

CCD Time-Series Photometry of White Dwarf Stars

by

Paul Robin Brian Chote

A thesis
submitted to the Victoria University of Wellington
in fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics.

Victoria University of Wellington
2014

Abstract

This thesis describes a practical programme that focused on CCD photometry of pulsating white dwarf (WD) stars. The first part of this thesis describes the development of two high-speed CCD photometer instruments and their data reduction pipeline, while the remainder describes the observation and analysis of several pulsating WDs and other targets.

The two photometers (Puoko-nui North and South) share a common hardware design that is optimized for acquiring efficient photometry with integration periods of milliseconds through to minutes. The design integrates a commercial CCD (Charge-Coupled Device) camera and GPS (Global Positioning System) receiver with custom timing electronics and control software. The reduction and visualization software developed for these instruments provide detailed real-time information to the observer, and a streamlined data reduction pipeline.

EC 04207-4748 is a pulsating helium atmosphere WD that shows significant non-sinusoidal intensity variations. We show that the pulsation spectrum of this WD can be described by four independent pulsation eigenmodes plus linear combinations that arise from non-linear energy transport through a sub-surface convection layer. Our results are consistent with similar analyses that have been made for similar stars, and add an additional data point to the growing catalogue of these convection measurements. We argue that the convection layer depth may form a useful substitute for the effective temperatures of these WDs.

GW Librae is the class prototype of the accreting WD pulsators. These stars exist in cataclysmic variable (CV) systems, and show a mix of CV and pulsating WD-related phenomena. Our observations of GW Librae four – six years after its 2007 outburst show signs of quasi-stable intensity modulations that we believe may be caused by non-radial pulsations, but these are not convincingly explained by existing WD or CV models.

L19-2 is a hydrogen atmosphere WD pulsator that shows extremely stable pulsation behaviour. We combine new observations with archival observations dating back to the mid 1970's, and derive a preliminary estimate of the period rate of change \dot{P} for two of the pulsation modes in this target. We show a clear result for the main 192 s pulsation mode $\dot{P}_{f_2} \lesssim 10^{-14} \text{ s s}^{-1}$, and discuss the improvements that we plan to make in order to convincingly improve this constraint by an additional order of

magnitude.

Observations of other rapidly variable targets include two extremely low mass (ELM) WDs, which exhibit variability due to their orbital motion (J0751) as well as non-radial pulsations (J1518); the 33 ms optical period of the Crab Pulsar; the helium atmosphere WD pulsators EC 05221-4725 and EC 20058-5234; the stable hydrogen atmosphere pulsator G117–B15A; and the eclipsing sdB binary system PG 1336-018.

Contents

Acknowledgments	xiii
List of Abbreviations	xv
1 Pulsating White Dwarf Stars	1
1.1 White Dwarfs	1
1.2 Pulsations	3
1.3 Applications	8
1.4 Thesis Structure	10
2 Instrumentation	13
2.1 Time-series Photometry	14
2.2 History	15
2.2.1 Photomultiplier Instrumentation	15
2.2.2 CCD Instrumentation	17
2.3 The Puoko-nui Photometer System	21
2.3.1 Science Camera	23
2.3.2 Timing Hardware	30
2.3.3 Control Computers	33
2.3.4 Acquisition Software	34
2.3.5 Autoguiding	37
2.4 Lightbox and Star Simulator	38
3 Data Reduction	41
3.1 Photometry Methods	45
3.2 Reduction Procedure	46
3.2.1 Calibration Frames	46
3.2.2 Frame Registration	50
3.2.3 Synthetic Aperture Photometry	51
3.2.4 Full-Width at Half-Maximum	54
3.2.5 Relative Photometry	55
3.2.6 Fourier Analysis	57

3.2.7	Aperture Size Optimization	60
3.2.8	Time Conversions	60
3.2.9	False Alarm Probability	61
4	EC 04207-4748	63
4.1	Background	63
4.2	Observations and DFT Analysis	64
4.3	Convection Theory	68
4.4	Convection Analysis	72
5	GW Librae	75
5.1	Accreting Pulsating White Dwarfs	77
5.2	Observations	80
5.3	The 19 Minute Periodicity	84
5.3.1	The Combined DFT	85
5.3.2	Mode Stability	88
5.4	The 4–6 Minute Periodicity	91
5.5	Discussion	92
6	L19-2	95
6.1	Evolution of Pulsation Periods	97
6.1.1	The O–C diagram	99
6.2	Observations	101
6.3	Analysis	102
6.3.1	Proper Motion	104
6.3.2	Direct Period Comparison	105
6.3.3	O–C Analysis	105
7	Other Observations	109
7.1	ELM White Dwarfs	109
7.1.1	SDSS J075141.18-014120.9	110
7.1.2	SDSS J151826.68+065813.2	112
7.2	Crab Pulsar	113
7.3	EC 05221-4725	116
7.4	EC 20058-5234	117
7.5	PG 1336-018	119
7.6	G117–B15A	120
8	Conclusions	121
A	Puoko-nui Instrument Manual	123

A.1	Hardware Setup	123
A.2	Collecting Data	125
A.2.1	Preparing the Acquisition	125
A.2.2	Autoguiding	125
A.2.3	Starting the Acquisition	126
A.2.4	Configuring the Online Reduction	129
A.2.5	Additional Acquisition Software Settings	133
A.3	Example <code>tsreduce</code> Commands	134
A.4	Troubleshooting	135
A.5	Software Updates	136
B	Timer Firmware Description	141
C	Finding Charts	151
D	Observations & Publications	155

List of Figures

1.1	Propagation diagram of a pulsating WD	6
1.2	Surface luminosity variations of a pulsating WD	7
2.1	The VUW single-channel photometer	16
2.2	The VUW two- & three-channel photometers	18
2.3	Frame corruption examples	20
2.4	Puoko-nui instrument block diagram	21
2.5	The Puoko-nui South instrument	22
2.6	The Puoko-nui North instrument	22
2.7	CCD quantum efficiency plots	26
2.8	CCD dark current vs temperature relationship	27
2.9	CCD dark current noise relationship	27
2.10	CCD transfer curve	28
2.11	Star-simulator CCD parameter comparison	29
2.12	The Puoko-nui timer unit	31
2.13	Acquisition software block diagram	35
2.14	Screenshot of the acquisition and online reduction software	35
2.15	Acquisition software configuration options	36
2.16	Offset guider schematic	37
2.17	Puoko-nui bench-test station	39
2.18	Star simulator data comparison	39
2.19	High-speed bench-test observations	40
3.1	Master dark and flat frames	46
3.2	Calibrated frame comparison	47
3.3	Frame registration example	50
3.4	Background intensity histograms	52
3.5	Aperture integration geometry	53
3.6	Aperture integration geometry (excluded cases)	54
3.7	FWHM calculation demonstration	55
3.8	EC04207-4748 online reduction plot	57

3.9	L19-2 pre-whitening demonstration	59
3.10	False alarm probability demonstration	62
4.1	EC 04207-4748 observation data	65
4.2	EC 04207-4748 pulsation residuals	67
4.3	EC 04207-4748 convection model fit	73
4.4	DAV and DBV τ_0 versus T_{eff} relationships	73
5.1	GW Librae 1998 MJUO observations	76
5.2	GW Librae observation data	82
5.3	GW Librae observation data (May 2012)	83
5.4	Frequency determination improvement	86
5.5	Combined may12 DFT	87
5.6	Frequency multiplet model comparison	88
5.7	GW Librae may12 spectrograms	89
5.8	Monte Carlo modelling procedure	91
5.9	GW Librae folded lightcurves	92
5.10	GW Librae 4000 – 5000 μ Hz spectrograms	93
6.1	L19-2 xCOV12 DFT	96
6.2	Comparison of astronomical clocks	97
6.3	L19-2 proper motion	104
6.4	L19-2 \dot{P} measurements	106
6.5	L19-2 \dot{P} cycle count verification	107
6.6	L19-2 Monte Carlo phase uncertainty estimation	108
7.1	WD J0751 observations	111
7.2	WD J0751 folded lightcurve	111
7.3	WD J1518 observations	112
7.4	Crab Pulsar observations	115
7.5	Crab Pulsar folded lightcurve	115
7.6	EC 05221-4725 observations	116
7.7	EC 20058-5234 observations	118
7.8	PG 1336-018 observations	119
7.9	G117–B15A observations	120
A.1	Puoko-nui South MJUO mounting notes	123
A.2	Puoko-nui South cable connections	124
A.3	Guide star offset calculation	126
A.4	Puoko-nui acquisition screenshot	127
A.5	Puoko-nui acquisition configuration	128
A.6	Example <code>config.sh</code> script	129

A.7	Aperture selection illustration	130
A.8	Timer unit circuit board	139
B.1	Timer firmware logic: main loop	145
B.2	Timer firmware logic: USB processing	146
B.3	Timer firmware logic: GPS serial processing	147
B.4	Timer firmware logic: 1 Hz interrupt	148
B.5	Timer firmware logic: camera monitoring	149
B.6	Timer firmware logic: 1 ms timer interrupt	150

List of Tables

1.1	Pulsating WD classifications	4
1.2	Table of targets	11
2.1	MicroMax camera CCD properties	26
2.2	Timer operating modes	32
4.1	EC 04207-4748 observation log	64
4.2	EC 04207-4748 frequencies	66
4.3	EC 04207-4748 convection model parameters	72
5.1	Accreting pulsator systems	78
5.2	GW Librae observation log	81
5.3	GW Librae May 2012 frequencies	90
6.1	L19-2 pulsation frequencies	96
6.2	L19-2 historical observation log	103
6.3	L19-2 period measurements	105
7.1	EC 20058-5234 frequencies	118
A.1	Advanced <code>tsreduce</code> parameters	132
A.2	Advanced acquisition software parameters	133
D.1	Observation log: References key	155
D.2	Observation log: March 2011	156
D.3	Observation log: July 2011 (a)	156
D.4	Observation log: July 2011 (b)	157
D.5	Observation log: November 2011	157
D.6	Observation log: March 2012	157
D.7	Observation log: April 2012	158
D.8	Observation log: May 2012	158
D.9	Observation log: December 2012	158
D.10	Observation log: January 2013	158

D.11 Observation log: March 2013 159

D.12 Observation log: April 2013 159

D.13 Observation log: May 2013 160

D.14 Observation log: June 2013 160

Acknowledgments

First and foremost I would like to thank my supervisor Denis Sullivan (DJS) for his guidance and support over the last several years. I have learned many things from Denis, not in the least that the study of physics is as much about understanding and (where appropriate) presenting the underlying processes as it is in presenting the results, and also that a well prepared plot trumps almost every other method of communicating scientific data. I have enjoyed our discussions on many topics, from general physics and technology through to politics and english grammar. His invaluable editorial comments and advice have significantly improved the presentation of this thesis.

I would also like to acknowledge Rod Brown for his expert work with the electronics described in this thesis, and for guidance when I was finding my feet in the area of embedded software. Also, to the University of Canterbury for providing access to the Mt John facilities, and in particular to Alan Gilmore and Nigel Frost for being available to help whenever technical issues struck at the telescope.

I am grateful for the opportunity that the White Dwarf group at the University of Texas at Austin provided to extend our instrumentation to their new camera system, as this added a lot of depth to my instrumentation work. In particular, I would like to thank Sam Harrold from UT, and Dean Chandler and Willie Strickland from the Central Texas Astronomical Society who were available and happy to help with the hands-on tasks that couldn't be done via a remote internet connection.

I'd finally like to acknowledge and thank the friends and family, both local and remote, who helped keep me stay grounded during the last three years. This has been a long and, at times, stressful journey, and it has been extremely satisfying to see everything converge into the form that is presented in the following chapters.

This research was supported by the Marsden Fund Council from Government funding, administered by the Royal Society of New Zealand.

List of Abbreviations

A number of acronyms are used throughout this text. Most of these are defined at the point they are first introduced, but they are also listed here for quick reference.

BJD	Barycentric Julian Day: Time measured from the solar system barycentre.
CCD	Charge Coupled Device: A type of 2D imaging detector.
CV	Cataclysmic Variable: A binary star system that features an accreting WD. These systems undergo periodic bright outbursts.
DAV	Stellar classification: Degenerate A Variable (H atmosphere).
DBV	Stellar classification: Degenerate B Variable (He atmosphere).
DFT	Discrete Fourier Transform.
DJS	Denis J. Sullivan. My PhD supervisor.
DS9	Astronomical image display software, developed by the Smithsonian Astrophysical Observatory.
EC04	EC 04207-4748: A DBV WD.
FAP	False Alarm Probability: A statistical test for identifying a noise threshold in the Fourier domain.
FFT	Fast Fourier Transform: An efficient DFT algorithm.
IRAF	Image Reduction and Analysis Facility: A software suite for reducing astronomical images and spectroscopic data.
M_{\odot}	Solar mass: the mass of our Sun (2×10^{30} kg).
MJVO	Mount John University Observatory, Lake Tekapo, NZ (operated by the University of Canterbury).
MLT	Mixing length theory: A model for convection in stellar atmospheres.
mma	Milli-modulation amplitude: A relative amplitude measurement, $10 \text{ mma} = 1\%$ of the total amplitude.
mmi	Milli-modulation intensity: A relative intensity (time-domain) measurement, where $10 \text{ mmi} = 1\%$ intensity change.
PMT	Photomultiplier Tube: A type of photo-electric detector.

PROMPT	Panchromatic Robotic Optical Monitoring and Polarimetry Telescopes: A collection of 0.41 m telescopes at Cerro Tololo Inter-American Observatory in Chile.
SAAO	South African Astronomical Observatory, Cape Town, South Africa.
sdB	Stellar classification: Subdwarf B star.
SOAR	Southern Astrophysical Research Telescope. A 4.1 m telescope at Cerro Pachón, Chile.
T_{eff}	Effective temperature: A temperature measurement obtained from spectral fitting.
UT	University of Texas at Austin.
VUW	Victoria University of Wellington.
WD	White Dwarf.
WET	Whole Earth Telescope: A collaboration of observers distributed by longitude for obtaining extended coverage observations of selected pulsating targets.
XCOV	The code prefix for WET campaigns, short for eXtended COverage.

Chapter 1

Pulsating White Dwarf Stars

There is a lot that could be said about the white dwarf stars and their pulsations, more than I could reasonably do justice to here. The following discussion briefly outlines our understanding of these interesting objects, and the ways in which they can be used as stellar laboratories. For more details I recommend an interested reader to see the comprehensive review articles by Althaus et al. (2010); Fontaine and Brassard (2008); Winget and Kepler (2008).

1.1 White Dwarfs

Stars below about 10 solar masses on the main sequence – around 98% of all stars (Winget and Kepler, 2008) – are believed to ultimately form a white dwarf (WD). The evolutionary processes that form WDs produces a peaked mass distribution, with the majority of WDs having a mass of $0.6 \pm 0.1 M_{\odot}$ (Kepler et al., 2007) and a radius comparable to the Earth. This makes them extremely dense, with strong surface gravity ($\log g \sim 8$ in cgs units, i.e. $g \sim 10^8 \text{ cm s}^{-2}$) and a degenerate core. The majority of WDs feature a degenerate carbon / oxygen core, which is covered by a thin stratified atmosphere with non-degenerate layers of helium and (for about 80% of WDs) hydrogen.

The WDs (and pre-WDs) that we are potentially interested in span a wide range of temperatures, from $> 10^5 \text{ K}$ down to $\sim 4000 \text{ K}$ for the oldest detected stars (Althaus et al., 2010). They are intrinsically faint despite these high temperatures: the insulating atmosphere and small surface area restricts the rate at which energy can be radiated away. The main energy source is thermal energy stored in their core – nuclear fusion is negligible or entirely absent in most WDs.

The details of how a given star evolves from the main sequence into a WD can be

complicated, and depends on the initial conditions. The most common evolutionary track has the WD progenitor evolve along the red giant branch, where it fuses helium to carbon and oxygen via the triple- α process (Althaus et al., 2010). A strong stellar wind strips away the (relatively cooler) outer layers, and leaves the hot central core which contracts into the WD.

Stars with high metallicity or in close binary systems are able to lose mass at a much higher rate than isolated stars, which may prevent the triple- α process from igniting (Bergeron et al., 1992). These stars are believed to evolve into a low-mass helium core WDs, which are discussed further in Chapter 7.

The extreme pressure and density in the WD core ($\sim 10^6$ times that of the sun) ionizes the chemical constituents into a plasma of ions and a degenerate electron gas. This degenerate gas exhibits a pressure (Fowler, 1926) that resists the inward gravitational contraction, giving the WD its stable mechanical structure. The Fermi energy of this degenerate gas is much greater than the thermal energy, so the mechanical properties are essentially independent of temperature. The ions in the core can be treated as a classical gas, which acts as a thermal reservoir. This separation between mechanical and thermal properties greatly simplifies the physical modelling of WD interiors.

The electron degeneracy leads to the counter-intuitive white dwarf mass-radius relationship, where the radius *decreases* with increasing mass. A more massive WD has a stronger inward gravitational force, which must be balanced by a corresponding increase in the restoring force. This occurs at a smaller radius, as compressing the electrons closer together will increase their velocity and subsequent degeneracy pressure. The mass-radius relationship becomes even more curious when one adds special relativity, which shows that the relationship between density and pressure must weaken as the electron velocities approach the speed of light. Chandrasekhar (1931) showed that this implies that there is a maximum stable mass that can be supported by electron degeneracy alone, and this has now been well established to be $1.4 M_{\odot}$. Stars that are more massive than the Chandrasekhar limit will continue to contract until a different supporting mechanism can act – for neutron stars this is provided by the (much stronger) neutron degeneracy pressure. The most massive stars have no known supporting mechanism, and collapse to form black holes.

WD evolution is dominated by cooling. The basic cooling model by Mestel (1952) relates the cooling rate to the star’s mass, luminosity, and composition. More detailed models include secondary energy effects, such as the production of plasmon neutrinos at temperatures $> 25,000$ K, and the liberation of gravitational potential energy as elements stratify within the core.

A reasonable intuitive model for the cooling is a hot metallic core covered with an insulating blanket. Thermal energy is transported efficiently through the core via conduction, but the non-degenerate atmosphere is significantly less efficient, and restricts the rate at which energy is radiated away. This leads to very long cooling times ($\sim 10^9$ years) despite there being no active energy generation.

WDs can be classified into one of several groups using a system defined by Sion et al. (1983). These categories types start with D (for degenerate) followed by one or two letters that characterize the spectral type. WDs with special properties receive an additional letter, such as V for variable stars, P for those that show polarization.

The majority (around 75%) of WDs are classified as DA, as their spectra feature only hydrogen absorption lines. These stars have a thin surface layer of hydrogen ($< 10^{-6}$ by mass) which masks the underlying helium layer ($\sim 1\%$ by mass).

The remaining 25% of WDs show no detectable hydrogen, implying that this outer layer is completely absent. The majority of these WDs are classified as DB, based on their strong helium (HeI) absorption lines. The DO / planetary nebula stars, although technically pre-WDs, have also been studied intensively. They are characterised via their HeII absorption lines, and are believed to be the precursors of the DB and/or DA stars (although the mechanism that governs this transformation is not well understood). A third relevant class are the hot DQ stars (Dufour et al., 2008), which show predominantly carbon atmospheres.

1.2 Pulsations

The first pulsating WD to be discovered was HL Tau 76 (Landolt, 1968; Warner and Nather, 1970). Targeted searches following the discovery of this strange hydrogen atmosphere object revealed a number of additional variable WDs with similar temperatures and surface gravities. It was clear that this heralded the discovery of a new class of variable stars (oddly, named after the *second* star to be discovered, R548 / ZZ Ceti), but the source of this variability was a mystery – there were no known mechanisms that could result in multi-periodic variability with a timescale of \sim minutes for these stars.

The pulsation mechanism was eventually identified as non-radial g-mode pulsations driven within a layer of partially ionized hydrogen (Chanmugam, 1972; Warner and Robinson, 1972). This led to the famous prediction and subsequent discovery of similar pulsations driven by the partial ionization of helium in the DB WDs (Winget et al., 1983b). Similar non-radial pulsations were later discovered in DO / planetary

Class name	Ionization Species	T_{eff} (kK)	$\log g$	Discovery Paper
DOV / PNNV (GW Vir)	C/O	80 – 180	5.5 – 7.5	McGraw et al. (1979)
DBV (V777 Her)	HeII	22 – 29	7.6 – 8.2	Winget et al. (1982)
DQV	C?	15 ~ 20	8 ~ 9	Montgomery et al. (2008)
DAV (ZZ Ceti)	HI	10.5 – 12.5	7.8 – 8.8	Landolt (1968)
ELM	HI?	8 ~ 11	5.5 ~ 7.5	Hermes et al. (2012b)

Table 1.1: Non-radial pulsations occur within several types of white dwarf. The specified temperature ranges define the ‘instability strips’ where the specified partially ionized constituent can drive pulsations. The DQV and ELM stars are recent discoveries, and their properties have not been precisely identified.

nebula stars, DQ (carbon atmosphere) WDs, and in the helium-core extremely low mass (ELM) WDs. These different groups are listed with their basic properties and discovery papers in Table 1.1.

There are two basic requirements for driving pulsations in these stars (outlined by Montgomery et al., 2008): (a) a sub-surface partial ionization zone, and (b) an oscillatory mechanism where the maximum pressure lags the maximum density.

Partial ionization occurs when the thermal energy of the atmosphere matches the ionization potential of one of the electron orbitals. The ionized and non-ionized atoms exist in an equilibrium, which allows for energy to be absorbed into ionizing an electron, and then later released during a subsequent recombination.

There are two competing models for the pulsation mechanism in WDs. Early theories assumed that this was provided by the $\kappa - \gamma$ mechanism (where κ measures the opacity or absorptive properties of a medium, and γ is the ratio of specific heats) that is known to drive pulsations in other stars such as Cepheid variables. A more recent model, first proposed by (Brickhill, 1991a), explains the pulsations in terms of ‘convective driving’, which is claimed (e.g. Montgomery et al., 2008) to solve some internal inconsistencies that arise when applying the $\kappa - \gamma$ model to these WDs.

Both convective driving and $\kappa - \gamma$ models correctly describe the observed result that the pulsations become overstable (driven) once the thermal timescale associated with partial ionization matches the inherent pulsation eigenmode periods, which depend on the stellar structure. This produces a well-defined ‘blue edge’ (upper temperature limit) to the pulsation instability region. The red edge remains poorly defined, as it is believed to result from non-linear effects that are not completely handled by either model. One definite advantage of the convective driving model is that it naturally explains the non-sinusoidal lightcurves that are observed in WDs near the red edge of the instability strips. This is discussed in more detail in Chapter 4.

The general operation of the pulsation mechanism can be illustrated in terms of the $\kappa - \gamma$ model:

1. Consider a region of partially ionized material (e.g. hydrogen or helium) that is initially out of equilibrium, and pulled inwards by gravity. The material in the region is compressed as it moves deeper into the star.
2. For an ideal gas, this compression would have a corresponding increase in temperature. This energy is instead absorbed into ionizing electrons, so the temperature does not change.
3. The temperature gradient between the region and the (hotter) surrounding material causes additional heat to flow into the region, causing even more ionization to occur.
4. The effect of temperature and density on the opacity is commonly approximated by Kramers law: $\kappa \approx \sigma T^{-7/2}$. The increase in density and constant temperature leads to an increase in opacity, and restricts the rate at which heat can pass through the region.
5. The reduced heat flow through the region causes the temperature (and therefore pressure) to build up beneath it, until it reaches a level sufficient to push the layer outward (i.e. maximum pressure lags maximum density).
6. If the recombination timescale of the partial ionization process matches the oscillation period, then the energy released from recombination can maintain a constant temperature as the region moves outwards and expands. The opacity decreases, and the cycle can repeat. If the timescales do not match, then the oscillation will be damped, and the pulsation will not be driven.

The behaviour of the stellar material in response to perturbations caused by the pulsations can be described using two key parameters: the Lamb frequency (L_l) is related to the local speed of sound, and the Brunt-Väisälä frequency (N) describes the natural oscillation frequency of a fluid element in the local gravitational field. These parameters vary with radius in response to changes in density, temperature, chemical composition, and other parameters. A plot of N^2 and L_l^2 against radius (illustrated in Figure 1.1) is known as a propagation diagram, and encodes essentially all of the important information about the pulsation physics and stellar properties.

A more detailed analysis of the underlying physics shows that pulsations can be driven at particular discrete frequencies (σ) within two different regimes on the propagation diagram. Radial p-mode pulsations can propagate in regions where $\sigma^2 > N^2, L_l^2$, and have short periods (\sim seconds) in WDs. These pulsations rely

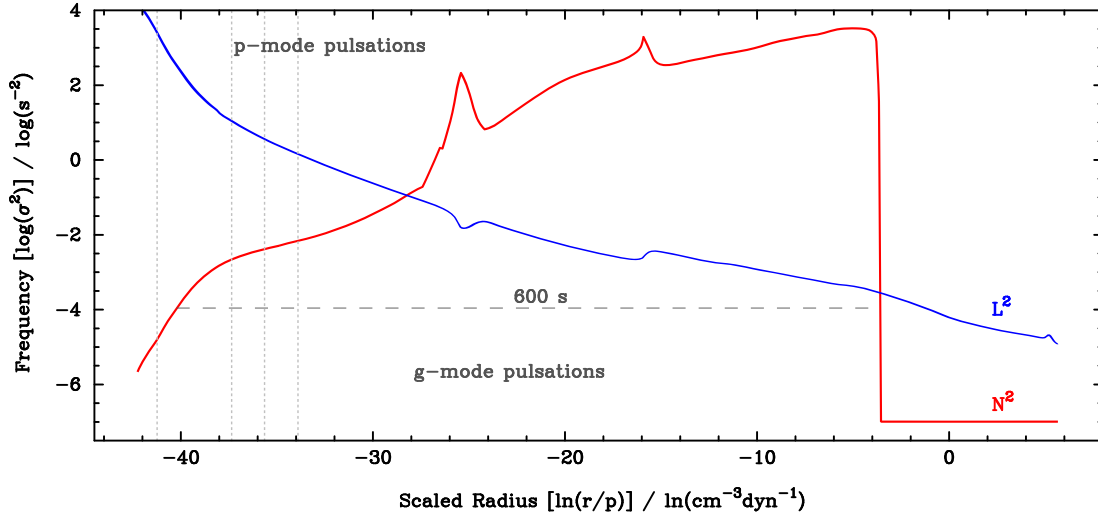


Figure 1.1: An example of a propagation diagram for a particular DAV white dwarf (BPM 37093), duplicated from Winget and Kepler (2008). These diagrams show the regions within the WD where pulsations of different frequencies can propagate. The shape of the (squared) Lamb and Brunt-Väisälä frequency curves encode many important stellar parameters, and can (in principle) be recovered from observations of many independent pulsation modes.

on pressure as the restoring force. The amplitudes of these modes are strongly suppressed by the strong surface gravity of the WD, and any subsequent luminosity variations have yet to be detected despite several targeted searches using large telescopes (e.g. Silvotti et al., 2011; Kilkenney et al., 2014).

The pulsations that we observe are g-modes, which propagate in regions where $\sigma^2 < N^2, L_t^2$. This corresponds to periods of $100 \sim 1000$ seconds within WDs. These non-radial modes have gravity as a restoring force, and the motion of the stellar material is largely horizontal, so these are not suppressed by the large surface gravity.

The g-modes manifest as temperature and luminosity variations on the surface of the star, featuring relatively hotter and cooler patches that can be described by a sum of spherical harmonics. Most of these patches cancel each other out, but our ability to detect light only from the visible disk leaves a small differential that we can detect. The time-dependence of these variations are largely sinusoidal, which leaves a signature in the intensity that we can measure. This is demonstrated for a single spherical harmonic mode in Figure 1.2.

The spherical harmonic analysis describes the three-dimensional standing wave pattern inside the star using linearly-independent normal modes. Each mode has a well defined eigenfrequency, and can be identified by counting the number of nodes or stationary points laterally (ℓ) and longitudinally (m) around the surface, and radially within the star (n or k). This formalism is identical to the description of

electron orbitals around a hydrogen atom. Real WDs typically show multiple simultaneous eigenmodes, which can be individually resolved by examining the intensity measurements in the Fourier domain.




Figure 1.2: The g-mode pulsations manifest as surface luminosity variations which can be modelled as a sum of spherical harmonics. The detected light intensity is given by the integral over the visible disk, which leads to a sinusoidal time variation with a period that corresponds to the eigenfrequency of the spherical harmonic mode (shown for $\ell = 2$, $m = 1$).

The goal of asteroseismology is to identify the specific eigenmode for each pulsation frequency, and then use these to obtain a propagation diagram. Each pulsation mode samples a different slice from the diagram, and so one could in principle (given enough modes) invert these to directly recover profiles for N and L_l . This procedure has so far only been achieved for the sun (see e.g. Basu and Antia, 2008), which has many thousands of identified eigenmodes.

We cannot detect enough WD pulsation modes to calculate the propagation diagram by inversion. There are multiple factors that contribute towards this: modes may not be excited, or may be excited with undetectably low amplitudes. In particular, the geometric cancellation significantly improves with increasing ℓ , and so modes with $\ell > 2$ are unlikely to have visible amplitudes in most pulsating WDs.

The standard procedure when faced with this situation (having a calculation that is simple in one direction, but near-impossible in the other) is to throw large amounts of computational power at the problem: a large number of trial propagation diagrams can be calculated and then compared against the known pulsation modes. If the input physics are realistic, then the diagrams that provide the best agreement with observation are likely to provide an accurate description of the star. The main theoretical modelling groups have traditionally maintained independent codes for this (e.g. the UT Austin group has evolved their WDEC code over many years), but

the MESA project (Paxton et al., 2013) looks promising as a future package that incorporates WD pulsations into a more general calculation environment.

In rough terms, if we assume that we know the eigenmode associated with each frequency (which is not a simple task in itself), then we can relate the spacing of the pulsation periods to physical quantities: the mean period spacing between adjacent k modes encodes the total stellar mass (more massive stars are more dense, which changes the Brunt-Väisälä frequency); differences in the individual periods away from the mean spacing encodes information on chemical and structural boundaries such as the H–He transition region or crystallization within the core; and the spacing between components in $\ell \geq 1$ multiplets encodes information on the rotation and magnetic properties of the star. Each measured eigenfrequency contributes an independent constraint on these parameters, and so the amount of information that we can reliably constrain varies from star to star. Several examples of asteroseismic analyses of stars include PG 1159-035 (Winget et al., 1991), GD 358 (Bradley and Winget, 1994), BPM 37093 (Kanaan et al., 2005), and EC 20058 (Sullivan et al., 2008).

At the other end of the spectrum, many WDs do not have enough modes for this technique to be useful. One workaround for this problem is to group multiple stars with similar properties, and analyse the ensemble group properties (e.g. Clemens, 1995; Romero et al., 2013). Different modes are likely to be excited in different stars within the sample, which can produce a large number of modes to match against the numerical models.

1.3 Applications

Theory and observation both strongly support the idea that *all* WDs will pulsate when they pass through the instability strips listed in Table 1.1. This means that the information which we obtain using asteroseismology applies to white dwarfs in general, and makes this a particularly powerful tool for studying the interior structure of these stars. The interior is where most of the interesting physics occur, but we cannot directly observe them using standard photometric or spectroscopic techniques. Asteroseismology is similar in many ways to the way in which seismologists infer the structure of the Earth by analyzing earthquake data.

White dwarfs contribute important measurements towards the understanding of our galaxy and universe. The coolest WDs place an important constraint on the age of our local galactic disk, and so it is important to understand the many contributions to the WD cooling rate if we want to relate WD temperature to an age distribution.

Fontaine et al. (2001) provide a comprehensive overview of the potential of WD ‘cosmochronology’.

White dwarfs in close binary systems can show other interesting phenomena, with examples such as the cataclysmic variable (CV) systems and Type Ia supernovae. These supernovae are important standard candles in measuring distances and the expansion rate of the universe, and so improving our understanding of the structure and internal physics is critical if we want to increase the precision of these measurements. Chapter 5 examines a CV system that contains a pulsating WD, and Chapter 7 includes observations of two ELM binary WD systems.

WDs are the end point of evolution through the red giant branch, and provide important constraints on this evolution, including the significant amount of mass that is ejected into the interstellar medium.

Information obtained by modelling convection in WD atmospheres (see Chapter 4) is expected to provide important constraints on the standard ‘mixing length’ models, which has important consequences on the accuracy of spectroscopic models (see e.g. Bergeron et al., 1995).

A small number of WDs have been observed with extremely stable pulsations. The evolutionary cooling of these stars can be detected by the small phase drifts that accumulate over several decades of observation (see Chapter 6). This can be used to constrain the WD cooling rate (important for age estimates), and also to constrain additional sources of cooling such as the emission of plasmon neutrinos (Winget et al., 2004) or axions (e.g. Mukadam et al., 2013). These phase drifts may also contain signatures of orbiting planets that have survived the violent transition of the host star from the main sequence into a white dwarf (e.g. Mullally et al., 2008).

1.4 Thesis Structure

This thesis describes the development of a time-series CCD photometer and its application to the observation of pulsating white dwarfs and other variable stars. Chapters 2 and 3 focus on the development of the two Puoko-nui photometers and the `tsreduce` reduction software. Chapter 4 presents observations that were used to identify the pulsation behaviour of the DBV WD EC 04207-4748, and discusses the effect of convection on the observed lightcurves. Chapter 5 describes the complicated accreting pulsating WD systems and investigates the semi-coherent periodicities that were observed in GW Librae. In Chapter 6, new and archival data on the DAV WD L19-2 are used to place preliminary constraints on the period rate of change of the primary pulsation mode. Chapter 7 presents observations of several targets and discusses their significance without presenting an in-depth analysis. Chapter 8 concludes the main body of the thesis by summarising the key results and some closing remarks.

Each chapter commences with a brief preface that provides a relatively informal introduction to the chapter that follows, including the historical background and motivation for undertaking the work presented in the chapter. The main body of the chapters then proceed as normal in the standard formal tone.

Four appendices provide an instrument manual for the Puoko-nui photometers, a detailed explanation of the firmware logic inside the timer units, a list of finding charts and observing details for several of the targets presented in earlier chapters, and a full observing log.

Table 1.2 summarises the details of the targets that are included in this thesis.

Name	RA	Dec	V-Mag	Type	Motivation	Section
EC 04207-4748	04 ^h 22 ^m 11 ^s	-47°22'19"	14.3	DBV WD	Characterisation, Convection	Chapter 4
GW Librae	15 ^h 19 ^m 55 ^s	-25°00'25"	18.5	Cataclysmic Variable	Characterisation	Chapter 5
L19-2 (MY Aps)	14 ^h 32 ^m 18 ^s	-81°20'09"	13.8	DAV WD	Pulsation stability	Chapter 6
SDSS J075141.18-014120.9	07 ^h 51 ^m 41 ^s	-01°41'21"	17.5	ELM eclipsing binary	Characterisation, Pulsation stability	Section 7.1.1
SDSS J151826.68+065813.2	15 ^h 18 ^m 27 ^s	+06°58'13"	17.6	ELM pulsator	Characterisation, Convection	Section 7.1.2
Crab Pulsar (CM Tau)	05 ^h 34 ^m 32 ^s	+22°00'42"	N/A	Optical pulsar	Instrument test	Section 7.2
EC 05221-4725	05 ^h 23 ^m 31 ^s	-47°22'19"	16.6	DBV WD	Characterisation	Section 7.3
EC 20058-5234 (QU Tel)	20 ^h 09 ^m 40 ^s	-52°25'24"	14.9	DBV WD	Pulsation stability	Section 7.4
PG 1336-018 (NY Vir)	13 ^h 38 ^m 48 ^s	-02°01'49"	13.3	sdB eclipsing binary	Instrument test	Section 7.5
G117-B15A (RY LMi)	09 ^h 24 ^m 15 ^s	+35°15'51"	15.5	DAV WD	Instrument test	Section 7.6

Table 1.2: The targets discussed in this thesis cover a range of topics. Finding charts and offset-guider coordinates are available in Appendix C.

V-Mag represents the apparent stellar magnitude as measured through a V (visual) filter, with every 2.5 units of magnitude representing a factor of 10 decrease in apparent brightness.

Chapter 2

Instrumentation

Instrumentation is a particular interest of mine, and it has formed a significant part of this thesis work. The focus of these efforts were directed towards fixing and improving an existing CCD photometer instrument, called Puoko-nui. This instrument had been (mostly) operational since 2007, but had several major flaws that made it frustrating to operate, and rendered it completely inoperable at times.

I became familiar with the symptoms of these problems during discussions with DJS over the course of my MSc. My MSc work (involving theoretical and numerical aspects of Gravitational Microlensing) was several steps removed from hands-on data collection, which I found somewhat unsatisfying. The idea of a more practical PhD project was quite appealing – working directly on the instrument, plus a generous amount of observing. It was clear from the outset that this would involve a significant amount of electronics and low-level software engineering, which fitted well within my interest and skill sets.

Solving these problems ultimately required an almost complete rebuild of the instrument. This was very successful as it allowed a re-evaluation of earlier design decisions, and resulted in a simpler, more robust and more flexible design which now reliably produces excellent quality photometry of faint variable targets.

This flexibility made it relatively straightforward to integrate a different camera system that was based with the white dwarf group at the University of Texas at Austin (UT). DJS has had close links with the UT group for many years, and earlier generations of instrumentations at VUW have been heavily inspired by their designs. It has therefore been satisfying to share our developments back to them, and this work has resulted in a second instrument using our Puoko-nui design.

This ‘Puoko-nui North’ photometer saw first light in August 2012, and it has since seen regular use with the 0.6 m robotic telescope at Meyer Observatory (PJMO) in

central Texas, and the 0.9 m and 2.1 m telescopes at McDonald Observatory.

Work on this second instrument spurred an additional set of improvements that we had not originally considered for our NZ-based instrument: this included dramatically improving the time resolution (the UT system can now acquire scientifically useful data at rates of up to 200 frames per second) and designing a much simpler user control interface. A new undergraduate physics experiment was also created using the SBIG guide camera and the ‘star simulator’ hardware that was originally developed for testing the Puoko-nui systems.

Instrumentation development is an iterative process, and is rarely done in isolation. For this reason, it is important to build some context before jumping into the details of Puoko-nui. I begin below with a few words on high-speed time-series photometry, and then outline the history of instrumentation development at VUW. This is followed by a discussion of the two Puoko-nui photometers, which expands on the details from Chote et al. (2014). A reasonably comprehensive manual for Puoko-nui South is included as Appendix A, and Appendix B provides a detailed description of the timer unit’s internal firmware logic.

2.1 Time-series Photometry

CCD time-series photometry is not inherently complicated: all one needs is a camera and a sufficiently accurate clock to measure the time at the start and/or end of each integration. The intensity of the target(s) of interest can then be extracted from each frame and plotted versus time (this plot is known as a lightcurve) or frequency (by taking a Fourier transform) in order to reveal systematic variations in the stellar intensity. These intensity changes may be intrinsic to the star (e.g. pulsations) or introduced by external effects such as eclipses. They may be periodic (e.g. pulsations, eclipses) or aperiodic (e.g. outbursts, occultations). The variability timescales span several orders of magnitude: from milliseconds for quickly rotating pulsars and occultations, through to hundreds of days or more for slowly pulsating objects like the Mira variables.

Our primary interest is the group of pulsating white dwarf stars, which have pulsation periods between roughly 100 – 1000 s. These objects are intrinsically faint, but the need for high time-resolution limits the maximum integration length that we can use. Long, uninterrupted observation sequences are required to resolve the often-seen complicated multi-periodic frequency structure, and it can be difficult (but not impossible) to obtain the required amount of time on large telescopes.

The optimum balance between telescope size and availability seems to be a 1 – 2 m

class telescope paired with a photometer instrument that has been specifically designed to maximize the photon collection during high-speed operation. This combination can observe objects as faint as $17^{\text{th}} \sim 19^{\text{th}}$ magnitude using exposures lengths of \sim tens of seconds. From the 1970's \sim 2000's this role was filled by multi-channel photometers, but these have now been made almost entirely obsolete by CCD cameras (which are discussed in more detail in Section 2.3.1).

Other short-period variable stars that require high-speed photometry include hot subdwarf B stars (with periods $10^2 \sim 10^3$ s; e.g. Fontaine et al., 2006), eclipsing short period double-degenerate systems ($10^1 \sim 10^2$ s; e.g. Hermes et al., 2012a) planetary occultations ($\sim 10^{-1}$ s), and optical pulsars ($\sim 10^{-2}$ s; e.g. Cocke et al., 1969).

High-speed photometry has a relatively rare requirement of stable and accurate timing over long continuous periods. Measurements of long period variability (tens of days) could be sufficiently covered by a single observation each night, each with a timing accuracy of \sim minutes. To obtain similar coverage on a fast pulsator (hundreds of seconds), one must sample almost continuously with a timing accuracy of a second or better. Factors such as clock drift and light-travel time become significant at this level of precision, and must be accounted for in the design of the instrument and its reduction pipeline. Furthermore, collaborations such as the Whole Earth Telescope (Nather et al., 1990) have made it routine to coordinate observations from multiple sites, which extends the problem of time synchronization to multiple geographical locations.

The advent of the Global Positioning System (GPS) in the 1990's has largely solved this problem: a modern GPS disciplined clock can provide timing accurate to a microsecond or better from almost anywhere on Earth. The systematic offsets that are introduced by the orbital motion of the Earth can be modelled and corrected to an accuracy of milliseconds or better during the data reduction phase (see Section 3.2.8). The uncertainty in the reduced time measurements can therefore be completely neglected for most of our observations.

2.2 History

2.2.1 Photomultiplier Instrumentation

The origins of the astrophysics programme at VUW go back to the early 1970's, with a group of instrumentally-minded physicists (David Beaglehole, DJS, and Joe Trodahl) who thought it would be interesting to dabble in astronomical observations. The first successful instrument was a photomultiplier tube (PMT) based monochro-

mator (Trodahl et al., 1973) that used a rotating semi-circular interference filter to scan through a fixed wavelength range. This instrument was used for low resolution spectrophotometry studies using the 0.4 m telescope at Carter Observatory in Wellington and the original 0.6 m Optical Craftsman telescope at the University of Canterbury's Mt John Observatory (now Mt John University Observatory; MJUO) at Lake Tekapo.



Figure 2.1: DJS and Mike Clark with the VUW single-channel photometer mounted on the 0.6 m Optical Craftsman telescope at MJUO during the late 1970's.

The first version of this instrument featured a simple motor and microswitches, but these were replaced in a second version of the instrument (Bringans et al., 1974, pictured in Figure 2.1) which used a PDP-8 computer-controlled stepper motor and employed single photon counting techniques with the PMT. Software changes (machine code in a 4k ferrite core memory system!) allowed the instrument to operate as a high-speed single channel photometer, which was used to observe lunar occultations and other phenomena during the late 1970's and early 1980's. During this period Beaglehole and Trodahl moved on to other interests, but DJS continued developing this instrumentation theme.

A significantly improved two-channel photometer (Van der Peet, 1987) was completed in the mid-1980's, and used extensively with the (then) new 1.0 m McLellan telescope at MJUO. The second channel was used to monitor a second (non-variable) star in parallel with the target, and provided a measurement of atmospheric transparency variations, which were a significant problem for the single-channel instrument. These variations could be corrected during subsequent analysis to produce a relative measurement of the intrinsic target variability. The main component of this instrument was a mirror unit which redirected the outer field of view of the

telescope to a second PMT, which could be positioned using a 2D slide mechanism. This ‘offset guider’ unit (as we now call it; see Section 2.3.5) was designed and built by the University of Texas at Austin, and is still in use today.

DJS became involved with the Whole Earth Telescope (WET) collaboration in the early 1990’s. This solidified our NZ connection with the white dwarf research community, and the UT group in particular. This close collaboration led to a further improved three-channel photometer (Sullivan, 2000), which copied some of the innovations developed by the UT group (Kleinman et al., 1996), and saw first light in 1997. The offset-guider box was retained as the basis of the instrument, but it underwent substantial modifications. The single large PMT in the primary light path was replaced by two small PMTs mounted side by side (one for the target, the other for measuring the sky intensity). A dichroic filter was added to the secondary light path, with the blue wavelengths being reflected to a third small PMT, and the red wavelengths passing through to a small CCD camera. The CCD (initially a SBIG ST-4) was integrated with the telescope control system, and automatically adjusted the telescope pointing (‘autoguiding’ the telescope) to keep the target star centred on the primary PMT aperture. The two and three channel photometer systems are shown in Figure 2.2.

The VUW astrophysics group developed an interest in precision timing as part of the programme of observing lunar occultations of stars. This was originally satisfied by building a precision clock using a commercial oven-controlled crystal. The clock could be aligned to radio time signals (VNG from Australia, and WWV from Hawaii) using digital controls and an oscilloscope display, and maintained time to an absolute accuracy of better than a millisecond. The advent of the GPS network made it straightforward to maintain precision timing using a suitable GPS receiver. The VUW astrophysics group has routinely used GPS for instrument timing since the mid 1990’s.

2.2.2 CCD Instrumentation

The Argos CCD photometer (Nather and Mukadam, 2004) was a significant departure from earlier PMT-based instruments, and this instrument pioneered efficient high-speed CCD photometry of pulsating WD stars. DJS had first-hand experience with this instrument during research visits to UT and McDonald Observatory in 2002 and 2003.

The move from PMTs to CCDs offers significant scientific and practical advantages, and the prospect of being able to adopt a copy of the Argos system with minimal changes for use in NZ was very attractive. The original plan was to retrofit the




Figure 2.2: DJS with the VUW two and three channel photometers mounted on the 1.0 m McLellan telescope during the late 1980's and in 2004.

3-channel photometer with the same CCD camera used by Argos, and then adopt the existing timing and instrument control systems that were developed for Argos. Funding was secured to purchase the necessary hardware in 2006, and UT provided the Argos control software and timer dongle.

The physical modifications to the instrument were straightforward, and completed without any trouble: the CCD replaced the two PMTs in the primary light path, and the third PMT was removed from the secondary light path – leaving it as a dedicated channel for the guide camera. The field of view incident on the CCD was increased by using the $f/7.7$ configuration of the 1.0 m telescope (the earlier instruments were operated at $f/13.5$), which increases the number of suitable comparison and guide stars available on the primary and secondary CCDs respectively.

The software side of the instrument was another matter entirely. The Argos software was tied closely to a specific version of the Linux kernel (presumably one of the 2.4 releases), which lacked support for the newer hardware in the acquisition PC that was purchased for the VUW instrument. Initial investigations found that a significant rewrite of the `Quilt` software would have been required in order to fix this, and that it would be more efficient to make a clean break than to try to adapt `Quilt`. We have heard that this rewrite was eventually required for Argos, and that it was indeed a difficult task!

The Argos timer dongle (which attached to a now-legacy parallel port) was seen as a liability for future-proofing, and so this was also abandoned in favour of a separate microcontroller-driven unit that cleanly separated the instrument timing logic away from the acquisition PC.

By mid 2009 work was well underway on this new approach. Johnny McClymont

(the School of Chemical and Physical Sciences’ electronics technician) had designed and built the first microcontroller timer unit, and Mike Reid (a previous PhD student of DJS’) had created acquisition software. This development suffered significantly from being developed part-time, and from the disconnect between those building the system (Johnny and Mike) and those using it (Denis). When I joined the project in late 2010 the instrument, which had been named Puoko-nui by this point, was functional but had several critical problems.

The worst of these issues would see the camera ignore the input trigger, and continuously readout frames every 1.4 seconds. We were unable to identify a specific cause for this issue, and there were no known workarounds except to wait for the problem to resolve itself (which could take anywhere between 5 minutes and several days). We never found an adequate explanation for this problem, but we did eventually find a reliable workaround: inverting the trigger polarity so that the normal line level is high, and pulls low to initiate a frame transfer.

A lot of dialog was exchanged with the manufacturer (Princeton Instruments) over this issue, but it was largely unhelpful – they were not able to reproduce the problem using their proprietary software, and so claimed that the issue must be in our timing hardware or software. We even went as far as sending the camera to them for investigation, but they were unable to identify any faulty components. This led us to the conclusion that this problem was caused by a design flaw in the camera’s trigger detection circuitry.

A second issue would result in corrupted frame data being returned from the (closed source) camera interface library. This manifested in two ways: the first would produce an image that was clearly corrupted (see Figure 2.3); the second would produce correct-looking frame data that was acquired from an earlier exposure. This was much more subtle and dangerous, as it could invisibly corrupt the data by producing photometry with incorrect times, or in some cases runs of repeated intensity measurements (also shown in Figure 2.3).

We eventually isolated the cause of this issue to a flaw in the closed-source camera control library. The issue is triggered by calling the camera API command to halt an exposure sequence while the camera was actively transferring data to the acquisition PC. It appears that a pointer into the frame buffer is reset when the halt command is received, but that the remaining frame data continued to be written after this. The overflow data is then treated as the first bytes in the next exposure sequence, and produces the corruption that is visible in Figure 2.3.

The manufacturer again claimed to be unable to reproduce the problem, so we were left with our own ingenuity to find a workaround. The camera library did not provide

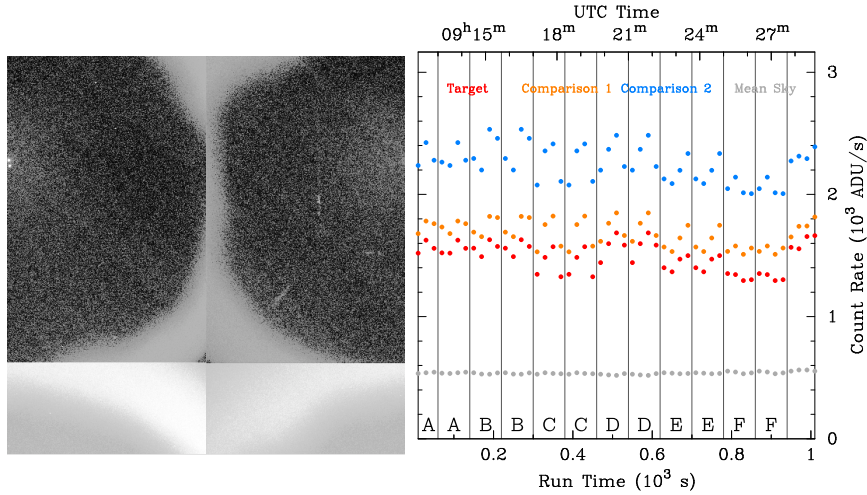


Figure 2.3: An issue with the camera control library could result in corrupted frames or stale frame-data being returned for multiple exposures (see text).

a function for querying the camera status directly, but it did include an option for setting a hardware output signal when the camera is actively reading a frame. The timer unit was modified to monitor this level and relay the status changes to the acquisition PC, which could then wait for a frame readout to complete before issuing the shutdown command.

There were a number of other issues (unrelated to the camera) that led to similar frustrations, but these are not individually interesting enough to mention. The lack of functional online reduction and sub-optimal control interface also made the instrument somewhat frustrating to operate. Most of these were resolved by the end of 2011. By mid-2012 we had an instrument that we could be proud of, and a collection of observations that proved its worth. The state of the instrument at this point was detailed in a poster presentation at the 2012 European White Dwarf Workshop (EUROWD) in Krakow, Poland (Chote and Sullivan, 2013b).

The return route from Poland took me via Texas, where I spent two weeks working with the UT WD group. One of the goals of this trip was to integrate our Puoko-nui timing and control systems with a new ProEM camera system that had been purchased following the accidental destruction of one of the Argos cameras at McDonald. The UT group were at the time investigating options for a successor to Argos (which was still operational, using a backup camera), and we felt that the Puoko-nui design would fit well. We were able to create a working proof of concept instrument during this time, which was installed on the 0.6 m robotic telescope at Meyer Observatory in central Texas before I returned to New Zealand.

This work spurred a second wave of instrumentation developments, working closely with a UT graduate student, Sam Harrold, as well as members of the Central Texas Astronomical Society who operate Meyer Observatory. These changes were targeted

mainly at improving the ProEM instrument, but many of the changes also improved our original NZ Puoko-nui. After some confusion over naming, we eventually settle on Puoko-nui South for our NZ instrument, and Puoko-nui North for the UT ProEM variant. The final results of these instrumentation efforts are described in the following sections.

2.3 The Puoko-nui Photometer System

Puoko-nui is a high-speed CCD photometer system that combines a commercial camera system and GPS receiver with a custom timing unit and open-source instrument control and reduction software. The instrument is optimized for acquiring continuous exposure sequences with integration periods of tens of milliseconds through to hundreds of seconds.

Figure 2.4 gives a schematic overview of the main instrument components. The design is highly modular, with a clean separation between the timing, imaging, reduction, and auto-guiding functions (in the case of Puoko-nui South). This provides significant flexibility towards future hardware changes, either for upgrading the current instruments, or for the creation of new ones.

Two instruments are currently in operation. The original Puoko-nui (South) photometer (Figure 2.5) is based at VUW in Wellington, and used primarily with the 1.0 m McLellan telescope at MJUO. The newer Puoko-nui North system (Figure 2.6) is based at the University of Texas at Austin, and has been used with the 0.6 m robotic telescope at Meyer Observatory in central Texas, and the 0.9 m and 2.1 m telescopes at McDonald Observatory.

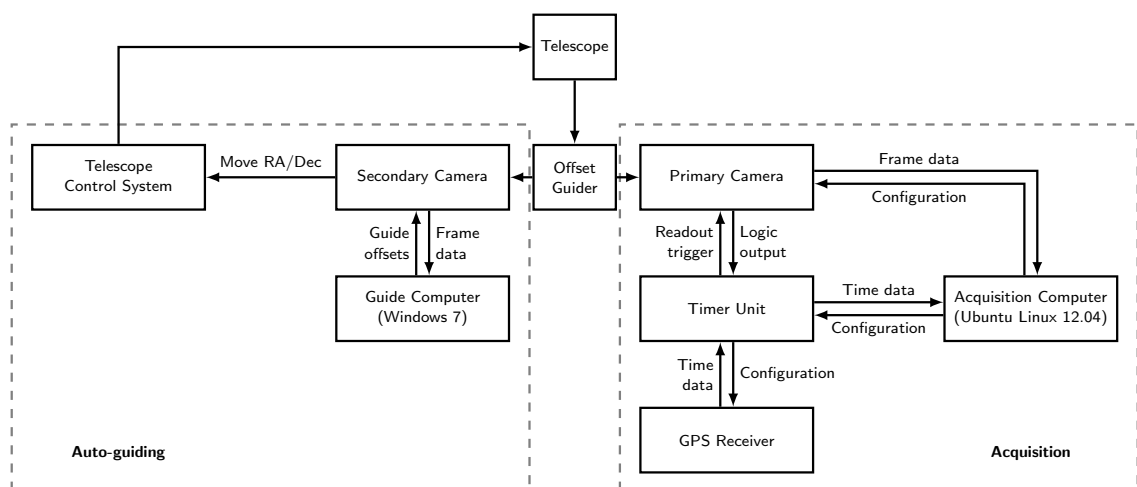


Figure 2.4: A block diagram depicting the main components of Puoko-nui South. The Puoko-nui North instrument mounts the science camera directly on the telescope and auto-guides using the science frames.



Figure 2.5: Puoko-nui South mounted on the 1.0 m McLellan telescope at MJUO.




Figure 2.6: Puoko-nui North mounted on the 0.6m robotic telescope at Meyer Observatory, and the 2.1m Otto-Struve telescope at McDonald Observatory (photo by Sam Harrold). The timer and acquisition PC are situated in the control room outside the dome.

2.3.1 Science Camera

Our original instrument (Puoko-nui South) uses a Princeton Instruments MicroMax camera system, which features a thermoelectrically cooled 1024×1024 px backside-illuminated frame-transfer CCD. The camera hardware is separated into a CCD head and a controller box that includes the digitization and communication electronics. The controller connects to the CCD head using a shielded analogue cable, and to the acquisition PC by USB 2.0 (using a manufacturer-provided upgrade card – the original system used a proprietary serial protocol called **TAXI**). The manufacturer provides a software interface for controlling the camera, which consists of a closed source software library (**PVCAM**) and an open source (but abandoned) USB driver (**rspiusb**). The Puoko-nui software includes a custom version of **rspiusb** that has been modified to support the Linux 3.x kernel versions and to fix incorrect memory accesses that would otherwise crash the USB subsystem when the computer was under heavy load.

The UT system (Puoko-nui North) uses a newer generation ProEM camera, also manufactured by Princeton Instruments. The ProEM camera features a similar 1024×1024 px (EM)CCD, but includes an additional electron multiplying (EM) readout mode that can pre-amplify the integrated signal by as much as $1000\times$ before readout. This decreases the effective readout noise (which is often the dominant noise source), and makes it possible to acquire photon-noise limited images at much shorter exposures. The EM readout mode supports faster readout rates, which reduces the minimum exposure length that can be acquired by the frame-transfer CCD. The ProEM camera includes all of its electronics in one unit, and communicates with the acquisition PC using gigabit ethernet. Proprietary software and drivers expose a similar software API (called **PICAM**) to the earlier generation of cameras.

CCD detectors have revolutionized astrophysical observations since they were introduced during the 1980's, and they are a remarkable example of solid-state physics and engineering. The Charge Coupled Device was invented by Willard Boyle and George E. Smith at Bell Labs as an electronic analogue of magnetic bubble memory. The fundamental principle behind the CCD, which earned Boyle and Smith the 2009 Nobel prize in Physics, is its ability to move charge between discrete wells. This is usually achieved by layering an oxide layer and metal electrodes on a suitably doped silicon substrate to create a collection of MOSFETs (Metal Oxide Semiconductor Field Effect Transistors). The gate voltage applied to the MOSFET can be used to control the depth of the depletion region, which creates a potential well that can be manipulated to trap and/or release the mobile charge carriers. Charge can be moved freely between adjacent wells by raising the potential of one well while lowering the

potential of the adjacent well.

Photons incident on the semiconductor will excite electron-hole pairs, with one component (typically the electron) captured in the depletion region to be read out at the end of the exposure. The simplest CCDs (originally used in devices such as fax machines) are linear devices, and consist of a single column of transistors with readout / digitization circuitry attached at one end. The entire N-transistor column can be digitized in turn by repeating the shuffle / readout process N times.

A two-dimensional detector device can be created by arranging multiple columns to shuffle charge into a perpendicularly oriented ‘readout register’. The readout register can be filled by one shuffle of the imaging area, and read in the same manner as the linear example above. The register is then filled with the next row from the imaging area, and the process repeats until the entire image has been read.

Traditional CCDs have a significant drawback which limits their usefulness for high-speed photometry: the readout process is serial, and can take a significant amount of time if there are a large number of pixels (e.g. it takes ~ 11 s to digitize a 1024×1024 px device at 100 kHz). The primary contribution to the readout time is the delay between the two measurement samples in the Correlated Double Sampling (CDS) procedure that is used to measure the electron count. This time can be reduced, but this comes with a penalty of increased readout noise.

This problem was largely solved with the invention of the frame-transfer CCD. These devices include twice as many pixels as an equivalent regular CCD, but permanently mask half of the pixels from incident light for use as a charge storage buffer. The integrated frame can be quickly transferred into this region at the end of an exposure (taking $1 \sim 10$ ms for a 1024×1024 px CCD), and then digitized in parallel with the next exposure. The charge transfer is fast enough to remove the need for a mechanical shutter, which allows for a completely solid-state operation. Frame transfer CCDs have two drawbacks: they are significantly more expensive than regular CCDs, and the readout time limits the minimum exposure length that can be used.

To underline just how much of an improvement this brings, consider 20 s exposures of a 1024×1024 px image using a 100 kHz digitization rate. A regular CCD would require about 11 s to digitize the CCD, for a total exposure cycle time of 31 s. 35% of the photons incident on the telescope will remain undetected in each exposure (assuming a 100% quantum efficiency). A frame transfer CCD has an exposure cycle time of 20.001 s, with only 0.005% undetected photons. This is equivalent to increasing the telescope diameter from 0.8 m to 1.0 m!

Electron multiplying CCDs (Mackay et al., 2001) include an additional amplification

stage between the readout register and ADC circuitry. This is achieved by clocking charge through dedicated multiplication register using a very large clock voltage, which accelerates the electrons to a high velocity as they move between wells. A small number of the transferred electrons ($\sim 1\%$) will impact-ionize an additional electron from the substrate, which compounds along the length of the register to create a significant overall gain factor (up to $1000\times$ for the ProEM camera in Puoko-nui North). The precise gain factor can be calibrated and configured by adjusting the applied clock voltages.

This gain is applied before the digitization step, and so it effectively reduces the readout noise by the gain factor. The CDS procedure can then be operated at faster rates without appreciably increasing the noise, which decreases the minimum useful exposure time achievable with a frame-transfer EMCCD.

A detailed study (see Mackay et al.) shows that EMCCDs can be operated in one of three regimes: 1) with unity gain, featuring identical properties to a standard CCD; 2) with moderate gain, featuring sub-electron effective readout noise but halved effective quantum efficiency; 3) with high gain, in a 2D photon-counting mode. The ProEM camera in Puoko-nui North can be operated in either of the first two modes.

A further improvement in CCD technology that we take advantage of is backside illumination. Limitations in semiconductor fabrication during the 1970's meant that the CCD had to be oriented with the metal gate layer in the light path, which reduces the detection efficiency – particularly at shorter wavelengths (such as those from hot white dwarfs). Backside illuminated devices have a mechanically thinned substrate layer, which allows them to be oriented with the gate layer at the rear. This dramatically improves the quantum efficiency (the percentage of incident photons that are detected by the CCD). This is illustrated in Figure 2.7, which compares the quantum efficiency of two frame-transfer CCDs options for the PI MicroMax camera system.

There are a small number of parameters that dictate the relationship between the number of incident photons and the measured signal. CCDs are semi-conductor devices, and so they feature thermally excited electron-hole pairs. These thermal electrons produce a roughly uniform background signal (known as the dark current) that increases exponentially with temperature. The dark current generation is a statistical process, which introduces a source of random noise. The CCD digitization procedure also adds a bias voltage to each measurement, which adds an additional uniform background signal and an additional source of noise. The relationship between the measured number of electrons and the arbitrary output units (known as Analogue to Digital units - ADU) is determined by the CCD gain. Table 2.1 and

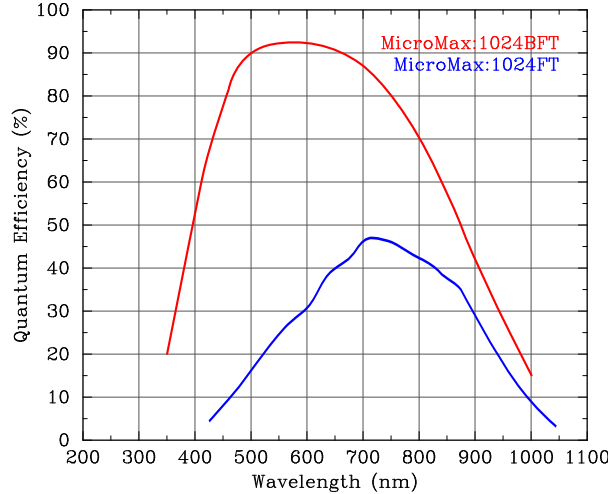


Figure 2.7: The backside illuminated CCD in Puoko-nui South (red) has a detection efficiency above 90% for photons between 500 – 650 nm, and is 5 – 10 times more sensitive at short wavelengths than the equivalent frontside illuminated CCD (blue). QE data was extracted from the MicroMax camera specifications.

Setting	Gain	Read Noise	Read Noise
	e^- / ADU	ADU	e^-
1 MHz, Low Gain	4.4	2.6	11.6
1 MHz, Med Gain	2.2	4.4	9.9
1 MHz, High Gain	1.0	7.6	7.9
100 kHz, Low Gain	4.5	1.6	6.9
100 kHz, Med Gain	2.0	3.3	6.7
100 kHz, High Gain	1.0	5.0	5.0

Table 2.1: Measured CCD properties for the MicroMax camera system. The uncertainty in these figures is half the final significant figure.

Figures 2.8 – 2.10 show some of these measured properties as measured for our MicroMax Camera.

The PI cameras we use are cooled using multi-stage thermo-electric coolers. The Puoko-nui South CCD is factory-configured with a temperature set-point of -40°C . Our standard acquisition configuration lowers this to -50°C for a significant reduction in dark current and thermal noise (see Figures 2.8 and 2.9). We believe this will have no meaningful impact in the lifetime of the cooling system as the default temperature will have been chosen for operation in a room-temperature environment. The temperature in the dome is usually below 10°C while observing, so the total cooling power should remain within the manufactured limits.

Extrapolating Figures 2.8 and 2.9 to -90°C (roughly what liquid nitrogen cooled cameras tend to be operated at) shows that the dark current and thermal noise will be completely negligible compared to the bias level and readout noise. The convenience of thermo-electric cooling offsets the significant, but not huge, improvement




Figure 2.8: Lab-based measurements of the CCD dark current show an approximately exponential increase with temperature, with the dark current doubling every $\sim 6^\circ\text{C}$.

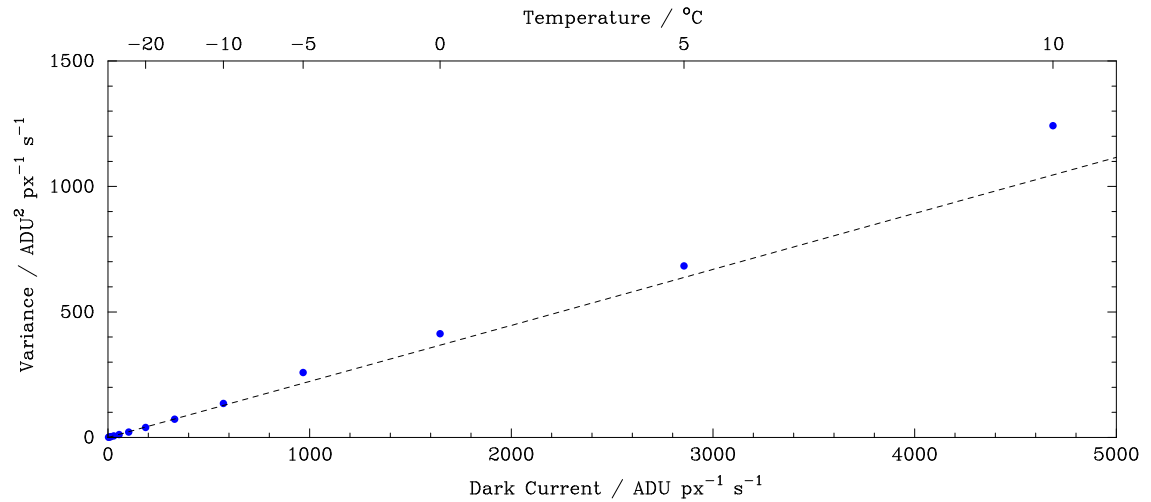


Figure 2.9: Lab-based measurements of the CCD dark frame variance approximately follow the expected linear relationship between the variance and mean dark signal. The proportionality constant is given by the inverse CCD gain (4.42 e^- for these low-gain measurements), which is indicated with the dashed line. The deviation away from this line at temperatures above -10°C is of no practical concern, as the high dark current and variance at these warm temperatures preclude useful operation.




Figure 2.10: A plot of image variance over a range of uniform intensities is known as a transfer curve, and can be used to measure the CCD linearity and gain (see Section 3.2.1). We did not have a light source suitable for measuring this under controlled conditions, but a large set of sky flat frames (690 images in total) provides a useful substitute. The main trend produces a gain measurement of $2.01 \text{ e}^- / \text{ADU}$ for the medium gain setpoint, and shows that the CCD is linear for pixel intensities up to at least 44 kADU. The scatter toward larger variances is caused by star images and/or streaks which are visible in the individual frames.

in data that would come with liquid nitrogen cooling.

The MicroMax camera has a single readout port which can be operated at 100 kHz or 1 MHz, each with three gain options. Faster readout rates have greater noise, so we use the 100 kHz rate whenever it is practical. The slower digitization time at 100 kHz limits the minimum exposure time to 10.75 s (3.2 s at 2×2 binning) vs 1.2 s (0.5 s at 2×2 binning) for the 1 MHz configuration.

Figure 2.11 shows the improvement that can be obtained by changing the acquisition parameters. Two acquisition runs were made using the star-simulator lightbox environment (see Section 2.4) under identical conditions except for the specified camera parameters. The blue points show the default camera parameters that were used for early Puoko-nui observing runs, while the red show the improved camera parameters that are now used.

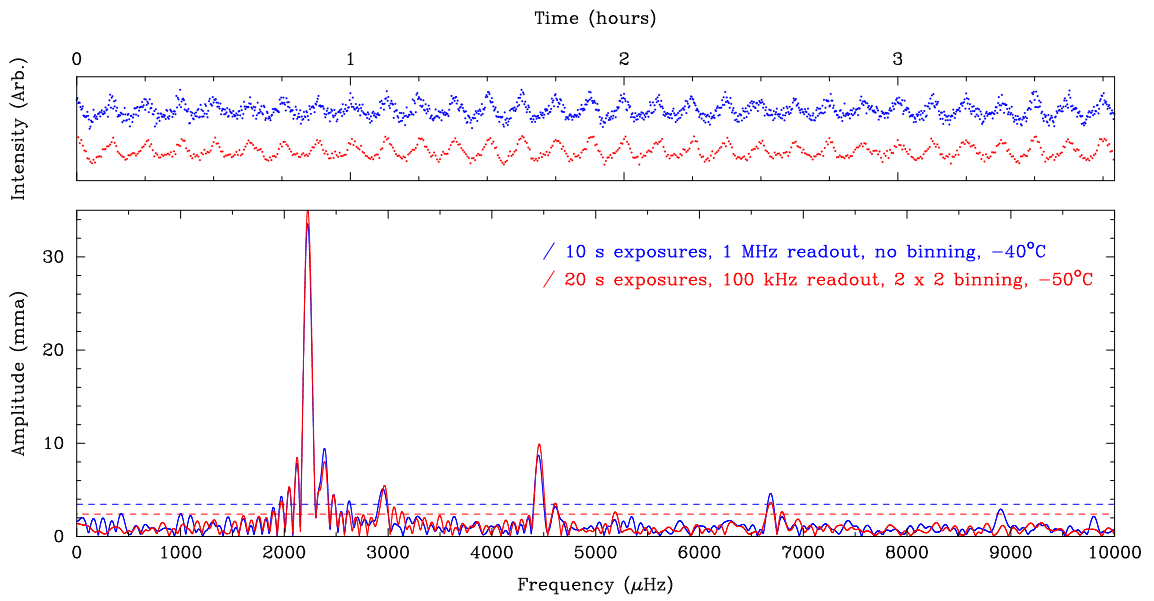


Figure 2.11: A comparison of lightbox observations for a simulated EC 04207 target demonstrates how a careful choice of acquisition parameters can dramatically improve the resulting photometry. Our standard observing configuration (red) is visibly less noisy than the default camera configuration (blue). The 0.1% false-alarm probability threshold (dotted lines; see Section 3.2.9) improves from 3.5 mma to 2.4 mma.

2.3.2 Timing Hardware

Precise exposure timing is provided by a custom timer/counter unit based around an Atmel AVR microcontroller. The unit communicates with the acquisition PC via USB, which offers significant flexibility over the PCI bus (as used by the Agile photometer (Mukadam et al., 2011a)), or the now-legacy parallel-port used by Argos. USB 2.0 is ubiquitous, and allows the instrument to be operated using a small-form-factor ‘netbook’ PC or laptop when paired with a USB or ethernet connected camera.

The timer, shown in Figure 2.12, connects to a GPS receiver (currently supporting both a Trimble Thunderbolt and a legacy Magellan OEM receiver) via a serial RS232 connection and coaxial cables carrying 1 Hz and (optional) 10 MHz pulse trains. The RS232 output provides absolute time information each second via an ASCII-coded string, and both coaxial signals provide edge-aligned pulses disciplined to GPS time. Coaxial cables to the camera provide an output for frame-transfer triggers and an input for status monitoring. A display on the front panel shows the current time and acquisition status.

All inputs contain high-impedance buffers and all outputs contain high-current drivers. This mitigates noise induced in the long cable runs. Electrical noise used to be a significant problem at MJUO, where the 1.0 m dome was driven using using a powerful AC motor and a rope drive. The dome motor (and also the floor hydraulics) generated a significant amount of noise, which could trigger spurious CCD readouts or artificially advance the timer pulse count. Adding the buffers and drivers largely eliminated this problem. The old dome motor was replaced with three direct-drive motors in mid-2013, and we expect that these will generate much less noise.

The timer firmware logic is written in the C programming language using the open source `avr-gcc` toolchain, and can be upgraded remotely using the unit’s USB connection. A thorough description of the firmware logic is included as Appendix B.

Two timing modes are supported, and are summarized in Table 2.2. Both modes operate by counting pulses on either the 10 MHz pulse input (high resolution mode) or the 1 Hz pulse input (low resolution mode), and trigger the camera after the desired exposure interval has elapsed. This limits the minimum trigger period to 1 ms and 1 s respectively. In practice, the shortest exposure is limited by the camera readout time, or more commonly by the need to collect sufficient photons to obtain a reasonable signal-to-noise ratio.

An alternative approach to timing, which is used in instruments like ULTRACAM (Dhillon et al., 2007), runs the camera asynchronously (using its own internal timing), and then records the GPS-determined time at the start and/or end of each



Figure 2.12: The Puoko-nui timer unit provides precise and flexible timing control. The unit connects to a GPS receiver (currently supporting both a Trimble Thunderbolt and a legacy Magellan receiver) via a serial RS232 connection and coaxial cables carrying 1 Hz and (optional) 10 MHz pulse trains. The RS232 output provides absolute time information each second via an ASCII-coded string, and both coaxial signals provide edge-aligned pulses synchronised to GPS time. Coaxial cables to the camera provide an output for frame-transfer triggers and an input for status monitoring. A selector switch allows the microcontroller clock to be switched between the 10 MHz pulse train and an internal oscillator.

Mode	High-res.	Low-res.
Exp. range	1 ms - 65.5 s	1 s - 18.2 h
Exp. resolution	1 ms	1 s
Trigger delay	$< 1 \mu\text{s}$	18 ms
Trigger instability	$< 10 \mu\text{s}$	$< 10 \mu\text{s}$
Required inputs	1 Hz, 10 MHz	1 Hz

Table 2.2: The timer unit features two operating modes. The high-resolution mode allows exposures as short as 1 ms, but in practice the camera digitization speed limits exposures to > 5 ms and telescope-size currently limits useful white-dwarf observations to > 1 s for the NZ-based instrument.

exposure. This simplifies some aspects of the instrument (there is no need to duplicate the triggering logic externally), but complicates others (the instrument must quickly record the time when an exposure finishes). There should be little practical difference in the timing precision of the two approaches, but we consider our approach to be more elegant: it cleanly separates the timing, camera, and storage functionality, and the time stamps recorded in each frame remain aligned to the UTC second to better than millisecond precision.

The key concept behind Puoko-nui’s high-resolution timing mode is that the microcontroller clock is driven directly from the GPS-disciplined 10 MHz signal. This guarantees that each clock cycle is 100 ns long, and means that we can count clock cycles to measure time at a much greater resolution than the 1 Hz rate that Argos and the original Puoko-nui implementation were limited to.

This mode is also significantly more robust against glitches in the GPS-clock signals. The 1 Hz pulse train is used as a consistency check on the integrated pulse count: we expect an integrated count of exactly 1000 milliseconds to have elapsed between each 1 Hz pulse. If this check fails, then we know that something has gone wrong (e.g. an intermittent cable fault is causing pulses to be lost), and a warning is raised that the observer can react to. The high-resolution mode can also correct for up to 65 consecutive missing serial time stamps. This is useful, as the Trimble receiver has been observed to miss time stamps at a rate of one or two per night.

The trigger delay figure in Table 2.2 was measured by comparing the 1 Hz GPS pulse train against the generated trigger pulses using an oscilloscope. The camera response time between the trigger and frame transfer beginning was similarly measured at $< 1 \mu\text{s}$. In the case of the high-resolution mode, the phase of the trigger pulse can be advanced or delayed in order to minimize the difference against the UTC-second aligned 1 Hz input. We have explicitly verified the absolute timing accuracy of Puoko-nui South to better than 20 ms by imaging the time display of the older three-channel photometer master clock. Observations of the Crab Pulsar using

Puoko-nui North (see Chapter 7) further verifies the timing accuracy to within a few ms, limited by the numerical precision of `tsreduce`. The trigger instability figure was calculated based on the possibility of a hardware interrupt being delayed while another is serviced, which may introduce a delay of up to 100 clock cycles ($10\ \mu\text{s}$). This situation is rare, and the measured jitter usually remains below $1\ \mu\text{s}$.

The main downside of the high-resolution mode is that it forces the microcontroller to be operated at 10 MHz, which is inconvenient for serial communication. The fastest baud rate that can be derived (with an acceptably low error rate) from a 10 MHz clock rate is 9600 baud. This is fast enough for ‘normal’ operation with many-second long exposures, but provides insufficient bandwidth for transferring trigger time stamps to the acquisition PC for exposures shorter than 100 ms. Shorter exposures are supported by sending only 1 in N time stamps to the PC, where N is chosen to keep the data transfer similar to 0.5 s exposures. The intermediate time stamp values are synthesized by the acquisition PC. One convenient side-effect is that the 9600 baud rate matches the default configuration of the Trimble and Magellan GPS receivers. This simplifies the operation of the unit’s relay mode.

The relay mode provides direct two-way communication between the GPS receiver and the acquisition PC. The timer will also send a readout trigger to the camera for each pulse received on the 1 Hz input line when in this mode. This mode is useful for running the proprietary GPS software without switching cables, and allows the ProEM camera to be operated using the proprietary Lightfield acquisition software (supplied by PI) in a rudimentary GPS-triggered mode.

A secondary function of the timer’s standard mode is to monitor the camera readout status via a logic output. This is necessary to work around a bug in the closed source camera firmware and/or the PVCAM library for the MicroMax camera: terminating an exposure sequence mid-readout can corrupt the internal state of the ‘black-box’ software and cause subsequent exposure sequences to be corrupted. The software API does not allow the camera status to be queried, and so a hardware work-around was implemented using the timer to monitor the programmable logic output on the camera and report state changes back to the acquisition software.

2.3.3 Control Computers

Acquisition and auto-guiding tasks for the NZ system are run on two small net-top computers (ASUS eeeBox EB1012). These computers are situated on an equipment trolley in the dome, along with the timer and GPS receiver. The main advantage of this type of computer is their extreme portability (they are only $222\text{ mm} \times 178\text{ mm} \times 26.9\text{ mm}$, about the size of a double-width DVD case). These

computers have relatively low performance by modern standards, but offer enough computing power for the purposes of instrument control and online reduction.

The computers are controlled remotely from the telescope control room (or anywhere else in the observatory or world) using the VNC screen sharing protocol.

The main control machine (**pico**) runs the acquisition software and online reduction on the Ubuntu 12.04 operating system. Data is saved to the internal hard drive, and then copied to network storage at VUW at the end of each night.

The auto-guiding computer (**nano**) runs the SBIG acquisition software (CCDOPS 5) on the Windows 7 operating system. This machine also has a second partition duplicating the acquisition PC, allowing a quick fallback for acquiring without auto-guiding if the acquisition computer were to unexpectedly fail during a run.

The UT system uses a single control computer (**ARGOS-DEV**), which is a powerful custom-built tower running Windows 7.

2.3.4 Acquisition Software

The acquisition and control software (also called **puokonui**) is written primarily in C, with small amounts of C++ and bash scripting. It is portable across 32- and 64-bit versions of Linux, Windows and Mac OS X, but the availability of usable camera drivers restricts operation to 32-bit Linux for Puoko-nui South, and 64-bit Windows 7 for Puoko-nui North.

The acquisition software performs three main tasks: configuring acquisition parameters, matching frame data with time stamps during acquisition, and notifying the user of the current hardware status. Each of the major components of the software (shown in Figure 2.13) runs in its own thread, and communicates with other components by passing messages asynchronously. This ensures that unexpected delays in one system (e.g. saving a frame to disk) won't impact other systems.

The user interface, shown in Figure 2.14, is designed to be as simple and uncluttered as possible. The configuration options are separated into two groups: those that can be changed while an acquisition is in progress (frame metadata and output parameters), and those that require the acquisition be halted (camera and timer parameters). Both sets of options are shown in Figure 2.15. Appendix A includes a list of advanced options that aren't exposed in the user interface: these options aren't useful during standard operation, but are provided for exceptional circumstances that may require them.

The run configuration parameters are sent to the camera and timer when the acqui-

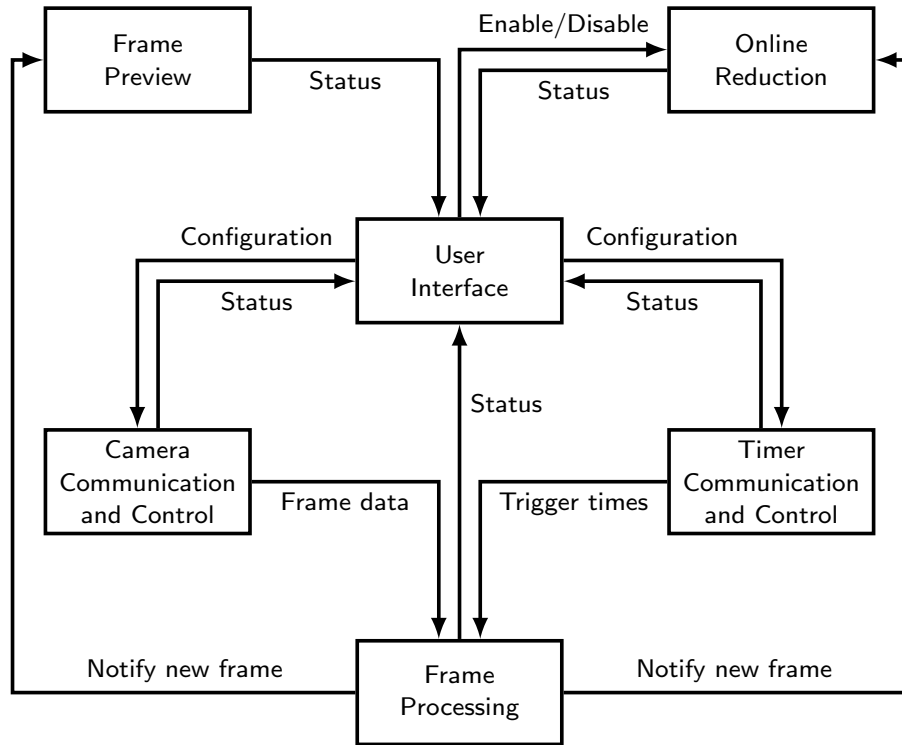


Figure 2.13: Block diagram of the Puoko-nui acquisition software threads.

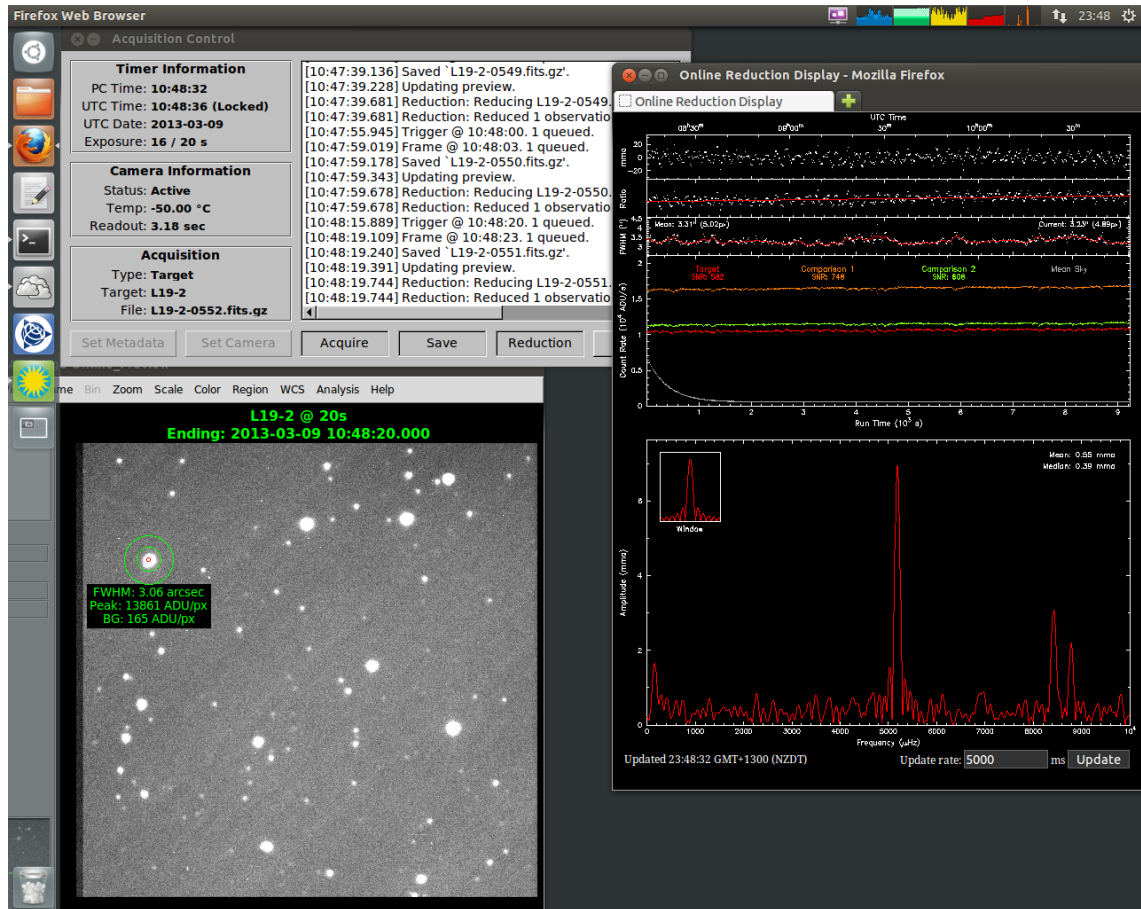


Figure 2.14: A screenshot of the Puoko-nui acquisition and online reduction software during an observation run.

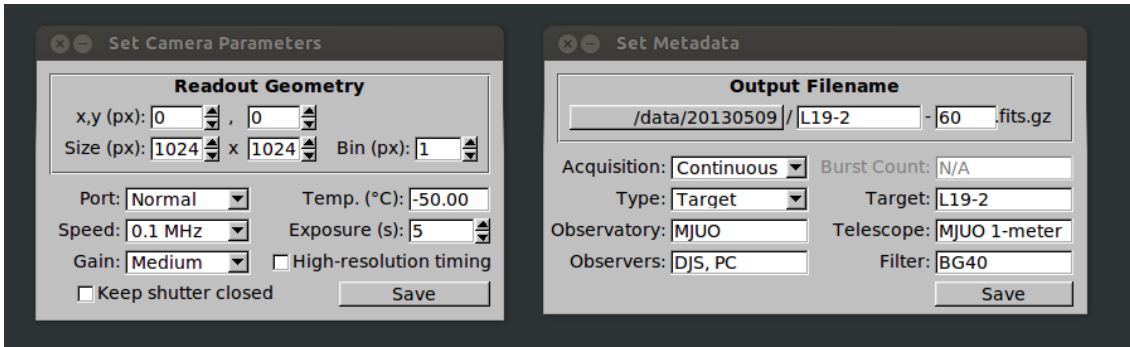


Figure 2.15: The instrument configuration is split into two groups: run configuration parameters that can only be changed when the instrument is idle, and frame metadata that can be changed mid-acquisition.

sition begins, and the camera thread will begin polling the camera driver for new frames (the timer thread continuously polls for data, even when idle). Each frame is tagged with the current GPS time as soon as it is received, before being passed to the frame processing thread. The timer thread also forwards trigger time stamps to this thread. The processing thread waits until at least one frame and time stamp is available, and then performs a simple consistency check: the trigger time plus readout time should match the time that the frame was received by the camera thread. A generous tolerance of 1.5s is included to avoid false-positives that would otherwise be triggered by delays from the asynchronous operation. Most observations have exposures > 5 s, so this tolerance remains short enough to catch invalid time stamp matches from occurring. The frame and trigger queues will be discarded if this check fails, and a warning is displayed in the observation log.

Successfully matched frames are written to disk as compressed FITS images, with the time stamp and other metadata included in the file headers. The preview and reduction threads are then notified that a new frame is available.

The reduction thread invokes a bash script (`reduction.sh`), which runs `tsreduce` to update the online display. This design was chosen to cleanly separate the `puokonui` and `tsreduce` programs, and to allow other reduction steps or packages to be cleanly integrated into the system.

The preview thread invokes another bash script (`preview.sh`), which runs `tsreduce` to update the `ds9` frame preview and (optionally) overlay information about the intensity and FWHM of selected stars. It can also optionally output star coordinates to a file that is read by the auto-guiding software for the 2.1 m telescope at McDonald observatory (see Appendix A for more information).




Figure 2.16: A schematic diagram of the offset guider box. The red arrows indicate photon paths that pass through to the science camera, while the blue arrows are redirected out the side of the box. The guide camera can be positioned in alignment with a suitably bright star for guiding the telescope.

2.3.5 Autoguiding

The auto-guiding functionality in Puoko-nui South is inherited from the earlier VUW three-channel photometer. The only modification from the earlier instrument was the replacement of the previous control computer (a laptop running Windows 2000, situated in the control room via a long USB cable) with the eeeBox situated in the dome and controlled remotely via VNC.

The offset-guider box forms the main body of the instrument, as it did for the two and three channel photometers that preceded it. The box contains a square mirror mounted at 45° , with a central hole that allows light from the central field of view to pass unimpeded to the science camera. The outer field of view is redirected by 90° out the side of the instrument, where the guide camera is mounted on a 2D slide mechanism. The mirror can be moved into a second position, where it redirects the central field by 90° and blocks light from entering the science camera. This is useful for inspecting the field of view visually, using a wide field eyepiece (adapted from a military surplus tank periscope) in place of the guide camera. It is also useful for baffling external light when acquiring dark frames for the (shutterless) main camera. A schematic view of the offset-guider arrangement is illustrated in Figure 2.16.

The auto-guiding functionality is provided by a SBIG ST402ME CCD camera and its bundled CCDOPS 5 software. The camera includes a RJ11 socket with four output pins which are switched high by the guide software to signal that motion is required in each of \pm RA or Dec. A microcontroller unit (provided by MJUO) translates these signals into the serial string commands that are used by the telescope control system.

The field of view of the guide camera is relatively small (3×2 arcmin on the MJUO 1.0 m telescope), but this is more than sufficient for tracking a single bright star. The slide mechanism allows the camera to be positioned within a much larger area: defined by an inner radius of approximately 6.5 arcmin (limited by the hole in the mirror), and an outer square of approximately 17 arcmin (limited by the slide mechanism). The image quality degrades significantly at larger off-axis radii due to the telescope optical configuration, and so it is best to choose a bright star with an offset of ~ 10 arcmin from the centre of the main field. The autoguiding algorithms operate by tracking the image centroid, and so the moderate distortions within this annulus do not have a major impact on the guiding capabilities.

Appendix A includes some notes on calculating the offset guider coordinates from a finding chart of the target field.

2.4 Lightbox and Star Simulator

Most of the instrument development was done in the lab at VUW. This was supported by simulating a realistic observing environment using a lightbox with LEDs positioned at the camera's focal plane. A microcontroller unit was developed to drive up to four LEDs with arbitrary intensity characteristics, which allows for accurate simulations of pulsating WDs other arbitrary signals with periods of \sim seconds or greater. The LEDs can be driven at higher frequencies using a standard signal function generator. Figure 2.17 shows a photo of this test environment.

Each of the four output channels in the star simulator features a software-selectable constant-current source (5, 50, 500 μ A, 5 mA) which is rapidly switched by a pulse-width-modulated (PWM) signal to provide fine-grained control of the LED intensity. The PWM and current-select signals are generated by an AVR ATmega128 microcontroller, which can be programmed using an onboard USB port.

Figures 2.18 and 2.19 demonstrate the capabilities of this bench testing environment. Having access to a controlled and repeatable light source is invaluable for testing different observing configurations, as was shown earlier in Figure 2.11.




Figure 2.17: The Puoko-nui test station at VUW, with the star simulator and lightbox.




Figure 2.18: The bench-testing environment can produce intensity variations that are virtually indistinguishable from real observations. A comparison is shown between photometry acquired on the DBV pulsator EC 04207-4748 (see Chapter 4) from MJUO (bottom panel) and photometry acquired in the lightbox driven by the star simulator (top panel).




Figure 2.19: The high-speed capabilities of the NZ Puoko-nui were tested by driving the lightbox with a standard signal generator. A sequence of 57 ms frames were acquired of a nominally 6 Hz sine wave. The raw time series (top-left panel) is not well enough sampled to reveal the oscillation shape, but the DFT (bottom-left panel) clearly reveals that it is periodic. Folding the timeseries at the period obtained from the DFT clearly recovers the oscillation shape, which is a truncated sine-wave caused by the LED switching off when the driving voltage falls below the minimum required forward-voltage.

Chapter 3

Data Reduction

The process of transforming raw data from a detector into a form suitable for analysis can be split into two functional blocks: (a) transforming the raw data into a series of individual intensity measurements, and (b) processing these intensities to remove systematic effects (due to e.g. atmospheric transparency changes) and assemble a list of normalized intensity measurements as a function of a consistent time standard (e.g. barycentric Julian day).

Part (a) was simple in the era of photo-multiplier tube instrumentation: the instrument outputs were direct intensity measurements for the target, a comparison star, and (on a three-channel instrument) the sky intensity. The advent of the CCD detector significantly improved the practical aspects of observing, but came with the (relatively minor) trade-off of requiring a more complicated reduction procedure.

Three main techniques (discussed below) exist for extracting stellar intensities from a CCD frame. Of these, synthetic aperture photometry is generally considered the best technique for our pulsating WD targets, which are typically faint and seen in relatively sparse stellar fields (a selection effect due to their faintness).

As one might expect, software does exist for performing synthetic aperture photometry, but these tools did not completely satisfy the requirements for this project. I would like to highlight three tools in particular: the automatic reduction built into the **Quilt** acquisition software for the Argos CCD photometer, **ccd_hsp** which is a collection of **IRAF** scripts developed by Antonio Kanaan (an ex-UT grad student) and used (at the time) widely by the WET collaboration, and the new (at the time) reduction software **MAESTRO** developed by James Dalessio (a grad student at the University of Delaware) on top of the **MATLAB** runtime environment.

As discussed in the last chapter, the original plan was to adopt the Argos system for

Puoko-nui, which featured online reductions as part of its control software *Quilt*. This software was tightly coupled to the rest of the Argos system, and so we were hesitant to try and extract this one component. We were therefore faced with the task of finding or building our own online reduction solution.

The `ccd.hsp` scripts were originally developed using the facilities provided by IRAF as a stop-gap offline reduction procedure for the WET until a more permanent solution could be found. It appears to have worked well enough to become the de facto standard (of the time) among many pulsating white-dwarf observers for long after its originally intended shelf-life. IRAF is a well established software project for astronomical data reduction, having initially started in the late 1980's. It features a wide variety of tools and routines for processing CCD images, flux calibration, and spectroscopic tasks. It includes an integrated scripting language, which was used to create the streamlined `ccd.hsp` routines. Unfortunately, the long history of IRAF has left it with a lot of legacy baggage, and we unable to successfully install or configure this software on our acquisition PC.

The MAESTRO software was still under development in 2010, and still had a number of rough edges. Despite this, it appeared to work nicely (and has improved with time) in the role it was designed to fill. Unfortunately, this role was not online reductions. MAESTRO is designed to operate in a batch mode for offline processing, requiring input files listing calibration and data frames, and then stepping through a number of aperture sizes and generating a (nominally) optimum reduction for the run. This reduction process takes several minutes, and updating the reduction with new data as it comes in required the full calculation to be repeated from scratch. While it would have been possible to write a script front-end to use this tool to periodically update the reduction (say once every 10 minutes), we felt that this would have been quite unsatisfactory – quick (ideally real-time) feedback is an important part of observing.

Given that my skills and interests were already aligned in this direction, we decided that it would be a better approach to create a ‘simple’ reduction procedure in the vein of *Quilt* which could generate a ‘good enough’ reduction in real-time to provide instant observer feedback. Our original intention was then to use an offline tool like MAESTRO for our final reductions. The original version of `tsreduce` was built in a few days using Python, and later rewritten in C to improve performance.

The original versions of the software would output data in a format compatible with DJS's existing photometric reduction pipeline (part (b) above). This worked well enough originally, but this pipeline was not designed for automatic operation: it required manual text input for certain commands, and had multiple intermediate data formats. This procedure was frustrating, and so these features were eventually

incorporated into **tsreduce** to simplify operation and maintain a single data format for our online display.

Over time, **tsreduce** accreted additional functionality and improvements until we found the quality of its reductions to be essentially indistinguishable from **MAESTRO**. At this point we decided to adopt **tsreduce** for our offline reductions, which gave us a simple and uniform data reduction pipeline.

The **tsreduce** software has grown from a simple synthetic aperture photometry script to a robust and flexible set of routines for online and offline reductions, and common analysis tasks. These routines include:

1. Online (real-time) calculations:
 - (a) Extraction of raw photometry and noise estimates from CCD frames with up to millisecond time stamp resolution.
 - (b) Normalization of target and comparison lightcurves, and polynomial fitting to isolate target pulsations in milli modulation-intensity (mmi).
 - (c) FWHM (seeing) estimates.
 - (d) Calculation of the discrete Fourier transform (DFT) and window function.
 - (e) Real-time plots of the above information.
2. Offline (after the fact) calculations:
 - (a) Generation of lightcurves and noise estimates over a range of aperture sizes to find the optimum size.
 - (b) Combination of multiple observation files and conversion to a uniform BJD_{TDB} (Barycentric Julian Day, measured using Barycentric Dynamical Time) time basis.
 - (c) Generation of DFT and window data files suitable for plotting.
 - (d) Identification and pre-whitening of peaks in the DFT.
 - (e) Generation of model light curves from identified peak frequencies.

3. Miscellaneous tasks:

- (a) Listing FITS header keys in a saved CCD frame.
- (b) Displaying the motion of stars over the CCD frame during a run.
- (c) Conversion of UTC time stamps to BJD_{TDB} .
- (d) Playback of a previously acquired run as a high-speed animation.

Instructions for installing and using **tsreduce** are included as part of the Puoko-nui manual in Appendix A.

3.1 Photometry Methods

A typical observing session will generate between a few hundred and a few thousand frames per object per night. Each of these frames must be calibrated to remove systematic instrumental effects, the stars of interest must be identified, and then integrated to extract a set of intensities to display to the observer.

Most photometry algorithms can be classified as one of three basic types: aperture photometry, profile fitting, or difference imaging.

Aperture photometry is the simplest technique, and it works well with the moderately bright stars in sparse fields where the specific advantages of the other techniques aren't needed. The algorithm simulates in software what one would do with a physical aperture: a circle is positioned over the star of interest, and the intensity of the pixels within the circle are added. The size and position of the aperture affects the quality of the data obtained, and there are various techniques (discussed in the following sections) to reduce noise in the result.

Profile fitting is useful for crowded fields where the target star image may not be uniquely resolved from nearby stars. A Gaussian (or Gaussian-like) profile is fitted to the point spread function of each stellar image, and the intensity is then evaluated analytically. This reduces the inherent per-pixel noise that affects aperture photometry, and allows stars with overlapping images to be individually resolved. This technique is seen by some as overly complex for observations of white dwarfs, but the results of the profile fitting used for the FWHM estimates in `tsreduce` look promising. It would be an interesting experiment to re-evaluate this technique using a modern profile-fitting code and data from a modern CCD instrument such as Puoko-nui.

Difference imaging is the most complex technique, and is most effective when searching for variable stars in crowded fields. It is absolutely indispensable in areas such as gravitational microlensing. A good quality frame is chosen as a reference, and then compared with each of the other frames in the acquisition sequence. A numerical kernel is calculated which encodes the seeing and intensity changes between the frames, and this is used to convert the two images into a common image space where the difference can be meaningfully calculated. This difference image can then be searched for stars that have changed between frames, and the delta-flux integrated to find the relative intensity change.

3.2 Reduction Procedure

3.2.1 Calibration Frames

The first step in the reduction process (and independent from the photometry technique chosen) is to characterize the systematic instrumental effects that impact each image equally. This can be done by acquiring two sets of calibration frames (typically at the start of the run) known colloquially as ‘darks’ and ‘flats’. Dark frames are typically made with the same exposure time as the science frames, and characterize the thermal noise of the CCD. A *master dark* frame can be created from a set of darks, and then subtracted from each science frame to remove this static background signal. Flat frames capture a uniformly lit background (either the twilight sky, or a screen inside the telescope dome), and characterize the combined instrument sensitivity as a function of CCD position. Each science frame can be divided through by a *master flat* frame to compensate for this variation, producing a uniform intensity measurement across the frame.

Figures 3.1 and 3.2 illustrate this procedure, showing example master frames and a science frame before and after calibration.

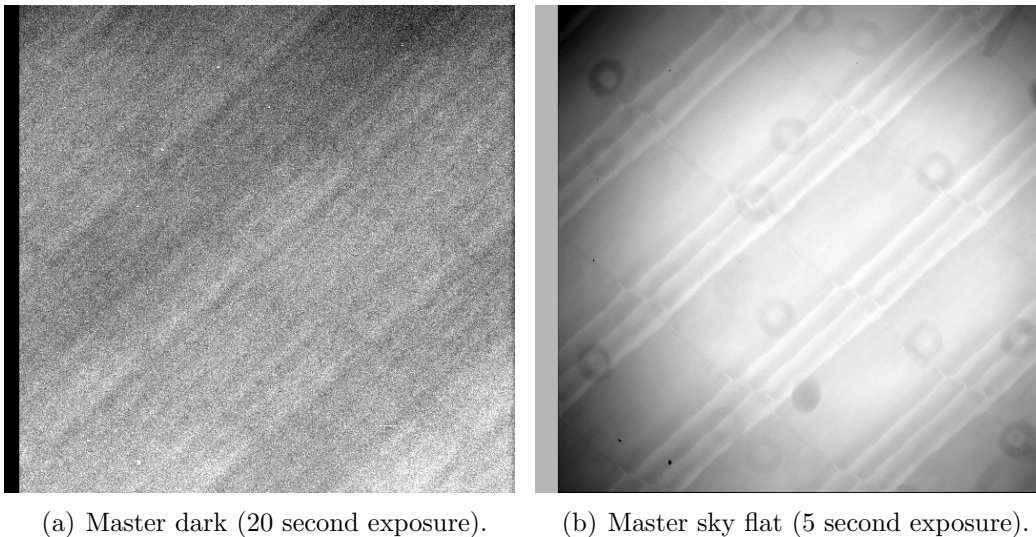


Figure 3.1: Puoko-nui calibration frames acquired on 2013-05-13 using the MJUO 1 m telescope. The structure of the master flat is a combination of three main effects: vignetting by the telescope optics, defocused images of dust particles on the optical surfaces, and sensitivity changes caused by the physical CCD structure.

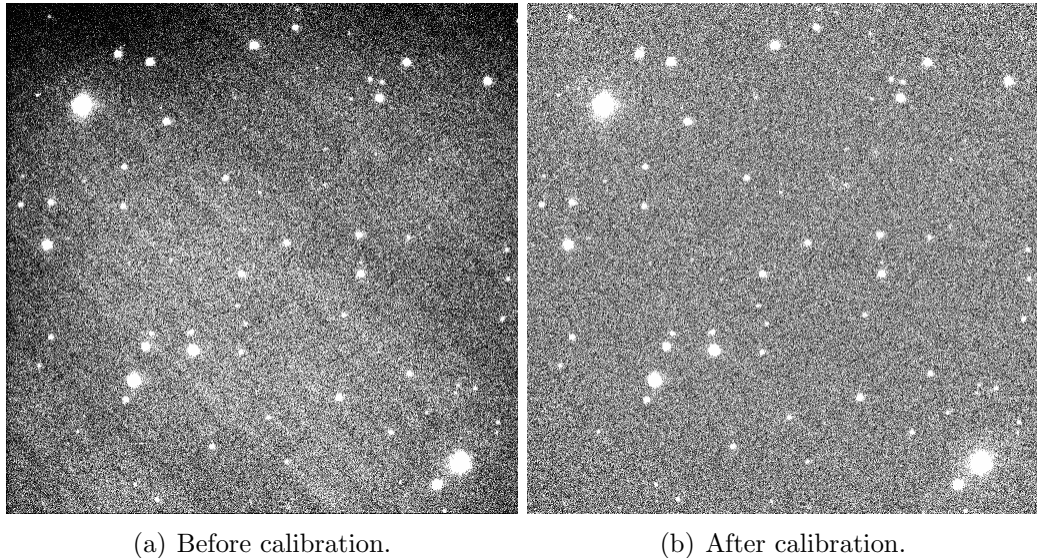


Figure 3.2: A science frame of GW Librae from 2011-08-01 before and after calibration. The dynamic range of the frames have been reduced in order to highlight the fixed structure that is visible in the background of the unreduced frame. This structure is removed by the calibration procedure, leaving only Gaussian noise.

Dark and bias subtraction

Recalling the discussion about CCDs in Chapter 2, we know that each CCD pixel will have a readout contribution due to dark current and bias. The bias is applied during readout, so will have no significant pixel-to-pixel variation. The dark current will increase linearly with distance from the readout port, as the more distant pixels will spend more time on the CCD (and so accumulate more thermal electrons) before digitization. This effect is relatively minor in practice, as in most situations the readout time (3.2 s using our standard configuration) is much less than the exposure time (typically tens of seconds).

The ‘standard recipe’ for CCD photometry is that observers with liquid-nitrogen cooled CCDs should take bias frames (zero-second darks) and not worry about darks, while those with thermo-electrically cooled CCDs can use dark frames and not be concerned with biases.

This is a logical approach, as the dark frames do include the constant bias level, but this removes the possibility of using the bias measurements as a diagnostic tool (the Argos CCD had known issues with bias fluctuations), and prevents one from acquiring a single set of long-exposure dark frames and then scaling the master dark to match the observation length (as the bias contribution is a fixed contribution, regardless of exposure). The best of both worlds can be achieved by configuring the CCD controller to digitize more columns than are physically present on the CCD:

these ‘overscan’ columns contain no integrated charge, and so they directly measure the bias signal (plus a tiny amount of dark current accumulated while the virtual row is shuffled through the digitization register).

The dark current is measured by taking a number (typically 30) of exposures with the CCD shuttered from external light. `tsreduce` calculates a master dark frame by loading these into a 3 dimensional cube (x,y,frame) and then sorting the frame axis by pixel intensity. The middle 50% (this is configurable) of the pixel values are averaged to give the final pixel value. This median-mean combination rejects outliers and dramatically increases the signal-to-noise ratio compared with a single dark frame.

Flat-fielding

There are several contributions to the observed sensitivity variation across the CCD frame. Vignetting by the telescope aperture reduces the number of photons collected around the edge of the frame, and dust or other obstructions in the optical path show up as dark halos. Variations in the physical structure of the CCD introduces small per-pixel variations in the quantum efficiency, and can cause fringing effects for longer light wavelengths.

These effects can be quantified by acquiring a collection of frames (typically 30) of a uniformly lit background. Our preference is to acquire ‘sky flats’: exposures are made of a relatively empty patch of sky during astronomical twilight while shifting the telescope at low speed. Another choice is dome flats: taking images of an illuminated panel on the inside of the telescope dome.

Each of these flat frames are bias/dark subtracted, and then scaled to have a mean intensity of unity. This scaling step is important for sky flats, as the background illumination changes continuously between frames (becoming darker in the evening, or brighter before dawn). The same median-mean process described for the dark frames is then applied, which rejects the pixels that are contaminated with additional starlight (moving the telescope ensures that each frame will have different pixels illuminated), and produces a master flat frame which is effectively an image of the uniformly lit sky, free of stars. Dividing a dark-subtracted science frame by this master flat will undo the effects of the optics and recover a uniform frame suitable for analysis.

These flat frames can also be used to estimate the read noise and gain of the CCD; these values are calculated by `tsreduce` and stored in the master flat image header for use in later reduction steps.

A histogram of the pixel values in the overscan strip shows a normal distribution with a mean equal to the bias, and a standard deviation equal to the readout noise of the detector. Subtracting the master dark from each flat frame removes the bias component, and so the variance can be calculated by summing the (square of the) overscan pixel residuals in all of the flat frames. The square root of the variance gives the RMS read noise in ADU units.

The CCD gain relates the arbitrary ADU (analog-digital-unit) counts of a pixel to the number of physical electrons measured by the CCD. The gain G is usually defined in units of e^-/ADU , and can be measured by comparing the variance of the flat field frames with the theoretically expected value.

The static structure in each flat frame is removed by subtracting the bias, master dark and master flat frames; the remaining signal again shows a normal distribution, with a standard deviation that is the sum of several uncorrelated sources. The dominant noise sources are photon shot noise σ_p , thermal noise from the dark current σ_t , and readout noise σ_{rd} . The thermal and shot noise are both Poissonian in nature, and the readout noise is constant. We can therefore write the total noise (for an intensity measured in electrons) as

$$\begin{aligned}\sigma_{e^-}^2 &= \sigma_{p,e^-}^2 + \sigma_{t,e^-}^2 + \sigma_{rd,e^-}^2 \\ &= (\sqrt{P_{e^-}})^2 + (\sqrt{D_{e^-}})^2 + \sigma_{rd,e^-}^2 \\ &= P_{e^-} + D_{e^-} + \sigma_{rd,e^-}^2 ,\end{aligned}$$

where P_{e^-} and D_{e^-} are the number of electrons associated with detected photons and thermal noise.

We can convert these quantities to ADU by introducing the gain,

$$(G\sigma)^2 = G(P + D) + (G\sigma_{rd})^2 ,$$

and then rearrange to obtain

$$G = \frac{P + D}{\sigma^2 - \sigma_{rd}^2} . \quad (3.1)$$

The quantities P and D are the mean intensity (in ADU) of the pre-processed flat and master dark frames, σ_{rd} is the read noise measured from the overscan region, and σ^2 the variance measured above. This calculation can be repeated for each of the 30 flat frames, and the median taken as the best estimate of the ‘true’ gain.

3.2.2 Frame Registration

The final pre-processing step is to ensure that the synthetic apertures are positioned correctly on the frame. Positional errors can accumulate over time due to imperfect telescope tracking, and can reach the order of tens of pixels. The integration algorithm can correct for small shifts, but requires an approximately correct position (roughly within an aperture radius) as a starting point. The aperture positions in `tsreduce` are specified relative to a reference frame, so we can measure and correct these drifts by comparing the frame of interest against the reference frame.

Finding the linear transform that relates two images is a standard image recognition task, and there are many software packages and sophisticated algorithms for solving this (and many other) problems in general. In our situation most of these capabilities are unnecessary - we are only interested in observations from equatorially mounted telescopes, so we know that only translational offsets will be present (altitude-azimuth telescopes without a field rotator will also feature rotation). This allows us to adopt a simpler, and more intuitive approach in `tsreduce`.

The pixel intensities along each row and column are summed into individual x and y intensity profiles. The mean level is calculated and subtracted from each profile, and any negative values are set to zero; the strong peaks from the star intensities ensure that the mean intensity is sufficiently high for this to truncate the background signal. The x and y profiles of the two frames are then cross-correlated to find the offset that maximizes the similarity between the frames, as illustrated in Figure 3.3.




Figure 3.3: The translational shift between two frames can be calculated by independently cross-correlating the x and y intensity profiles. The target field surrounds the DAV white dwarf L19-2, with two frames offset by $(-1, -17)$ px.

3.2.3 Synthetic Aperture Photometry

The actual synthetic aperture integration is quite simple, with one caveat: the signal-to-noise ratio of the integration depends strongly on the chosen aperture size. If the aperture is too small then useful signal will be lost, but if the aperture is too large then the noise (which increases as r^2) will begin to dominate. The process of selecting an optimum aperture size is deferred until Section 3.2.7; for now, we will assume that we have chosen a fixed aperture size.

The required aperture size can be reduced somewhat by centring the aperture over the star image to sub-pixel accuracy each frame, removing the positional jittering due to tracking errors. This is implemented in `tsreduce` using the modified moment algorithm described by Stone (1989). This algorithm relies on having an accurate estimate of the background sky intensity, and suggests using the mode of the histogram of the surrounding background.

In the interests of having a simple and transparent algorithm, `tsreduce` takes the pragmatic approach of requiring the user to select the sky aperture when marking the target stars during the reduction configuration. This is achieved by presenting the user with a calibrated reference frame (usually the first science frame in an acquisition sequence) in the SAOImage `ds9` display software, where they can mark stars for reduction using the included annulus tool.

The background intensity within each annulus is found by generating a histogram of the pixel intensities, then discarding pixels more than 10σ above the mean level to remove any defective CCD pixels. The statistical mode (most common value) is then estimated using a rule of thumb (attributed to Karl Pearson) for a nearly-symmetric unimodal distribution:

$$\text{mode} = 3 \cdot \text{median} - 2 \cdot \text{mean} .$$

This method works well for the sparse fields that we have been interested in, but it is very sensitive to starlight contamination. It would be straightforward to replace this calculation with a more robust approach, but this has not been done to date because our targets have not required anything more complex. Figure 3.4 demonstrates this process for targets with sky apertures with and without background star contamination.

The stellar intensity can then be integrated by adding the pixel intensities within the aperture, each weighting the fractional overlay of the pixel with the aperture. The background intensity is then corrected by subtracting πr^2 times the estimated sky intensity.




Figure 3.4: The background calculation estimates the mode of the background intensity using a simple calculation. The estimated mode (red arrow) matches the intensity histogram well for an uncontaminated sky aperture (top panels), but fails if the aperture contains significant flux from a bright star (bottom panels - the pixels associated with the stars are well outside the range of the histogram).

The weights for each pixel are evaluated analytically (other routines use a polygon approximation), using a straightforward algorithm based on checking the number of corners within the aperture. The weight will be 1 if all four corners are in the aperture, otherwise each edge is checked against the aperture to find the two intersection points \vec{p}_1 and \vec{p}_2 .

Figure 3.5 demonstrates how the areas of interest are then broken down into individual cases of a polygon and circle segment. The area of a polygon can be shown (e.g. Chote, 2011) to be

$$A_{\text{poly}} = \frac{1}{2} \sum_{i=1}^M (x_{i-1}y_i - x_iy_{i-1}) ,$$

and, the area of the segment can be derived as

$$A_{\text{segment}} = r^2 \sin^{-1} \left(\frac{b}{r} \right) - b\sqrt{r^2 - b^2} ,$$

where $b = |\vec{p}_2 - \vec{p}_1| / 2$ is half the length of the line joining the intersection points.




Figure 3.5: The fractional overlap between the circular aperture and square pixel can be broken down into the areas of a polygon (grey) and a circular sector (red).

There are two cases, shown in Figure 3.6, where the aperture may overlap a pixel while all the corners remain outside. We can reject Case 1 outright as it would require an aperture radius smaller than $\sqrt{2}$ px. Case 2 can affect at most four pixels in the entire integration, each contributing a negligible weight to the integration, so this case can also be safely ignored.

The noise contribution from each pixel can be found by rearranging Equation 3.1:

$$\sigma^2 = W (P + D + \sigma_{\text{rd}}^2) ,$$




Figure 3.6: There are two cases where the aperture may intersect with a pixel while all four corners remain outside. Both cases can be safely ignored.

where W is the fractional area weight of the pixel and the remaining terms are as defined for Equation 3.1.

3.2.4 Full-Width at Half-Maximum

Most stars are sufficiently distant that their incoming light can be treated as plane-waves originating from a point source. An ideal detection system would focus these incoming waves into points on the detector; however, the real world is hardly ideal: the atmosphere refracts and distorts the incoming plane waves with a variability timescale of tens of milliseconds, and the telescope optics introduces additional static distortions. These effects blur the point-image into a Gaussian-like profile, known as the point spread function (PSF). The full-width at half-maximum (FWHM) of this profile provides a useful estimate of the atmospheric conditions, known as the *seeing*, which is typically measured in arcseconds. Poor seeing conditions spread the flux over a larger area and produces noisier photometry than an otherwise equivalent night with improved seeing.

The conversion between pixels and arcseconds is given by the plate scale, which can be calculated using the telescope geometry and CCD pixel size. The MJUO telescope (in the configuration we use for observing) has an effective focal length of $f/7.7$ and a diameter of 1 m, so from the geometry each arcsecond of view is projected to $37\ \mu\text{m}$ on the CCD at the focal plane. The CCD has $13\ \mu\text{m}$ pixels, and so the plate scale for this configuration is $0.35\ \text{arcsec} / \text{px}$. Our standard observation procedure uses 2×2 pixel binning, giving an effective plate scale of $0.70\ \text{arcsec} / \text{px}$.

The PSF can be estimated by integrating synthetic apertures with increasing radii of sizes, and then taking differences to plot a mean intensity profile versus radius.

This can then be fit with a Gaussian to extract the FWHM. Figure 3.7 shows this process for the same object on two nights with different seeing. These plots show a slight reduction in intensity at larger radii, which is caused by a slight over-compensation in the background intensity, and is not a significant factor in measuring the FWHM.

One can re-cast a Gaussian equation into a polynomial form by taking the logarithm of the profile, which produces (in theory) a simple form for linear-least-squares fitting. This technique fails in practice, because the logarithm suppresses the bright central region of the PSF and amplifies the contribution of the noisy wings. A simple and robust approach for our purposes (where we don't need high accuracy) is to step through a range of FWHM values and identify the one which best fits the measured profile.




Figure 3.7: A demonstration of the FWHM calculation on nights with different seeing conditions. The target star is the ELM white dwarf WDJ1518.

3.2.5 Relative Photometry

The physical quantity of interest in our pulsating WD targets is not the absolute intensity; it is the relative intensity changes over time, which can then be converted into the Fourier domain. A convenient measurement unit for these small-scale intensity changes is the milli-modulation intensity (mmi), which is defined by $10 \text{ mmi} = 1\%$ change. By limiting ourselves to relative photometry we are able to simply and effectively remove any atmospheric effects that uniformly affect the field of view by point-wise dividing the measured target intensity by one or more constant comparison star intensities.

The most dramatic effect on the lightcurve is cloud, which (depending on the type)

can introduce smooth changes in intensity, rapid point-to-point changes, or completely wipe out the visibility. The first two cases can be almost completely corrected by the comparison division as long as the cloud obscuration isn't too large. Data with as much as 75% cloud attenuation can often be salvaged into a usable form.

The other significant effect is a smooth change in intensity as the star rises and/or sets, which is caused by a change in the path-length that light travels through the atmosphere. Light from a star near the horizon has a much longer path through the atmosphere, which introduces additional attenuation and refraction, decreasing intensity and increasing FWHM. A reasonable model for this is given by a plane-parallel atmosphere, with an 'airmass' defined as $\sec(z)$ at an angle z from the zenith. An airmass of two has double the path length than the zenith, and occurs about $\sim 30^\circ$ above the horizon, Airmass three is about $\sim 20^\circ$ above the horizon, and is roughly the limit at which we can obtain useful white dwarf photometry from MJUO.

The attenuation (more commonly called extinction in this context) of the atmosphere depends on wavelength, so is not fully corrected by the target/comparison division (this is also true for clouds, but is a much smaller effect with no clear solution). White dwarfs are generally bluer than other stars, so there will usually be a residual extinction signature left in the data after the light-curve division. This is mitigated to a certain extent by observing through a blue BG40 filter, but not enough to avoid needing corrections. This could be avoided if a suitable airmass correction were made to the lightcurves before division, however this introduces unnecessary complication as it depends on the spectral type of each individual star. The intrinsic variability that we are usually interested in has timescales of tens of minutes or less, so we lose no information by simply fitting and subtracting a low-order polynomial (usually a quadratic, but `tsreduce` allows a higher order if necessary) to remove the the long-period extinction changes.

This polynomial correction can be combined with the mmi conversion, using the relationship

$$\text{mmi} = \frac{1000(\text{ratio} - \text{fit})}{\text{fit}}$$

evaluated for each measurement.

Outlying measurements can unduly influence this calculation, and so this procedure is applied twice: the first time is used to calculate a mean and standard deviation, which are used to discard outliers (defaulting to those outside 5σ , but this is again configurable) before repeating the calculation to obtain an improved value.

The results of all these calculations are plotted in real-time for the observer, and include panels with the raw intensity measurements; ratio, and corrected mmi values; and the mean FWHM of the PSF. An example online photometry plot is shown in Figure 3.8. A second plot (not shown) shows the Fourier transform and window function for the data, which are discussed in the next section.

The reduced data are stored in a simple plain-text format, which consists of a file header containing the run parameters and columns listing the source image, observation time, and raw intensities, FWHM, and corrected aperture position for each star. The ratio and mmi values are generated from this raw data when they are needed.




Figure 3.8: An online photometry plot of observations taken on 2011-08-01 of the DBV white dwarf EC04207-4748. The target lightcurve is divided by the comparisons and then fitted with a quadratic polynomial to remove differential extinction. Clean photometry can be extracted through cloud attenuation of as much as 75% for a bright target like EC0407. Data segments irreparably affected by cloud can be marked for rejection by `tsreduce`.

3.2.6 Fourier Analysis

The Discrete Fourier Transform (DFT) is one of the most useful techniques in a physicist's mathematical toolbox. The Fourier transform decomposes a time domain signal into its component frequencies; signals that are complicated functions of time may resolve themselves to a relatively simple combination of sinusoidal variations. Periodic but non-sinusoidal functions will contain harmonic frequencies in

their DFT, with the time-domain shape encoded in the relative amplitude and phase of the harmonics.

The frequency resolution of the Fourier transform is limited by the overall data length, and its maximum measurable frequency by the sampling rate. The Nyquist-Shannon sampling theorem states that the maximum frequency that can be uniquely resolved is half the sampling frequency, which is intuitively reasonable: you need to sample at least the positive and negative lobes of the sinusoid to be able to distinguish it. For pulsating WDs with periodicities $\gtrsim 100$ s, sampling rates of $10 \sim 30$ s give no problems, so the maximum frequency resolution is rarely an issue.

The primary contribution to the appearance of the signal in the frequency domain is the total run length. A finite segment of data can be thought as an infinitely long signal multiplied by a square pulse of width T , which translates to the Fourier domain as a convolution with a sinc function with a central peak width of $2\pi/T$. Gaps within the time series data (due to cloud or combining multiple runs) will introduce additional alias peaks, which can significantly complicate the frequency identification procedure.

These aliases can be quantified by the *window function*, which is the DFT of a noise free sinusoid sampled at the same times as the original time series (note that this is different to the definition used in other fields, where the window function is a tapering function applied to data in the time domain in order to minimize these aliases – the two are inherently related).

The window function described here allows us to visually separate sampling effects from noise. For a multi-night single-site set of observations the window function will be dominated by a series of peaks separated by $11.6 \mu\text{Hz}$ ($1/24$ h) introduced by the daily gaps in the observations.

An effective method for reducing these window aliases is to use the frequencies identified in the Fourier domain to generate and fit sinusoids in the time domain. The χ^2 differences between fits can help break the degeneracy between adjacent alias frequencies, and the fitted sinusoids can be subtracted (‘pre-whitened’) to remove the frequency and its aliases from the DFT. This may uncover lower amplitude frequencies that were previously hidden by the window aliases. An example of this is shown in Figure 3.9; the $5191.8 \mu\text{Hz}$ frequency in the DAV white dwarf L19-2 is split into three components separated by $\sim 13 \mu\text{Hz}$ due to stellar rotation.

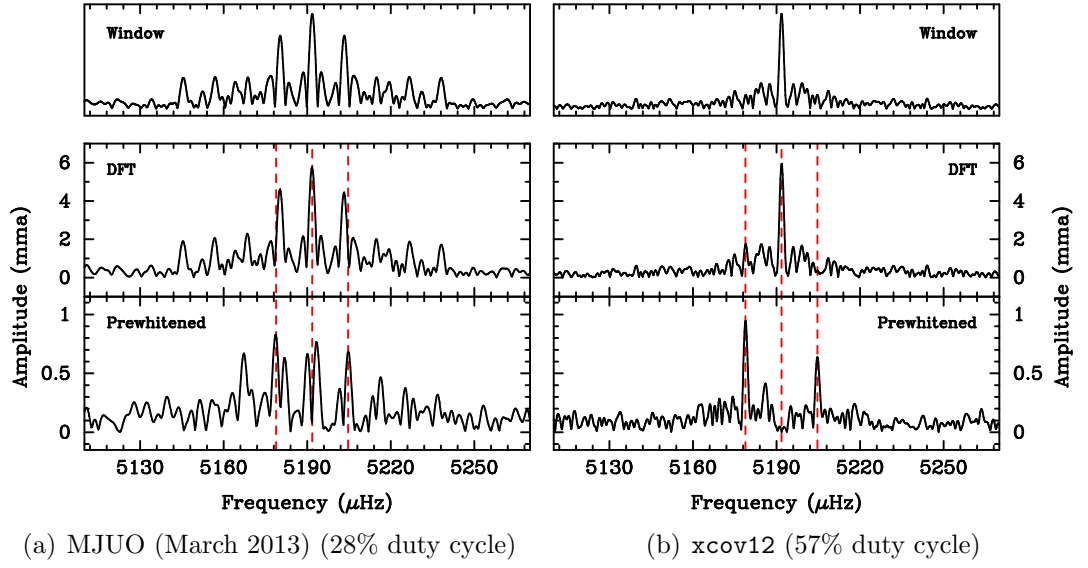


Figure 3.9: The dominant frequency in the DAV WD L19-2 is split into a triplet (indicated with dotted lines) by stellar rotation. The low-amplitude side peaks are hidden in the window pattern of the strong central frequency, but become clearly visible after pre-whitening. This is demonstrated for our best single-site MJUO data and the `xcov12` WET campaign; we see that even the best multi-site observations require pre-whitening in order for these closely spaced multiplets to be resolved. The $11.6 \mu\text{Hz}$ frequency aliases that dominates the single-site window function are highly suppressed in the multi-site data, which does not feature the regularly spaced gaps enforced by sunrise at a single site.

3.2.7 Aperture Size Optimization

Finding an aperture size that gives the best signal-to-noise ratio is a relatively tricky problem. The ‘best’ size may be different for different stars in the image, and it will typically vary over an acquisition run as the seeing changes. Varying the aperture size during the reduction introduces a number of issues, and is generally considered to be an unnecessary complication. We can dramatically simplify the task by using a single fixed aperture size for all the stars in a run.

The pragmatic approach for finding the optimum radius is to repeat the reduction with a range of sizes, and then pick the value that produces the lowest noise in the Fourier domain. This works well for offline reductions, and `tsreduce` includes two routines for this purpose (`reduce-range` and `plot-range`; see Appendix A for more details). This approach does not work for the online reductions, and so this requires an algorithm that can estimate a ‘good enough’ (but not necessarily perfect) value from the first one or two frames acquired at the beginning of a run.

Howell (2006) shows that the theoretically optimum aperture size is given by the FWHM of the PSF, which (being a function of the atmosphere and instrumentation) is independent of the stellar intensity and exposure time. This is problematic for faint stars, however, as their faintness means that the wings of their PSFs are lost in the background noise, which gives them a narrower apparent PSF than brighter stars. One could, in principle, estimate the true FWHM by analyzing the PSF of the brightest non-saturated stars in the field, but this would add unnecessary complexity.

A simpler approach is implemented in `tsreduce`: the radial profile (see Section 3.2.4) is calculated for the target and comparison stars in the selected reference frame, and the radius is chosen from either three times the apparent FWHM (suggested by Howell as a reasonable compromise for faint targets), or the radius where the radial profile falls to within 5σ of the background (an arbitrary choice that was found to give good results in most cases). This results in a unique radius measurement for each star, from which the maximum is taken as the best reduction value.

3.2.8 Time Conversions

The time that we measure on Earth contains systematic offsets associated with leap seconds and the secular motion of the Earth around the sun. We need to ensure that any frequency variations that we observe in a target are intrinsic, and not effects of our moving observation platform.

These effects can be broken into two parts: those intrinsic to the time measurement

(e.g. leap second offsets, gravitational time dilation), and those associated with light travel time as the observer moves about the solar system barycentre. This is surprisingly difficult to do rigorously, but we are fortunate in that there are well established models and software libraries that can do most of the hard work.

Eastman et al. (2010) provide a good discussion of the complexities involved in this procedure, as does the “SOFA Time Scale and Calendar Tools” documentation of the SOFA library (IAU SOFA Board, 2010) which we use to implement these conversions in `tsreduce`.

The net result of these conversions is a time measured in BJD_{TDB} (Barycentric Julian Day, measured using Barycentric Dynamical Time), which can be loosely considered as the time at which an observer positioned at the solar system barycentre would have seen the same event, assuming the mass of the solar system were removed. The SOFA documentation specifically mentions that this is a dubious and unnecessary comparison, but I think it is better to have a hand-waving explanation that conveys the key ideas than to hide behind rigorous but opaque mathematical terminology.

3.2.9 False Alarm Probability

The concept of the false alarm probability (FAP) was first introduced by Scargle (1982), and provides a simple intuitive test for differentiating between random noise peaks and low amplitude coherent signals in the Fourier transform. The technique that we use (from Sullivan et al., 2008) differs somewhat in form from Scargle’s approach, but performs the same function. The speed of modern computers makes it feasible to apply numerical Monte Carlo techniques to work with the measured time-domain noise instead of relying on an assumed theoretical noise model.

The FAP procedure calculates a threshold amplitude A that corresponds to a $1/N$ probability of a peak of amplitude A or greater being caused by incoherent noise fluctuations. By making N large, we can choose an arbitrarily strong confidence threshold for classifying peaks as signal or noise.

We start by pre-whitening all the large-amplitude periodicities from the time series data (the technique can be applied without pre-whitening, but that tends to then over-estimate the noise level), and then shuffle the times associated with each measurement to destroy any remaining coherent signals. The incoherent noise is unaffected by this process, so we can claim that any peaks in the resulting DFT are exclusively due to the incoherent noise. The amplitude of the highest peak in the DFT is recorded, and this shuffle-and-DFT process repeated N times. A histogram

of these maximum peak amplitudes can be plotted (but isn't strictly necessary), and the largest amplitude of any incoherent peak is taken as the FAP threshold.

It should be clear that for N repeats (assuming N is large enough to avoid issues with small number statistics), the highest overall amplitude A defines a threshold where there is a $1/N$ probability of that amplitude being due to caused by incoherent noise. We typically choose $N = 1000$, which sets a 0.1% probability that any peaks larger than A are caused by incoherent noise.

Figure 3.10 shows the FAP peak-amplitude histogram for a typical observing run. These distributions tend to be asymmetric, with a large tail extending to higher amplitudes. A histogram of the peak frequencies is roughly flat, as we should expect for incoherent (white) noise.

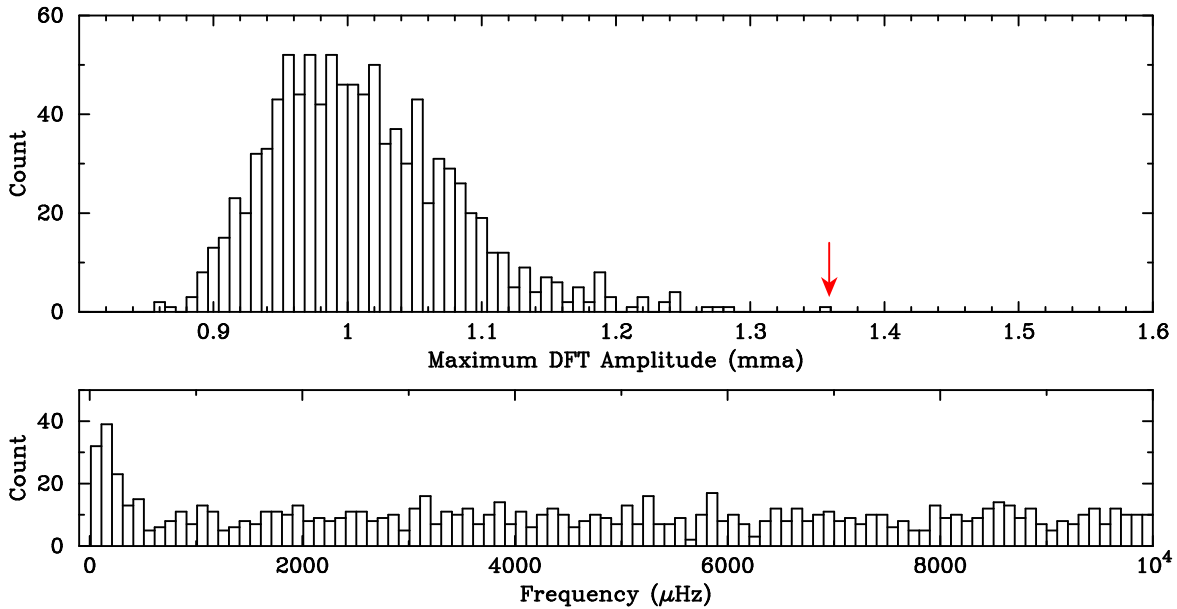


Figure 3.10: The FAP histogram calculated for the November 2011 observation run on EC04207. The maximum amplitude at 1.36 mma is marked with an arrow.

Note that there is nothing special about the time shuffling in this procedure; it is simply an efficient way of generating a random time-domain signal that has the same incoherent noise properties as the original data. This is equivalent to the procedure that will be described in Chapter 6, which fits the pre-whitened noise distribution with a Gaussian and then randomly samples that distribution to obtain random noise samples.

As with any Monte Carlo simulation, this calculation requires a good source of random numbers to ensure unbiased statistics. The Mersenne Twister algorithm (Matsumoto and Nishimura, 1998) is used in `tsreduce`, as it is computationally efficient and offers good quality pseudo-random numbers.

Chapter 4

EC 04207-4748

EC 04207-4748 (subsequently EC04) was the first object that we devoted significant time to observing at the start of my PhD, and so it became the benchmark that I used to test developments on the photometer system, reduction code, and analysis procedures. In many ways EC04 was a perfect target for this: it is bright (V magnitude 15.3) and has strong pulsation amplitudes, which makes it relatively forgiving to observe, and the pulsation spectrum is neither overly complex nor boringly simple.

We devoted four seven-night observing runs to EC04 during 2011. The final session in November was as part of the xCOV28 WET run, which were coordinated with additional observations from Chile (using the PROMPT and SOAR telescopes) and South Africa (using the 1.0 m telescope at the SAAO’s Sutherland site).

Our single-site observations and analysis were published as Chote et al. (2013), and although the additional WET data did produce a more extended lightcurve, nothing new of significance was visible in the DFT of the extended dataset.

4.1 Background

EC04 was first discovered by the Hamburg European Southern Observatory survey for type Ia supernova candidates and named HE 0420-4748. Spectroscopic analysis as part of this programme showed that it is a typical DB WD, with an effective temperature and surface gravity of $T_{eff} = 25000$ K, $\log g = 8.2$ (Koester et al., 2001) or $T_{eff} = 27300$ K, $\log g = 7.8$ (Voss et al., 2007). The discrepancy between these fits is an inherent problem with models for DBV WDs, and is discussed in more detail later.

The EC name prefix and its discovery as a pulsator came from the Edinburgh-Cape

Run Name	Start Time [BJD]	N	ΔT [hr]	Duty Cycle
mar11	2455622.8752867	1972	11.0	8%
jul11a	2455744.2439362	2497	13.9	10%
jul11b	2455770.0678829	3713	26.6	18%
nov11	2455883.9212131	3373	18.7	13%

Table 4.1: Files containing the combined photometry for each of the four one week observing runs. Column 1 gives the run name, 2 the start time of the first exposure in barycentric Julian day (BJD), 3 and 4 the number of useful exposures and effective observation times, and the final column states the duty cycle for each of the individual sessions.

Southern hemisphere survey for faint blue objects (Stobie et al., 1997). From this, Kilkenny et al. (2009) identified EC04 as a pulsator, with four potential frequencies, based on ten hours of photometry acquired in 2002 and 2003 using the SAAO 1.0 and 0.75 m telescopes. This paper also identified a second new southern DBV, EC 05221-4725; we obtained a few short observation runs on this object (see Chapter 7), but did not pursue it further as other targets had higher priority.

4.2 Observations and DFT Analysis

We acquired a total of 56 h of useful integrations during our four 2011 observation runs, which are presented visually in Figure 4.1. The weather and seeing were poor during the first run in March; this can be seen in the large noise baseline of the March DFT and the correspondingly large false alarm probability in Table 4.1. The weather during the July and November runs was unremarkable by MJUO standards, with moderate seeing and occasional cloud.

The CCD frames were reduced using our standard procedure (described in Chapter 3) to create a combined photometry file for each week-long run, which are summarised in Table 4.1. All observations were acquired using a 20 s integration time, with the CCD configured for 100 kHz readout, medium gain, and 2×2 pixel binning.

The analysis features of `tsreduce` were then used to identify the frequencies and calculate the false alarm probability (FAP) for each run. A total of 15 significant frequencies were found, and are listed in Table 4.2. These frequencies can be classified into two groups: $f_1 - f_{11}$ are the main contributions to the non-sinusoidal lightcurves, and are well established across the four sessions. $f_a - f_d$ pass our FAP criteria, but are not nearly as well established.

The frequency identification of modes between the four runs was found to agree to better than $1 \mu\text{Hz}$ after accounting for the degeneracy between adjacent $11.6 \mu\text{Hz}$




Figure 4.1: A total of 56 h of useable integrations were obtained on EC 04207-4748 using the Puoko-nui South photometer on the MJUO 1.0m telescope. The left hand panels show the time-series data for each night, which were aggregated into the four week-long runs. The right hand panels show the DFTs and window functions of each combined run.

Item	Frequency [μHz]	Period [s]	Amplitude [mma]				Assumed Combinations
			mar11.ts [2.7]	jul11a.ts [1.5]	jul11b.ts [1.4]	nov11.ts [1.4]	
f_a	736	1358.9	-	-	1.7	1.6	-
f_b	1478	676.5	4.1	-	(1.3)	-	-
$\star f_1$	1669	599.1	4.7	2.6	1.9	-	-
$\star f_2$	2236	447.2	33.1	38.8	38.4	38.3	-
$\star f_3$	2361	423.5	8.7	8.1	8.1	9.2	-
f_c	2918	342.7	(1.6)	-	1.7	2.1	-
$\star f_4$	2973	336.4	5.8	4.7	5.6	5.1	-
f_d	3906	256.0	(2.1)	1.6	-	-	$f_2 + f_1$
f_5	4473	223.6	7.7	9.0	10.7	10.4	$2f_2$
f_6	4598	217.5	(2.1)	3.2	2.6	2.0	$f_2 + f_3$
f_7	5209	192.0	2.7	2.1	3.1	2.5	$f_2 + f_4$
f_8	6709	149.1	2.9	2.8	4.1	4.2	$3f_2$
f_9	6831	146.4	-	(1.4)	(1.3)	1.5	$2f_2 + f_3$
f_{10}	7445	134.3	-	-	1.7	1.4	$2f_2 + f_4$
f_{11}	8945	111.8	-	-	1.4	1.8	$4f_2$

Table 4.2: Table of the detected frequencies in the four individual observing runs. The frequency and period values are rounded to their nearest certain digit (assume ± 0.5 uncertainty in the final digit). The amplitudes are specified in millimodulation amplitude (mma) units and those given in parentheses are below the 0.1% false alarm probability (FAP) for the particular run, but which we believe are likely to be real (see text). The run FAP values are specified in the column headers in square brackets. Frequencies $f_1 - f_{11}$ labelled by numbered suffixes are well determined across all four observing runs, whereas the four items labelled with letter suffixes a – d exhibit significant power in at least one run, but we consider these only marginal detections of real power. Frequencies deduced to be associated with independent eigenmodes of the star are labelled with a ‘ \star ’.

alias peaks. The frequency uncertainty was therefore empirically determined to be $< 1 \mu\text{Hz}$, and this was later confirmed by running the Monte Carlo technique that will be described in Section 5.3.2 on one of these runs.

Figure 4.2 shows the DFT of each run after pre-whitening the $f_1 - f_{11}$ frequencies. It should be clear from this that f_a and f_b aren't really believable; they were only given a second look because they are above the 0.1% FAP threshold during one of the runs. Without any additional supporting evidence to the contrary, we believe that these happen to be caused by the one in a thousand chance of a 'noise conspiracy'. f_c looks to be relatively convincing in the final two runs, and could perhaps be a variable-amplitude mode that remained below the detection amplitude during the first two runs. Similarly, we believe that f_d is most likely real, corresponding to the linear combination $f_2 + f_1$, but we cannot firmly justify it based on the data that we have. However, even if it is real, it is clearly a combination frequency and not an independent pulsation mode.

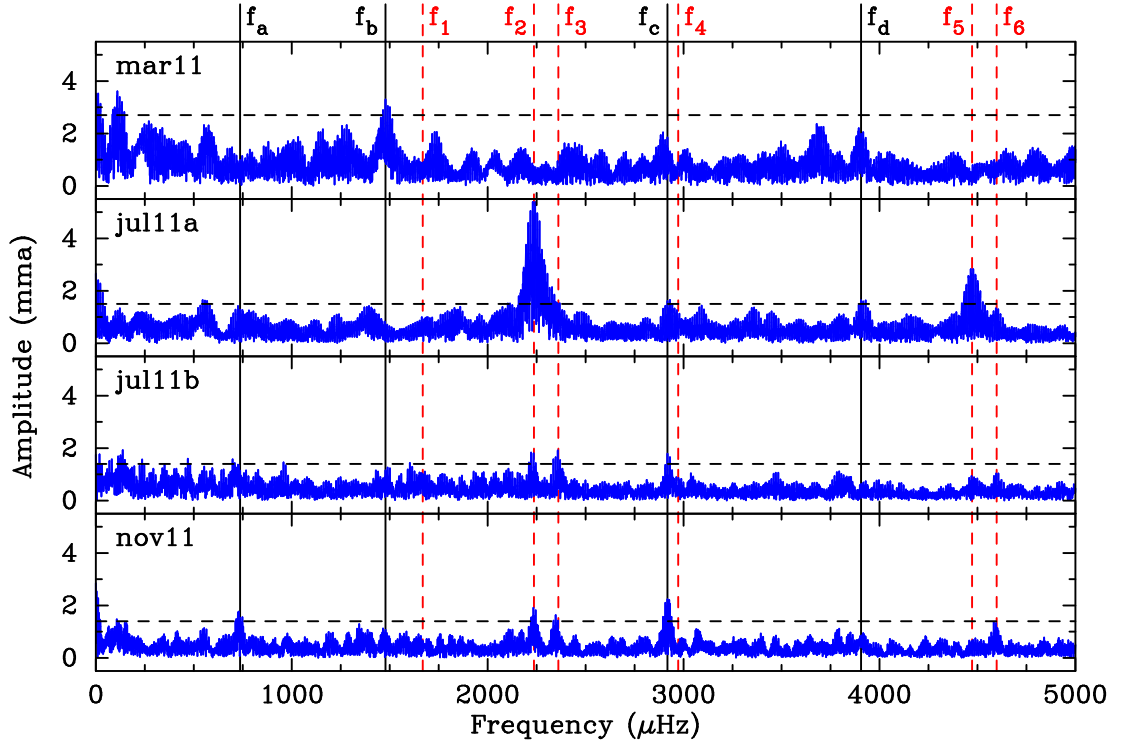


Figure 4.2: A more detailed look at the residuals that remain after pre-whitening the frequencies $f_1 - f_{11}$. The location of these subtracted frequencies are marked with dotted red lines. Four frequencies that formally pass the FAP threshold are marked, and discussed in the text along with the residuals left after subtracting f_1 and f_3 . The peaks below $200 \mu\text{Hz}$ are long-period variations that are associated with imperfect atmospheric corrections and are not associated with the WD. Note the changed frequency scale: $f_7 - f_{11}$ did not show any residual power after pre-whitening, so they are not included in this plot.

Another feature worth noting in Figure 4.2 is the power which remains after pre-whitening the three main frequencies (f_2 , f_3 , and f_5) in each data set. These are

quite significant in the first July run, but remain above the FAP threshold in the second July and November runs. These could possibly be caused by weak amplitude modulation in these modes, or by unresolved multiplet components.

The first July run was impacted by an instrument problem: certain conditions would cause the camera drivers to return buffered frame data from an earlier exposure. This resulted in incorrect time tags during some exposure sequences. We were careful to identify and correct or discard these frames (the delay was always equal to four 20 second exposures), but an imperfect correction could lead to DFT artefacts like the ones we see. Selectively discarding the remaining segments did not appreciably improve the situation, and the artefacts are relatively minor compared to the primary frequencies, so we do not consider them to have any meaningful effect on our results. This problem with the camera driver was exposed by changes made to the acquisition software between the March and July runs, and the problem was resolved before the second July run. We are confident that this was the only run affected by this issue.

Table 4.2 shows that most of the identified frequencies are linear combinations of other modes. Such linear-combination frequencies are seen in many pulsating WDs (particularly near the cooler end of the instability strips), and are known to be artefacts caused by non-sinusoidal lightcurve shapes. They do not correspond to physical pulsation modes in the WD, and so it appears that EC 04207-4748 has only four detectable eigenmodes. This makes EC04 poorly suited for traditional asteroseismological analysis, but the strongly non-sinusoidal lightcurve (shown in Figure 4.3) makes it an ideal candidate for studying convection effects.

4.3 Convection Theory

The combination frequencies that are visible in the DFTs of many pulsating WDs are a clear sign of non-linearities in the transformation from the internal flux variations generated by the pulsations into the photospheric flux variations that we observe. One source of non-linearities is the T^4 relationship between temperature and luminosity (the Stefan-Boltzmann equation). Another, more interesting, contribution is the effect of a sub-surface convection zone within the stellar atmosphere.

Convection was largely neglected in the original theoretical treatments of WD pulsations, and its effects weren't seriously considered until Brickhill introduced a model for *convective driving* (Brickhill, 1983, 1990, 1991a,b, 1992). One of the attractive features of this model is that it provides a natural explanation for the observed non-linearity: the outgoing flux passes through a sub-surface convection zone which

has a heat capacity that varies in response to local temperature changes (Brickhill, 1992). The small temperature variations that occur during a pulsation cycle cause energy to be absorbed by the convection zone when the flux increases, and then released after a delay: this results in a lightcurve that has relatively sharper peaks and flatter troughs.

Aspects of this early numerical work were explored analytically by Goldreich and Wu (1999), and Wu (2001). Wu derived an analytical relationship between the thickness of the convection zone (described using a timescale, τ) and the amplitude of the combination modes. Montgomery (2005) adopted a modified parameterisation of this model, and created a hybrid (semi-analytical) code that could generate lightcurves suitable for fitting to observational data. His original code was limited to mono-periodic pulsators, but was subsequently extended with support for the multi-periodic variety (Montgomery et al., 2010).

The convective region is a relatively thin layer between the photosphere and radiative interior of the WD. This layer is believed to overlap with the partial ionization region in pulsating WDs, and the exchange of energy between them provides a mechanism for driving the pulsations (which is the origin of the term convective driving), as was described in Chapter 1.

Energy transport within the star is dominated by radiation at high temperatures, so convection in hotter WDs is limited to a shallow layer near the surface. Cooler WDs have a thicker layer, with more mass and a longer convective timescale. This produces larger distortions and more non-sinusoidal lightcurves. This relationship between temperature and lightcurve shape has been well established observationally, and Fontaine and Brassard (2008) show nice plots of this trend within the DAV and DBV instability strips.

A small number of assumptions make this behaviour tractable for modelling: that the convection layer reacts instantaneously to temperature changes (this is reasonable because the thermal response timescale is ~ 1 second, compared with hundreds of seconds for the g-mode pulsations); that the behaviour of the convection zone is purely radial (the convection zone is thin compared with the angular variation of the pulsations); that geometrical effects associated with pulsation are ignored (the extreme temperature sensitivity of the convection zone means that temperature and flux effects will dominate); and that flux perturbations beneath the convection zone are sinusoidal in time and spherical harmonic in angle.

One can then define a simple equation relating the flux entering the base of the

convection zone to the flux exiting into the photosphere,

$$F_{\text{phot}} = F_{\text{base}} - \frac{d\tilde{Q}}{dt} , \quad (4.1)$$

where $d\tilde{Q}/dt$ is the energy per unit time absorbed by the convection layer. Montgomery (2005) explains that this absorption factor can be computed from static atmosphere models over a range of effective temperatures, which allows this to be rewritten as

$$F_{\text{phot}} = F_{\text{base}} + \tau_C \frac{dF_{\text{phot}}}{dt} . \quad (4.2)$$

The parameter τ_C describes the magnitude of the change in heat capacity of the convection layer as the local photometric flux and temperature vary. The full calculation of τ_C depends on the underlying physics in the convection layer, which is described by mixing length theory (MLT) (Böhm-Vitense, 1958; Böhm and Cassinelli, 1971), but can be approximated using a simple power law over a restricted range of temperatures (e.g. the DAV or DBV instability strip):

$$\tau_C = \tau_0 \left(\frac{T_{\text{eff}}}{T_{\text{eff},0}} \right)^{-N} . \quad (4.3)$$

The convective response timescale, τ_0 , is a function of the mixing length parameter α , and is highly sensitive to the effective temperature of the WD. It can be thought of as the average time that energy is trapped inside the convection layer, or equivalently as the layer depth or mass.

T_{eff} is the instantaneous effective temperature, which is related to the observed flux, and $T_{\text{eff},0}$ is its time-averaged value. N describes the strength of the power law, and is very high: $90 \sim 95$ over the DAV instability strip, and $20 \sim 25$ across the DBV strip.

The assumed behaviour of the flux entering the base of the convection zone is given by the sum of spherical harmonics,

$$F_{\text{base}} = F_{\text{base},0} + \text{Re} \left\{ \sum_{j=1}^M A_j e^{i(\omega_j t + \phi_j)} Y_{\ell_j m_j}(\theta, \phi) \right\} , \quad (4.4)$$

which introduces the familiar parameters of eigenmode frequency, phase, amplitude (here specified at the base of the convection zone, which is not directly observable), ℓ and m .

The final step for calculating a model lightcurve is to integrate Equation 4.2 over the visible disk of the WD, which introduces an additional parameter, θ_i , for the relative inclination between the pulsation axis and line of sight. This integration

can be done numerically, and takes into account the T^4 relationship, limb darkening, and the observation passband.

An important measure of a model is in how well it can describe observed behaviour with a limited number of parameters, and in its ability to make predictions. This convection model succeeds on both counts: Figure 4.3 shows a model fit to one of the EC04 lightcurves that we acquired in 2011, and fits to other modelled WDs show similarly good results (see e.g. Provencal et al., 2012, and references therein). The underlying convection theory provides a self-consistent explanation for the driving of the g-mode pulsations, and a simple explanation for the observed relationship between temperature and lightcurve shape across the instability strip.

The main predictive power of this model comes from the measured convective response timescale, τ_0 , and its relationship to T_{eff} and the MLT α parameter. α defines the relative size of the convection cells, and plays an important role in the atmospheric models that are used to extract physical quantities from WD spectra. Bergeron et al. (1995) showed that varying the value of α significantly changes the temperature, mass, and radii measurements extracted from their spectroscopic fitting procedure.

Bergeron et al. (1995) and Tremblay et al. (2010) found values of $\alpha = 0.6$ and 0.8 produced the best self-consistent results with their respective spectroscopic models, but Fontaine and Brassard (2008) found that $\alpha \approx 1$ is required to correctly describe the observed blue edge of the DAV instability strip. Tremblay et al. (2010) notes that α is expected to vary significantly as a function of depth, and that the spectroscopic models are only sensitive to its value near the surface. One should therefore expect a certain amount of disagreement between α values obtained using different techniques.

This convection model also provides an independent constraint on the spherical harmonic degree (ℓ) and azimuthal number (m) identifications made via asteroseismology; however, these fits have been found to be relatively insensitive to these values, so one should be cautious about drawing too many conclusions from these results. Additionally, the pulsation angle θ_i may be useful for constraining effects such as magnetic fields. Montgomery et al. (2010) discusses how an apparent modulation of θ_i could be described in terms of a relative inclination between the rotation and pulsation axes. This inclination could, in turn, be interpreted in terms of a magnetic field.

Mode	Frequency [μHz]	Period [s]	Amplitude [mma]	ℓ	m	Parameter	Value
f_2	2236.24	447.18	23.8	1	0	T_{eff} (K)	25800 ± 1500
f_3	2361.47	423.50	9.8	1	1	τ_0 (s)	148 ± 5
f_4	2972.66	336.40	4.0	1	0	θ_i ($^\circ$)	38.8 ± 5

Table 4.3: A list of input and output parameters from our best-fit EC 04207-4748 convection models.

4.4 Convection Analysis

Our November run was selected for convection fitting by Mike Montgomery, who used our lightcurve and identified mode frequencies $f_2 - f_4$ (f_1 was not visible in the November run), plus the previous spectroscopic temperature estimates as known inputs. The ℓ and m values were varied to find the optimum fit, and the remaining parameters left as free parameters for the model to optimize. The results of the fit are given in Table 4.3, and a segment of the fitted lightcurve (showing only one of the four nights) is shown in Figure 4.3. The two temperature estimates were found to give slightly different τ_0 values – this systematic uncertainty has been incorporated into the stated uncertainty of τ_0 .

This result adds an additional data point to the growing collection of τ_0 measurements that are being collected for DAV and DBV pulsators. Figure 4.4 plots τ_0 vs T_{eff} for all the modelled WDs, and shows the clear relationship that exists between τ_0 , T_{eff} , and the MLT α parameter.

If T_{eff} is well known then we can use this relationship to obtain a non-spectroscopic estimate of α , and the current results do appear consistent with $\alpha \approx 1$ as required by Fontaine and Brassard (2008).

The traditional spectroscopic measurements of the DBV white dwarfs do not yield accurate measurements of T_{eff} due to the fact that (a) visible spectral lines are relatively insensitive to temperature changes in the vicinity of the DBV instability strip, and (b) the derived temperatures are sensitive to undetectable trace levels of hydrogen (Beauchamp et al., 1999). However, it is clear that the value of τ_0 obtained from this convection fitting reliably tracks the convection zone depth within the DAV and DBV instability strips, and therefore this parameter can be used as a calibration or substitute for the effective temperature if we choose a specific value for α .




Figure 4.3: The model parameters from Table 4.3 provide an excellent fit to observations obtained on 2011-11-24 using Puoko-nui South on the MJUO 1.0 m telescope.




Figure 4.4: A clear relationship can be seen between τ_0 and T_{eff} for both DAV (left panel, using data from Provencal et al., 2012) and DBV (right panel, from Chote et al., 2013) WDs.

Chapter 5

GW Librae

GW Librae (GW Lib) is the class prototype for the group of accreting pulsating white dwarfs. These WDs exist in cataclysmic variable (CV) systems, where they are in a close binary orbit with another star. The secondary (often called the donor) star is disrupted by the strong gravity of the nearby WD, and matter – primarily hydrogen, but also heavier elements – streams onto the WD, usually via an accretion disk.

These systems are known for undergoing periodic outbursts (nova) where they dramatically brighten by several magnitudes, and then dim back to their original intensity. These are triggered by a large amount of energy being dumped into the WD atmosphere, which raises the surface temperature by several thousand degrees. There are two main mechanisms that can cause this: a flash of nuclear fusion when the accreted hydrogen passes a critical level; or the liberation of gravitational potential energy when changes in the accretion disk cause a large amount of material to be dumped onto the WD. These systems also show a range of shorter period, quasi-periodic variability associated with flickering and other behaviour within the accretion disk.

The story of GW Librae (GW Lib) begins with Maza and Gonzalez (1983), who first noticed GW Lib as a system undergoing a $\sim 9^{\text{th}}$ magnitude outburst. A comparison with archived photometric plates showed that the quiescent magnitude was around 18^{th} magnitude. Duerbeck and Seitter (1987) tentatively classified GW Lib as a nova or dwarf nova type CV based on spectroscopic observations, and Ringwald et al. (1996) later confirmed it to be a dwarf nova. Little further interest was shown until Warner and van Zyl (1998) discovered photometric variability that was consistent with the non-radial g-mode pulsations that had previously only been seen in isolated white dwarfs.

GW Lib was relatively well studied during the following several years, which re-




Figure 5.1: Observations of GW Librae acquired in 1998 using the MJUO CCD on the 1.0 m McLellan telescope show the quiescent pulsation behaviour of GW Librae. Three clearly visible periodicities are identified as non-radial g-mode pulsations in the accreting WD primary.

sulted in the measurement (or at least constraint) of several key system parameters. Thorstensen et al. (2002) used radial velocity measurements to find an orbital period of 76.79 ± 0.02 minutes and a binary inclination of $i \sim 11^\circ$ (i.e. almost face on). Szkody et al. (2002) identified solar-like metals in the UV spectrum (making GW Lib a DAZQ spectral type in the Sion et al. (1983) system), and found the best-fitting model temperature was provided by having $\sim 60\%$ of the surface at 13,300 K and the rest at 17,100 K. The best single-temperature model estimated a surface temperature of 14,700 K. All of these temperatures are hotter than the blue edge of the standard DAV instability strip.

A multi-site photometric observation campaign in 1998 (van Zyl et al., 2004) characterized the non-radial pulsation spectrum. The NZ contribution to this run included two 7 hour runs in May 1998 using the Mt John Top Hat photometer on the 1.0 m telescope (Sullivan et al., 2000). The original CCD data were recovered and re-reduced using `tsreduce`; the DFTs of these two runs are shown in Figure 5.1. The full observation campaign identified three pulsation frequencies in the vicinity of 650, 370, and 230 s, but the exact periods appeared to vary between observations.

Szkody et al. (2002) showed that these three pulsation modes were also visible in the UV wavelengths, with amplitudes ~ 16 times greater than in the optical regime.

This is consistent with the expected effects of limb-darkening on the non-radial WD pulsations.

An initial seismological analysis by Townsley et al. (2004) modelled a WD mass of $1.02 M_{\odot}$, with an accreted layer mass of $0.31 \times 10^{-4} M_{\odot}$, and an accretion rate of only $7.3 \times 10^{-11} M_{\odot} \text{yr}^{-1}$. This is consistent with the assumption that only very low-accretion CVs will show non-radial pulsations (see below).

GW Lib underwent another outburst in 2007 (Templeton et al., 2007; Vican et al., 2011), which lasted for 26 days and peaked at 8th magnitude. The increased surface temperature during and after the outburst caused the pulsations to cease as expected. Byckling et al. (2009) note that this outburst was somewhat unusual compared to other dwarf novae, in that the x-ray flux did not decrease during the outburst, but instead increased by three orders of magnitude. GW Lib is also taking longer than expected to cool back to its quiescent state.

Efforts to monitor GW Lib for the return of these pulsations have been underway since the outburst, with Copperwheat et al. (2009), Schwieterman et al. (2010) and Bullock et al. (2011) all reporting negative results for the return of non-radial pulsations between 2008 – 2010.

We became involved with this project in 2011, with our first observations in March just preceding scheduled Hubble Space Telescope (HST) observations in April 2011. The results of our MJUO observations during 2011 formed an important part of Szkody et al. (2012a), and our monitoring efforts have continued through to 2013.

5.1 Accreting Pulsating White Dwarfs

The cataclysmic variable (CV) classification is relatively broad, and covers a range of observationally-distinct systems. We are only interested in the specific subset that can show non-radial pulsations, which corresponds to the WZ Sagittae subgroup. These dwarf novae systems (Howell et al., 1995) lack the common short period outbursts that are seen in other types of CV, but instead show ‘superoutbursts’ that last longer (weeks vs days) and take longer (months – years) for the WD to cool into quiescence. These systems spend long periods of time ($\gtrsim 10$ years) in quiescence between superoutbursts, and the WD contributes a significant proportion of the visible flux (about half for GW Lib) during this time. This allows us to detect the presence of non-radial g-mode pulsations in the primary WD.

The most important property that allows these systems to exhibit pulsations is a very low accretion rate. The accreted material compresses and heats the surface of

Star	Temperature (kK)	Outbursts (year)	Pulsation periods (s)	Discovery paper
GW Lib	15.4	1983, 2007	650, 380, 240	Warner and van Zyl (1998)
V386 Ser	14.5	-	610, 350, 300, 220	Woudt and Warner (2004)
SDSS 0131-09	14.5	-	600, 340, 260	Warner and Woudt (2004)
SDSS 2205+11	15	2011	580, 480, 330	Warner and Woudt (2004)
HS 2331+3905	10.5	2007	400, 360	Araujo-Betancor et al. (2005)
PQ And	12	1938, 1967, 1988, 2010	1260, 630	Vanlandingham et al. (2005)
RE 1255+26		1994	1340, 1240, 670	Patterson et al. (2005)
SDSS 1339+48	12.5	2011	640	Gänsicke et al. (2006)
SDSS 1514+45		-	560	Nilsson et al. (2006)
SDSS 0745+45	15.1	2006	1200	Mukadam et al. (2007)
SDSS 0919+08	13.5		260	Mukadam et al. (2007)
SDSS 1507+52			1140, 660, 500	Patterson et al. (2008)
SDSS 0804+51		2006, 2010	1000	Pavlenko (2009)
VSX 0747+06		2008	680, 240	Woudt and Warner (2011)
SDSS 1457+51			1200, 600	Uthas et al. (2012)
BW Scl	~15		1240, 620	Uthas et al. (2012)

Table 5.1: Table of accreting pulsator systems. These properties have been collated from the literature, and may not be complete. The pulsation periods are given to the nearest 10 s, as many of these periods are observed to vary at this level.

the WD, which provides a heating source that can (if not sufficiently low) keep the WD temperature too hot for the partially-ionized layer that drives the pulsations to form. These WZ Sge systems also tend to have orbital periods near the theoretical minimum for CV systems, which is around 80 minutes. This makes them interesting for the CV community, but is not of particular concern here. They also have low magnetic fields, and so mass transfer occurs via an accretion disk.

GW Lib was the first of these systems to be discovered with variability that is consistent with non-radial pulsations (van Zyl et al., 2000), but there are now (at least) 15 other systems that show similar behaviour. These are presented in Table 5.1.

The accreting WD instability strip appears to extend to much higher temperatures than for the isolated DAVs. Arras et al. (2006) show that this can be attributed to differences in the atmospheric composition: these stars are expected to have a higher fraction of helium than the isolated DAVs due to the accreted material from the donor star. Their models reveal an additional instability strip associated with HeII partial ionization at around 15,000 K, which may merge with the HI instability strip at 10,500 K to create a single wider strip that extends up to around 16,000 K.

Accurate temperature measurements are difficult to obtain for these systems. The accretion disk significantly contaminates the visible spectra, and makes it difficult to precisely fit the spectroscopic WD models. One effective technique (see e.g. Szkody et al., 2010) is to measure the Lyman α lines in the ultraviolet region, where the hotter WD dominates the flux, but the Earth’s atmosphere limits these UV measurements to space-based instruments.

Time-tagged UV spectra obtained from instruments such as the Cosmic Origins Spectrograph on the Hubble Space Telescope can be transformed into a UV lightcurve

by integrating over a range of wavelengths. The WD signal can be enhanced by excluding the wavelengths known to be associated with the accretion disk emission, which further improves the signal quality. These lightcurves show the same pulsation frequencies, but with much higher amplitudes than in the visible wavelengths. This is expected behaviour for non-radial pulsations, as the increased limb darkening at UV wavelengths decreases the geometric cancellation across the visible disk of the WD (see e.g. Robinson et al., 1995).

These accreting systems are significantly more complicated than the isolated pulsating WDs, and the effects that this will have on their pulsations are still largely unknown. Millennia of accretion (and the associated angular momentum increase) spins the rotation of the WD up to very fast rates – periods of hundreds of seconds or less are not unusual. The accretion in these low mass-transfer systems occurs in an equatorial band, and may lead to compositional or temperature changes around the surface of the WD. These two effects combine to break the spherical symmetry that is explicitly assumed in the WD pulsation theory, although it is not clear how much of an effect this would have on the resulting pulsation modes. Some discussion of these effects are included in Szkody et al. (2012b, 2013).

There are still a number of unexplained phenomena associated with these WDs. Long-baseline observations of GW Lib, SDSSJ1610-01 and VSX J0747+06 during quiescence all show apparent frequency multiplets with a spacing of $\sim 1 \mu\text{Hz}$ (van Zyl et al., 2004; Mukadam et al., 2010; Woudt and Warner, 2011). The presence of such multiplets in an isolated WD is usually treated as a sign of rotation, but the spacings measured for these objects would imply a rotational period of \sim days, which directly contradicts the \sim hundreds-of-second periods from radial velocity measurements.

Other systems, such as SDSS 1457+51 and SDSS 0745+4538 show broadly stable pulsation modes, but with periods that change by relatively large amounts between nights. Several additional systems with parameters that should place them within the accreting WD instability strip have not shown any visible pulsations (e.g. Mukadam et al., 2007; Gänsicke et al., 2005). However, one of these systems (BW Scl) *was* later found to pulsate (Uthas et al., 2012).

Even with these unknowns, there are compelling reasons for wanting to study these pulsators. The periodic outbursts heat their surface layers outside the instability strip, which is followed by the system cooling back to its quiescent temperature. This cooling happens on a timescale of years, which is much more accessible to human observers than the evolutionary cooling timescales of isolated WDs. This opens a window into observing the effects of cooling on the WD structure via its effect on the pulsation periods (see Chapter 6 for more details). Changes in the pulsations

before and after an outburst may also provide information on the accretion event itself, such as the total accretion mass (e.g. Mukadam et al., 2011b).

Small number statistics limit our ability to understand these systems in the immediate future. They are extremely complicated, and it is unreasonable to expect that we can accurately model all the behaviour seen in each system. Matters may become clearer in the future once more systems have been discovered, and more outbursts observed, and the ‘normal’ behaviour of these systems become known.

5.2 Observations

Our new observations of GW Lib span three years, from March 2011 through to March 2013. We obtained approximately 39 hours of useable integrations in 2011, which were published in Szkody et al. (2012a). An additional 77 hours were obtained in 2012, and a further 9.5 hours in 2013. The details of the individual acquisition runs are presented in Table 5.2, and the lightcurves and DFTs for each run are shown in Figures 5.2 and 5.3.

We found no sign indicating the return of the pre-2007 quiescent pulsation frequencies. However, we did find several recurring periodicities at other frequencies (periods). The runs in Table 5.2 are labelled with these types, but omit the specific period length: this is partly to reduce unnecessary clutter, but also because many of these frequencies are seen to vary between runs. The long period (> 1 hour) frequency determinations are expected to be distorted slightly by the amplitude distortions introduced by differential extinction and the low-order polynomial subtraction.

The periodicity identifications that we believe are present in these data are:

- A** A double-humped, 80 minute ($200 \mu\text{Hz}$ and harmonics) ‘late superhump’. This has been suggested to be associated with the precession of an elliptical accretion disk, which gives it a period slightly longer than the 77 minute orbital period. This signal has been reported by several other groups in both the visible and UV regimes (e.g. Bullock et al., 2011; Kato et al., 2008).
- B** Periods in the range of 4 – 6 minutes ($3000 - 4000 \mu\text{Hz}$), which are examined in detail within Szkody et al. (2012a). These peaks are seen in the ultraviolet as well as the visible wavelengths, which is a strong sign that this periodicity originates from the WD. However, the extremely variable frequency and amplitude between nights makes it unclear as to the exact nature of this signal.
- C** A ~ 19 minute ($\sim 900 \mu\text{Hz}$) periodicity, which is examined in detail below.

D A ~ 3 h ($90 \mu\text{Hz}$) periodicity that is present at the same time as **C**. The DFT shows additional peaks at 30 and $60 \mu\text{Hz}$, which may indicate that this is a harmonic of a longer period phenomenon.

The primary focus of this chapter is the May 2012 (may12) observation run. The lightcurve during this period was dominated by a strong (~ 50 mma, or 5% of the total intensity) non-sinusoidal intensity modulation that was not seen in our 2011 or 2013 observations. Our first run on May 17 appeared to indicate a partially-resolved multiplet structure, which we interpreted as a potential indication of an underlying mode structure. We decided to change our focus for this observing run exclusively to GW Lib, with the goal of quantifying this (apparently) new behaviour. An early analysis of this data was presented in Chote and Sullivan (2013a), and a journal paper that concisely reports on the investigation below is underway.

The coordinates of GW Lib ($15^{\text{h}}19^{\text{m}}55^{\text{s}}$, $-25^{\circ}00'25''$) allow for it to be observed all night from MJUO in May. This, combined with uncharacteristically good observing conditions for Mt John, allowed us to acquire 61 h of excellent quality photometry over a total of 6 nights – an impressive 42% observing duty cycle, out of a potential $\sim 50\%$ for May at Mt John.

Date	Start UTC	End UTC	Exposure (s)	Useable (hours)	Observed frequency types
2011-03-02	13:09	17:23	20	4	A
2011-03-04	14:33	16:19	20	1.5	A
2011-07-01	07:02	12:06	20	5	B
2011-07-02	06:23	11:00	20	4.5	B
2011-07-04	06:36	14:02	20	7.5	B
2011-07-06	06:35	09:28	20	3	(A) B
2011-07-27	06:48	12:47	20	5.5	(A)
2011-08-01	06:39	12:53	20	5	A
2011-08-02	06:35	10:04	20	3	
2012-03-24	13:02	17:49	30	3.5	(A)
2012-03-25	12:40	17:56	30	5	A
2012-04-23	11:22	18:24	30	7	(B) (C) D
2012-05-17	06:53	18:06	30	11	B C D
2012-05-18	06:44	17:53	30	11	B C
2012-05-19	06:33	17:55	30	11.5	B C D
2012-05-20	07:40	18:06	30	6	C D
2012-05-21	07:32	17:46	30	10	C D
2012-05-22	06:29	18:04	30	12	C D
2013-03-08	13:22	15:24	30	1.5	(A)
2013-03-10	12:57	17:26	30	2.5	(A)
2013-03-13	12:46	15:24	20, 30	5.5	A

Table 5.2: Table of GW Lib observation runs. The frequency type descriptions are included in the main text. Labels in parenthesis indicate probable identifications.




Figure 5.2: Observations (left) and DFTs (right) of GW Librae between 2010 and 2013 show a prominent 80 minute periodicity that is believed to originate from the accretion disk, but no sign of the expected pre-outburst pulsation frequencies. Several runs show relatively coherent 4 – 6 minute periodicities, which may be associated with non-radial pulsations. The vertical plot scales are equal for all plots, and only labelled once to reduce visual clutter. The 0.1% false alarm probability for each run is indicated by the red line in the DFT panels, and the red lines in the time series panels indicate the double-humped 80 minute periodicity (type A) that is visible in many of the runs.

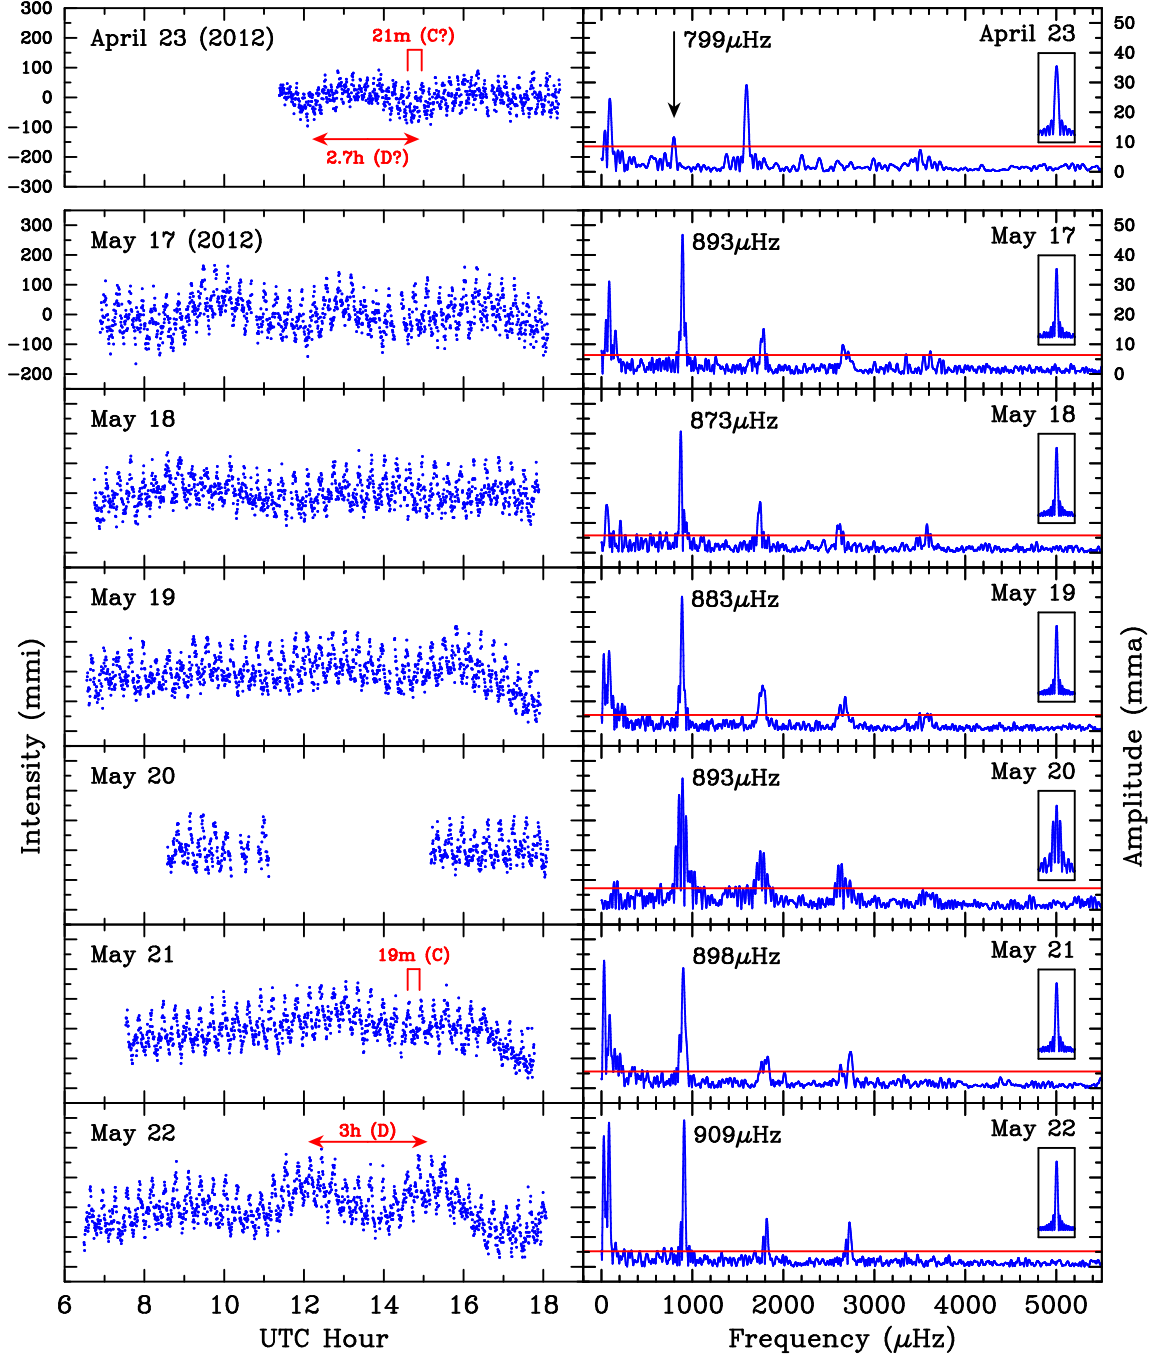


Figure 5.3: Observations (left) and DFT (right) of GW Librae in May 2012 were dominated by a ~ 19 minute periodicity that was not visible in our 2011 or 2013 observations. The vertical plot scales are equal for all plots, and only labelled once to reduce visual clutter. The two main periodicities are marked on the time series panels, and the 0.1% false alarm probability for each run is indicated by the red line in the DFT panels. The April 23 run is also included in this figure, as it raises interesting questions as to the origin of this periodicity (see text).

Our next observation of GW Lib was not until March 2013, and these runs showed no signs of the 19 minute signal. The time span of nearly 10 months between these two observations makes it difficult to constrain the length of time that it was present, and so further data on this signal may remain elusive.

5.3 The 19 Minute Periodicity

It became clear early into our analysis of the may12 data set that this signal couldn't be described using a simple coherent pulsator model. Each nightly run showed slightly different frequencies, and the apparent structure in the DFT was found to behave more like artefacts than as unresolved multiplets, which indicated that the stability timescale is shorter than even 12h. This posed a puzzle: the models for g-mode pulsations, based on isolated DAV pulsators, do not allow for such large frequency changes on this short timescale.

A search through the literature revealed that a very similar periodicity was seen over a period of several months in 2008 (see Copperwheat et al., 2009; Schwieterman et al., 2010; Bullock et al., 2011; Vican et al., 2011). Vican et al. comment that the signal could have been present for anywhere between four months and two years. These observations showed the same apparent instability that we see in our may12 data, which led Bullock et al. to attribute its origin to a quasi-periodic oscillation (QPO; see e.g. Warner and Woudt, 2008) associated with the accretion disk.

Vican et al. note that this would be unusual, as QPOs are usually only observed in systems with much larger accretion rates than GW Lib. They explain the common-sense fact that these accreting systems are sufficiently different from the isolated DAVs that we must be careful about drawing too-close of a comparison to them. Essentially, these authors note that the cause of the observed signal is uncertain, and we cannot rule out non-radial WD pulsations.

The case for non-radial pulsations is strongly supported by observations of another accreting pulsator, SDSS 0745+4538 (Mukadam et al., 2007). The lightcurve of SDSS 0745+4538 is remarkably similar to what is seen in our may12 data, with a non-sinusoidal oscillation near 20 minutes and a DFT that shows a strong peak with multiple harmonics. This period and its harmonics is observed to vary between nights in a similar fashion to what is seen in our may12 data. Given this observational similarity, it seems reasonable to attribute the same driving mechanism to the periodicities in both objects.

SDSS 0745+4538 underwent an outburst in October 2006, and the variability ceased. It later returned, with essentially identical properties, once the system returned to

quiescence in 2010 (Mukadam et al., 2011b). This is a strong sign that the signal may indeed be caused by non-radial pulsations: the pulsation frequencies are determined primarily by the interior structure of the WD, and the outburst heating simply removes the partial-ionization driving mechanism. The same periods are expected to return when the temperature at the base of the partial ionization zone returns to its quiescent level.

Mukadam et al. argue that the apparent frequency variability in SDSS 0745+4538 can be explained in terms of amplitude modulation across an underlying stable set of pulsation eigenmodes. They propose that only one mode is excited at any time, and that the apparent frequency shifts are caused by power shifting between these modes on a timescale of less than a day.

We find an additional complication arises when we consider our sole observation in the April 2012 run. This showed a strong 672 s periodicity, which is tantalizingly close to the 650 s pre-outburst pulsation. On more careful inspection, however, we see that this is actually the first harmonic of a double-humped 21 min periodicity. Given the short time span between the April and May observations, it seems quite plausible that these are related. There is a sign of power at 21 min in our March 2012 observations, but the signal is too weak in these relatively poor quality data to be able to make any quantitative statements about it. The spread of frequencies of the main 19 min periodicity in the may12 run was $\sim 50 \mu\text{Hz}$; the April observation suggests that this might vary over an even wider range.

Both the April and May observations show a long period intensity modulation: 2.7 h for the April run, and 3 h for the May runs. The ratio of period lengthening ($2.7 / 3 = 90\%$) in this signal matches the ratio of period decrease that is seen in the 19 min signal ($19 / 21 = 90\%$). The significance of this is not clear, but it could potentially suggest that (a) the 19 and 21 minute signals are indeed related, and (b) that they are somehow coupled to this longer period modulation.

Given this uncertainty over the origin of this signal, we try to approach our may12 data set with a view to characterize its phenomenological behaviour, without imposing a specific underlying physical model for its origin.

5.3.1 The Combined DFT

The simplest explanation for apparent frequency modulations in a stable pulsator is that there are multiple coherent but unresolved pulsation modes that undergo beating. The apparent instability can be easily explained once the observation baseline increases sufficiently to uniquely resolve the underlying modes. This baseline is usu-




Figure 5.4: The frequency determination of a stable pulsation can be greatly improved by calculating the overall DFT of several nights of single-site data. However, the daily data gaps introduce significant side (alias) peaks that are separated from the real frequency peak by $\pm 1 \text{ day}^{-1}$ ($\pm 11.6 \mu\text{Hz}$). The DFTs and corresponding window functions in the two panels on the left illustrate the improvements obtained for the very stable 251 s ($3559 \mu\text{Hz}$) pulsation mode in the DBV white dwarf EC 20058–5234 (Sullivan et al., 2008) when moving from a single night to multiple nights. The improvement in the DFT resolution as data is added matches the improvement of the window function. This is in stark contrast to the 19 min period DFTs for GW Lib shown in the right panels: the single night DFT peak is essentially destroyed when multiple nights are combined. Note that the two data sets for EC20058 and GW Lib were comparable in duty cycle and length, as should be apparent by inspecting the window functions.

ally longer than a single night, and so one is usually required to combine multiple nights of single-site observations together. This introduces gaps in the total data set, which introduce artefact peaks with a spacing of 1 day^{-1} ($11.6 \mu\text{Hz}$). If the underlying pulsation is stable, then these aliases will remain enveloped inside the lower resolution peaks from each of the individual run segments.

This is demonstrated in Figure 5.4 for a known stable pulsator (EC 20058-4725; Sullivan et al., 2008), and contrasted against our may12 GW Lib observations. We see that, in contrast to EC20058, the GW Lib DFTs do not show an improved frequency identification as the baseline increases, and that the DFT behaviour when adding additional nights does not match that of the window function.

Figure 5.5 presents an expanded view of $900 \mu\text{Hz}$ region of the may12 DFT. An optimistic frequency identification finds five distinct frequencies on top of a relatively incoherent ‘hump’ of power. The DFT frequencies around the first and second harmonics show even less coherent behaviour than the main peak. This kind of smearing in the Fourier domain is characteristic of an underlying incoherency in the signal of interest.




Figure 5.5: A combined DFT of the full may12 dataset does not resolve a coherent underlying frequency structure. The majority of the frequency structure can be fitted with five frequencies, but a large amount of power remains. The harmonics (with frequency scales of $2\times$ and $3\times$ the bottom panel) that are clearly visible in the individual runs become smeared over a wider range of frequencies and lack a clear structure.

We can further demonstrate that this identification is a poor description of the data by generating a simultaneous linear least squares fit of these five sinusoids to the time series data. The harmonic frequencies act to change the peak shape, but won't have any significant effect on the overall phase and amplitude characteristics so they can be safely neglected. The result, shown in Figure 5.6, does a reasonable job of fitting the phases over many (but not all) of the nights, but does a poor job of fitting the amplitudes. In particular, it requires amplitude changes within the runs that are not reflected in the data, and forces the observed variations to disappear entirely at times that are coincidentally spaced during gaps in our observations.

These results make it quite clear that these variations cannot be described by a set of stable pulsation modes. This indicates that the pulsation mode (or modes) must be unstable in either frequency or amplitude. The next simplest model is the one proposed by Mukadam et al. (2011b), where there are multiple stable frequencies that have unstable amplitudes. This is potentially consistent with the results of our combined DFT – the five peaks could potentially correspond to the underlying frequencies, and the amplitude modulation introduces the incoherent frequency artefacts. The next step in our analysis should therefore be to constrain this behaviour further.




Figure 5.6: A five-frequency sinusoidal fit to the full may12 dataset requires strong beating effects in order to adequately describe the observed frequencies and phases. This is an improbable model at best.

5.3.2 Mode Stability

A useful technique for visually investigating the stability of a signal is to generate a spectrogram. Spectrograms show a running plot of the signal intensity (using colour) over a range of frequencies as a function of time. They are commonly used in audio signal analysis, but have also been applied to WD pulsations (see e.g. Provencal et al., 2012).

Five of the six nights (we exclude May 20, which was affected by cloud) contain unbroken runs of 10+ hours, which is enough to produce a meaningful spectrogram for each night. Each vertical slice in the spectrogram measures the DFT intensity of a 1.5 h lightcurve segment centred at the specified time. This segment length is the shortest that we can use while maintaining a reasonable frequency resolution, and so this defines the minimum time resolution of the spectrogram.

The long period variability (D) was first removed from the data using the spline fitting tool `ts3fix` (Sullivan et al., 2008), and then spectrograms were calculated for each run. These plots, shown in Figure 5.7, appear to show similar behaviour in all five runs: the main period remains stable for several hours, before transitioning to a different stable frequency. We can constrain the upper limit of the transition time to within the 1.5 h length of our spectrogram window. The harmonics appear to track the main period, which indicates that they are indeed not independent modes. The spectrogram from the 18th (Figure 5.7, top-centre panel) appears to show the harmonics growing from a small amplitude after the main frequency shifts.

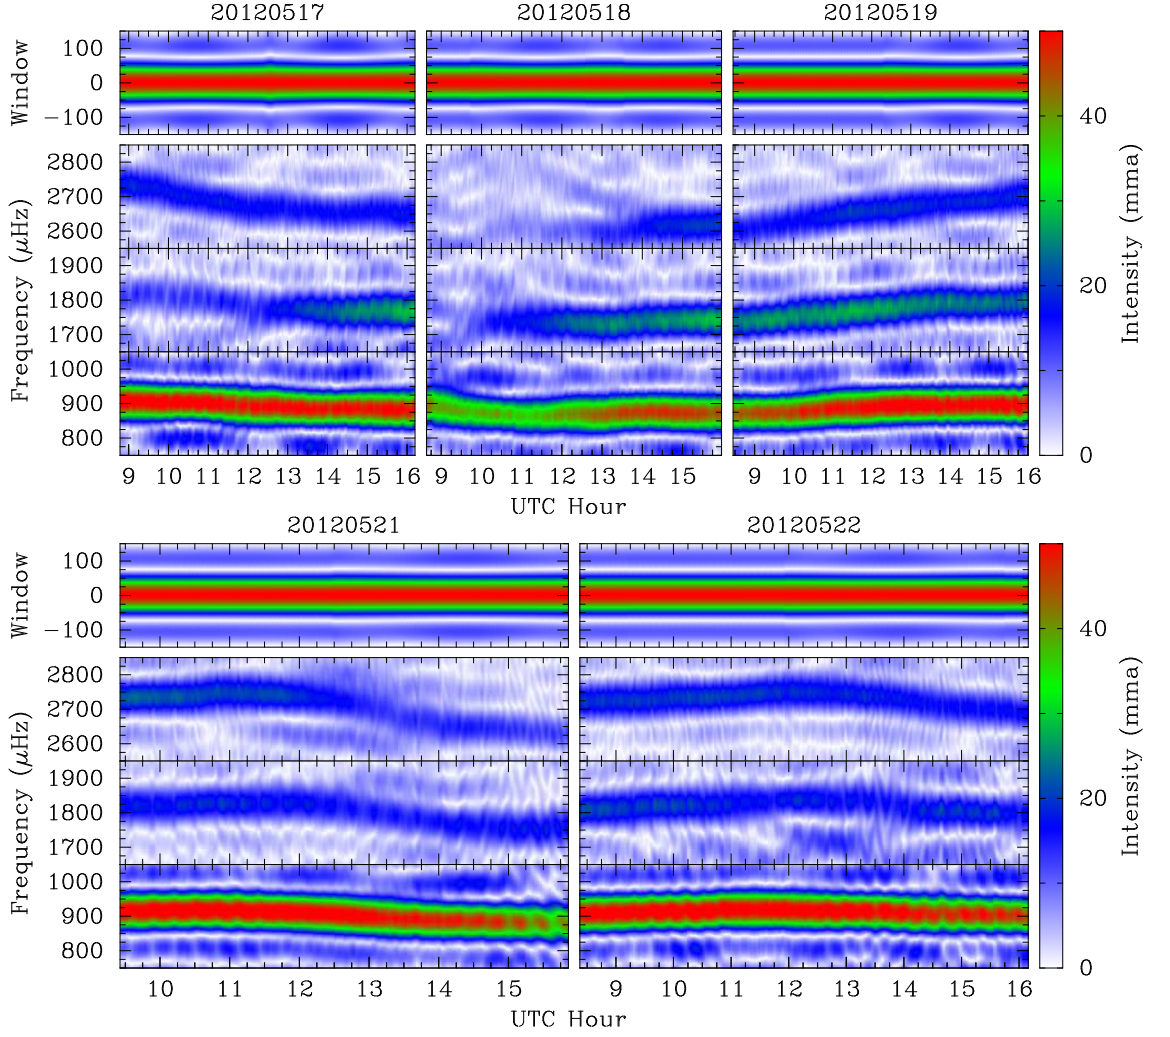


Figure 5.7: Spectrograms of the 19 minute GW Lib periodicity and its harmonics.

Each night was then split at the mid-point of each of the apparent transitions so that each apparently-stable sub-run could be analysed separately. The two short data segments acquired on the 20th were also included as separate runs. These sub-run segments, detailed in Table 5.3, feature much cleaner DFTs (see Figure 5.3; top panel), and show none of the artefacts that were previously visible in the full nightly runs. The harmonic peaks are also artefact-free, with frequencies of exactly $2, 3, 4\times$ the base frequency. The frequency (period) change between adjacent sub-runs is quite significant – $25\,\mu\text{Hz}$ (30 s) or more, compared with the $\sim\mu\text{Hz}$ formal uncertainty in the frequencies within each sub-run.

The individual frequencies within the sub-runs appear to fall into a regularly spaced sequence, with the most common frequencies being $900\,\mu\text{Hz}$ (four sub-runs), $871\,\mu\text{Hz}$ (three sub-runs), and others near $880, 885, 895, 905,$ and $915\,\mu\text{Hz}$. The physical significance of this is not clear - if this was a rotationally split multiplet it would imply an $\ell = 5$ mode, which should be highly suppressed by geometrical cancellation effects. It would be an unprecedented discovery if this were true, but we do not

Date (UTC)	Time (UTC)	Run Name	Primary Frequency ($\mu\text{Hz} \pm 3\sigma$)	Primary Amplitude ($\mu\text{Hz} \pm 3\sigma$)	Additional Frequencies (μHz)
17 May	06:53 - 12:00	20120517a	903.4 ± 1.9	50 ± 5	
	12:00 - 18:06	20120517b	884.5 ± 1.3	47 ± 4	3615
18 May	06:43 - 09:00	20120518a	900 ± 13	47 ± 8	
	09:00 - 11:00	20120518b	850 ± 53	30 ± 7	3555
	11:00 - 17:52	20120518c	870.1 ± 1.1	45 ± 4	3569
19 May	06:33 - 12:00	20120519a	872.3 ± 1.9	45 ± 7	3609
	12:00 - 17:55	20120519b	894.5 ± 1.0	50 ± 3	
20 May	06:06 - 11:07	20120520a	902.8 ± 3.4	52 ± 6	
	15:10 - 18:06	20120520b	873.7 ± 2.5	42 ± 4	
21 May	07:32 - 13:00	20120521a	912.5 ± 0.9	51 ± 3	
	13:00 - 17:46	20120521b	879.4 ± 2.1	42 ± 4	
22 May	06:28 - 10:00	20120522a	906.2 ± 2.3	51 ± 5	
	10:00 - 14:00	20120522b	915.3 ± 1.8	54 ± 4	
	14:00 - 18:03	20120522c	898.4 ± 2.8	45 ± 5	

Table 5.3: Best-fit frequencies for the ~ 19 minute periodicity seen during our may12 run. The non-sinusoidal nature of the lightcurve was accounted for by fitting up to the third harmonic for the 17–20 May data, and the second harmonic for the 21–22 May data. Additional frequencies were identified above the FAP level in each plot, and those that showed believable behaviour in the phase analysis were included in the ‘additional frequencies’ column.

believe this to be a realistic interpretation.

The uncertainties in Table 5.3 were calculated using a Monte Carlo procedure. The time-series data were first pre-whitened by the best-fit primary frequency and harmonics, and the remaining signal was fitted with a Gaussian (Figure 5.8; panel A) to characterize the noise statistics. A synthetic lightcurve was then generated using the modelled frequencies and random noise sampled from the Gaussian parameters, and a new frequency determination was made using a non-linear least-squares fit. This procedure was repeated 1000 times, and a histogram of the best-fit frequencies (panel B) and amplitudes (panel C) were fitted with Gaussians to obtain a statistical uncertainty estimate for our original frequency and amplitude measurements.

These result do appear to be consistent with the proposal of power shifting between underlying stable modes, and appears to show us that the timescale for power shifting between modes is very fast (< 2 hours).

Pre-whitening the 19 min periodicity and its harmonics from each night’s data reveals the presence of the higher frequency oscillation (B) in several of the sub runs; these are indicated in Table 5.3, and shown for one of the sub-runs in Figure 5.8.




Figure 5.8: The Monte Carlo procedure used to calculate the values in Table 5.3 is demonstrated for the 20120517b sub-run. See text for details.

These accurate sub-run period determinations allow us to fold the sub-run lightcurves at the base period to obtain the mean pulsation shape. Figure 5.9 plots these for the two May 17 sub-runs, and these are representative of the shapes found for the other sub-runs. We see a double-humped profile, with the secondary peak offset significantly in phase from the first. Also note the large difference in period (24 s) between the two May 17 sub-runs.

The folded lightcurve of the April run shows a similar double-humped profile, but the secondary peak shifts to an earlier phase. If this longer period 21 min periodicity is indeed related to the 19 min periodicity then we can rule out eclipses as the cause of the double-humped modulation (but this was unlikely in any case, given the low inclination of the GW Lib system).

5.4 The 4–6 Minute Periodicity

Several of our runs in 2011 and 2012 show significant peaks in the region of 3000 – 4000 μHz . Szkody et al. (2012a) show that these were visible in additional optical observations from Apache Point observatory and Kitt Peak, and in UV observations from the Hubble Space Telescope.

Much like the 19 min period, these shorter (4 – 6 min) periodicities were seen to vary between nights, and the DFT peaks again showed signs of messy artefacts. We can apply the same spectrogram technique to this region of the DFT in order to visualize the stability of these modes; however, many of these runs are much shorter than our may12 nights, and so these plots are somewhat less useful here.




Figure 5.9: The lightcurves from the two sub-runs on the 17th of May are folded at the primary pulsation period to produce a mean pulse shape. This double-humped profile is compared with the folded lightcurve from the April 23 observations.

Figure 5.10 shows the spectrograms of several selected nights, and show the (completely unsurprising) result that these frequencies shift in a similar manner to the 19 min periodicity. This seems to support the idea that there is a common mechanism behind these frequency shifts, and that it is indeed associated with the WD.

5.5 Discussion

GW Librae has been an interesting target to observe from a phenomenological viewpoint, but attempting to extract anything quantitative from our data has been challenging. I have instead aimed to present the data, and to try to explore what it can tell us rather than prescribing any specific model.

The biggest question raised by these observations is the origin of the driving mechanism for these quasi-stable intensity modulations. It is our belief (which is undoubtedly biased by our interest in the regular WD pulsators) that these are surely related to the WD, and not the accretion disk. WDs are known to be stable oscillators, and it is not unrealistic to assume that the additional complexities introduced




Figure 5.10: Spectrograms of the 4000 – 5000 μHz region of the DFT of our 2011 – 2013 data sets appear to reveal similar frequency transitions as were seen in the 19 min periodicity. The may12 runs (labelled with *) were pre-whitened of the 19 min period and harmonics in order to remove the $\sim 3600 \mu\text{Hz}$ harmonic from this region.

by accretion act to reduce this stability in a way in which the details are not yet clear. This viewpoint appears to be supported by observation: the 4 – 6 min signal is visible in the UV wavelengths normally associated with the WD luminosity, and SDSS 0745+4538 saw its similarly-behaving signals disappear and reappear following an outburst in the same manner that is expected for true non-radial pulsations.

The other important result that we take away from this is that the standard analysis tools and procedures that we take for granted with the isolated WD pulsators do not necessarily apply to these systems. They appear to be unstable on short timescales (even for the ‘stable’ non-radial pulsations; see the earlier comments from the introduction), and so increasing the observational baseline by combining multiple consecutive nights will not be useful in identifying any underlying frequency structure. It is clear that undertaking a multi-site WET campaign, for example, would not be productive.

Since we have so few well-studied examples of this type of system at the current time, it is uncertain whether the behaviour shown by GW Lib is normal for the accreting pulsators, or whether it is some strange outlier that we may never adequately explain. We hope that matters will become clearer in the future once more of these systems have been observed, and more detailed models have been produced.

Chapter 6

L19-2

L19-2 (also known as MY Aps or WD 1425-811) is a bright southern DAV white dwarf. It has a declination of -81° , which makes it circumpolar from MJUO (-43° latitude). This combines with its relative brightness ($V = 13.8$) to make L19-2 a convenient year-round backup target that can be usefully observed during inclement weather conditions. We acquired several runs during 2011 and 2012, and then shifted our focus primarily to this target during our 2013 observing sessions.

Our goal for L19-2 was to connect our new observations with the existing archival data set (which contains observations going back to the mid-1970's) in order to track the phase of the main pulsation modes over time. These phase measurements are expected to contain the signature of a lengthening period caused by evolutionary cooling. This picks up a project that DJS started during the 1990's (Sullivan, 1998a, 1999; Sullivan and Sullivan, 2000), but which had dropped off the priority list after it became clear that the primary target of interest (EC 20058-4725) would not reveal its secrets easily.

The main pulsation properties of L19-2 were identified in the 1980's through a targeted photometric observation campaign from the South African Astronomical Observatory (SAAO) (O'Donoghue and Warner, 1982, 1987). The xcov12 WET run in 1995 confirmed and improved these earlier results (Sullivan, 1998a). L19-2 pulsates with ten detectable eigenmodes, which are grouped into a triplet with periods near 193s, two doublets near 143 and 119s, and an isolated 350s mode. These modes are presented in Table 6.1, and shown graphically in Figure 6.1.

Asteroseismic models of these eigenmodes (e.g. Bradley, 2001; Romero et al., 2013) broadly agree with independent spectroscopic models (e.g. Koester and Allard, 2000), and show that L19-2 is a relatively massive WD near the hot edge of the DA instability strip: with $M \approx 0.75 M_\odot$, and $T_{\text{eff}} \approx 12100$ K.

Label	Frequency (μHz)	Period (s)	Amplitude (mma)
f_1	2855.9 (2855.9)	350.15 (350.16)	0.77 (0.79)
f_8	5178.8 (5178.9)	193.09 (193.09)	0.95 (1.00)
$f_2 \star$	5191.8 (5191.8)	192.61 (192.61)	5.92 (5.74)
f_9	5204.7 (5205.2)	192.13 (192.12)	0.61 (0.67)
f_3	6972.5 (6973.0)	143.42 (143.41)	0.49 (0.41)
f_4	6991.4 (6990.7)	143.03 (143.05)	0.41 (0.39)
f_5	8426.3 (8426.3)	118.68 (118.68)	1.34 (1.41)
f_6	8437.4 (8437.5)	118.52 (118.52)	1.72 (1.51)
$f_7 \star$	8789.1 (8789.1)	113.78 (114.78)	2.05 (2.14)
f_{10}	8828.7 (8827.9)	113.27 (113.28)	0.43 (0.32)

Table 6.1: Ten pulsation modes have been identified in L19-2. The periods and amplitudes values are shown for the 1995 xCOV12 observations, and compared against our combined 2013 determinations (in parentheses). The modes marked with a \star were selected for analysis using the O–C technique.




Figure 6.1: A Fourier transform of the xCOV12 data clearly shows the five distinct frequency groups in L19-2. Pre-whitening these main frequencies reveals additional modes that are introduced by rotational splitting.

The pulsations exhibited by L19-2 are notable for two reasons: the first is the asymmetric splitting of the 192 s mode (f_2). This asymmetry is small ($\sim 1\%$), but not negligible. This has been interpreted as a sign of second-order rotational effects (O'Donoghue and Warner, 1987), but it is not clear whether the models for these effects are physically realistic. The second notable aspect is that the two strongest modes (f_2 and f_7) show extremely stable pulsation periods and amplitudes. This makes L19-2 an important laboratory for studying the effects of evolutionary cooling.

6.1 Evolution of Pulsation Periods

A small number of WDs have been observed to feature extremely stable pulsation modes: these WDs rank with pulsars as being the most stable astrophysical clocks in the sky. Figure 6.2 compares the relative pulsational stability of several WDs and pulsars. The period of these WD pulsation modes are expected to show extremely small secular changes as the star evolves – primarily due to cooling as the WD radiates away its stored thermal energy.




Figure 6.2: A comparison of the pulsation stabilities of certain astronomical clocks. The cross-hatched regions indicate the theoretically expected regions for the (pre-)WD pulsators. Reproduced from Sullivan (1998b), and updated with recent \dot{P} values.

The earliest constraints on these evolutionary period changes were made by Stover et al. (1980) for R548, and Kepler et al. (1982) for G117-B15A. More recent mea-

surements have improved these constraints to definite measurements (not just upper limits) that are of the order of $10^{-15} \text{ s s}^{-1}$ (Kepler et al., 2005; Mukadam et al., 2013). Several additional stable pulsators have been discovered (see e.g. Mullally et al., 2008), but the observational baseline for these objects is not yet long enough to definitively measure their period rate of change (\dot{P}).

These most stable pulsators appear to be restricted to the hot DAV stars (Winget and Kepler, 2008). Theoretical considerations based on ‘simple’ WD models give a picture that *all* pulsations should be stable up to evolutionary timescales, but observations show that this isn’t the case. The decreased stability is likely to be associated with higher order effects that are not included in the models, such as weak magnetic fields, turbulent convection, diffusion, or other poorly-constrained parameters. The recent discovery of a hot DAV star WD 0111+0018 (Hermes et al., 2013a), with a measured $\dot{P} > 10^{-12}$ is also puzzling, and appears to indicate the presence of an as-yet-unknown evolutionary timescale.

In any case, we can avoid difficult questions about these unknown effects by restricting our attention to the few objects that don’t appear to show them. The dominant effect on these DAV and DBV pulsators is evolutionary cooling: the core degeneracy increases as the WD cools, which causes the Brunt-Väisälä frequency to decrease. This frequency governs the propagation of the g-mode pulsations in the star, and so the pulsation periods will correspondingly lengthen ($\dot{P} > 0$).

Contraction can play an important role in the DOV pulsators, and contributes towards a period *decrease* ($\dot{P} < 0$) as the radial wavelength decreases in step with the radius. The relative size of the cooling and contraction terms change across the extended DOV instability strip (Winget et al., 1983a): the hot DOV pulsators are expected to show a negative \dot{P} dominated by contraction, and cooler DOVs a positive \dot{P} dominated by cooling.

The possibility for mode trapping (Kawaler and Bradley, 1994) adds additional complications to all three classes (DA, DB, DO) of pulsator. A mode can become trapped when the radial wavelength matches the thickness of a compositional layer (e.g. H or He). The mode periods then track the evolution of the compositional layer it is trapped in, with a different rate (and even a different sign, see Kawaler and Bradley (1994)) to untrapped modes.

The simple Mestel (1952) cooling model shows that the cooling rate, and therefore \dot{P} , depend strongly on the mass and chemical composition of the core. More massive WDs have a larger total heat capacity (as there are more ions to store energy), and a smaller surface area through which to radiate away energy. This decreases the surface cooling rate and \dot{P} relative to a less massive WD. The ratio of carbon /

oxygen (and heavier elements) has a similar but opposite effect – increasing mean atomic mass *decreases* the heat capacity of the core (as there are less ions per unit mass), and *increases* \dot{P} . This means that a firm measurement of \dot{P} can be used to infer information about the core composition, which can provide an important constraint on the previous stellar evolution, and factors such as the $^{12}\text{C}(\alpha, \gamma)\text{O}^{16}$ nuclear cross section.

The detection of cooling rates that exceed the expected surface radiative flux can constrain additional interesting physics. A specific example of this is the neutrino flux that is expected to be generated within the cores of hot WDs: photons propagating through the core can couple to the plasma, which allows them to decay into a neutrino and antineutrino via the creation of intermediate virtual electron-positron pairs (Winget et al., 2004). These plasmon neutrinos will then escape the star without further interaction, carrying away energy. This process is expected to dominate the cooling of WDs with $T_{\text{eff}} \gtrsim 25,000 \text{ K}$, and should in principle be detectable in the \dot{P} of hot stable DBV pulsators. Córscico et al. (2012) and Mukadam et al. (2013) applied a similar argument along with the measured \dot{P} for G117–B15A and R548 to constrain the mass of the theoretical axion particle (which is a candidate for dark matter), and also to place an upper limit on the secular variation of the gravitational constant G (Córscico et al., 2013).

Figure 6.2 indicates the expected magnitude of \dot{P} in the three instability strips. Except for the DOV stars, these \dot{P} values are far too small to measure as direct period changes. A significantly more sensitive technique is required.

6.1.1 The O–C diagram

A well established technique for measuring small \dot{P} values is the $O - C$ (observed minus calculated phase) diagram (Kepler et al., 1991), which tracks the *phase* of the pulsation mode.

The ephemeris for a stable pulsation mode can be defined by the equation

$$T_{\text{max}} = T_{\text{max}}^0 + PE,$$

where P is the pulsation period, T_{max}^0 is a reference epoch (time of peak maximum at some known time), and E is the number of pulsation cycles between the current and reference epochs.

The effect of a slowly varying period can be found by expanding this in a Taylor series:

$$T_{\text{max}} = T_{\text{max}}^0 + PE + \frac{1}{2}P\dot{P}E^2 + \mathcal{O}(E^3).$$

The E^3 and higher terms depend on \ddot{P} and higher derivatives, and can be neglected for most stars.

We can recast this into a more useful form by defining the quantities:

$$\begin{aligned} O &= T_{max} = \text{Observed peak maximum (measured),} \\ C &= T_{max}^1 + P_1 E = \text{Calculated peak maximum for a fixed period } P_1, \\ \Delta P &= \text{Error in the estimated period } P, \\ \Delta E_0 &= \text{Error in the estimated } E_0, \end{aligned}$$

and rearranging to produce the standard $O - C$ equation:

$$(O - C) = \Delta E_0 + \Delta P E + \frac{1}{2} P \dot{P} E^2.$$

The $O - C$ term measures the total phase drift (measured as a time) away from the expected value, which we can see increases quadratically with with cycle count (i.e. time). A plot of $O - C$ versus E (called an $O - C$ diagram) is therefore significantly more sensitive than than a direct period calculation, which improves only linearly with time.

The first step in constructing an $O - C$ diagram is to divide the full data set into shorter runs. The natural choices for dividing single-site data are either ‘per run’ (combining multiple adjacent nights to improve the phase fit), or ‘per night’ (giving less accurate phases, but many more points). Our preference is the former choice, but others (e.g. Kepler et al., 1991) appear to prefer the latter.

The local time of maximum, O , can be measured directly from the photometry for each run. The period P can be measured from one of the runs, and is assumed to vary slowly enough that it can be considered constant over the full data set. The initial (relatively poor) estimate of P can be iteratively improved by adding the ΔP term from the $O - C$ fit, and repeating the process until this term vanishes.

The cycle count E associated with each O measurement can be found by dividing the time difference $(O - T_{max}^0)$ by the period, and then rounding to the nearest integer. T_{max}^0 is an arbitrary time that corresponds to the ‘zero’ of the $O - C$ diagram. The $O - C$ offset is then just the fractional cycle offset multiplied by the period:

$$(O - C) = P [(O - T_{max}^0) - \lfloor O - T_{max}^0 \rfloor] .$$

The $O - C$ diagram is sensitive to additional phenomena that may impact the pulsation phase in other ways. The most interesting of these is the reflex motion of the

WD due to orbiting planets, which will show a sinusoidal phase variation over time. Long-term surveys of a number of WDs are underway (e.g. Mullally et al., 2008) with the aim of detecting this motion, and offer the possibility of finding planets that have survived the violent red giant evolutionary phase.

The O–C diagram is a useful tool, but we do need to treat its results with some caution. There have been a number of missteps associated with over-optimistic interpretations of these results and under-estimates of uncertainties. The most well known example is for the DOV prototype star PG 1159-035. The first \dot{P} measurement for the 516 s mode in this star indicated two possible values, one positive and the other negative (Winget et al., 1985). The χ^2 statistics heavily favoured the negative solution, and additional observations appeared to further confirm this result. This became a problem once the models for these stars improved, and showed that a negative \dot{P} should not be possible. An explanation was eventually found by invoking mode trapping effects (Kawaler and Bradley, 1994). The observation baseline eventually became long enough for a direct period measurement to be made (Costa et al., 1999), which showed that the original positive \dot{P} was correct all along! More recent studies (e.g. Costa and Kepler, 2008) show that PG 1159-035 has a mix of positive and negative \dot{P} for the various pulsation modes, and so the mode trapping model is still necessary for this complicated pulsator.

The discovery of the (apparently) stable DBV WD EC 20058-4725 in the Edinburgh-Cape survey provided the opportunity to test the neutrino emission prediction for these stars. Unfortunately, after 15 years of observations on this target it became clear that EC 20058-4725 was not as stable as it originally appeared. The pulsation modes were found to wander cyclically in phase by ± 0.5 cycles or more over this observation baseline, which masks any potential cooling signature in \dot{P} . A recent O–C diagram showing this behaviour is presented in Dalessio et al. (2013).

6.2 Observations

We obtained 145 h of new observations on L19-2 using Puoko-nui South on the MJUO 1.0 m telescope between 2011 and 2013. These new data add to the archival data set on this object, which consists of:

- 75 h of CCD photometry that was obtained by DJS using the original Puoko-nui instrument between 2008 and 2010. These frames were reduced using a modified version of `tsreduce`.
- 2 h of CCD photometry acquired by DJS using the Raptor/“Argos B” instrument on the 40” telescope at Siding Spring Observatory in May 2002. These

frames were reduced using a modified version of `tsreduce`.

- 162 h of PMT photometry acquired from MJUO, SAAO, and Pico dos Dias observatory in Brazil during the 1995 XCOV12 WET run.
- 588 h of PMT photometry from the 0.75 and 1.0 m telescopes at SAAO between 1976 and 1984, and 13 h of PMT photometry acquired from the 0.9 m telescope at Cerro Tololo Inter-American Observatory in August 1976 (O’Donoghue and Warner, 1982, 1987; Hesser et al., 1977). These data were supplied in their original format, which consisted of a plain text file for each observation sequence, with times specified in modified Julian Days. These were transformed into our standard ‘.ts’ data format by grouping the files into longer aggregate runs, and calculating new BJD observation times based on the UTC start times specified in the header of each run file.

Table 6.2 lists the final set of data files that were used in the following analysis. An additional large data set of two and three channel photometry acquired by DJS from MJUO exists, but was not available in time to include in this analysis. This set consists of 11 runs between 1993 and 1999 and less frequent runs between 1999 and 2008. This will fill the large gap between the XCOV12 and CCD observations, and so the results below should be considered preliminary for now.

Seven of the 32 runs (marked with \star) contained insufficient data to adequately constrain the phase of the two modes. They were excluded from further analysis.

6.3 Analysis

The first estimates of \dot{P} for L19-2 were made by O’Donoghue and Warner (1987), who found $\dot{P}_{f_2} = (1.8 \pm 3.0) \times 10^{-14} \text{ s s}^{-1}$ and $\dot{P}_{f_7} = (1.5 \pm 2.0) \times 10^{-14} \text{ s s}^{-1}$ based on 8 years of observations. The other pulsation modes are too weak to reliably track the phases at the precision required for the O–C diagram (Figure 6.6 includes a comparison of the relative phase uncertainty for the 10 modes). Bradley (2001) predicted theoretical values of $\dot{P}_{f_2} = (1.4 \sim 2.3) \times 10^{-15}$ and $\dot{P}_{f_7} = (0.9 \sim 1.4) \times 10^{-15}$ using a seismological model.

L19-2 is a challenging target to apply the O–C technique to, because the primary pulsation mode is a triplet with a spacing of $13 \mu\text{Hz}$ between the rotational components. This is close to the $11.6 \mu\text{Hz}$ (1 day^{-1}) aliases that are unavoidable in single-site observations, which makes it difficult to uniquely resolve the modes and their respective phases.

Sullivan (1998b) showed that these side frequencies, if unresolved, will introduce

Data File	Start Time (BJD - 2440000)	Length (days)	N (points)	Site (%)	Instrument
saao_jul76.ts	2987.25719	6.06 (8%)	4398	SAAO 0.75 m	PMT
saao_aug76.ts	3011.48444	3.17 (17%)	4768	CTIO 0.91 m	PMT
saao_mar79.ts	3954.42743	5.17 (13%)	5718	SAAO 0.75 m	PMT
saao_apr79.ts	3987.31089	6.27 (16%)	8573	SAAO 0.75 m	PMT
saao_jun79.ts	4051.23441	6.33 (25%)	13518	SAAO 1.0 m	PMT
saao_dec79.ts	4219.32736	6.18 (14%)	7675	SAAO 0.75 m	PMT
saao_feb80.ts	4290.45357	5.18 (12%)	5478	SAAO 1.0 m	PMT
saao_mar80.ts	4312.33966	4.30 (36%)	13479	SAAO 0.75 m	PMT
saao_jun80.ts	4394.33560	6.32 (40%)	21968	SAAO 0.75 m	PMT
saao_sep80.ts	4481.23033	17.41 (20%)	30266	SAAO 0.75 m	PMT
saao_oct80.ts *	4516.56070	1.97 (4%)	647	SAAO 1.0 m	PMT
saao_jul83.ts	5522.26745	3.05 (8%)	2192	SAAO 0.75 m	PMT
saao_sep83.ts *	5608.28549	1.06 (14%)	1315	SAAO 0.75 m	PMT
saao_mar84.ts	5760.48848	8.09 (4%)	2747	SAAO 0.75, 1.0 m	PMT
saao_may84.ts	5822.42861	3.13 (8%)	2053	SAAO 0.75 m	PMT
xcov12.ts	9833.24863	8.54 (57%)	58358	WET	PMT
bargos.ts *	12407.87839	1.98 (4%)	742	SSO 41"	Raptor
2008-06.ts	14619.75362	2.07 (15%)	2678	MJUNO 1.0 m	Puoko-nui *
2009-05.ts *	14981.75362	1.08 (7%)	677	MJUNO 1.0 m	Puoko-nui *
2009-06.ts	15008.75350	1.15 (27%)	2662	MJUNO 1.0 m	Puoko-nui *
2009-10.ts	15122.95731	3.17 (10%)	2635	MJUNO 1.0 m	Puoko-nui *
2010-04.ts	15292.79386	4.16 (20%)	7056	MJUNO 1.0 m	Puoko-nui *
2010-05.ts *	15330.75341	2.93 (3%)	814	MJUNO 1.0 m	Puoko-nui *
2010-07.ts	15382.75334	6.08 (10%)	5275	MJUNO 1.0 m	Puoko-nui *
march2011.ts *	15625.42305	0.05 (100%)	214	MJUNO 1.0 m	Puoko-nui
july2011.ts	15745.75340	4.37 (12%)	2229	MJUNO 1.0 m	Puoko-nui
march2012.ts *	16015.22679	0.20 (96%)	559	MJUNO 1.0 m	Puoko-nui
april2012.ts	16041.81541	5.18 (19%)	2802	MJUNO 1.0 m	Puoko-nui
march2013.ts	16359.83420	5.15 (28%)	6236	MJUNO 1.0 m	Puoko-nui
april2013.ts	16390.79398	7.18 (12%)	3861	MJUNO 1.0 m	Puoko-nui
may2013.ts	16420.75331	6.41 (24%)	6509	MJUNO 1.0 m	Puoko-nui
july2013.ts	16453.75367	4.32 (24%)	4460	MJUNO 1.0 m	Puoko-nui

Table 6.2: A summary of the observations that were available for the L19-2 O–C analysis. The length column specifies the multi-night baseline of each data set, and includes the observation duty cycle in parenthesis. Runs marked with * contained insufficient data to constrain the mode phases, and were excluded from the analysis. Puoko-nui * refers to early versions of Puoko-nui from before the redesign work described in Chapter 2.

periodic phase differences into the O–C plot of up to ± 8 s. O’Donoghue and Warner (1987) did not explicitly account for this, but instead relied on the large number of individual phase measurements (one per run) and the constrained parabolic fit to average out these uncertainties. We have chosen to take a more direct approach in our analysis - by combining multiple nights of observations into a single phase measurement we are able to individually resolve the triplet components, and largely avoid this uncertainty.

The f_7 mode is a doublet, but has a larger separation ($39.6 \mu\text{Hz}$), and the secondary component (f_{10}) is much weaker. The effect of this mode on the phase of f_7 is not expected to be a problem.

6.3.1 Proper Motion

An interesting observational detail is that L19-2 has a large proper motion of 0.45 arcsec/year . Figure 6.3 shows the noticeable shift in apparent position over the 37 years that it has been actively monitored.




Figure 6.3: L19-2 has noticeably shifted over the last 40 years. The changing geometry introduced by this large proper motion is expected to contribute towards \dot{P} at the $10^{-15} \text{ s s}^{-1}$ level. Panel (a) shows the position near the time of its discovery, as captured by the Digitized Sky Survey in 1976. Panel (b) shows the position in 2002, observed by DJS using the UT Argos B (Raptor) instrument on the $40''$ telescope at Siding Spring Observatory. Panel (c) shows the position in April 2012, observed using Puoko-nui South at MJUO.

This proper motion also has an implication for our \dot{P} measurements: Pajdosz (1995) first showed that the changing line of sight to the moving WD introduces a positive \dot{P} contribution of $10^{-14} \sim 10^{-15}$. Pajdosz also calculated proper motion \dot{P} values for several stars, including L19-2: with expected rates of $2.2 \times 10^{-15} \text{ s s}^{-1}$ for f_2 and $1.3 \times 10^{-15} \text{ s s}^{-1}$ for f_7 . These values define a lower bound for our expected \dot{P} values.

Run	BJD (days)	Frequency ($\mu\text{Hz} \pm 1\sigma$)	Amplitude (mma $\pm 1\sigma$)	Period (s $\pm 1\sigma$)
saao_jun79	2444054	5191.824 ± 0.012	6.17 ± 0.09	192.6105 ± 0.0004
saao_sep80	2444481	5191.893 ± 0.008	4.98 ± 0.09	192.6080 ± 0.0002
xcov12	2449836	5191.835 ± 0.007	5.96 ± 0.05	192.6101 ± 0.0003
2010-04	2455295	5191.79 ± 0.03	5.72 ± 0.12	192.612 ± 0.001
2013-05	2456424	5191.814 ± 0.013	5.77 ± 0.07	192.6109 ± 0.0004

Table 6.3: Best-fit period determinations for the f_2 mode at different times.

6.3.2 Direct Period Comparison

Direct measurements of a changing period have been achieved for the DOV pulsators (e.g. Costa et al., 1999), but aren't practical for the much smaller \dot{P} DAV pulsators. This can be demonstrated by way of a simple calculation.

Table 6.3 compares period measurements from the best quality data sets spanning 34 years. The periods and uncertainties were calculated using the Monte Carlo procedure described in Chapter 5.

The expected period change across this 34 year baseline (assuming $\dot{P} \sim 10^{-15}$) is only $\sim 10^{-6}$ s – two orders of magnitude smaller than the period uncertainties. It is unlikely that any group could significantly improve the period determination over what was achieved by the 1995 WET run, and so we must instead resolve ourselves to waiting another ~ 3000 years to increase our time baseline by the required two orders of magnitude.

6.3.3 O–C Analysis

Each of the runs in Table 6.2 were individually fitted with sinusoids matching the ten pulsation modes in Table 6.1, and used to determine $O - C$ values for f_2 and f_7 . T_{max}^0 was chosen to be the first maxima in the XCOV12 data set, which conveniently falls near the middle of the full data set range.

Figure 6.4 presents the O–C plot of our incomplete data set. We find formal linear least squares fits of $\dot{P}_{f_2} = (3.8 \pm 2.3) \times 10^{-15}$ and $\dot{P}_{f_7} = (0.07 \pm 0.8) \times 10^{-15}$, but are somewhat skeptical of these values (as discussed below). We consider a more appropriate limitation provided by the plotted constraint curves, which show $\dot{P}_{f_2} \lesssim 1 \times 10^{-14}$ and $\dot{P}_{f_7} \lesssim 0.5 \times 10^{-14}$.

The uncertainty in the phase / time is quantified using a Monte Carlo simulation. The lightcurve from each run is first pre-whitened by the fitted modes to extract the noise signal for the run. A histogram of the amplitudes shows the expected Gaussian




Figure 6.4: The preliminary O–C diagram for the f_2 and f_7 modes confirms the stability of the f_2 mode, and shows a potential glitch in the f_7 mode. The solid red curves show the fitted $\dot{P}_{f_2} = (3.8 \pm 2.3) \times 10^{-15} \text{ s s}^{-1}$ and $\dot{P}_{f_7} = (0.07 \pm 0.8) \times 10^{-15} \text{ s s}^{-1}$. A more realistic constraint on \dot{P} is provided by the dashed grey curves, which indicate $\dot{P}_{f_2} < 2.5 \times 10^{-14}$ and $\dot{P}_{f_7} < 10^{-14}$.

distribution for each run. A synthetic lightcurve is created with the same sampling times as the original data, using the fitted mode amplitudes and phases plus a random noise contribution sampled from the noise distribution for each measurement time. If one assumes that our original model is correct, then this procedure amounts to randomly shuffling the noise while leaving the coherent signal undisturbed. The synthetic lightcurve is then fit with a new set of sinusoids to obtain a (slightly different) set of amplitudes and phases. This process is repeated 1000 times to obtain reasonable statistics, and the 1σ uncertainty is found by fitting the phase distribution with a Gaussian curve. This procedure is demonstrated for the xCOV12 data set in Figure 6.6.

Note that this is a statistical uncertainty that only quantifies our ability to fit a set of fixed frequency sinusoids to the data. It doesn't account for any external uncertainties in the measurement times, or for correlated errors between modes. This is therefore a significant underestimate of the total error, and this can be clearly seen by the large scatter relative to the error bars in the SAAO f_2 measurements presented in Figure 6.4. This scatter appears to come largely from the format of the SAAO observing programme, which consisted of short nightly runs (~ 2 h) that poorly resolve the main triplet. This will have introduced a correlated uncertainty between the phases of f_2 , f_8 , and f_9 which is not explicitly accounted for in this estimate of the uncertainty.

It is straightforward to demonstrate that these data do not suffer from the cycle counting errors that contributed to the original incorrect \dot{P} determinations of




Figure 6.5: Cycle counting errors can be unambiguously ruled out from our incomplete data set on L19-2. The frequency determination from the xCOV12 DFT is sufficiently precise to allow the 2008 – 2013 data sets to be unambiguously connected. The best-fit periods near ± 1 cycle / 13 years ($\pm 9 \times 10^{-5}$ s) introduce clear linear trends in the SAAO and MJUO observations, which rule these out as valid solutions.

PG 1159-035. This is shown in Figure 6.5.

The 2008 – 2010 MJUO data appears to show a strange glitch in the phase of f_7 (Figure 6.4; bottom panel). We need to be cautious about this: if it is real, then it would invalidate the \dot{P} estimate for this mode. Most of this apparent feature comes from only two data points (2008-06 and 2009-06). These were both acquired using the original Puoko-nui instrument, which had a known issue that could result in frame times that were delayed by one exposure period (10 s in this case). The f_2 data is also consistent with a 10 s delay, and so this ‘glitch’ may simply be a positive detection of an instrument timing error. The MJUO PMT data will help resolve this question by filling in the gap between the 1995 and 2008 data. If this phase measurements during this period show the expected gradual change, then we will be able to correct or discard these two suspect data points.

If we take the linear least squares fits at face value, then our measurement of \dot{P}_{f_2} is completely consistent with the expected proper motion, and this would imply that the evolutionary \dot{P} is $\sim 1 \times 10^{15}$ or less. The fitted \dot{P}_{f_7} is inconsistent with proper motion by an order of magnitude, but this is quite clearly due to insufficiencies in the fit rather than in the theory.

The best conclusion that we can confidently present from this preliminary analysis is that L19-2’s f_2 mode is indeed stable, and that our measured O–C results clearly show $\dot{P}_{f_2} \lesssim 1 \times 10^{-14}$. We will be able to further improve this estimate once the full data set is available and a more careful consideration of the phase uncertainties has been made. We intend on finishing this analysis and publishing the final results in the coming months.

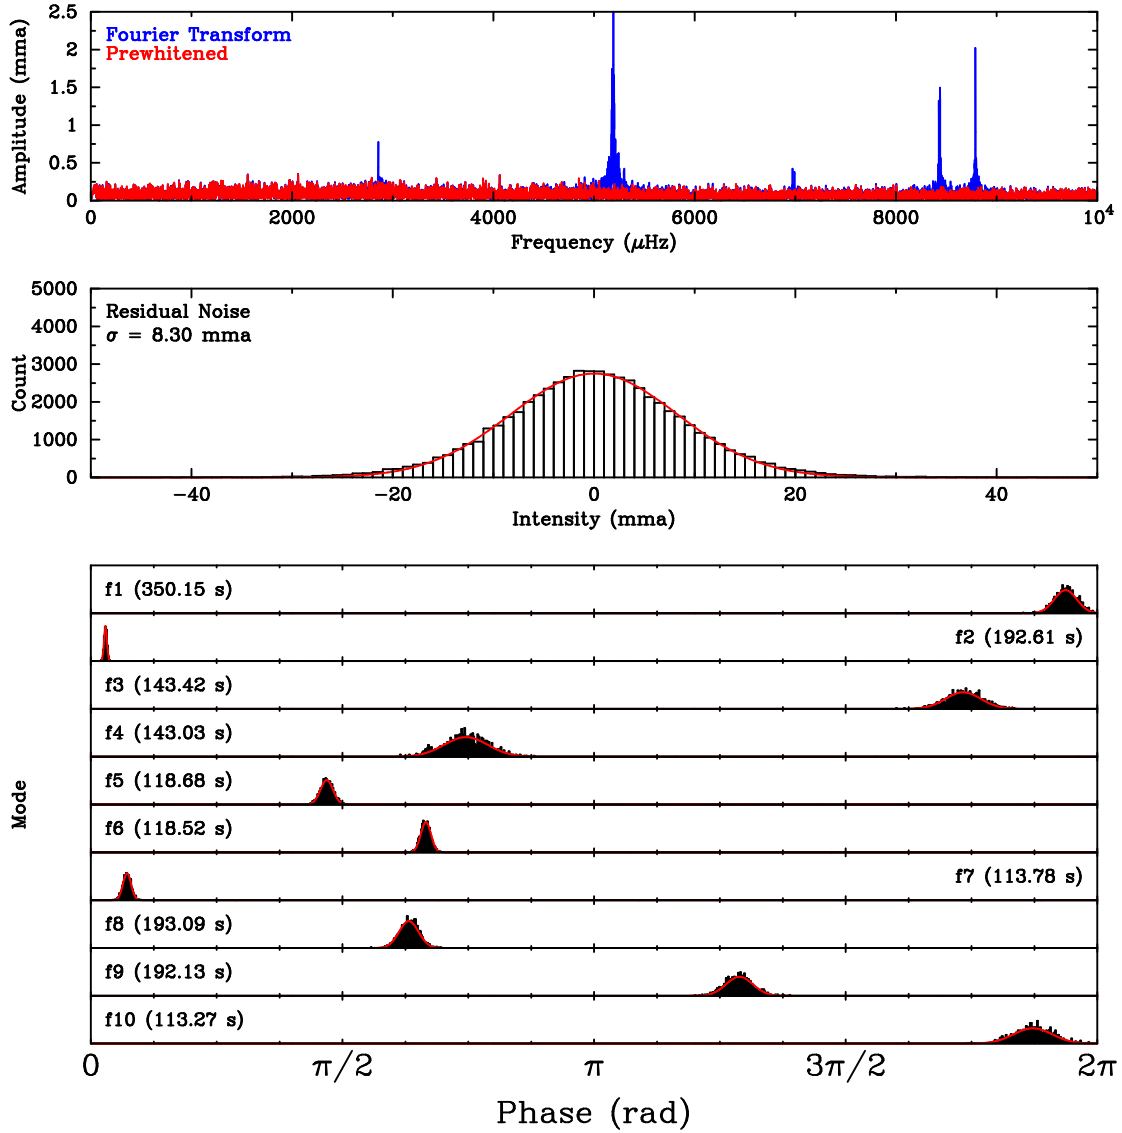


Figure 6.6: A demonstration of the Monte Carlo procedure used to estimate the phase uncertainty of the xCOV12 data set (see text). The top panel shows the DFT before and after pre-whitening. The middle panel shows that the pre-whitened lightcurve is consistent with white noise. The bottom panels show the phase distributions obtained from 1000 simulations.

Chapter 7

Other Observations

This chapter collates a number of other observations which we have made using the Puoko-nui instruments, and which did not require a detailed analysis on my part. These were either contributions towards larger projects, self contained instrument verification tests, or observations that we started but then chose not to continue.

7.1 ELM White Dwarfs

Extremely Low Mass (ELM) white dwarfs are a relatively rare subset of white dwarfs, with $M < 0.25M_{\odot}$. A recent survey for these objects (Brown et al., 2010, and subsequent papers) has identified a number of new WDs in this group, and this has opened the way for observing and understanding these objects in more detail.

ELM WDs are believed to form from a more massive progenitor in a binary system. The companion causes strong mass loss to occur at the start of the red giant branch, which prevents the star from fusing helium into heavier elements. The resulting star is left with a helium core covered with a thick hydrogen atmosphere (Iben and Tutukov, 1985). While isolated low mass stars are theoretically capable of evolving into WDs, their evolutionary timescale is too long for this to have occurred within our galactic neighbourhood.

There are two types of ELM WD that are of particular interest. The double degenerate binary systems (ELM with another, more massive WD or neutron star) have small orbital radii and periods, but remain detached. This can lead to tidal distortions that introduce periodic variations at half the orbital period (caused by the changing area of the WD disk). The intensity of these variations can be related to the mass and radii of the binary system, which can be combined with spectro-

scopic measurements to extract radius measurements for the individual WDs. A small number of systems have been observed to eclipse, which allows additional information to be extracted.

These systems are also expected to lose angular momentum via gravitational radiation at a significantly faster rate than wider binaries, and this should lead to a detectable period change over a span of around a decade.

ELM WDs can also sustain g-mode pulsations within a particular instability strip (Steinfadt et al., 2010). These pulsations have much longer periods than the more massive WDs, and probe deep into the core of the star. These pulsation modes should therefore place strong constraints on the core mass and composition. The first pulsating ELM WD was discovered by Hermes et al. (2012b), and several additional pulsators have since been detected.

We have contributed observations towards two ELM systems.

7.1.1 SDSS J075141.18-014120.9

SDSS J075141.18-014120.9 (Brown et al., 2013), subsequently J0751, is double-degenerate binary system, with masses of 0.19 and 0.97 M_{\odot} . It does not exhibit g-mode pulsations, but shows ellipsoidal variations with a primary period of 1.9 h.

We were approached about observing this target in 2012 after earlier observations (both optical and radial velocity) had appeared to show poor phase coherency over time. One theory to explain this was that there may have been an undetected third body in the system; near-continuous observations over a week or two could in theory constrain the presence and/or properties of this potential companion.

We acquired 17.9 h of photometry towards this goal during two observing sessions in December 2012 and January 2013. This data is shown in Figure 7.1. Nothing conclusive was visible from the O–C phase diagram, but our data supplemented the 45.3 h of photometry acquired from McDonald Observatory as part of a larger study on the properties of these double degenerate systems. Excerpts of our observations and the total combined data set are shown in Figure 7.2.

Two particularly interesting results came from these combined observations: the system was in fact an eclipsing binary (the fifth known eclipsing ELM system), and the properties of this system make it very likely to be one of the first discovered progenitors of the AM CVn systems (Kilic et al., 2014). J0751 is expected to undergo a type Ia or .Ia (as defined by Bildsten et al., 2007) supernova in $\sim 10^8$ years.




Figure 7.1: 17.9h of photometry were acquired on the double-degenerate ELM WD J0751.




Figure 7.2: The full WD J0751 data set is folded at the 1.9h orbital period and binned by phase. The fitted curve includes ellipsoidal variations, eclipses, and Doppler beaming. Plot is adapted from Kilic et al. (2014).

7.1.2 SDSS J151826.68+065813.2

SDSS J151826.68+065813.2 (Hermes et al., 2013b), subsequently J1518, is a pulsating ELM WD which shows signs of multiple pulsation frequencies as well as strong non-linear effects due to convection. This makes the object particularly interesting - the pulsation eigenmodes offer the chance of constraining the interior structure via asteroseismology, and a similar non-linear analysis that was applied to EC 04207 could offer new insights into convection within WD atmospheres.

J1518 was the focus of the xCOV29 WET campaign in April and May 2013. Our contributions towards this campaign was 27.6 h of photometry which is shown in Figure 7.3. The analysis of the full data set is underway by the WET group, with a journal paper expected in the near future.




Figure 7.3: 27.6 h of photometry were acquired on the pulsating ELM WD J1518 for the xCOV29 Whole Earth Telescope campaign.

7.2 Crab Pulsar

The Crab Pulsar and its surrounding nebula (see Hester (2008) for a thorough review) are the result of a supernova which was observed and recorded by Chinese astrologers on July 4, 1054 AD. Pulsations in the optical regime were first discovered in early 1968 by Cocke et al. (1969), with confirmation by Nather et al. (1969) using the 2.1 m telescope at McDonald Observatory.

Nather went on to become a key figure in the WD community, creating both the Whole Earth Telescope group and the Argos photometer from which we derived Puoko-nui. We therefore considered it particularly meaningful to be able to demonstrate the return of true high-speed photometry capabilities (which were lost in the initial transition from PMT detectors to CCDs) by repeating this earlier observation using the same 2.1 m telescope equipped with Puoko-nui North.

The Crab pulsar is a particularly challenging target to observe: the primary pulsation period is only 33 ms, and so an exposure period < 10 ms is needed in order to resolve the pulsation shape (which is far from sinusoidal).

We determined that a 5 ms exposure stride (4 ms exposure plus 1 ms for frame-transfer) would be possible by reading a 336×128 px window of the ProEM EMCCD with 8×8 binning and the 10 MHz EM readout mode. The column shift time was also decreased from $1 \mu\text{s}$ to $0.8 \mu\text{s}$ in order to keep the frame-transfer time below 1 ms. The camera was operated with the high gain readout plus $1000 \times$ EM gain in order to maximize the camera sensitivity and reduce the readout noise.

An initial attempt at this observation in April 2013 was unsuccessful, primarily due to large amounts of dust in the air (which can be a serious problem at McDonald Observatory) and high airmass (the Crab sets shortly after dark at that time of year). A second attempt in November 2013 was successful, and a sequence of 5000 frames were acquired during a 25 s acquisition run.

A small segment of the raw photometry, shown in the upper panel of Figure 7.4, is relatively disappointing; however, the DFT (lower panel) reveals a strong periodicity at 33.689 ms. Folding the time series at this period fills the gaps between the 5 ms observations and reveals the characteristic pulsar intensity profile. Figure 7.5 shows this folded lightcurve plus composite images of the observation field at three representative phases.

These data conclusively demonstrate the excellent relative timing of Puoko-nui North, as even a small amount (say 0.5 ms) of timing jitter would wipe out the details of the lightcurve. This, in theory, also provides a stringent test on the absolute timing of the instrument. This is actually quite difficult in practice, with the

limiting factor being the UTC \rightarrow BJD time conversion in `tsreduce`. Calculating this time conversion at sub-ms accuracies requires a careful consideration of additional effects such as the observer’s position on the Earth ($\Delta t \sim 20$ ms) and changes in gravitational time dilation as the Earth moves through the gravitational wells of (primarily) Jupiter and Saturn ($\Delta t \sim 2$ ms).

The BJD time calculation was repeated by hand using the Eastman et al. (2010) routines and the JPL Earth ephemeris, which offer much more precision than the analytic approximations provided by SOFA. The time of the main peak was then compared against the Jodrell Bank radio ephemeris (Lyne et al., 1993, 2014), and found to agree to within 2.5 ± 5 ms. The relatively large uncertainty accounts for the numerical precision limitations of 64-bit floating point numbers, and covers the range of times that were found by adjusting the calculation steps in `tsreduce` while attempting to remove several known precision limitations.




Figure 7.4: The raw time-series measurements of the crab pulsar are too sparse for the pulsations to be clearly visible in the time domain. The DFT, however, reveals that there is a strong 33.689 ms periodicity hidden within this data.

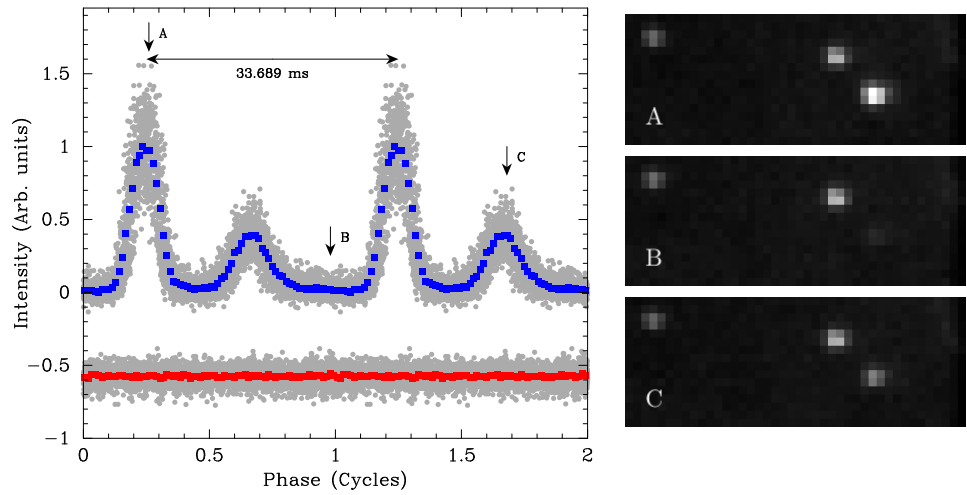


Figure 7.5: Re-plotting the time-series data as a phase of the 33.689 ms periodicity recovers the characteristic pulsation shape of the pulsar. The raw intensities (grey) are then grouped into 75 phase bins for the Crab (blue) and comparison (red). The comparison intensity has been offset by -0.75 units to improve clarity. Images A – C show the average of 100 frames around the indicated phases.

7.3 EC 05221-4725

EC 05221-4725 (subsequently EC05) is a southern DBV white dwarf, discovered by the Edinburgh-Cape survey and presented by Kilkenney et al. (2009). Like EC 04207-4748 (EC04), this target had not been studied in detail following its original discovery, and was too far south to be observed from northern observatories. This made it an ideal potential target to focus on from MJUO.

EC05 appears close in the sky to EC04, and both targets are optimally visible during the summer months. These two targets therefore competed for the limited observing time in which they were visible. We considered EC04 to be a higher priority, and so observations of EC05 were delayed until that investigation (see Chapter 4) was completed.

We obtained 32 h of photometry during two observing sessions in March and April 2012, which are shown in Figure 7.6. The lightcurves from both sessions show significant beating effects, and there were no common frequencies visible in the DFTs of the two runs; this made it clear that we would need to devote a significant amount of observing effort over the summer months in order to resolve this complicated frequency structure. These follow-up observations conflicted with the J0751 observing project, and so we decided to remove EC05 from our observing schedule.

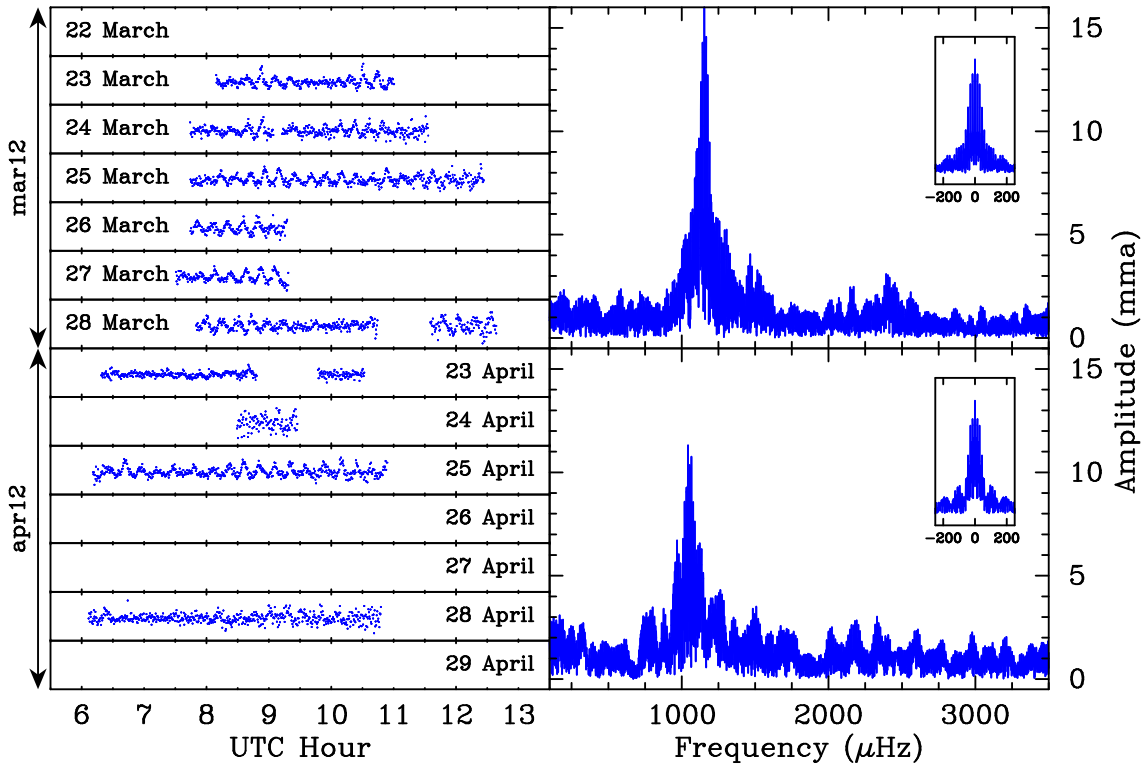


Figure 7.6: Observations of the DBV EC05221-4725 hint towards a number of closely spaced pulsation modes. Further observations conflicted with other priorities, so we made the decision to not continue with this target.

7.4 EC 20058-5234

EC 20058-5234 (EC20058) is a well studied southern DBV white dwarf. The xCOV15 Whole Earth Telescope campaign in 1997 (Sullivan et al., 2008) characterized the pulsation spectrum and many key physical properties of this star. A total of 11 eigenmodes were identified, plus a number of additional linear combination frequencies.

The two strongest pulsation modes (with periods of 281 and 257 s) are remarkably stable, which made EC20058 a potential laboratory for detecting plasmon neutrino emission by way of an additional cooling measured from the O - C diagram (Sullivan, 2005, see also Chapter 6). The results of this study have so far shown that there are additional unknown effects that cause periodic phase variations in these periods, which dominate any changes which may be associated with plasmon neutrino emission (Daleissio et al., 2013).

We obtained 19.4 hours of photometry on EC20058 during July 2011, which are shown in Figure 7.7. This target was never a high priority, but it did fill an observing gap between our two primary targets (GW Librae and EC 04207-4748) in July 2011. Table 7.1 shows the pulsation modes that we detected during these two runs, using the nomenclature and frequencies from Sullivan et al. (2008).

These data, plus additional CCD observations from 2008 – 2010, add to the ongoing long-term period monitoring of this object, but we have no immediate plans for doing any additional work with these data.




Figure 7.7: 19.4 hours of data were acquired on the DBV EC 20058-5234 in July 2011.

Item	Frequency [μHz]	Period [s]	Amplitude [mma]	
			jul11a.ts [0.9]	jul11b.ts [1.4]
f_2	1903.5	525.35	1.6	1.7
f_3	2852.5	350.57	(0.9)	(0.9)
f_4	2998.7	333.48	2.7	2.8
f_5	3489.0	286.62	1.3	1.6
f_6	3559.0	280.98	7.2	7.2
f_8	3893.2	256.86	6.4	7.2
f_{10}	4887.8	204.59	2.1	2.5
f_{11}	4902.2	203.99	1.8	1.5
f_{12}	5128.6	194.98	2.4	2.2
f_{14}	7452.2	134.19	1.2	1.7
f_{15}	9021.7	110.84	(0.8)	(0.8)
a	3564.8	280.52	1.7	-
b	3887.6	257.23	2.0	-
c	4881.0	204.88	1.3	-

Table 7.1: Table of the detected frequencies in the two observing runs. The 0.1% FAP threshold for each run is indicated in square brackets, and expected peaks that fall below this threshold are included in parentheses. The peaks labelled $f_2 - f_{15}$ match the frequencies identified by Sullivan et al. (2008), while a – c were not previously identified, and are not considered to be significant.

7.5 PG 1336-018

PG 1336-018 (Kilkenny et al., 2003) is an eclipsing binary system featuring a pulsating sdB star. Its position (-02° declination) and dramatic lightcurve made it a good candidate for concurrent observation using the two Puoko-nui instruments.

Puoko-nui North was operated remotely by Sam Harrold on the 0.6 m robotic telescope at Meyer Observatory (PJMO), and Puoko-nui South was used on the MJUO 1.0 m telescope. In total, we acquired 2.5 h using Puoko-nui North and 4 h using Puoko-nui South on April 12 2013, with 1.5 h of concurrent observations.

The analysis of our MJUO data was complicated by the lack of a suitable comparison star within our field of view. We were fortunate, however, to have near-perfect observing conditions which meant that the ‘single channel’ photometry was free from any short period transparency variations.

The difficulty came in matching the two lightcurves: our MJUO data were difficult to correct for extinction given the dramatic nature of the lightcurve, and the PJMO data featured residual extinction from the large colour difference between the target and comparison stars. The two data sets were brought into agreement by first applying a linear fit to the MJUO data to align the bottom of the primary eclipses, and then fitting a plot of relative PJMO intensity - relative MJUO intensity with a low order polynomial in order to remove the residual extinction effects.

The results of this comparison are presented in Figure 7.8.




Figure 7.8: Concurrent observations of PG 1336-018 provided a useful test for the two Puoko-nui instruments. The bottom panel compares the full lightcurves, with the PJMO intensity offset vertically to improve clarity. The top panel focuses on the primary eclipse, which shows the good agreement between the two data sets. The MJUO photometric uncertainties are given by the point size, and the PJMO uncertainties are a factor of two larger.

7.6 G117–B15A

Two runs were acquired on the northern DAV pulsator G117–B15A (Kepler et al., 2005) during the Puoko-nui North testing run on the McDonald Observatory 2.1 m telescope by Sam Harrold and Dean Chandler in April 2013. G117–B15A is an extremely stable pulsator (see Chapter 6), and discrepancies between the observed and calculated phases had been previously used to catch a timing error on the Argos photometer. These observations were, like the Crab pulsar, considered an important test for demonstrating the capabilities of Puoko-nui North to the UT group.

Figure 7.9 shows an excerpt of the time series data and the DFT of one of the two runs. The $O-C$ plot, using archival phase data provided by S. O. Kepler (private communication), shows that our 2013 phase measurements are entirely consistent with the timing prediction from the model of evolutionary cooling. This provides a simple check on the absolute timing accuracy of the instrument, and extends the catalogue of phase measurements on this target by an additional year.

It is interesting to note the general improvement in measurement precision that was obtained by moving from photomultiplier instruments (blue points) to frame transfer CCD measurements (Argos, red points) in 2002. Puoko-nui measurements (pink) should continue this trend.

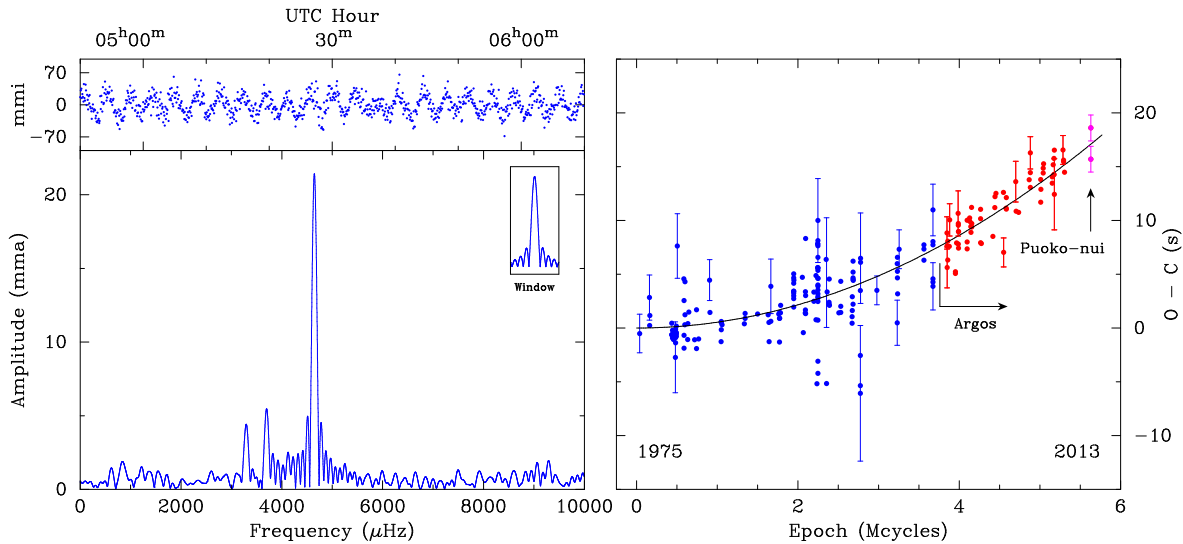


Figure 7.9: Observations of the very stable DAV pulsator G117–B15A were taken on two nights in April 2013 using Puoko-nui North on the 2.1 m McDonald telescope. The left plot depicts a short lightcurve segment and the DFT from one of the two runs, and the right panel compares the measured phase of the 215 s mode against 38 years of archival data using an $O-C$ diagram (see Chapter 6). Only a representative subset of the measurement uncertainties have been included in the plot to minimise clutter.

Chapter 8

Conclusions

This instrumentation- and observation-based thesis has described the development of two CCD photometer instruments that are designed to efficiently acquire high-speed photometry of variable stars, and presented observations of several white dwarf targets of contemporary interest.

The Puoko-nui instruments build on the principles previously established over several generations of instrumentation at VUW and the University of Texas at Austin, and mark the return of ‘true’ high-speed photometry, which was lost in the initial shift to CCD-based photometers.

The key outcomes of this instrumentation work are (a) two CCD photometers that offer a 200-fold improvement in the achievable time resolution compared with the previous generation of instruments; (b) a simple, accurate, and efficient online data visualization system that integrates well with these instruments; (c) a streamlined offline reduction pipeline that is competitive with other software packages; (d) a bench-testing environment that can simulate realistic observing conditions for use with instrument verification and for improving observing strategies.

Feedback from users of the Puoko-nui North instrument with telescopes at McDonald and Meyer Observatories has so far been overwhelmingly positive, and this instrument appears to have a strong future ahead of it. Our Puoko-nui South instrument is now stationed at MJUO, where it will remain available for other users of the Mt John facilities for the foreseeable future.

Observations of the DBV white dwarf EC 04207-4748 have significantly improved our understanding of this star’s pulsation characteristics. We identified 11 distinct frequencies in the pulsation spectrum (including seven that were previously unknown), and classified these as four eigenmodes plus combinations that arise due to convective mixing in the sub-surface convection zone. A non-linear fit to our data

was used to extract a characteristic convection timescale $\tau_0 = 148 \pm 5$, which is consistent with expectations. We argued that this τ_0 parameter provides a useful proxy for the effective temperature in the DBV WDs, and that this will become more useful once additional DBV τ_0 measurements have been made.

Observations of the pulsating WD in the cataclysmic variable GW Librae extended the post-outburst monitoring by an additional three years, and identified recurring but apparently unstable pulsations with periods of around $4 \sim 6$ and 19 minutes. We argue that these are most likely due to non-radial WD pulsations in spite of the lack of extended frequency coherency, but that we cannot be certain because we do not have a good theoretical understanding of these complicated binary star systems.

Observations of the DAV WD L19-2 were combined with archival data, and a preliminary $O - C$ diagram was constructed in order to identify the effects of evolutionary cooling on the pulsation stability. Linear least squares fits to this diagram found $\dot{P}_{f_2} = (3.8 \pm 2.3) \times 10^{-15}$ and $\dot{P}_{f_7} = (0.07 \pm 0.8) \times 10^{-15}$. The f_2 rate of change is remains consistent within uncertainties to the expected geometrical (proper motion) value, and the f_9 pulsation mode analysis is complicated by a potential signature of a non-parabolic phase change. Additional archival data that will be included in the final analysis will reduce the uncertainties in the $O - C$ fits and should resolve these issues with the fits.

Our observational contributions on the double-degenerate ELM system J1518 helped identify this system as one of the first known progenitors of the AM CVn star systems. Contributions towards the xCOV29 WET campaign on J0751 are expected to help extend our understanding of the pulsating variety of ELM WDs.

Additional observations of several ‘classic’ targets provided useful tests of the Puokoni instruments’ capabilities. These included the Crab Pulsar, the eclipsing binary sdB system PG 1336-018, and the stable hot DAV pulsator G117-B15A.

The topics presented in this thesis have covered a relatively wide area within the pulsating WD field. The topics of convection, \dot{P} , and the accreting pulsators contribute toward active areas of current research, and the new instrumentation looks to provide a solid foundation for continued observational work in the future.

Appendix A

Puoko-nui Instrument Manual

The following is meant to provide a basic overview for mounting, operating, and maintaining the Puoko-nui photometers.

A.1 Hardware Setup

Figure A.1 illustrates how the Puoko-nui South instrument should be mounted on the MJUO 1 m telescope. The offset-guider box is oriented with the guide camera on the southern face, and the top edge of the CCD should face the control electronics. The cables between the CCD, camera electronics, and instrument cart should be looped around the telescope strut and cable tied where they cross to provide strain relief. Figure A.2 shows a close up of the cable connections on the camera and timer units.

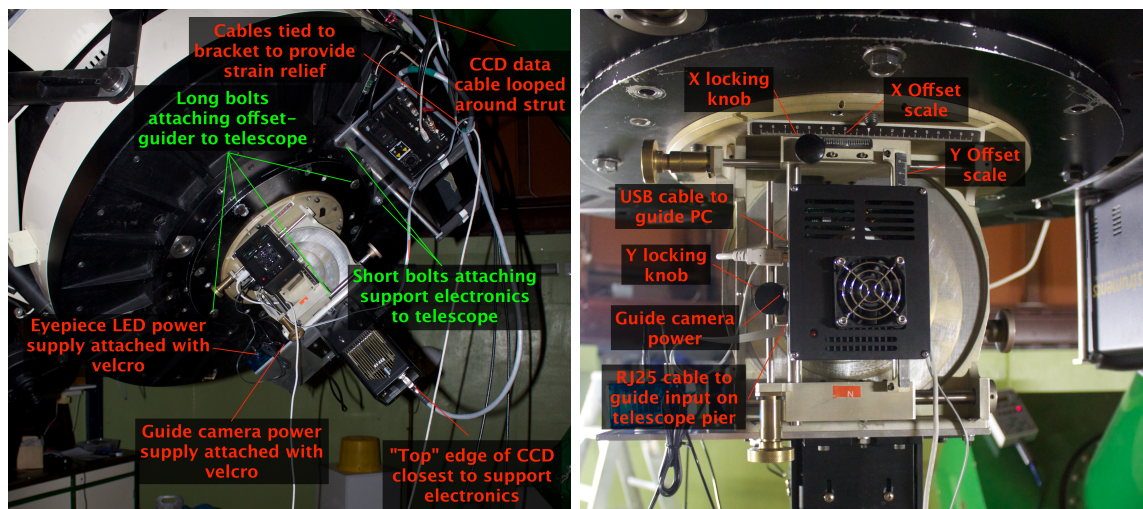


Figure A.1: Illustrated notes for mounting the instrument on the MJUO 1 m telescope.

The optimum guide camera focus can be found by screwing the focus ring completely




Figure A.2: Illustrated notes on the cable connections on the timer unit and CCD support electronics. Note that the support electronics unit is mounted upside-down on the telescope.

in, and then screwing it out by 6.5 turns. Ensure that both locking screws are tightened and that the (unused) dichroic filter is moved out of the light path.

A.2 Collecting Data

A.2.1 Preparing the Acquisition

A useful task each night should be to (re-)enable internet access for the acquisition and guide PCs. Run `ienabler.exe` on the guide PC using the standard MJUO user-name/password. Use `telnet` on the acquisition PC to connect to `ienabler.canterbury.ac.nz` port 259 and follow the prompts.

The next step is to synchronize the acquisition PC clock using NTP (Network Time Protocol). The default timeserver is not accessible from MJUO, but you can specify to connect to `ntp.vuw.ac.nz` or `1.nz.pool.ntp.org` directly:

```
$ sudo ntpdate ntp.vuw.ac.nz
```

The PC time isn't used in the reduction process, but it is stored in the acquired frames so it is a good idea to keep it roughly correct. The acquisition PC clock can drift by several seconds per day, so it pays to synchronize it manually at the start of each night.

A.2.2 Autoguiding

The guide camera is controlled using the `CCDOPS 5` software, which can be launched from the desktop or start menu. Once the software has started, select **Camera** → **Setup...** from the menu and enable temperature regulation. The optimum temperature set point depends on the ambient temperature: the cooler should operate at 50–70% power, which is usually around -10°C.

The auto-guider needs to be calibrated for each target field. The guide star is usually bright enough to use for this, but if not you may need to shift the telescope slightly to line up a brighter star that isn't too far away from the target field.

The continuous focus mode can be used for aligning the guide star and optimizing focus. This can be activated from the menu by selecting **Camera** → **Focus...** Adjust the telescope pointing so the star is in the centre of the field, and then select **Track** → **Calibrate...** to calibrate the calibration. The best x and y movement times will depend on the target field: targets the equator can use 5 s for both, while those near the pole will need a larger RA time: as much as 20 or 30 s.

When the run is ready to start you can begin auto-guiding by selecting **Track** → **Autoguide...** from the menu.

A suitable guide star for each target can be identified visually by replacing the guide camera with an eyepiece, or by calculating the offsets from a printed finding chart (obtained from e.g. <http://archive.eso.org/dss/dss>) and the approximate slide calibration of 4.43 arcmin / cm and offsets of (0.2, 0.1) cm. Figure A.3 demonstrates the calculation for a particular star (BPM37093). Note that the relationship between linear distances and right ascension changes with declination, and so it is simplest to work with pixel or printed distances instead of using the right ascension and declination directly.




Figure A.3: A demonstration of calculating the offset-guider coordinates for a specific target (BPM37093). Note that there is both a scaling factor and individual x,y offsets.

A.2.3 Starting the Acquisition

The acquisition software can be launched from the shortcut on the launcher (Ubuntu) or desktop (Windows). Right-clicking the Ubuntu launcher provides options for running with simulated hardware for testing purposes.

The main software interface is shown in Figure A.4. The left column of the acquisition software shows the camera and timer status, and the right column shows an information log. This log information is also saved as a text file in the **logs** subdirectory within the application directory.

Frame metadata can be changed in the **Set Metadata** panel (Figure A.5). Metadata can be changed while a run is in progress, provided that saving is disabled. The

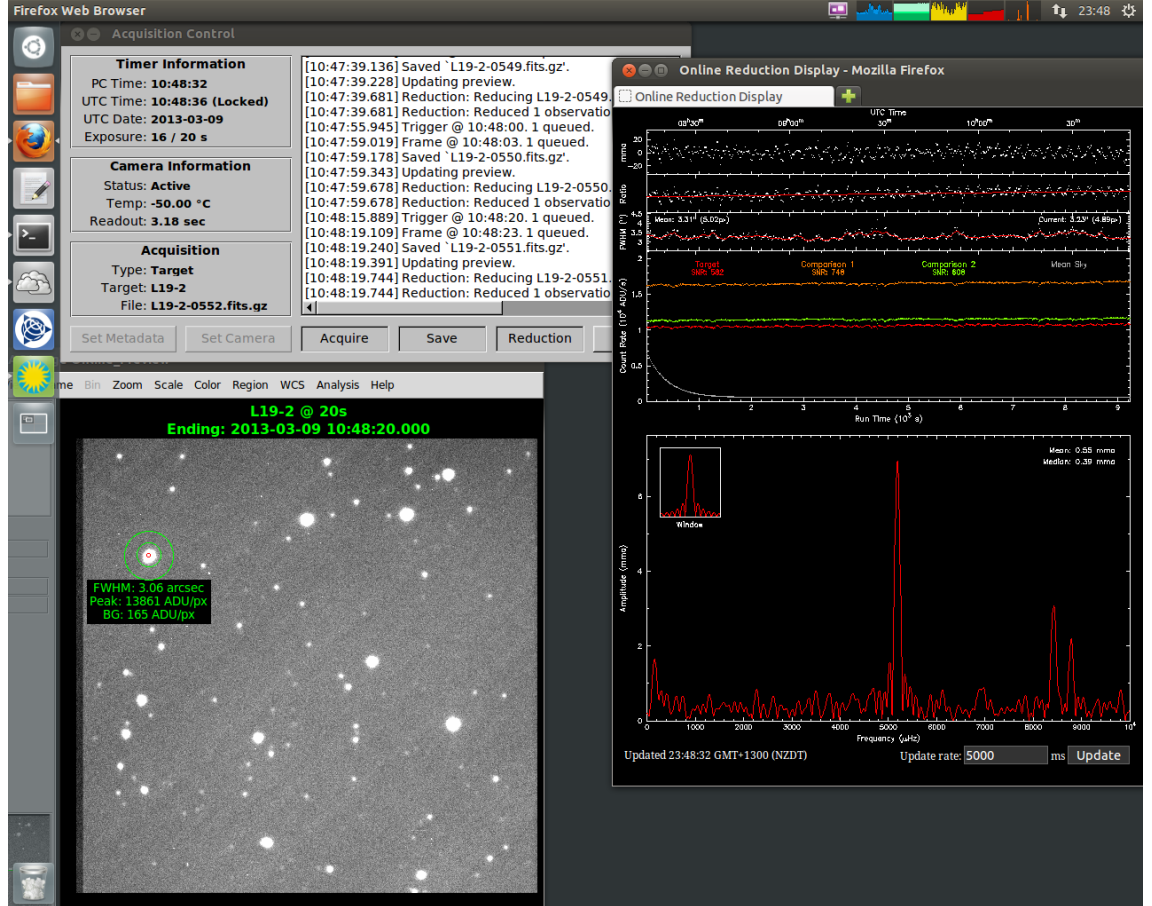


Figure A.4: The acquisition software interface as it appears during a typical acquisition run. The acquisition control window controls the instrument and provides status information. The frame preview window shows the latest acquired frame, optionally overlaid with star information. The online reduction window shows the evolving lightcurve and Fourier transform as data is acquired.

frame prefix is used by `tsreduce` to distinguish between dark/flat/object frames, so ensure that a unique prefix for each acquisition sequence. The acquisition type can be changed between **Continuous** for regular observations, or **Burst** for saving a specified number of frames (set in the **Burst Count** field). The frame type can be set to **Target** allowing a target name to be entered. The other options pre-set the target name to standard values.

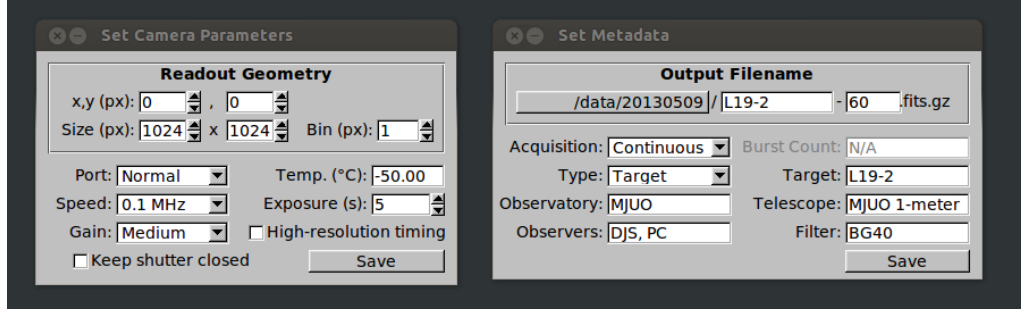


Figure A.5: The acquisition camera and metadata settings panels.

Camera properties can be changed in the **Set Camera** panel when the camera is idle. The readout geometry can be changed to read a restricted window of the CCD or to bin pixels during readout. The x and y settings specify the corner of the window closest to the readout port and its size. The width can be increased beyond 1024 to read the overscan columns if they have been enabled in `preferences.dat`.

The readout port, speed, and gain can also be configured to optimize the signal / noise of the observations. The ProEM camera includes a checkbox to keep the integrated shutter closed while acquiring dark frames.

The timing resolution can be configured between pulse counting and high resolution modes. The high resolution mode changes the exposure time resolution to 1ms, up to a maximum of 65.535s.

Acquisition can be started and stopped with the **Acquire** button. Acquired frames will only be saved if the **Save** button is toggled. The filename that the next acquired frame will use is displayed in the left sidebar.

The intensity and FWHM of stars in the target field can be monitored by dragging annuli in the `ds9` window. Target intensity and FWHM can be monitored in the frame preview by dragging an annulus around the star, which is used (see Section A.2.4 for tips about selecting apertures). These statistics are recalculated for each acquired frame, and do not require the frames to be saved to disk. This can be seen in action in Figure A.4.

The **Reduction** button enables the online reduction. The reduction will proceed if `<run prefix>.dat` (created using `tsreduce create` - see Section A.2.4) exists in

the frame directory, and the reduction preview (opened from the launcher [Linux] or desktop [Windows]) will automatically update to display the most current information.

A.2.4 Configuring the Online Reduction

The acquisition software interacts with external software (`tsreduce` and `ds9`) by executing a bash script. This makes it relatively straightforward to customize the interaction between tools: either by editing the script directly, or by defining variables that the script can query.

The acquisition software will execute `reduction.sh` each time a frame is saved to disk while the **Reduction** button is toggled. The default behaviour of the script is to search for `<run prefix>.dat` in the frame directory. If found, `tsreduce update` will be run on this file, and then `tsreduce plot` is called with custom parameters that make it save two images (`online_ts.gif` and `online_dft.gif`) to the acquisition software directory. These images are displayed in a web browser, which loads a web page (`preview.html`) that refreshes at a set interval. It is safe to run the acquisition when the reduction time is greater than the frame exposure time - the reduction preview will simply lag behind the data and update in large chunks.

There are two parameters that can be overridden by creating `config.sh` in the acquisition software directory. An example `config.sh` explaining these options is shown in Figure A.6.

```
# Override the filename controlling the reduction
# If unspecified, this will default to <run prefix>.dat
REDUCTION_FILENAME="object.dat"

# Size parameter for the online reduction plots
# If unspecified, this will default to "9"
REDUCTION_PLOTSIZE="6"
```

Figure A.6: An example `config.sh` script which overrides the reduction file-name and plot size.

The reduction file can be created as soon as the first target frame has been saved (for best results, ensure that dark and flat frames have also been acquired, but this is not required).

`tsreduce` is a command-line program, so open a bash prompt from the launcher (Linux) or desktop (Windows) and run:

```
$ cd <path to frame directory>
$ tsreduce create <run prefix>.dat
```

where `<run prefix>` is the frame prefix used for the target frames. If the run prefix is inconvenient then any other filename may be specified in `config.sh`.

`tsreduce` will prompt for the reduction parameters and, if required, create the master calibration frames. It will then open a `ds9` window showing the calibrated reference frame and prompt for the target apertures to be selected. Click and drag a generous annulus around the target star and one or more comparisons, making sure that the sky annulus is not contaminated by other stars in the field. `tsreduce` will centre the annulus about the mean moment of the contained flux and calculate the background intensity from the intensity within the annulus. The best results are obtained by picking a small number of comparison stars that are bright but remain within the linear CCD regime. Figure A.7 shows the apertures selection for a particular target field (EC 04207-4748).




Figure A.7: Demonstration of apertures selected during `tsreduce create`. The red circles define the integration aperture, and the dotted green circles define the sky annuli.

The aperture size is calculated from the selected reference frame using one of three techniques:

1. Radius where the radial profile of the brightest star falls to within 5σ of the background.
2. Diameter of 3 times the largest estimated full-width at half-maximum.
3. Manual (prompting the user for a value).

The `ds9` preview will be updated with the final aperture selection, and the user prompted for confirmation. Once the settings are confirmed, a reduction file is written to disk and the online reduction can be enabled by toggling the **Reduction** button in the acquisition software.

There are several hidden options that can be configured by editing the `.dat` file with a text editor. These options are listed in Table A.1.

Relay Mode

The timer unit features a relay mode that allows direct two-way communication between the acquisition PC and the GPS receiver. The timer can be prepared for relay mode by running the `timerutil` program:

```
$ timerutil <port> relay
```

where `<port>` is the OS handle of the timer (e.g. `/dev/ttyUSB0` under Linux, `/dev/tty.usbserial-00001004` under OSX, or `COM5` under Windows).

The next program that opens a serial connection to the timer will communicate directly with the GPS receiver. The normal timer operation will return for subsequent connections, and so this command must be run each time that an external program is to be used with the GPS receiver.

Note that the Trimble Studio software has been found to corrupt the timer firmware in some situations, and should not be used. The older Trimble Monitor software is recommended for checking or configuring the Trimble Thunderbolt status.

Keyword	Default	Description
# RatioFitDegree: <degree>	2	Polynomial degree to fit to the ratio data.
# PlotMaxRaw: <max>	5000	Maximum intensity to display in the raw data plot.
# PlotMaxDft: <max>	(automatic)	Maximum intensity to display in the Fourier spectrum plot.
# PlotMinUhz: <min>	0	Minimum frequency to display in the Fourier spectrum plot.
# PlotMaxUhz: <max>	10000	Maximum frequency to display in the Fourier spectrum plot.
# PlotNumUhz: <num>	10000	Number of points to display in the Fourier spectrum plot.
# BlockRange: (<min>, <max>)	–	Exclude a time range (in seconds) from ratio and Fourier calculations (heavy cloud, etc). Can be included multiple times with different ranges.
# PlotErrorBars: <num>	0	If non-zero, display error bars in online ratio and MMA plots.
# MMAFilterSigma: <num>	5	Exclude data points outside this range. Increase this if peaks are clipped.
# PlotFwhmSmooth: <num>	7	The number of points to average for the FWHM running estimate.
# RA: <HH:MM:SS>	}	Target coordinates used when converting UTC to BJD.
# DEC: <DD:MM:SS>		

Table A.1: Additional `tsreduce` keywords which can be specified in the reduction file header.

A.2.5 Additional Acquisition Software Settings

The acquisition software includes several advanced settings that may be configured by manually editing `preferences.dat`. Ensure that the acquisition software is closed before opening this file, or any manual changes will be overwritten.

The CCD overscan region can be enabled on the MicroMax camera by configuring `CameraOverscanColumns` with a non-zero value (24 is recommended). This increases the CCD width available in the geometry settings beyond 1024; the first several columns are masked physical pixels that measure the integrated dark current, and the remaining columns are overscan with a signal that measures the CCD bias. Two FITS header keys (`IMAG-RGN` and `BIAS-RGN`) are saved in each frame to indicate the image and bias regions.

Key	Default	Description
<code>Instrument</code>	Puoko-nui	Instrument name
<code>TimerMonitorLogicOut</code>	1	Enable monitoring of CCD logic out to when starting/stopping a run. Set to 0 to use fixed time delays.
<code>TimerSerialPort</code>	<code>/dev/ttyUSB0</code>	USB emulated serial port handle (e.g. <code>COM9</code> under Windows or <code>/dev/ttyusb0</code> under Linux)
<code>TimerBaudRate</code>	9600	Communication baud rate (requires timer firmware modification if changed).
<code>CameraOverscanColumns</code>	0	Number of overscan columns to digitize.
<code>CameraPlatescale</code>	0.33	Camera plate scale in arcsec/px.
<code>CameraFrameBufferSize</code>	5	Number of frames to allocate space for in the internal frame buffer used by the camera driver.
<code>ProEMExposureShortcut</code>	5	Number of milliseconds to shorten exposure times to allow the camera to complete frame transfer and prepare for subsequent triggers.
<code>ProEMEMGain</code>	1	Electron-multiplication gain.
<code>ProEMShiftMode</code>	1	Frame-transfer speed mode index.
<code>ValidateTimestamps</code>	1	Enable time stamp validation check.
<code>FrameFlipX</code>	0	Mirror image along x axis before saving.
<code>FrameFlipY</code>	0	Mirror image along y axis before saving.
<code>FrameTranspose</code>	0	Transpose image (flip x and y) before saving.
<code>PreviewRateLimit</code>	500	Minimum number of milliseconds between frame preview updates.

Table A.2: Advanced acquisition software options in `preferences.dat`.

A.3 Example `tsreduce` Commands

This section contains a list of examples of the commonly used `tsreduce` commands. Additional commands are available for offline analysis, but they are considered to be advanced functionality - refer to the `tsreduce` source code for details on these commands. Items in square brackets indicate optional parameters, and would not be included when entering these commands into a terminal. A single backslash at the end of a line indicates a line-break to fit the single-line command onto the page.

```
$ tsreduce create ec04207.dat
```

Create a reduction file `ec04207.dat` in the current directory. `tsreduce` will prompt for input regarding the reduction.

```
$ tsreduce update ec04207.dat
```

Update the reduction file `ec04207.dat` with any new frames.

```
$ tsreduce plot ec04207.dat [ts.ps/cps dft.ps/cps 10]
```

Generate an online reduction plot from the reduction file `ec04207.dat`. The optional parameters define the PGPLOT devices and size for the plots: saving to two postscript files with a PGPLOT scale of 10. The default plot devices are X-windows 5 and 6 under Linux/OS X and `ts.gif` and `dft.gif` under Windows.

```
$ tsreduce playback ec04207.dat 100 5 [ts.ps/cps dft.ps/cps 10]
```

Replay the acquisition of data at an accelerated rate. `tsreduce plot` is called repeatedly with a minimum wait of 100 ms between updates. Each update advances the data by 5 observations. The optional parameters match those for `tsreduce plot`

```
$ tsreduce tracer ec04207.dat
```

Displays the latest raw observation frame in ds9, overlaid with a line indicating the movement of the target and comparison stars. This function is useful to monitor shifts in the field of view over a run.

```
$ tsreduce display ec04207.dat ec04207-1234.fits.gz
```

The frame `ec04207-1234.fits.gz` is bias/dark subtracted, flat fielded, and displayed in ds9. The apertures used for the reduction are overlaid. The requested frame must have been previously reduced via `tsreduce update`.

```
$ tsreduce reduce-range template.dat 10 15 0.5 20110703
```

Reduce the reduction file `template.dat` with apertures between 10 px and 15 px in 0.5 px increments. Output files will be saved as `20110703-<aperture size>.dat`.

```
$ tsreduce plot-range 20110703-[0-9]+.[0-9]+.dat
```

Display timeseries and DFT plots of the data created with `tsreduce reduce-range`. The displayed aperture size can be changed using the `p` and `n` keys with the input window selected. This command only works under Linux and OSX.

```
$ tsreduce create-ts 2011-07-03 06:00:00 20110703.dat 20110704.dat \
  20110705.dat [...] july2011.ts
```

Combine a set of reduction files into a single timeseries containing relative BJD_{TDB} time, mmi intensity, and noise estimate. The first two arguments specify a reference date and time to use as time zero. The remaining arguments list (any number of) input filenames, followed by a single output filename. The first input file is required to define the `RA` and `Dec` keywords specifying the (J2000) target coordinates.

```
$ tsreduce dft july2011.ts 100 10000 1 july2011.dft [july2011.freq]
```

Generate a Discrete Fourier Transform of `july2011.ts`, with frequencies sampled between 100 and 10000 μHz in 1 μHz increments. The output data will be saved to `july2011.dft`. The optional parameter specifies a list of frequencies to pre-whiten before calculating the DFT (see `tsreduce` source code for details).

```
$ tsreduce window july2011.ts 1000 800 1200 0.01 july2011.win
```

Generate a window function for `july2011.ts` using 1000 μHz sinusoid between 800 and 1200 μHz . Data is saved to `july2011.win`.

A.4 Troubleshooting

The online frame preview exits and restarts if the frame is updated while selecting an aperture:

This appears to be an unavoidable bug in `ds9`. You can avoid this by ensuring that you aren't dragging an aperture when the frame preview updates.

`tsreduce create` cannot find frames with special prefix characters:

`tsreduce` converts the frame prefix into a (POSIX compatible) regular expression in order to find the files on disk. Characters in the set `[]\^$.|?*(\){}` have special meaning in regular expressions, and must be escaped with a `'\'`. `tsreduce create` will automatically escape these characters.

Cannot configure online reduction if the first frame in a run is cloudy:

The reference frame queried during `tsreduce create` is used to determine the aperture positions and time-zero. You can specify a more suitable reference frame, and then manually change the reference time in the `.dat` if you want to avoid negative time values.

The mouse jumps around the screen when running Relay Mode:

The timer unit acts as a USB-to-Serial adaptor for the attached GPS when running in relay mode. Windows may misinterpret the incoming serial data as input from a serial mouse and cause erratic cursor behaviour.

You can deselect the serial enumerator in the FTDI advanced driver properties to disable this behaviour.

A.5 Software Updates

The three main software components of the Puoko-nui instruments are managed using the `git` version control system, and are currently hosted on Github.

```
Acquisition software: https://github.com/pchote/Puoko-nui.git
tsreduce: https://github.com/pchote/tsreduce.git
Timer firmware: https://github.com/pchote/Karaka.git
```

`git` is an extremely powerful tool for managing source code across multiple locations (originally created by Linux Torvalds to simplify development of the Linux kernel), but we will focus on just the necessary commands to fetch and reset the local code to match the upstream version.

Github provides a good general introduction and tutorial at <http://learn.github.com/p/intro.html>, plus additional more specific topics at <https://help.github.com/>.

The code can be acquired for the first time using the `git clone` command, e.g.:

```
$ git clone https://github.com/pchote/Puoko-nui.git puokonui
```

This will create a subdirectory called `puokonui` that contains a copy of the acquisition software source code.

If you `cd` into an existing repository you can download any changes from the remote repository using `git fetch`:

```
$ git fetch -p
```


The `-p` flag tells fetch to *prune* any branches that have been deleted from the upstream server. These changes are downloaded, but not applied to your copy. The simplest way to do this, which will discard any local changes, is to *reset* the local state to match the remote state:

```
$ git reset --hard origin/master
```

where `origin/master` specifies the upstream branch to reset to. You can view a list of remote branches using the `git branch -r` command:

```
$ git branch -r
```

Once you have reset the local code to the latest version, you can **make** the updated executable:

```
$ make clean
$ make <project options>
```

The Makefile for each project starts with a list of options that can be configured. See each file for the options and their descriptions.

The overall process for updating the project will look something like:

```
$ cd <project directory>
$ git fetch
$ git reset --hard origin/master
$ make clean
$ make <project options>
```

For completeness, a list of the main **make** commands are included below:

Linux

```
Acquisition (Puoko-nui): $ make CAMERA_TYPE=PVCAM
Acquisition (ProEM):    $ make CAMERA_TYPE=PICAM
tsreduce:               $ make
```

Windows

```
Acquisition: $ make CAMERA_TYPE=PICAM
tsreduce:    $ make USE_READLINE=no
```

Timer Firmware:

The timer unit contains a boot-loader which will be run at startup and enter a firmware update mode if any data is received within four seconds of connecting via USB.

The upgrade process is controlled by the **Makefile** in the source code directory. Several configuration options are defined at the top of the file. Ensure that these are set correctly before running **make**: incorrect values put the unit into a nonfunctional state the requires physical access to fix.

The standard firmware can be updated via USB by the following:

```
$ make clean
$ make install PORT=<port>
```

where **<port>** is the OS handle of the timer (e.g. `/dev/ttyUSB0` under Linux or `/dev/tty.usbserial-00001004` under OSX).

The **PORT** parameter under windows needs to be escaped with several backslashes, which are mangled by the commandline. Instead, you can modify the **Makefile** to set the **PORT** value directly, e.g:

```
PORT := \\.\COM5
```

Timer Boot-loader:

The boot-loader can only be updated using the JTAG programmer. Connect the programmer to the port labelled in Figure A.8, and then upgrade the boot-loader and firmware together:

```
$ make clean
$ make jtag
```




Figure A.8: Illustrated notes on the timer unit circuit board.

Appendix B

Timer Firmware Description

The aim of this section is to give a detailed explanation of the logic implemented within the timer unit firmware. The diagrams that follow closely mirror the source code, but skip some of the unnecessary details involved in each block step.

Figure B.1 traces the main logic starting from the unit first receiving power, and covers initialization and the main logic loop.

The timer uses the DTR serial control line to trigger a hardware reset remotely from the Acquisition PC. This simplifies many details, but has one notable downside - the unit is forced to reset whenever software on the Acquisition PC first opens a connection via USB (this is strictly true only for Linux and OS X, but we force additional resets to ensure consistent behaviour on all platforms).

The unit will enter its firmware upgrade mode if any data is received via the USB connection within four seconds of powering up. This mode uses the AVR-109 protocol for flashing the firmware and performing other diagnostics. The update mode will exit if the command character ‘E’ is received, or if the unit is reset by toggling DTR.

A work-around is implemented to disable this mandatory four-second delay when the unit is in its relay mode. This is achieved by using the `timerutil` program (see Appendix A) to send a command packet which enables a one-time bypass of the boot-loader. The unit will restart into a one-time relay mode session the next time that the DTR line is toggled.

The main program begins after the boot-loader timeout (to avoid ambiguity, I’ll use timeout to refer to actions run via the microcontroller timer/counter subsystem), which initializes the hardware components and sends configuration messages for the Trimble and Magellan GPS receivers through the serial port.

Once initialization is complete, the main program loop will loop the following:

- Any message flags set by interrupts or the USB/GPS subroutines are read, and the appropriate data sent to the USB subsystem to be forwarded to the Acquisition PC.
- The camera logic input is polled. If a level change is detected a $500\ \mu\text{s}$ timeout is set to confirm that it wasn't triggered by a spurious noise pulse. Any actions that run as a consequence of this level change occur within the logic monitor timeout interrupt.
- USB commands are processed (see Figure B.2).
- GPS messages are processed (see Figure B.3).

Hardware or software interrupts may trigger any time during the main loop, and will take over the program flow until they complete. Interrupt sources include:

- GPS and USB serial reads/writes. These systems operate asynchronously, using interrupts to copy bytes into a buffer as soon as they are received, and a second buffer which data is copied to for transmission in parallel with the main loop.
- The timeout used to restore the camera trigger level after $512\ \mu\text{s}$.
- The timeout used to set the GPS status to `UNAVAILABLE` after 3s without data.
- The timeout used to verify and act on camera logic output changes (see Figure B.5).
- The timeout used to increment the high-speed exposure counter (see Figure B.6, and explained below).
- A hardware interrupt on the falling edge of each 1 Hz input pulse (see Figure B.4, and explained below).

The low resolution pulse-counting mode provides a 1 s resolution between 1 - 65535 s (18.2 hours). This mode decrements a counter in the 1 Hz hardware interrupt, and sends a readout trigger to the camera once the counter reaches zero. It also sets a message flag (`SEND_TRIGGER`), which will cause the main loop to send the next time stamp as the trigger time to the Acquisition PC.

It is assumed that the first time stamp received after a trigger pulse is sent corresponds to the trigger. An incorrect trigger time stamp may be sent in the (rare) event that the serial data for the time packet of the trigger is corrupted. The output trigger pulses are delayed by $11\ \mu\text{s}$ from the input pulse: this offset is consistent be-

tween exposures, and so it is assumed to be negligible compared to the 1 s exposure resolution.

The high-speed timing mode provides a 1 ms resolution for exposures between 1 - 65535 ms, but requires the timer to be operating in the external clock mode, using the GPS disciplined 10 MHz signal to drive the microcontroller clock (this is supported by the Trimble Thunderbolt and several other makes of GPS receivers, but not the Magellan receivers at VUW or MJUO).

A switch on the back of the unit selects between the internal oscillator and the external clock signal, which is then divided using the timer/counter subsystem within the microcontroller to produce an interrupt accurately aligned to each millisecond.

This millisecond-interrupt decrements the exposure counter and updates an internal estimate of the current millisecond time component. The millisecond counter is allowed to increment past 1000 ms, which adds a significant degree of robustness against glitches in the GPS serial stream: the timer will internally correct up to 65 continuous missing serial time stamps before timing accuracy is lost. The counter is reset modulo 1000 each time a serial time packet is received.

The 1 Hz signal interrupt is used to verify that the millisecond counter is an integer multiple of 1000 ms. A message flag (`TIME_DRIFT`) is set if this check fails, which causes a warning to be sent to the Acquisition PC.

The timer is started with an initial count value chosen to reduce the time offset between the 1 Hz input pulses and the output trigger to within 500 ns. A future improvement may make this delay customizable in order to correct an exposure phase shift introduced by the triggered operation of the ProEM camera.

Three input command packets are recognized by the timer. The `ENABLE_RELAY` command is sent by the `timerutil` utility to place the unit into relay mode. In relay mode, all data sent via USB is forwarded to the GPS receiver, and vice-versa, turning the timer unit into a USB-to-serial adaptor for the GPS receiver. The 1 Hz input pulses are also forwarded to the camera, but note that the polarity and pulse width are different to the outputs of the Trimble and Magellan receivers.

The `START_EXPOSURE` command is used to start an exposure sequence. The command packet includes the timing mode to use, whether to monitor the camera logic output or simulate the logic changes internally, and the exposure time. The unit will wait until the camera is ready to start exposing (signaled by the logic output going low), and will then wait a further period of time to ensure that exposure transitions will occur on the round minute (assuming an exposure period that evenly divides into 60 seconds). The first exposure will begin immediately if pulse-counting mode is active, otherwise an additional exposure period is required. This first exposure

(which is discarded by the acquisition software) may therefore be up to twice as long as the requested exposure time, and may cause bright objects in the field to oversaturate.

The `STOP_EXPOSURE` command immediately stops the exposure countdown, and waits for the camera output level to go high, signalling that a frame is *not* being read out. The camera monitor interrupt sets a message flag (`STOP_EXPOSURE`), signalling the Acquisition software that it is safe to shut down the camera. This two-step shutdown procedure is necessary to work around a serious memory-corruption bug in the PVCAM camera software library. This bug is triggered if an exposure sequence is terminated while the camera is reading out an exposure.

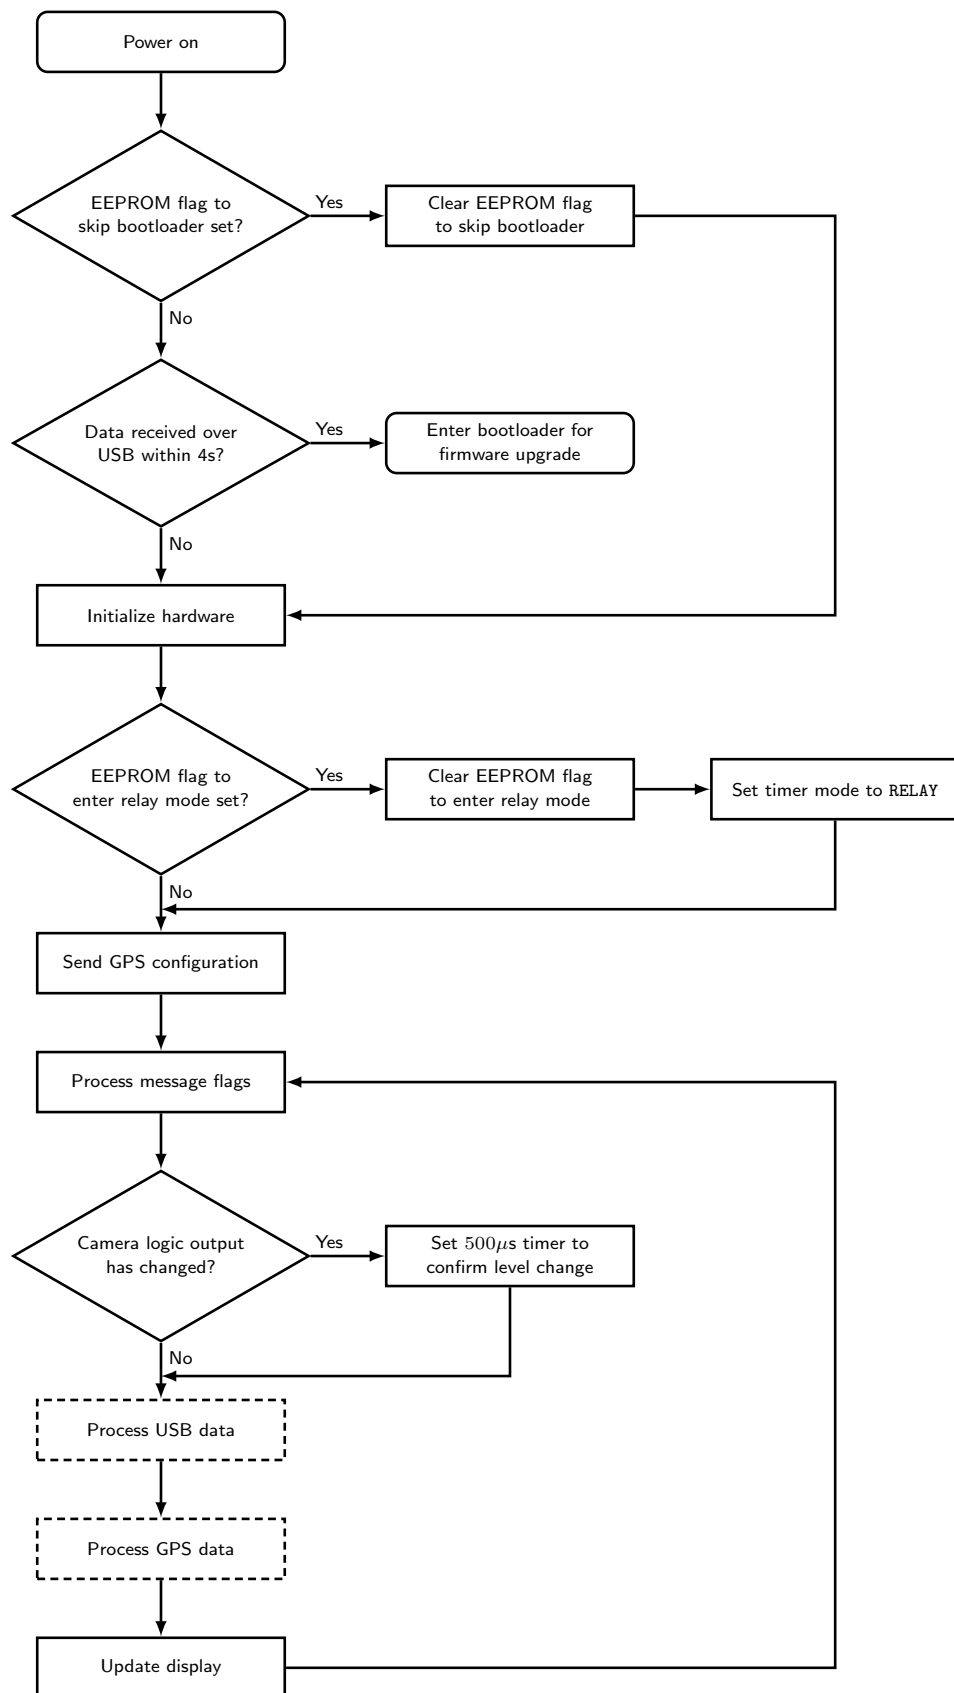


Figure B.1: The main loop handles initializes the hardware and handles non-time-critical functions. Each of the dotted boxes are expanded in additional figures.

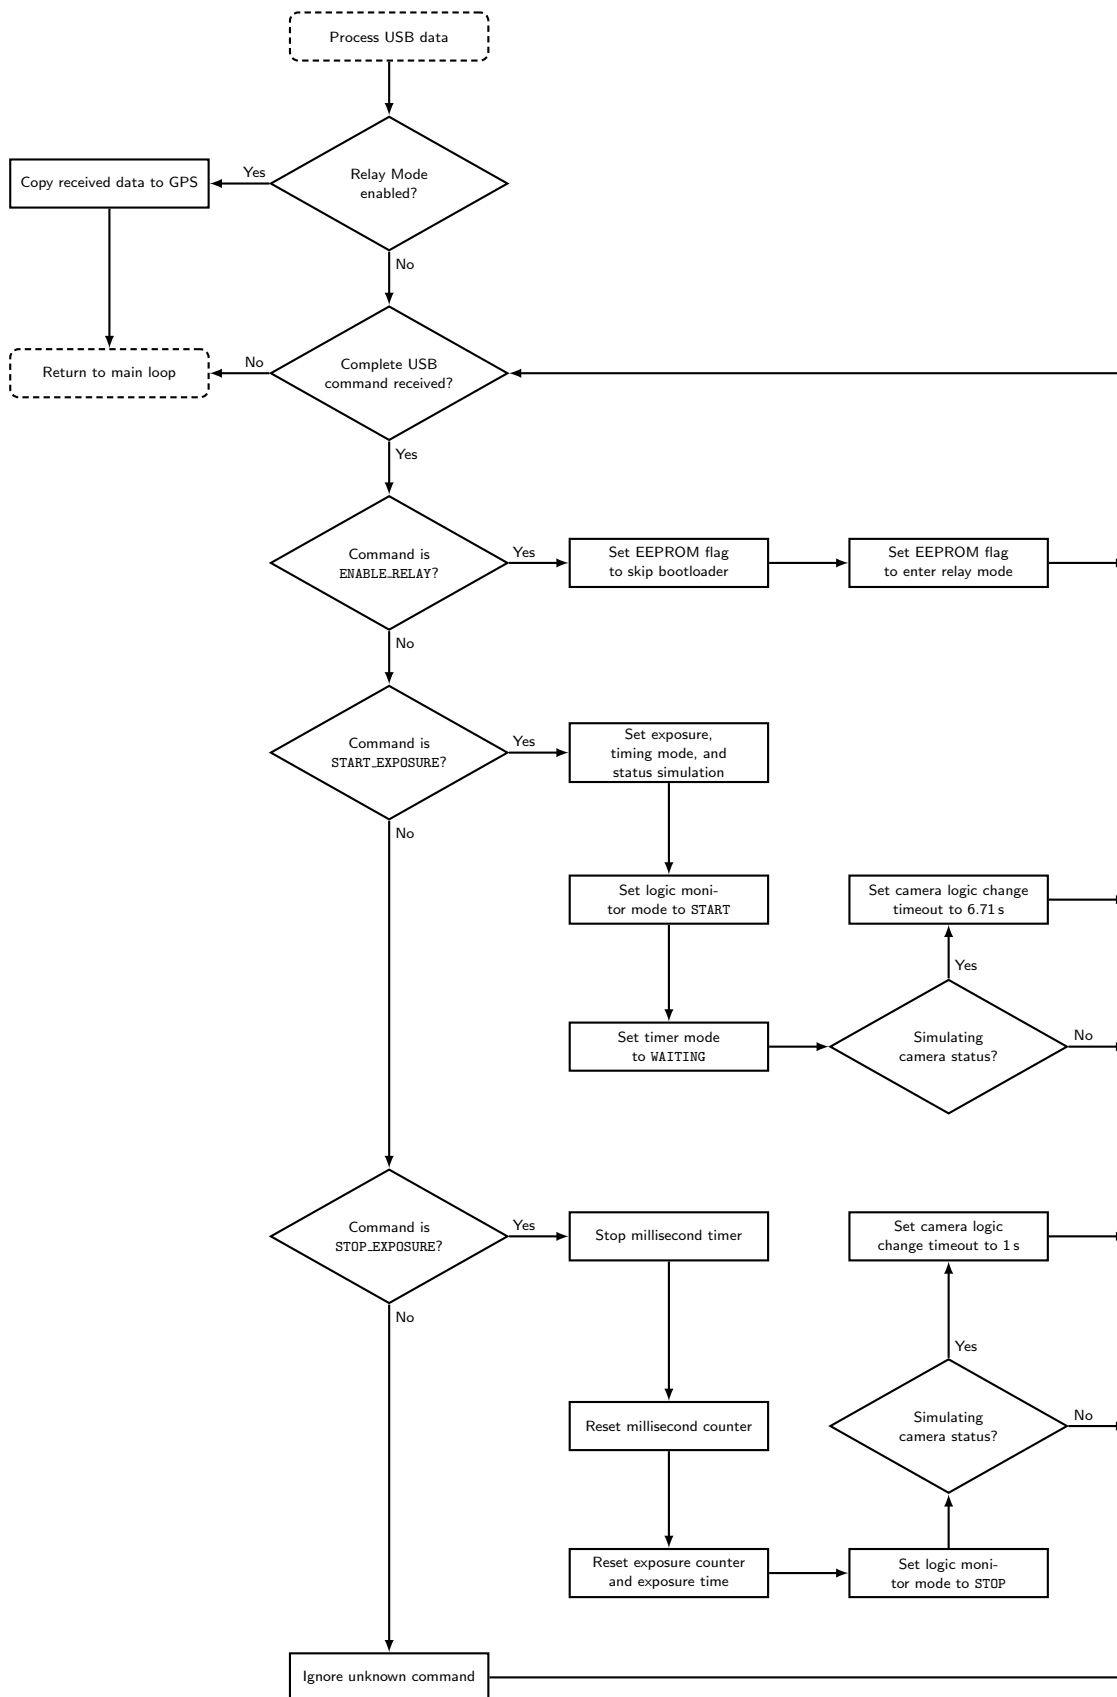


Figure B.2: Expanded detail of the USB command processing logic.

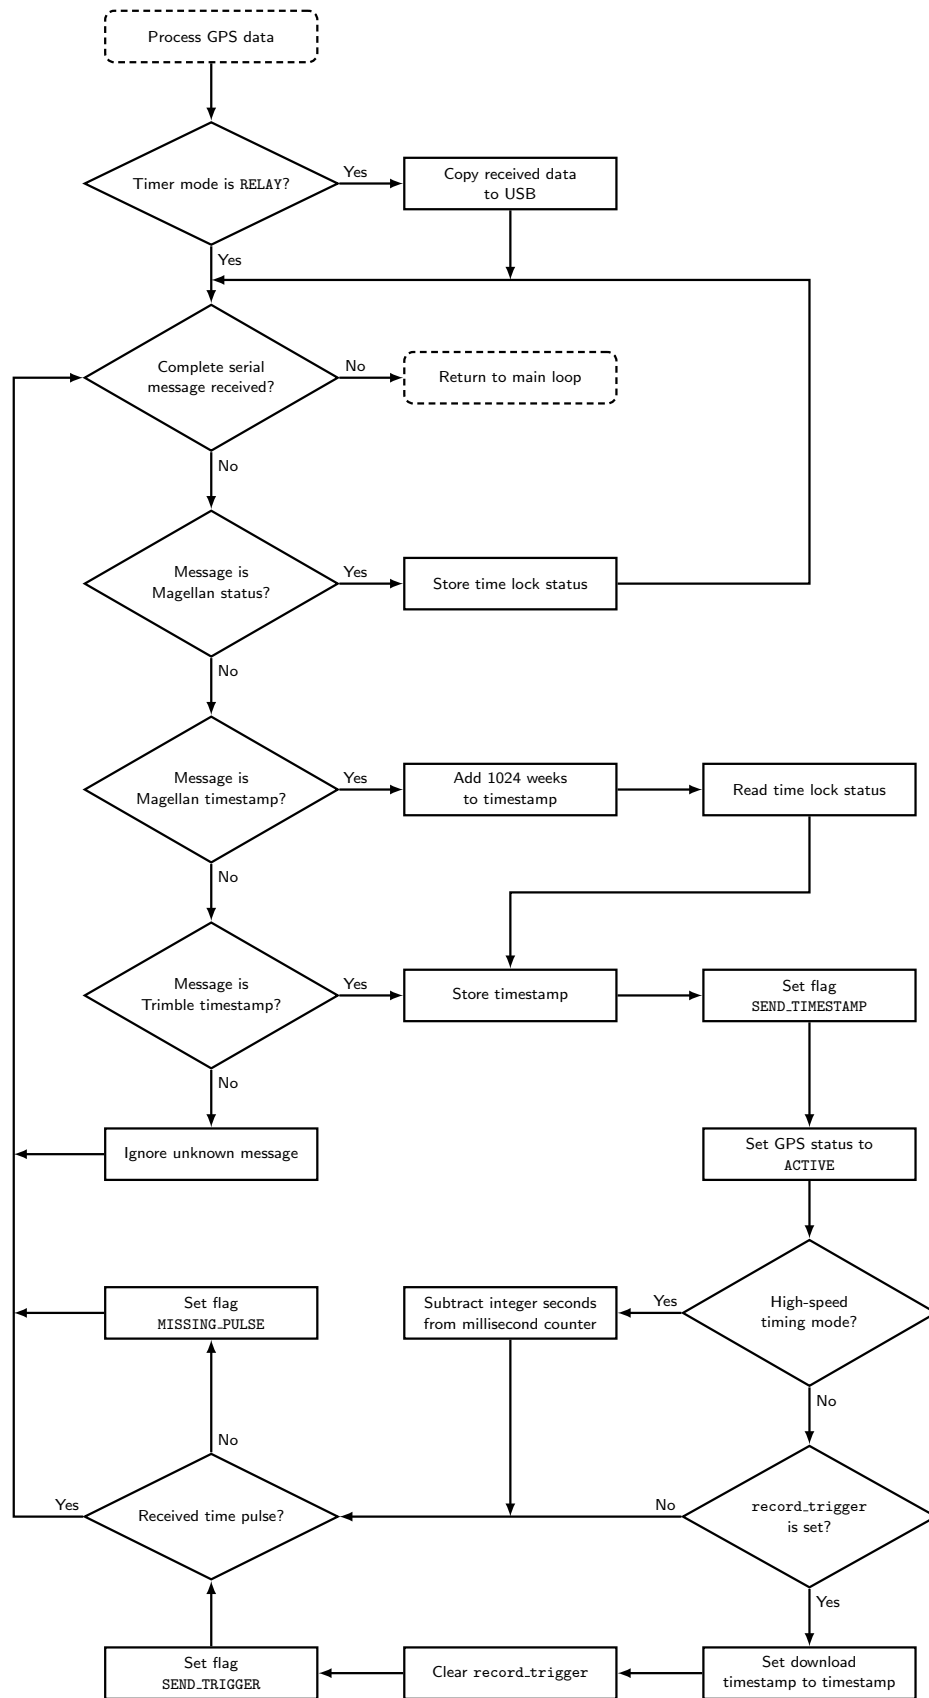


Figure B.3: Serial messages from the GPS are processed to extract the current time and satellite lock status. A check is in place to ensure that serial time stamps and time pulses are received in alternating order.

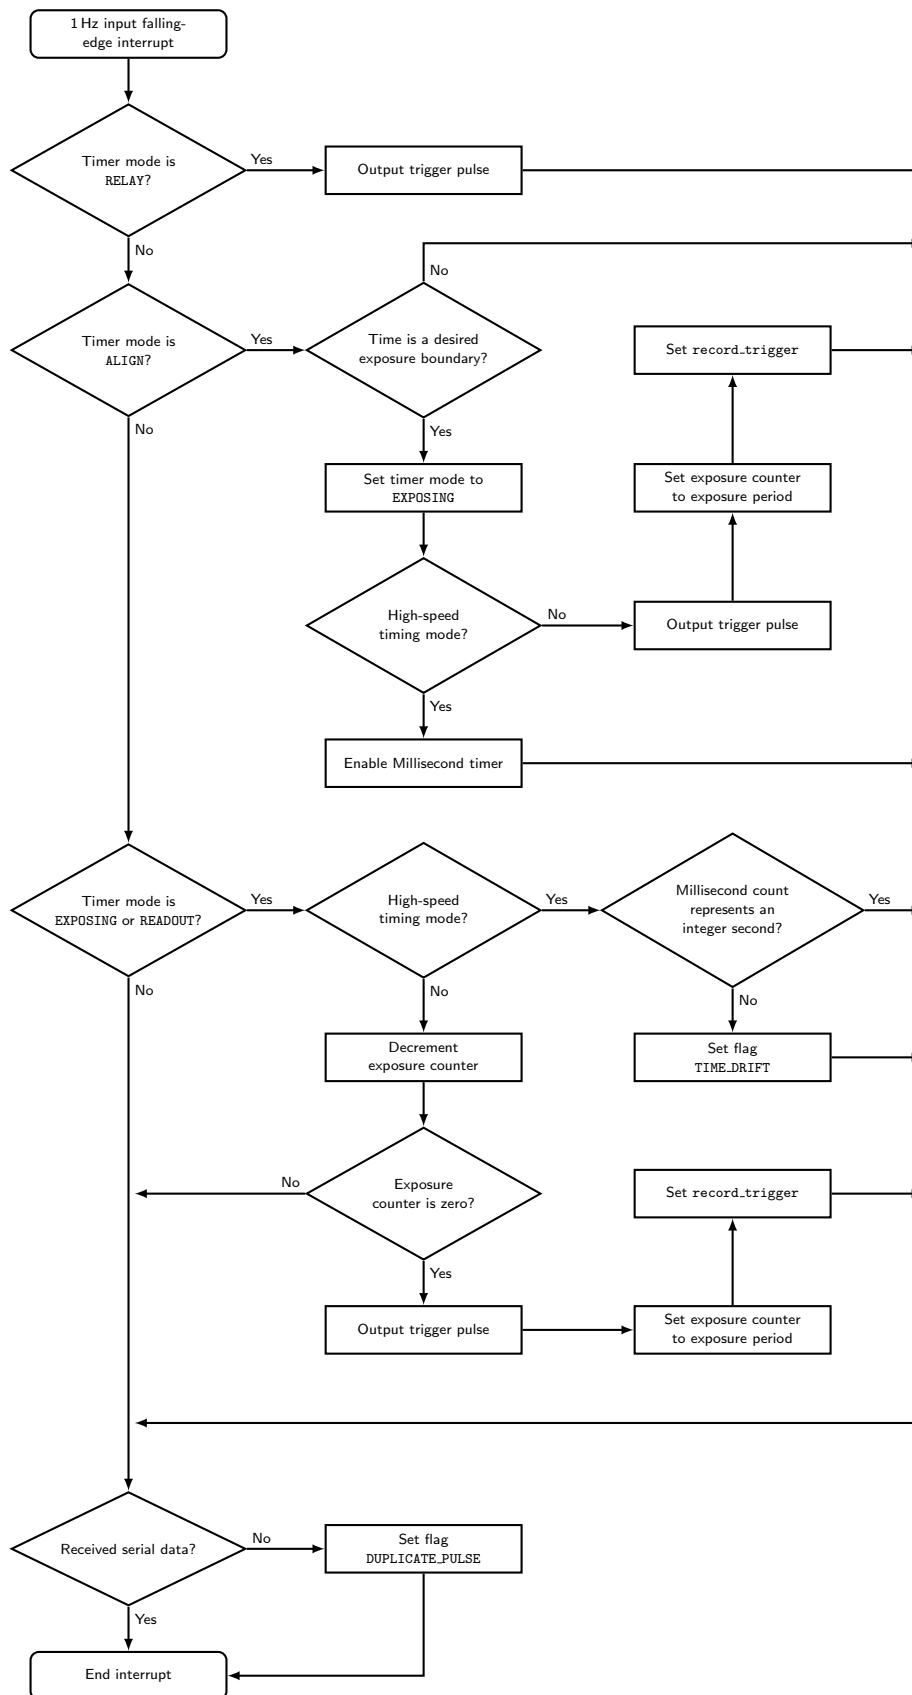


Figure B.4: A pulse on the 1 Hz input line triggers different actions depending on the current timer mode. See text for details.



Figure B.5: A block diagram showing the camera monitoring logic.

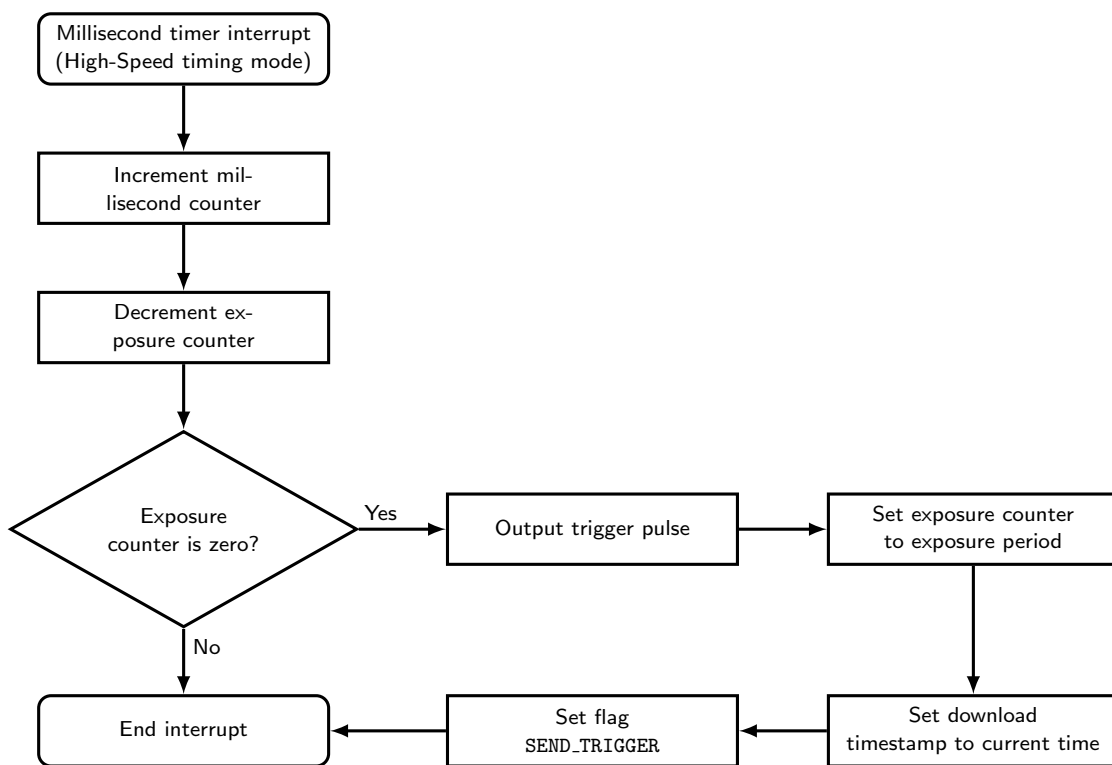



Figure B.6: The millisecond timer is used in the High-speed timing mode as a high-cadence equivalent to the 1 Hz input.

Appendix C


Finding Charts

This section chronicles the finding charts and offset-guider offsets for our targets of interest. Each chart shows a $20' \times 20'$ area around the target (indicated with an arrow). The approximate CCD alignment for the MJUO 1.0 m telescope at f/7.7 is indicated with a box.


EC 04207-4748



Coordinates (J2000)		Guide Offset	Information
RA:	04 ^h 22 ^m 11 ^s	X: -1.55	V-mag: 15.3
Dec:	-47°41'42"	Y: 2.72	Type: DBV


EC 05221-4725

Coordinates (J2000)	Guide Offset	Information
RA: 05 ^h 23 ^m 31 ^s	X: -1.70	V-mag: 16.6
Dec: -47°22'19"	Y: 4.05	Type: DBV

EC 20058-5234


Coordinates (J2000)	Guide Offset	Information
RA: 20 ^h 09 ^m 40 ^s	X: -1.90	V-mag: 14.9
Dec: -52°25'24"	Y: -2.75	Type: DBV

GW Librae




Coordinates (J2000)	Guide	Offset	Information
RA: 15 ^h 19 ^m 55 ^s	X:	2.50	V-mag: 18.5
Dec: −25°00′25″	Y:	2.80	Type: CV


L19-2



Coordinates (J2000)	Guide	Offset	Information
RA: 14 ^h 32 ^m 18 ^s	X:	−3.15	V-mag: 13.8
Dec: −81°20′09″	Y:	1.30	Type: DAV

J0751

Coordinates (J2000)	Guide Offset	Information
RA: 07 ^h 51 ^m 41 ^s	X: -1.80	V-mag: 17.5
Dec: -01°41'21"	Y: -1.40	Type: ELM

WDJ1518

Coordinates (J2000)	Guide Offset	Information
RA: 15 ^h 18 ^m 27 ^s	X: -2.25	V-mag: 17.6
Dec: 06°58'13"	Y: -0.10	Type: ELM

Appendix D

Observations & Publications

Observing was an important part of this thesis work. The following tables summarize the observations obtained during 13 sessions at Mt John University Observatory. Approximately 484 hours of new data were acquired over 92 nights using the Puokonui South photometer on the 1m McLellan telescope. Tables D.2 – D.14 present a log of these observations.

Much of this work has been undertaken as part of larger programmes, which has resulted in co-authorship of several publications (Szkody et al., 2012a; Kilic et al., 2014; Hermes et al., 2014). Our contributions towards the xCOV29 WET campaign on WDJ1518 should result in an additional publication.

I took the lead in our work on EC 04207-4748 (Chote et al., 2013), and we have additional publications in preparation presenting our results for L19-2 and GW Librae. Chote et al. (2014) provides a detailed description of the two Puokonui photometers.

Target	Ref	Citation
EC 04207-4748	A	Chapter 4
	B	Chote et al. (2013)
	C	EuroWD12 presentation
GW Librae	D	Chapter 5
	E	Szkody et al. (2012a)
	F	Chote and Sullivan (2013a)
L19-2	G	Chapter 6
	H	Kilic et al. (2014)
	I	Hermes et al. (2014)
	J	Chapter 7
	K	Chote et al. (2014)

Table D.1: Reference key for the observation log presented in Tables D.2 – D.14.

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2011-03-02	EC04207	09:00	12:00	20	3	Clear	A B C
	GW Librae	13:09	17:23	20	4	Cloudy	D E
2011-03-03	EC04207	08:46	11:24	20	1.5	Cloudy	A B C
2011-03-04	EC04207	08:34	12:47	20	2	Cloudy	A B C
	GW Librae	14:33	16:19	20	1.5	Cloudy	D E
	L19-2	13:01	14:12	20	1	Cloudy	G
2011-03-05	-	-	-	-	-	Overcast	-
2011-03-06	EC04207	08:48	11:00	20	1.5	Cloudy	A B C
	L19-2	14:40	15:00	20	-	Cloudy / Overcast	-
2011-03-07	EC04207	08:49	12:13	20	3	Cloudy	A B C
	L19-2	12:32	14:20	20	-	Cloudy / Overcast	-
2011-03-08	EC04207	08:09	12:03	20	4	Clear	A B C
	GW Librae	12:39	13:08	20	-	Cloudy / Overcast	D E

Table D.2: Observation log for March 2011

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2011-07-01	GW Librae	07:02	12:06	20	5	Clear	D E
	EC20058	12:38	16:29	20	4	Clear	J
	EC04207	17:53	19:27	20	1.5	Clear	A B C
2011-07-02	GW Librae	06:23	11:00	20	4.5	Clear	D E
	EC20058	11:20	15:41	20	4.5	Clear	J
	EC04207	16:16	19:31	20	2.5	Clear	A B C
2011-07-03	L19-2	06:39	11:46	20	5	Clear	G
	EC20058	12:06	15:28	20	3.5	Clear	J
	EC04207	15:50	19:31	20	3	Clear	A B C
2011-07-04	GW Librae	06:36	14:02	20	7.5	Clear	D E
	EC20058	14:31	15:27	20	1	Clear	J
	EC04207	15:40	19:22	20	3.5	Cloudy	A B C
2011-07-05	L19-2	14:00	15:26	20	1.5	Cloudy	G
	EC04207	15:57	17:47	20	0.75	Cloudy / Overcast	A B C
2011-07-06	GW Librae	06:35	09:28	20	3	Cloudy	D E
	L19-2	09:58	15:32	20	5.5	Clear	G
	EC04207	16:00	17:56	20	1.5	Cloudy / Overcast	A B C
2011-07-07	L19-2	14:53	15:29	20	0.75	Clear	G
	EC04207	15:41	19:21	20	3.5	Clear	A B C

Table D.3: Observation log for July 2011 (run 1)

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2011-07-27	GW Librae	06:48	12:47	20	5.5	Clear	D E
	EC04207	13:37	19:01	20	5.5	Clear	A B C
2011-07-28	EC20058	10:28	12:32	20	2	Cloudy	J
	EC04207	13:04	14:41	20	1	Cloudy	A B C
	EC20058	16:30	17:34	20	0.75	Cloudy / Overcast	J
	EC04207	17:59	18:57	20	1	Cloudy	A B C
2011-07-29	EC20058	08:53	10:32	20	1	Cloudy	J
	EC04207	13:58	14:50	20	0.8	Cloudy	A B C
2011-07-30	EC04207	13:16	19:12	20	5	Clear	A B C
2011-07-31	GW Librae	06:37	06:39	20	-	Cloudy/Overcast	E
	EC04207	14:27	18:52	20	0.5	Cloudy	A B C
2011-08-01	GW Librae	06:39	12:53	20	5	Cloudy	D E
	EC04207	13:19	18:44	20	4	Cloudy	A B C
2011-08-02	GW Librae	06:35	10:04	20	3	Clear	D E
	EC20058	10:18	14:02	20	3.5	Clear	J
	EC04207	14:18	19:01	20	4.5	Cloudy	A B C

Table D.4: Observation log for July 2011 (run 2)

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2011-11-18	EC04207	10:02	16:08	20	3	Cloudy	A B C
2011-11-19	-	-	-	-	-	Overcast	-
2011-11-20	-	-	-	-	-	Overcast	-
2011-11-21	EC04207	09:10	15:41	20	6	Cloudy	A B C
2011-11-22	-	-	-	-	-	Overcast	-
2011-11-23	EC04207	09:13	13:57	20	3.5	Cloudy / Overcast	A B C
2011-11-24	EC04207	09:28	16:04	20	6	Cloudy	A B C

Table D.5: Observation log for November 2011 (WET Run XCOV28)

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2012-03-22	-	-	-	-	-	Overcast	-
2012-03-23	EC05221	08:09	11:22	30	3.25	Clear / Overcast	J
	L19-2	13:13	13:54	30	-	Cloudy / Overcast	-
2012-03-24	EC05221	07:44	11:33	30	3.5	Cloudy	J
	GW Librae	13:02	17:49	30	3.5	Cloudy	D F
2012-03-25	EC05221	07:44	12:26	30	4.75	Clear	J
	GW Librae	12:40	17:56	30	5	Clear	D F
2012-03-26	EC05221	07:28	09:33	30	1.5	Cloudy	J
2012-03-27	EC05221	07:30	09:30	30	1.5	Clear / Overcast	J
2012-03-28	EC05221	07:50	12:39	30	4	Cloudy	J
	L19-2	12:58	17:50	30	4.75	Clear	G

Table D.6: Observation log for March 2012

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2012-04-23	EC05221	06:49	11:05	30	3.5	Cloudy	J
	GW Librae	11:22	18:24	30	7	Cloudy	D F
2012-04-24	EC05221	09:00	10:15	30	-	Cloudy	-
	L19-2	10:43	18:31	30	5.5	Cloudy	G
2012-04-25	EC05221	06:41	11:23	30	4.8	Clear	J
	L19-2	11:34	18:37	30	5.5	Cloudy	G
2012-04-26	L19-2	06:41	07:55	30	-	Cloudy	-
	EC05221	08:11	08:33	30	-	Cloudy / Overcast	J
	L19-2	08:47	09:20	30	-	Cloudy	-
2012-04-27	-	-	-	-	-	Overcast	-
2012-04-28	EC05221	06:37	11:18	30	4.5	Cloudy	J
	L19-2	11:34	18:25	30	6.75	Clear	G
2012-04-29	L19-2	13:51	16:43	30	3	Cloudy / Overcast	G

Table D.7: Observation log for April 2012

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2012-05-16	GW Librae	10:30	10:48	20	-	Cloudy / Overcast	-
	GW Librae	15:13	16:55	30	-	Cloudy / Overcast	-
2012-05-17	GW Librae	06:53	18:06	30	11	Clear	D F
2012-05-18	GW Librae	06:44	17:53	30	11	Clear	D F
2012-05-19	GW Librae	06:33	17:55	30	11.5	Clear	D F
2012-05-20	GW Librae	07:40	18:06	30	6	Cloudy / Overcast	D F
2012-05-21	GW Librae	07:32	17:46	30	10	Clear	D F
2012-05-22	GW Librae	06:29	18:04	30	12	Clear	D F

Table D.8: Observation log for May 2012

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2012-12-11	-	-	-	-	-	Overcast	-
2012-12-12	J0751	11:16	15:25	30	4.2	Clear	J H I
2012-12-13	-	-	-	-	-	Overcast	-
2012-12-14	J0751	12:22	15:25	45	2	Cloudy	J H I
2012-12-15	J0751	11:35	15:28	45	3.5	Cloudy	J H I
2012-12-16	J0751	11:22	15:04	45	3.75	Cloudy	J H I
2012-12-17	-	-	-	-	-	Overcast	-

Table D.9: Observation log for December 2012

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2013-01-11	J0751	10:18	10:47	45	-	Cloudy / Overcast	-
2013-01-12	J0751	09:51	11:15	45	0.75	Cloudy / Overcast	J H I
2013-01-13	J0751	14:46	15:52	45	1	Cloudy	J H I
2013-01-14	-	-	-	-	-	Overcast	-
2013-01-15	J0751	09:49	14:20	45	3.5	Cloudy	J H I
2013-01-16	-	-	-	-	-	Overcast	-
2013-01-17	J0751	10:48	15:50	45	4.5	Cloudy	J H I

Table D.10: Observation log for January 2013

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2013-03-08	L19-2	08:46	12:35	20	3.5	Cloudy / Overcast	G
	GW Librae	13:22	15:24	30	1.5	Cloudy / Overcast	-
2013-03-09	L19-2	08:14	17:20	20	9	Clear	G
2013-03-10	L19-2	08:00	12:43	20	4.3	Cloudy	G
	GW Librae	12:57	17:26	30	2.5	Cloudy / Overcast	-
2013-03-11	L19-2	08:05	17:28	20	5	Cloudy / Overcast	G
2013-03-12	L19-2	08:22	17:34	20	9	Clear	G
2013-03-13	L19-2	08:14	12:27	20	4	Clear	G
2013-03-13	GW Librae	12:46	15:24	20	2.5	Clear	-
2013-03-13	GW Librae	15:25	17:31	30	2	Clear	-
2013-03-14	-	-	-	-	-	Overcast	-

Table D.11: Observation log for March 2013

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2013-04-08	L19-2	07:28	10:16	20	0.3	Cloudy / Overcast	G
2013-04-09	L19-2	07:25	10:48	20	3	Cloudy	G
	WDJ1518	12:06	13:42	45	1.75	Cloudy	J
	WDJ1518	13:57	16:30	30	2.25	Cloudy	J
	L19-2	16:49	18:07	20	1.3	Clear	G
2013-04-10	L19-2	07:01	11:56	20	5	Clear	G
	WDJ1518	12:29	13:01	30	0.5	Cloudy	J
	WDJ1518	13:06	16:59	60	4	Cloudy	J
	L19-2	17:20	18:08	20	0.8	Clear	G
2013-04-11	WDJ1518	15:10	18:05	60	3	Cloudy	J
2013-04-12	L19-2	07:10	08:32	20	1.3	Clear	G
	PG1336	09:01	13:00	10	4	Clear	J K
	WDJ1518	13:23	17:48	60	4.5	Cloudy	J
2013-04-13	L19-2	07:03	11:39	20	4.5	Clear	G
	WDJ1518	12:18	17:57	60	7.6	Clear	J
2013-04-14	-	-	-	-	-	Overcast	-
2013-04-15	L19-2	07:03	11:44	20	4.75	Clear	G
	WDJ1518	12:04	18:00	60	6	Clear	J

Table D.12: Observation log for April 2013

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2013-05-09	L19-2	09:03	09:28	20	0.2	Cloud/Overcast	G
2013-05-10	L19-2	06:42	09:35	20	3	Clear	G
	WDJ1518	10:03	15:50	60	5.5	Cloudy	J
	L19-2	16:05	18:40	20	2.5	Clear	G
2013-05-11	L19-2	06:18	09:40	20	2.5	Cloudy	G
	WDJ1518	10:08	16:09	60	5	Cloudy	J
	L19-2	16:25	18:38	20	2	Cloudy	G
2013-05-12	L19-2	06:18	09:39	20	3.25	Clear	G
	WDJ1518	10:08	15:39	60	5.5	Clear	J
	L19-2	15:56	18:41	20	2.75	Clear	G
2013-05-13	L19-2	06:21	09:41	20	3.3	Clear	G
	WDJ1518	10:02	15:43	60	5.5	Clear	J
	L19-2	15:59	18:26	20	3.5	Clear	G
2013-05-14	L19-2	10:46	12:01	20	1.25	Cloudy	G
	WDJ1518	12:22	15:42	60	2.5	Cloudy	J
	L19-2	15:56	18:50	20	3	Cloudy	G
2013-05-15	L19-2	06:28	11:50	20	4.3	Cloudy	G
	WDJ1518	12:12	13:58	60	1.5	Cloudy	J
	L19-2	14:09	18:49	20	4	Cloudy	G

Table D.13: Observation log for May 2013 (WET Run XCOV29)

Date	Target	Start UTC	End UTC	Exposure (s)	Useable (hours)	Weather	Reference
2013-06-10	L19-2	06:05	17:54	20	9	Clear/Overcast	G
2013-06-11	L19-2	15:08	19:10	20	4	Cloudy	G
2013-06-12	L19-2	10:33	19:07	20	8.5	Clear	G
2013-06-13	L19-2	16:17	17:27	20	0.8	Cloudy/Overcast	G
2013-06-14	L19-2	06:02	08:30	20	0.8	Cloudy	G
2013-06-14	GW Librae	10:42	12:54	30	2	Cloudy	D
2013-06-14	L19-2	13:13	13:58	20	0.5	Cloudy	G
2013-06-15	-	-	-	-	-	Overcast	-
2013-06-16	-	-	-	-	-	Overcast	-

Table D.14: Observation log for June 2013

Bibliography

- Althaus, L. G., A. H. Córscico, J. Isern, and E. García-Berro
2010. Evolutionary and pulsational properties of white dwarf stars. *Astronomy & Astrophysics Reviews*, 18:471–566. doi:10.1007/s00159-010-0033-1.
- Araujo-Betancor, S., B. T. Gänsicke, H.-J. Hagen, T. R. Marsh, E. T. Harlaftis, J. Thorstensen, R. E. Fried, P. Schmeer, and D. Engels
2005. HS 2331+3905: The cataclysmic variable that has it all. *Astronomy & Astrophysics*, 430:629–642. doi:10.1051/0004-6361:20041736.
- Arras, P., D. M. Townsley, and L. Bildsten
2006. Pulsational Instabilities in Accreting White Dwarfs. *The Astrophysical Journal, Letters*, 643:L119–L122. doi:10.1086/505178.
- Basu, S. and H. M. Antia
2008. Helioseismology and solar abundances. *Physics Reports*, 457:217–283. doi:10.1016/j.physrep.2007.12.002.
- Beauchamp, A., F. Wesemael, P. Bergeron, G. Fontaine, R. A. Saffer, J. Liebert, and P. Brassard
1999. Spectroscopic Studies of DB White Dwarfs: The Instability Strip of the Pulsating DB (V777 Herculis) Stars. *The Astrophysical Journal*, 516:887–891. doi:10.1086/307148.
- Bergeron, P., R. A. Saffer, and J. Liebert
1992. A spectroscopic determination of the mass distribution of DA white dwarfs. *The Astrophysical Journal*, 394:228–247. doi:10.1086/171575.
- Bergeron, P., F. Wesemael, R. Lamontagne, G. Fontaine, R. A. Saffer, and N. F. Allard
1995. Optical and Ultraviolet Analyses of ZZ Ceti Stars and Study of the Atmospheric Convective Efficiency in DA White Dwarfs. *The Astrophysical Journal*, 449:258. doi:10.1086/176053.
- Bildsten, L., K. J. Shen, N. N. Weinberg, and G. Nelemans
2007. Faint Thermonuclear Supernovae from AM Canum Venaticorum Binaries.

The Astrophysical Journal, Letters, 662:L95–L98. doi:10.1086/519489.

Bohm, K. H. and J. Cassinelli

1971. Convective Envelopes and Acoustic Noise Generation in White Dwarfs. *Astronomy & Astrophysics*, 12:21.

Böhm-Vitense, E.

1958. Über die Wasserstoffkonvektionszone in Sternen verschiedener Effektivtemperaturen und Leuchtkräfte. Mit 5 Textabbildungen. *Zeitschrift fuer Astrophysik*, 46:108.

Bradley, P. A.

2001. Asteroseismological Constraints on the Structure of the ZZ Ceti Stars L19-2 and GD 165. *The Astrophysical Journal*, 552:326–339. doi:10.1086/320454.

Bradley, P. A. and D. E. Winget

1994. An asteroseismological determination of the structure of the DBV white dwarf GD 358. *The Astrophysical Journal*, 430:850–857. doi:10.1086/174456.

Brickhill, A. J.

1983. The pulsations of ZZ Ceti stars. *Monthly Notices of the Royal Astronomical Society*, 204:537–556.

Brickhill, A. J.

1990. The Pulsations of Zz-Ceti Stars - Part Two - Turbulent Convection and the Maximum Period. *Monthly Notices of the Royal Astronomical Society*, 246:510.

Brickhill, A. J.

1991a. The pulsations of ZZ Ceti stars. III - The driving mechanism. *Monthly Notices of the Royal Astronomical Society*, 251:673–680.

Brickhill, A. J.

1991b. The pulsations of ZZ Ceti stars. IV - The instability strip. *Monthly Notices of the Royal Astronomical Society*, 252:334–341.

Brickhill, A. J.

1992. The Pulsations of Zz-Ceti Stars - Part Six - the Amplitude Spectra. *Monthly Notices of the Royal Astronomical Society*, 259:529.

Bringans, R. D., D. J. Sullivan, and H. J. Trodahl

1974. Color Indices of HR 1653. *Publications of the Astronomical Society of the Pacific*, 86:693. doi:10.1086/129668.

Brown, W. R., M. Kilic, C. Allende Prieto, A. Gianninas, and S. J. Kenyon

2013. The ELM Survey. V. Merging Massive White Dwarf Binaries. *The Astrophysical Journal*, 769:66. doi:10.1088/0004-637X/769/1/66.

Brown, W. R., M. Kilic, C. Allende Prieto, and S. J. Kenyon

2010. The ELM Survey. I. A Complete Sample of Extremely Low-mass White Dwarfs. *The Astrophysical Journal*, 723:1072–1081. doi:10.1088/0004-637X/723/2/1072.

Bullock, E., P. Szkody, A. S. Mukadam, B. W. Borges, L. Fraga, B. T. Gänsicke, T. E. Harrison, A. Henden, J. Holtzman, S. B. Howell, W. A. Lawson, S. Levine, R. M. Plotkin, M. Seibert, M. Templeton, J. Teske, and F. J. Vrba

2011. GALEX and Optical Observations of GW Librae during the Long Decline from Superoutburst. *The Astronomical Journal*, 141:84. doi:10.1088/0004-6256/141/3/84.

Byckling, K., J. P. Osborne, P. J. Wheatley, G. A. Wynn, A. Beardmore, V. Braitto, K. Mukai, and R. G. West

2009. Swift observations of GW Lib: a unique insight into a rare outburst. *Monthly Notices of the Royal Astronomical Society*, 399:1576–1586. doi:10.1111/j.1365-2966.2009.15378.x.

Chandrasekhar, S.

1931. The Maximum Mass of Ideal White Dwarfs. *The Astrophysical Journal*, 74:81. doi:10.1086/143324.

Chanmugam, G.

1972. White dwarfs-Pulsation periods explained in terms of non-radial oscillations and gravity modes. *Nature*, 236:83–+.

Chote, P. and D. J. Sullivan

2013a. New Time-Series Observations of the Intriguing Object GW Librae. In *18th European White Dwarf Workshop.*, J. Krzesiński, G. Stachowski, P. Moskalik, and K. Bajan, eds., volume 469 of *Astronomical Society of the Pacific Conference Series*, P. 337.

Chote, P. and D. J. Sullivan

2013b. The Puoko-nui CCD Time-Series Photometer. In *18th European White Dwarf Workshop.*, J. Krzesiński, G. Stachowski, P. Moskalik, and K. Bajan, eds., volume 469 of *Astronomical Society of the Pacific Conference Series*, P. 249.

Chote, P., D. J. Sullivan, R. Brown, S. T. Harrold, D. E. Winget, and D. W. Chandler

2014. Puoko-nui: a flexible high-speed photometric system. *Monthly Notices of the Royal Astronomical Society*. doi:10.1093/mnras/stu348.

Chote, P., D. J. Sullivan, M. H. Montgomery, and J. L. Provencal

2013. Time series photometry of the helium atmosphere pulsating white dwarf

EC 04207-4748. *Monthly Notices of the Royal Astronomical Society*, 431:520–527. doi:10.1093/mnras/stt180.

Chote, P. R. B.

2011. A semi-analytical model for gravitational microlensing. Master’s thesis, Victoria University of Wellington.

Clemens, J. C.

1995. Perspectives of the Whole Earth Telescope Science. *Baltic Astronomy*, 4:142.

Cocke, W. J., M. J. Disney, and D. J. Taylor

1969. Discovery of Optical Signals from Pulsar NP 0532. *Nature*, 221:525–527. doi:10.1038/221525a0.

Copperwheat, C. M., T. R. Marsh, V. S. Dhillon, S. P. Littlefair, P. A. Woudt, B. Warner, D. Steeghs, B. T. Gänsicke, and J. Southworth

2009. ULTRACAM observations of two accreting white dwarf pulsators. *Monthly Notices of the Royal Astronomical Society*, 393:157–170. doi:10.1111/j.1365-2966.2008.14163.x.

Córsico, A. H., L. G. Althaus, E. García-Berro, and A. D. Romero

2013. An independent constraint on the secular rate of variation of the gravitational constant from pulsating white dwarfs. *Journal of Cosmology and Astroparticle Physics*, 6:32. doi:10.1088/1475-7516/2013/06/032.

Córsico, A. H., L. G. Althaus, M. M. Miller Bertolami, A. D. Romero, E. García-Berro, J. Isern, and S. O. Kepler

2012. The rate of cooling of the pulsating white dwarf star G117-B15A: a new asteroseismological inference of the axion mass. *Monthly Notices of the Royal Astronomical Society*, 424:2792–2799. doi:10.1111/j.1365-2966.2012.21401.x.

Costa, J. E. S. and S. O. Kepler

2008. The temporal changes of the pulsational periods of the pre-white dwarf PG 1159-035. *Astronomy & Astrophysics*, 489:1225–1232. doi:10.1051/0004-6361:20079118.

Costa, J. E. S., S. O. Kepler, and D. E. Winget

1999. Direct Measurement of a Secular Pulsation Period Change in the Pulsating Hot Pre-White Dwarf PG 1159-035. *The Astrophysical Journal*, 522:973–982. doi:10.1086/307655.

- Dalessio, J., D. J. Sullivan, J. L. Provencal, H. L. Shipman, T. Sullivan, D. Kilkenny, L. Fraga, and R. Sefako
2013. Periodic Variations in the O - C Diagrams of Five Pulsation Frequencies of the DB White Dwarf EC 20058-5234. *The Astrophysical Journal*, 765:5. doi:10.1088/0004-637X/765/1/5.
- Dhillon, V. S., T. R. Marsh, M. J. Stevenson, D. C. Atkinson, P. Kerry, P. T. Peacocke, A. J. A. Vick, S. M. Beard, D. J. Ives, D. W. Lunney, S. A. McLay, C. J. Tierney, J. Kelly, S. P. Littlefair, R. Nicholson, R. Pashley, E. T. Harlaftis, and K. O'Brien
2007. ULTRACAM: an ultrafast, triple-beam CCD camera for high-speed astrophysics. *Monthly Notices of the Royal Astronomical Society*, 378:825–840. doi:10.1111/j.1365-2966.2007.11881.x.
- Duerbeck, H. W. and W. C. Seitter
1987. A survey of nova remnants. *Astrophysics Space Science*, 131:467–473. doi:10.1007/BF00668127.
- Dufour, P., G. Fontaine, J. Liebert, G. D. Schmidt, and N. Behara
2008. Hot DQ White Dwarfs: Something Different. *The Astrophysical Journal*, 683:978–989. doi:10.1086/589855.
- Eastman, J., R. Siverd, and B. S. Gaudi
2010. Achieving Better Than 1 Minute Accuracy in the Heliocentric and Barycentric Julian Dates. *Publications of the Astronomical Society of the Pacific*, 122:935–946. doi:10.1086/655938.
- Fontaine, G. and P. Brassard
2008. The Pulsating White Dwarf Stars. *Publications of the Astronomical Society of the Pacific*, 120:1043–1096. doi:10.1086/592788.
- Fontaine, G., P. Brassard, and P. Bergeron
2001. The Potential of White Dwarf Cosmochronology. *Publications of the Astronomical Society of the Pacific*, 113:409–435. doi:10.1086/319535.
- Fontaine, G., E. M. Green, P. Chayer, P. Brassard, S. Charpinet, and S. K. Randall
2006. On the Empirical Instability Domains for Pulsating Subdwarf B Stars. *Baltic Astronomy*, 15:211–218.
- Fowler, R. H.
1926. On dense matter. *Monthly Notices of the Royal Astronomical Society*, 87:114–122.
- Gänsicke, B. T., P. Rodríguez-Gil, T. R. Marsh, D. de Martino, J. Nestoras, P. Szkody, A. Aungwerojwit, S. C. C. Barros, M. Dillon, S. Araujo-Betancor,

M. J. Arévalo, J. Casares, P. J. Groot, U. Kolb, C. Lázaro, P. Hakala, I. G. Martínez-Pais, G. Nelemans, G. Roelofs, M. R. Schreiber, E. van den Besselaar, and C. Zurita

2006. A ZZCeti white dwarf in SDSSJ133941.11+484727.5. *Monthly Notices of the Royal Astronomical Society*, 365:969–976. doi:10.1111/j.1365-2966.2005.09781.x.

Gänsicke, B. T., P. Szkody, S. B. Howell, and E. M. Sion

2005. Hubble Space Telescope STIS Observations of the Accreting White Dwarfs in BW Sculptoris, BC Ursae Majoris, and SW Ursae Majoris. *The Astrophysical Journal*, 629:451–460. doi:10.1086/431271.

Goldreich, P. and Y. Wu

1999. Gravity Modes in ZZ Ceti Stars. I. Quasi-adiabatic Analysis of Overstability. *The Astrophysical Journal*, 511:904–915. doi:10.1086/306705.

Hermes, J. J., W. R. Brown, M. Kilic, A. Gianninas, P. Chote, D. Sullivan, D. E. Winget, K. J. Bell, K. I. Winget, E. L. Robinson, P. Mason, S. T. Harrold, and M. H. Montgomery

2014. Radius constraints and relativistic effects from high-speed photometry of 20 merging low-mass white dwarf binaries. *in preparation*.

Hermes, J. J., M. Kilic, W. R. Brown, D. E. Winget, C. Allende Prieto, A. Gianninas, A. S. Mukadam, A. Cabrera-Lavers, and S. J. Kenyon

2012a. Rapid Orbital Decay in the 12.75-minute Binary White Dwarf J0651+2844. *The Astrophysical Journal*, 757:L21. doi:10.1088/2041-8205/757/2/L21.

Hermes, J. J., M. H. Montgomery, F. Mullally, D. E. Winget, and A. Bischoff-Kim

2013a. A New Timescale for Period Change in the Pulsating DA White Dwarf WD 0111+0018. *The Astrophysical Journal*, 766:42. doi:10.1088/0004-637X/766/1/42.

Hermes, J. J., M. H. Montgomery, D. E. Winget, W. R. Brown, A. Gianninas, M. Kilic, S. J. Kenyon, K. J. Bell, and S. T. Harrold

2013b. Discovery of Pulsations, Including Possible Pressure Modes, in Two New Extremely Low Mass, He-core White Dwarfs. *The Astrophysical Journal*, 765:102. doi:10.1088/0004-637X/765/2/102.

Hermes, J. J., M. H. Montgomery, D. E. Winget, W. R. Brown, M. Kilic, and S. J. Kenyon

2012b. SDSS J184037.78+642312.3: The First Pulsating Extremely Low Mass White Dwarf. *The Astrophysical Journal, Letters*, 750:L28. doi:10.1088/2041-8205/750/2/L28.

Hesser, J. E., B. M. Lasker, and H. E. Neupert

1977. High-frequency stellar oscillations. XII - L19-2, a low-amplitude ZZ Ceti variable with periods of 193 and 114 seconds. *The Astrophysical Journal, Letters*, 215:L75–L78. doi:10.1086/182482.

Hester, J. J.

2008. The Crab Nebula: An Astrophysical Chimera. *Annual Review of Astronomy & Astrophysics*, 46:127–155. doi:10.1146/annurev.astro.45.051806.110608.

Howell, S. B.

2006. *Handbook of CCD Astronomy*.

Howell, S. B., P. Szkody, and J. K. Cannizzo

1995. Tremendous outburst amplitude dwarf novae. *The Astrophysical Journal*, 439:337–345. doi:10.1086/175177.

IAU SOFA Board

2010. Iau sofa software collection. Available from: <http://www.iausofa.org>.

Iben, Jr., I. and A. V. Tutukov

1985. On the evolution of close binaries with components of initial mass between 3 solar masses and 12 solar masses. *The Astrophysical Journal, Supplement*, 58:661–710. doi:10.1086/191054.

Kanaan, A., A. Nitta, D. E. Winget, S. O. Kepler, M. H. Montgomery, T. S. Metcalfe, H. Oliveira, L. Fraga, A. F. M. da Costa, J. E. S. Costa, B. G. Castanheira, O. Giovannini, R. E. Nather, A. Mukadam, S. D. Kawaler, M. S. O'Brien, M. D. Reed, S. J. Kleinman, J. L. Provencal, T. K. Watson, D. Kilkenney, D. J. Sullivan, T. Sullivan, B. Shobbrook, X. J. Jiang, B. N. Ashoka, S. Seetha, E. Leibowitz, P. Ibbetson, H. Mendelson, E. G. Meistas, R. Kalytis, D. Ališauskas, D. O'Donoghue, D. Buckley, P. Martinez, F. van Wyk, R. Stobie, F. Marang, L. van Zyl, W. Ogloza, J. Krzesinski, S. Zola, P. Moskalik, M. Breger, A. Stankov, R. Silvotti, A. Piccioni, G. Vauclair, N. Dolez, M. Chevreton, J. Deetjen, S. Dreizler, S. Schuh, J. M. Gonzalez Perez, R. Østensen, A. Ulla, M. Manteiga, O. Suarez, M. R. Burleigh, and M. A. Barstow

2005. Whole Earth Telescope observations of BPM 37093: A seismological test of crystallization theory in white dwarfs. *Astronomy & Astrophysics*, 432:219–224. doi:10.1051/0004-6361:20041125.

Kato, T., H. Maehara, and B. Monard

2008. Late Superhumps in WZ Sge-Type Dwarf Novae. *Publications of the ASJ*, 60:L23.

Kawaler, S. D. and P. A. Bradley

1994. Precision asteroseismology of pulsating PG 1159 stars. *The Astrophysical Journal*, 427:415–428. doi:10.1086/174152.

Kepler, S. O., J. E. S. Costa, B. G. Castanheira, D. E. Winget, F. Mullally, R. E. Nather, M. Kilic, T. von Hippel, A. S. Mukadam, and D. J. Sullivan

2005. Measuring the Evolution of the Most Stable Optical Clock G 117-B15A. *The Astrophysical Journal*, 634:1311–1318. doi:10.1086/497002.

Kepler, S. O., S. J. Kleinman, A. Nitta, D. Koester, B. G. Castanheira, O. Giovannini, A. F. M. Costa, and L. Althaus

2007. White dwarf mass distribution in the SDSS. *Monthly Notices of the Royal Astronomical Society*, 375:1315–1324. doi:10.1111/j.1365-2966.2006.11388.x.

Kepler, S. O., R. E. Nather, J. T. McGraw, and E. L. Robinson

1982. The pulsation periods of the pulsating white dwarf G117-B15A. *The Astrophysical Journal*, 254:676–682. doi:10.1086/159779.

Kepler, S. O., D. E. Winget, R. E. Nather, P. A. Bradley, A. D. Grauer, G. Fontaine, P. Bergeron, G. Vauclair, C. F. Claver, T. M. K. Marar, S. Seetha, B. N. Ashoka, T. Mazeh, E. Leibowitz, N. Dolez, M. Chevreton, M. A. Barstow, J. C. Clemens, S. J. Kleinman, A. E. Sansom, R. W. Tweedy, A. Kanaan, B. P. Hine, J. L. Provencal, F. Wesemael, M. A. Wood, P. Brassard, J.-E. Solheim, and P.-I. Emanuelsen

1991. A detection of the evolutionary time scale of the DA white dwarf G117 - B15A with the Whole Earth Telescope. *The Astrophysical Journal, Letters*, 378:L45–L48. doi:10.1086/186138.

Kilic, M., J. J. Hermes, A. Gianninas, W. R. Brown, C. O. Heinke, M. A. Agüeros, P. Chote, D. J. Sullivan, K. J. Bell, and S. T. Harrold

2014. Found: the progenitors of am cvn and supernovae ia. *Monthly Notices of the Royal Astronomical Society: Letters*, 438(1):L26–L30. Available from: <http://mnrasl.oxfordjournals.org/content/438/1/L26.abstract>, doi:10.1093/mnrasl/slt151.

Kilkenny, D., D. O'Donoghue, L. A. Crause, N. Hambly, and H. MacGillivray

2009. Two new pulsating DB Stars from the EC survey, EC 04207-4748 and EC 05221-4725. *Monthly Notices of the Royal Astronomical Society*, 397:453–457. doi:10.1111/j.1365-2966.2009.14966.x.

Kilkenny, D., M. D. Reed, D. O'Donoghue, S. D. Kawaler, A. Mukadam, S. J. Kleinman, A. Nitta, T. S. Metcalfe, J. L. Provencal, T. K. Watson, D. J. Sullivan, T. Sullivan, R. Shobbrook, X. J. Jiang, S. Joshi, B. N. Ashoka, S. Seetha, E. Leibowitz, P. Ibbetson, H. Mendelson, E. Meištas, R. Kalytis, D. Ališauskas, P. Mar-

- tinez, F. van Wyk, R. S. Stobie, F. Marang, S. Zola, J. Krzesinski, W. Ogłóza, P. Moskalik, R. Silvotti, A. Piccioni, G. Vauclair, N. Dolez, M. Chevreton, S. Dreizler, S. L. Schuh, J. L. Deetjen, J.-E. Solheim, J. M. Gonzalez Perez, A. Ulla, R. Østensen, M. Manteiga, O. Suarez, M. Burleigh, S. O. Kepler, A. Kanaan, and O. Giovannini
2003. A Whole Earth Telescope campaign on the pulsating subdwarf B binary system PG 1336-018 (NY Vir). *Monthly Notices of the Royal Astronomical Society*, 345:834–846. doi:10.1046/j.1365-8711.2003.07007.x.
- Kilkenny, D., B. Y. Welsh, C. Koen, A. A. S. Gulbis, and M. M. Kotze
2014. A search for p-mode pulsations in white dwarf stars using the Berkeley Visible Imaging Tube detector. *Monthly Notices of the Royal Astronomical Society*, 437:1836–1839. doi:10.1093/mnras/stt2029.
- Kleinman, S. J., R. E. Nather, and T. Phillips
1996. The WET standard photometer. *Publications of the Astronomical Society of the Pacific*, 108:356.
- Koester, D. and N. F. Allard
2000. The ZZ Ceti Instability Strip Revisited. *Baltic Astronomy*, 9:119–124.
- Koester, D., R. Napiwotzki, N. Christlieb, H. Drechsel, H.-J. Hagen, U. Heber, D. Homeier, C. Karl, B. Leibundgut, S. Moehler, G. Nelemans, E.-M. Pauli, D. Reimers, A. Renzini, and L. Yungelson
2001. High-resolution UVES/VLT spectra of white dwarfs observed for the ESO SN Ia progenitor survey (SPY). I. *Astronomy & Astrophysics*, 378:556–568. doi:10.1051/0004-6361:20011235.
- Landolt, A. U.
1968. A New Short-Period Blue Variable. *The Astrophysical Journal*, 153:151. doi:10.1086/149645.
- Lyne, A. G., R. S. Pritchard, and F. Graham-Smith
1993. Twenty-Three Years of Crab Pulsar Rotational History. *Monthly Notices of the Royal Astronomical Society*, 265:1003.
- Lyne, A. G., M. E. Roberts, and C. A. Jordan
2014. Available from: <http://www.jb.man.ac.uk/~pulsar/crab.html>.
- Mackay, C. D., R. N. Tubbs, R. Bell, D. J. Burt, P. Jerram, and I. Moody
2001. Subelectron read noise at mhz pixel rates. Available from: <http://dx.doi.org/10.1117/12.426988>, doi:10.1117/12.426988.
- Matsumoto, M. and T. Nishimura
1998. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-

- random number generator. *ACM Trans. Model. Comput. Simul.*, 8(1):3–30. Available from: <http://doi.acm.org/10.1145/272991.272995>, doi:10.1145/272991.272995.
- Maza, J. and L. E. Gonzalez
1983. Novalike Object in Libra. *IAU Circulars*, 3854:2.
- McGraw, J. T., S. G. Starrfield, J. R. P. Angel, and N. P. Carleton
1979. PG1159-035: An MMT Discovery of a New Short-Period Variable Star. In *The MMT and the Future of Ground-Based Astronomy*, T. C. Weekes, ed., volume 385 of *SAO Special Report*, P. 125.
- Mestel, L.
1952. On the theory of white dwarf stars. I. The energy sources of white dwarfs. *Monthly Notices of the Royal Astronomical Society*, 112:583.
- Montgomery, M. H.
2005. A New Technique for Probing Convection in Pulsating White Dwarf Stars. *The Astrophysical Journal*, 633:1142–1149. doi:10.1086/466511.
- Montgomery, M. H., J. L. Provencal, A. Kanaan, A. S. Mukadam, S. E. Thompson, J. Dalessio, H. L. Shipman, D. E. Winget, S. O. Kepler, and D. Koester
2010. Evidence for Temperature Change and Oblique Pulsation from Light Curve Fits of the Pulsating White Dwarf GD 358. *The Astrophysical Journal*, 716:84–96. doi:10.1088/0004-637X/716/1/84.
- Montgomery, M. H., K. A. Williams, D. E. Winget, P. Dufour, S. De Gennaro, and J. Liebert
2008. SDSS J142625.71+575218.3: A Prototype for a New Class of Variable White Dwarf. *The Astrophysical Journal, Letters*, 678:L51–L54. doi:10.1086/588286.
- Mukadam, A. S., A. Bischoff-Kim, O. Fraser, A. H. Córscico, M. H. Montgomery, S. O. Kepler, A. D. Romero, D. E. Winget, J. J. Hermes, T. S. Riecken, M. E. Kronberg, K. I. Winget, R. E. Falcon, D. W. Chandler, J. W. Kuehne, D. J. Sullivan, D. Reaves, T. von Hippel, F. Mullally, H. Shipman, S. E. Thompson, N. M. Silvestri, and R. I. Hynes
2013. Measuring the Evolutionary Rate of Cooling of ZZ Ceti. *The Astrophysical Journal*, 771:17. doi:10.1088/0004-637X/771/1/17.
- Mukadam, A. S., B. T. Gänsicke, P. Szkody, A. Aungwerojwit, S. B. Howell, O. J. Fraser, and N. M. Silvestri
2007. Discovery of Two New Accreting Pulsating White Dwarf Stars. *The Astrophysical Journal*, 667:433–441. doi:10.1086/520700.
- Mukadam, A. S., R. Owen, E. Mannery, N. MacDonald, B. Williams, F. Stauffer,

- and C. Miller
2011a. High-Speed Time-Series CCD Photometry with Agile. *Publications of the Astronomical Society of the Pacific*, 123:1423–1433. doi:10.1086/663725.
- Mukadam, A. S., D. M. Townsley, B. T. Gänsicke, P. Szkody, T. R. Marsh, E. L. Robinson, L. Bildsten, A. Aungwerojwit, M. R. Schreiber, J. Southworth, A. Schwope, B.-Q. For, G. Tovmassian, S. V. Zharikov, M. G. Hidas, N. Baliber, T. Brown, P. A. Woudt, B. Warner, D. O’Donoghue, D. A. H. Buckley, R. Sefako, and E. M. Sion
2010. Multi-site Observations of Pulsation in the Accreting White Dwarf SDSS J161033.64-010223.3 (V386 Ser). *The Astrophysical Journal*, 714:1702–1714. doi:10.1088/0004-637X/714/2/1702.
- Mukadam, A. S., D. M. Townsley, P. Szkody, B. T. Gänsicke, D. E. Winget, J. J. Hermes, S. B. Howell, J. Teske, J. Patterson, J. Kemp, and E. Armstrong
2011b. First Unambiguous Detection of the Return of Pulsations in the Accreting White Dwarf SDSS J074531.92+453829.6 After an Outburst. *The Astrophysical Journal, Letters*, 728:L33. doi:10.1088/2041-8205/728/2/L33.
- Mullally, F., D. E. Winget, S. De Gennaro, E. Jeffery, S. E. Thompson, D. Chandler, and S. O. Kepler
2008. Limits on Planets around Pulsating White Dwarf Stars. *The Astrophysical Journal*, 676:573–583. doi:10.1086/528672.
- Nather, R. E. and A. S. Mukadam
2004. A CCD Time-Series Photometer. *The Astrophysical Journal*, 605:846–853. doi:10.1086/382676.
- Nather, R. E., B. Warner, and M. Macfarlane
1969. Optical Pulsations in the Crab Nebula Pulsar. *Nature*, 221:527–529. doi:10.1038/221527a0.
- Nather, R. E., D. E. Winget, J. C. Clemens, C. J. Hansen, and B. P. Hine
1990. The whole earth telescope - A new astronomical instrument. *The Astrophysical Journal*, 361:309–317. doi:10.1086/169196.
- Nilsson, R., H. Uthas, M. Ytre-Eide, J.-E. Solheim, and B. Warner
2006. New pulsating white dwarfs in cataclysmic variables. *Monthly Notices of the Royal Astronomical Society*, 370:L56–L60. doi:10.1111/j.1745-3933.2006.00188.x.
- O’Donoghue, D. E. and B. Warner
1982. Observations and analysis of the low amplitude ZZ Ceti variable, L19-2. *Monthly Notices of the Royal Astronomical Society*, 200:563–583.

O'Donoghue, D. E. and B. Warner

1987. The period structure and stability of the pulsating white dwarf L19-2. *Monthly Notices of the Royal Astronomical Society*, 228:949–955.

Pajdosz, G.

1995. Non-evolutionary secular period increase in pulsating DA white dwarfs. *Astronomy & Astrophysics*, 295:L17–L19.

Patterson, J., J. R. Thorstensen, and J. Kemp

2005. Pulsations, Boundary Layers, and Period Bounce in the Cataclysmic Variable RE J1255+266. *Publications of the Astronomical Society of the Pacific*, 117:427–444. doi:10.1086/429786.

Patterson, J., J. R. Thorstensen, and C. Knigge

2008. SDSS 1507+52: A Halo Cataclysmic Variable? *Publications of the Astronomical Society of the Pacific*, 120:510–522. doi:10.1086/588615.

Pavlenko, E.

2009. The white dwarf in dwarf nova SDSS J080434.20+510349.2: Entering the instability strip? *Journal of Physics Conference Series*, 172(1):012071. doi:10.1088/1742-6596/172/1/012071.

Paxton, B., M. Cantiello, P. Arras, L. Bildsten, E. F. Brown, A. Dotter,

C. Mankovich, M. H. Montgomery, D. Stello, F. X. Timmes, and R. Townsend
2013. Modules for Experiments in Stellar Astrophysics (MESA): Planets, Oscillations, Rotation, and Massive Stars. *The Astrophysical Journal, Supplement*, 208:4. doi:10.1088/0067-0049/208/1/4.

Provencal, J. L., M. H. Montgomery, A. Kanaan, S. E. Thompson, J. Dalessio, H. L.

Shipman, D. Childers, J. C. Clemens, R. Rosen, P. Henrique, A. Bischoff-Kim, W. Strickland, D. Chandler, B. Walter, T. K. Watson, B. Castanheira, S. Wang, G. Handler, M. Wood, S. Vennes, P. Nemeth, S. O. Kepler, M. Reed, A. Nitta, S. J. Kleinman, T. Brown, S.-L. Kim, D. Sullivan, W. P. Chen, M. Yang, C. Y. Shih, X. J. Jiang, A. V. Sergeev, A. Maksim, R. Janulis, K. S. Baliyan, H. O. Vats, S. Zola, A. Baran, M. Winiarski, W. Ogloza, M. Paparo, Z. Bognar, P. Papics, D. Kilkenny, R. Sefako, D. Buckley, N. Loaring, A. Kniazev, R. Silvotti, S. Galletti, T. Nagel, G. Vauclair, N. Dolez, J. R. Fremy, J. Perez, J. M. Almenara, and L. Fraga

2012. Empirical Determination of Convection Parameters in White Dwarfs. I. Whole Earth Telescope Observations of EC14012-1446. *The Astrophysical Journal*, 751:91. doi:10.1088/0004-637X/751/2/91.

Ringwald, F. A., T. Naylor, and K. Mukai

1996. The optical spectra of old novae. *Monthly Notices of the Royal Astronomical Society*, 281:192–210.
- Robinson, E. L., T. M. Mailloux, E. Zhang, D. Koester, R. F. Stiening, R. C. Bless, J. W. Percival, M. J. Taylor, and G. W. van Citters
1995. The pulsation index, effective temperature, and thickness of the hydrogen layer in the pulsating DA white dwarf G117-B15A. *The Astrophysical Journal*, 438:908–916. doi:10.1086/175132.
- Romero, A. D., S. O. Kepler, A. H. Córscico, L. G. Althaus, and L. Fraga
2013. Asteroseismological Study of Massive ZZ Ceti Stars with Fully Evolutionary Models. *The Astrophysical Journal*, 779:58. doi:10.1088/0004-637X/779/1/58.
- Scargle, J. D.
1982. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*, 263:835–853. doi:10.1086/160554.
- Schwieterman, E. W., M. A. Wood, D. Piwowar, J. Patterson, R. Rea, B. Monard, T. Krajci, G. Bolt, G. Roberts, J. Foote, and J. McCormick
2010. Time-Series Photometry of GW Librae One Year After Outburst. *Journal of the Southeastern Association for Research in Astronomy*, 3:6–10.
- Silvotti, R., G. Fontaine, M. Pavlov, T. R. Marsh, V. S. Dhillon, S. P. Littlefair, and F. Getman
2011. Search for p-mode oscillations in DA white dwarfs with VLT-ULTRACAM. I. Upper limits to the p-modes. *Astronomy & Astrophysics*, 525:A64. doi:10.1051/0004-6361/201015334.
- Sion, E. M., J. L. Greenstein, J. D. Landstreet, J. Liebert, H. L. Shipman, and G. A. Wegner
1983. A proposed new white dwarf spectral classification system. *The Astrophysical Journal*, 269:253–257. doi:10.1086/161036.
- Steinfadt, J. D. R., L. Bildsten, and P. Arras
2010. Pulsations in Hydrogen Burning Low-mass Helium White Dwarfs. *The Astrophysical Journal*, 718:441–445. doi:10.1088/0004-637X/718/1/441.
- Stobie, R. S., D. Kilkeny, D. O'Donoghue, et al.
1997. The Edinburgh-Cape Blue Object Survey - I. Description of the survey. *Monthly Notices of the Royal Astronomical Society*, 287:848.
- Stone, R. C.
1989. A comparison of digital centering algorithms. *The Astronomical Journal*, 97:1227–1237. doi:10.1086/115066.

- Stover, R. J., R. E. Nather, E. L. Robinson, J. E. Hesser, and B. M. Lasker
1980. Period stability of the pulsating white dwarf R548 (=ZZ Ceti). *The Astrophysical Journal*, 240:865–870. doi:10.1086/158300.
- Sullivan, D. J.
1998a. Frequency Stability of the DAV White Dwarf L19-2. In *A Half Century of Stellar Pulsation Interpretation*, P. A. Bradley and J. A. Guzik, eds., volume 135 of *Astronomical Society of the Pacific Conference Series*, P. 62.
- Sullivan, D. J.
1998b. The Pulsating White Dwarf L 19-2. *Baltic Astronomy*, 7:159–173.
- Sullivan, D. J.
1999. Period changes in the pulsating white dwarfs. In *11th European Workshop on White Dwarfs*, S.-E. Solheim and E. G. Meistas, eds., volume 169 of *Astronomical Society of the Pacific Conference Series*, P. 85.
- Sullivan, D. J.
2000. The New Zealand WET three channel photometer. *Baltic Astronomy*, 9:425.
- Sullivan, D. J.
2005. EC 20058-5234: A DBV White Dwarf and a Possible Plasmon Neutrino Detector. In *14th European Workshop on White Dwarfs*, D. Koester and S. Moehler, eds., volume 334 of *Astronomical Society of the Pacific Conference Series*, P. 495.
- Sullivan, D. J., T. S. Metcalfe, D. O'Donoghue, D. E. Winget, D. Kilkenney, F. van Wyk, A. Kanaan, S. O. Kepler, A. Nitta, S. D. Kawaler, M. H. Montgomery, R. E. Nather, M. S. O'Brien, A. Bischoff-Kim, M. Wood, X. J. Jiang, E. M. Leibowitz, P. Ibbetson, S. Zola, J. Krzesinski, G. Pajdosz, G. Vauclair, N. Dolez, and M. Chevreton
2008. Whole Earth Telescope observations of the hot helium atmosphere pulsating white dwarf EC20058-5234. *Monthly Notices of the Royal Astronomical Society*, 387:137–152. doi:10.1111/j.1365-2966.2008.13074.x.
- Sullivan, D. J., J. Pritchard, M. L. Reid, and T. Sullivan
2000. CCD Time-Series Observations of the DAV Cataclysmic Variable GW Librae. *Baltic Astronomy*, 9:223–230.
- Sullivan, D. J. and T. Sullivan
2000. Frequency Stabilities of Two Pulsating White Dwarfs. *Baltic Astronomy*, 9:81–86.
- Szkody, P., B. T. Gänsicke, S. B. Howell, and E. M. Sion
2002. Hubble Space Telescope Spectra of GW Librae: A Hot Pulsating White

- Dwarf in a Cataclysmic Variable. *The Astrophysical Journal, Letters*, 575:L79–L82. doi:10.1086/342916.
- Szkody, P., A. Mukadam, B. T. Gänsicke, A. Henden, M. Templeton, J. Holtzman, M. H. Montgomery, S. B. Howell, A. Nitta, E. M. Sion, R. D. Schwartz, and W. Dillon
2010. Finding the Instability Strip for Accreting Pulsating White Dwarfs From Hubble Space Telescope and Optical Observations. *The Astrophysical Journal*, 710:64–77. doi:10.1088/0004-637X/710/1/64.
- Szkody, P., A. S. Mukadam, B. T. Gänsicke, A. Henden, E. M. Sion, D. Townsley, P. Chote, D. Harmer, E. J. Harpe, J. J. Hermes, D. J. Sullivan, and D. E. Winget
2012a. HST and Optical Data Reveal White Dwarf Cooling, Spin, and Periodicities in GW Librae 3-4 Years after Outburst. *The Astrophysical Journal*, 753:158. doi:10.1088/0004-637X/753/2/158.
- Szkody, P., A. S. Mukadam, B. T. Gänsicke, E. M. Sion, D. Townsley, and A. Henden
2013. Enigmas of Accreting Pulsating White Dwarfs. In *18th European White Dwarf Workshop*, Krzesiński, J., G. Stachowski, P. Moskalik, and K. Bajan, eds., volume 469 of *Astronomical Society of the Pacific Conference Series*, P. 31.
- Szkody, P., A. S. Mukadam, B. T. Gänsicke, E. M. Sion, D. M. Townsley, A. Henden, D. J. Sullivan, and P. Chote
2012b. The instability strip for accreting pulsating white dwarfs as a probe for accretion heating/cooling. *Memorie della Societa Astronomica Italiana*, 83:513.
- Templeton, M., R. Stubbings, E. O. Waagen, P. Schmeer, A. Pearce, and P. Nelson
2007. GW Librae. *Central Bureau Electronic Telegrams*, 922:1.
- Thorstensen, J. R., J. Patterson, J. Kemp, and S. Vennes
2002. The Dwarf Novae of Shortest Period. *Publications of the Astronomical Society of the Pacific*, 114:1108–1116. doi:10.1086/342484.
- Townsley, D. M., P. Arras, and L. Bildsten
2004. Seismology of the Accreting White Dwarf in GW Librae. *The Astrophysical Journal, Letters*, 608:L105–L108. doi:10.1086/422411.
- Tremblay, P.-E., P. Bergeron, J. S. Kalirai, and A. Gianninas
2010. New Insights into the Problem of the Surface Gravity Distribution of Cool DA White Dwarfs. *The Astrophysical Journal*, 712:1345–1358. doi:10.1088/0004-637X/712/2/1345.
- Trodahl, H. J., D. J. Sullivan, and D. Beaglehole
1973. Monochromator Color Indices of δ Scuti and ρ Puppis. *Publications of the Astronomical Society of the Pacific*, 85:608. doi:10.1086/129513.

- Uthas, H., J. Patterson, J. Kemp, C. Knigge, B. Monard, R. Rea, G. Bolt, J. McCormick, G. Christie, A. Retter, and A. Liu
2012. Two new accreting, pulsating white dwarfs: SDSS J1457+51 and BW Sculptoris. *Monthly Notices of the Royal Astronomical Society*, 420:379–387. doi:10.1111/j.1365-2966.2011.20042.x.
- Van der Peet, A.
1987. *The VUW Two Channel Photometer*. Victoria University of Wellington.
- van Zyl, L., B. Warner, D. O'Donoghue, C. Hellier, P. Woudt, D. Sullivan, J. Pritchard, J. Kemp, J. Patterson, W. Welsh, J. Casares, T. Shahbaz, F. van der Hooft, and S. Vennes
2004. The non-radially pulsating primary of the cataclysmic variable GW Librae. *Monthly Notices of the Royal Astronomical Society*, 350:307–316. doi:10.1111/j.1365-2966.2004.07646.x.
- van Zyl, L., B. Warner, D. O'Donoghue, D. Sullivan, J. Pritchard, and J. Kemp
2000. GW Librae: An Accreting Variable White Dwarf. *Baltic Astronomy*, 9:231–246.
- Vanlandingham, K. M., G. J. Schwarz, and S. B. Howell
2005. Discovery of Nonradial Pulsations in PQ Andromedae. *Publications of the Astronomical Society of the Pacific*, 117:928–930. doi:10.1086/432456.
- Vican, L., J. Patterson, W. Allen, B. Goff, T. Krajci, J. McCormick, B. Monard, R. Rea, P. Nelson, G. Bolt, R. Koff, G. Roberts, M. Wood, and J. Kemp
2011. A Thousand Hours of GW Librae: The Eruption and Aftermath. *Publications of the Astronomical Society of the Pacific*, 123:1156–1168. doi:10.1086/662633.
- Voss, B., D. Koester, R. Napiwotzki, N. Christlieb, and D. Reimers
2007. High-resolution UVES/VLT spectra of white dwarfs observed for the ESO SN Ia progenitor survey. II. DB and DBA stars. *Astronomy & Astrophysics*, 470:1079–1088. doi:10.1051/0004-6361:20077285.
- Warner, B. and R. E. Nather
1970. Quasi-periodic outbursts in the whitedwarf Haro-Luyton Taurus No. 76. *Monthly Notices of the Royal Astronomical Society*, 147:21.
- Warner, B. and E. L. Robinson
1972. White dwarfs-More rapid variables. *Nature*, 239:2–7.
- Warner, B. and L. van Zyl
1998. Discovery of non-radial pulsations in the white dwarf primary of a cataclysmic variable star. In *New Eyes to See Inside the Sun and Stars*, F.-L. Deubner,

- J. Christensen-Dalsgaard, and D. Kurtz, eds., volume 185 of *IAU Symposium*, P. 321.
- Warner, B. and P. A. Woudt
2004. Pulsating white dwarfs in cataclysmic variables: The marriage of ZZ Cet and dwarf nova. In *IAU Colloq. 193: Variable Stars in the Local Group*, D. W. Kurtz and K. R. Pollard, eds., volume 310 of *Astronomical Society of the Pacific Conference Series*, P. 382.
- Warner, B. and P. A. Woudt
2008. QPOs in CVs: An executive summary. In *American Institute of Physics Conference Series*, M. Axelsson, ed., volume 1054 of *American Institute of Physics Conference Series*, Pp. 101–110. doi:10.1063/1.3002491.
- Winget, D. E., C. J. Hansen, and H. M. van Horn
1983a. Do pulsating PG1159-035 stars put constraints on stellar evolution? *Nature*, 303:781–782. doi:10.1038/303781a0.
- Winget, D. E. and S. O. Kepler
2008. Pulsating White Dwarf Stars and Precision Asteroseismology. *Annual Review of Astronomy & Astrophysics*, 46:157–199. doi:10.1146/annurev.astro.46.060407.145250.
- Winget, D. E., R. E. Nather, J. C. Clemens, J. Provencal, S. J. Kleinman, P. A. Bradley, M. A. Wood, C. F. Claver, M. L. Frueh, A. D. Grauer, B. P. Hine, C. J. Hansen, G. Fontaine, N. Achilleos, D. T. Wickramasinghe, T. M. K. Marar, S. Seetha, B. N. Ashoka, D. O'Donoghue, B. Warner, D. W. Kurtz, D. A. Buckley, J. Brickhill, G. Vauclair, N. Dolez, M. Chevreton, M. A. Barstow, J. E. Solheim, A. Kanaan, S. O. Kepler, G. W. Henry, and S. D. Kawaler
1991. Asteroseismology of the DOV star PG 1159 - 035 with the Whole Earth Telescope. *The Astrophysical Journal*, 378:326–346. doi:10.1086/170434.
- Winget, D. E., E. L. Robinson, R. D. Nather, and G. Fontaine
1982. Photometric observations of GD 358 - DB white dwarfs do pulsate. *The Astrophysical Journal, Letters*, 262:L11–L15. doi:10.1086/183902.
- Winget, D. E., E. L. Robinson, R. E. Nather, S. O. Kepler, and D. Odonoghue
1985. A measurement of secular evolution in the pre-white dwarf star. *The Astrophysical Journal*, 292:606–613. doi:10.1086/163193.
- Winget, D. E., D. J. Sullivan, T. S. Metcalfe, S. D. Kawaler, and M. H. Montgomery
2004. A Strong Test of Electroweak Theory Using Pulsating DB White Dwarf

Stars as Plasmon Neutrino Detectors. *The Astrophysical Journal, Letters*, 602:L109–L112. doi:10.1086/382591.

Winget, D. E., H. M. van Horn, M. Tassoul, C. J. Hansen, and G. Fontaine
1983b. The pulsation properties of DB white dwarfs - A preliminary analysis. *The Astrophysical Journal, Letters*, 268:L33–L36. doi:10.1086/184024.

Woudt, P. A. and B. Warner

2004. SDSS J161033.64-010223.3: a second cataclysmic variable with a non-radially pulsating primary. *Monthly Notices of the Royal Astronomical Society*, 348:599–602. doi:10.1111/j.1365-2966.2004.07369.x.

Woudt, P. A. and B. Warner

2011. VSX J074727.6 + 065050: A dwarf nova with a non-radially pulsating white dwarf primary. *Astrophysics Space Science*, 333:119–123. doi:10.1007/s10509-010-0573-x.

Wu, Y.

2001. Combination frequencies in the Fourier spectra of white dwarfs. *Monthly Notices of the Royal Astronomical Society*, 323:248–256. doi:10.1046/j.1365-8711.2001.04224.x.