Ap.J.*.

Kovetz, A., Prialnik, D., Shara, M. M., 1988, *Ap.J.*, **325**, 828.

Kraft, R.P., Matthews, J., Greenstein, J.L. 1962, *Ap.J.*, **136**, 312.

Kurtz, D.W., 1982, *MNRAS*, **202**, 761.

Kurtz, D.W., 1998, private communication.

Kurtz, D.W., 1999, in *Variable Stars as Important Astrophysical Tools*, C. Ibanoglu and C. Akan, eds., NATO-ASI, Kluwer Academic Press, in press.

Lamb, F.K., 1989, in *Timing Neutron Stars*, ed. H. Ögelman & E.P.J. van den Heuvel, Kluwer, p. 649.

Livio, M., 1993, *SAAS FEE Advance Course., Interacting Binaries*, ed. H. Nussbaumer & A. Orr, Springer-Verlag, Berlin, p.135.

Long, K.S., *et al.*, 1994, *Ap.J.*, **424**, L49.

Markiel, J.A., Thomas, J.H., van Horn, H.M., 1994, *Ap.J.*, **430**, 834.

Nather, R.E., Robinson, E.L., Stover, R.J., 1981, *Ap.J.*, **244**, 269.

Nather, R.E., *et al.*, 1990, *Ap.J.*, **361**, 309.

Nitta, A., 1996, MSc Dissertation, University of Texas at Austin.

Nitta, A., 1998, private communication.

Nitta, A., *et al.*, 1998a, submitted to *Ap.J.*.

Nitta, A., *et al.*, 1998b, submitted to *Ap.J.*.

Nitta, A., Winget D.E., 1997, *Baltic Astronomy*, **7**, 141.

O'Donoghue D., *South African Journal of Science*, 84, 1988.

O'Donoghue D., 1996, *Stellar Pulsation*, Honours course notes, UCT.

O'Donoghue D., *et al.*, 1991, *MNRAS*, **250**, 363.

O'Donoghue D., 1998, private communication.

Osaki Y., in *Theory of Accretion Discs II*, eds W.J.Duschl *et al.* , Kluwer, Dordrecht, 1994.

Paczynski, B., Sienkiewicz, R. 1981 *Ap.J.*, **248**, L27.

Patterson, J., 1980, *Ap.J.*, **241**, 235.

Patterson, J., Price, C. 1981, *Ap.J.*, **243**, L83.

Pringle, J.E., 1998, *MNRAS*, **230**, 587.

Provencal, J.L., 1994, Ph.D. Dissertation, University of Texas at Austin.

Rappaport, S., Joss, P.C., Webbink, R.F. 1982, *Astrophys. J.*, **254**, 616.

Ritter H., Kolb U., *A&AS*, 1997, submitted.

Robinson, E.L., Kepler, S.O., Nather, R.E., 1982, *Ap.J.*, **259**, 219.

Robinson, E.L., *et al.*, 1995, *Ap.J.*, **438**, 908

Schatzman, E.L., 1958, *White dwarfs*, Amsterdam, North-Holland Pub. Co.; New York, Interscience Publishers.

Shafter, A.W., *et al.*, 1985, *Ap.J.*, **290**, 707.

Shara, M.M., 1989 *PASP*, **101**, 5.

Shipman, H.L., 1989, in *IAU Colloq. 114, White Dwarfs*, ed. G. Wegner (Berlin:Springer), 220.

Sion, E.M., *et al.*, 1994, *Ap.J.*, **430**, L53.

Sion, E.M., Szkody, P., 1990, *IAU Colloq. No. 122*, p. 59.

Smak, J., 1993 *Acta Astr.*, **43**, 121.

Sparks, W.M., *et al.*, 1993, in Cataclysmic Variables and Related Physics, eds. O. Regev & G. Shaviv, Inst. Phys. Publ., Bristol, P. 96.

Szkody, P., *et al.*, 1999, Proceedings of the Warner Symposium, In Press.

Tassoul, J.-L., 1978, Theory of Rotating Stars, Princeton Univ. Press, Prinston p. 248.

Tassoul, M., 1980, *Ap.J.S.*, **43**, 469.

Unno, W., *et al.*, 1989, *Nonradial Oscillations of Stars*, (2nd ed.; Tokyo: Univ. Tokyo Press).

Van Kerkwijk, M., *et al.*, 1998, *MNRAS*, in press.

Vauclair, G. *et al.*, 1993, *A&A*, **267**, L35.

Verbunt, F., *Sp.Sci.Rev.*, **32**, 379, 1982.

Verbunt, F., Zwaan, C., 1981 *A&A*, **100**, L7.

Vuille, F., 1998a, private communication.

Vuille, F., 1998b, Proceedings of the Norway White Dwarf Conference, in press.

Warner, B., 1995, *Cataclysmic Variable Stars*, Cambridge University Press.

Warner, B., 1996, *Ap.Sp.Sci*, **241**, 263.

Warner, B., Robinson, E.L., 1972, *Nature Phys. Sci*, **234**, 2.

Warner, B., O'Donoghue, D., Fairall, A.P. 1981, *MNRAS*, **196**, 705.

Watson, T.K., 1992, IAU Circular 5603.

Weber, E., Davis, L., 1967 *Astrophys. J.*, **148**, 217.

Wedermann, V., 1990, *ARAA*, **28**, 103.

Wickramasinghe, D.T., Wu, K., 1993, *MNRAS*, **266**, L1.

Wilson, O.C., 1966, *Astrophys. J.*, **144**, 695.

Winget, D.E., Fontaine, G., 1982, in *Pulsations in Classical and Cataclysmic Variable Stars*, ed. J.P.Cox & C.J.Hansen (Boulder: University of Colorado Press), 46.

Winget, D.E, Hansen, C.J., Van Horn, H.M., 1983, *Nature*, **303**, 781.

Winget, D.E. *et al.*, 1987, *Ap.J.*, **315**, L77.

Winget, D.E. *et al.*, 1991, *Ap.J.*, **378**, 326.

Winget, D.E. *et al.*, 1994, *Ap.J.*, **430**, 839.

Winget D.E., Van Horn H.M., Hansen, C.J., 1981, *Ap.J.*, **245**, L33.

Winget D.E., and Van Horn H.M., *Sky & Telescope*, September 1982.

Winget D.E., in *Advances in Helio- and Asteroseismology*, ed. Kawaler S.D., Kluwer, Dordrecht, 1988.

Wood, M.A. 1992, *Ap.J.*, **386**, 539.

Wood, M.A., *et al.*1987, *Ap.J.*, **313**, 757.

Wood, J.H., *et al.*1986, *MNRAS*, **222**, 645.

Wood, J.H., *et al.*1993, *MNRAS*, **261**, 103.