

**Multi-wavelength observations of cataclysmic
variable stars**

by

Richard D.G. Hickman

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

May 2011

THE UNIVERSITY OF
WARWICK

Contents

Declarations	v
Abstract	vi
Abbreviations	vii
Chapter 1 Introduction	1
Chapter 2 Cataclysmic variables: a brief overview	3
2.1 Stellar evolution	3
2.2 The common envelope phase and Roche potential	4
2.3 The primary and the secondary	6
2.4 The disc and the bright spot	7
2.5 The boundary layer	9
2.6 Mass transfer and angular momentum	10
2.7 Classification of CVs	11
2.8 Dwarf novae outbursts	12
2.9 Viscosity	13
2.10 Elliptical discs	14
2.11 Observing CVs	15
2.12 Summary	16
Chapter 3 Detectors for modern astronomy	17
3.1 Detector technologies	18
3.1.1 Photographic plates	18
3.1.2 Photomultipliers	19
3.1.3 Digicons	19
3.2 Charge-coupled devices	20
3.2.1 Principles of operation	20

3.2.2	Buried channel potential well	23
3.2.3	Front and back illumination	23
3.2.4	Dark current and inverted operation	25
3.2.5	Noise sources	26
3.2.6	Signal-to-noise	28
3.2.7	CCD gain and dynamic range	28
3.3	Photometry	29
3.3.1	Aperture photometry	29
3.3.2	Bias frames	31
3.3.3	Dark frames	31
3.3.4	Flat fields	32
3.3.5	Flux calibration	32
3.4	Spectroscopy	33
3.4.1	Arc line calibration	34
3.4.2	Flux calibration	34
3.5	Summary	35
Chapter 4 Ultracam & Ultraspec: detectors for the study of CVs		36
4.1	Science goals	36
4.2	The common hardware stack	37
4.2.1	GPS timing	37
4.2.2	CCD controller & rack hardware	37
4.3	Ultracam	38
4.3.1	Collimators & beam-splitters	38
4.3.2	Ultracam CCDs	39
4.4	Ultraspec	41
4.4.1	EFOSC2	41
4.4.2	Ultraspec CCD	41
4.5	The common software stack	45
4.6	Summary	47
Chapter 5 Disc-rim absorption in CVs: the case for EX Dra		48
5.1	Observations and data reduction	49
5.1.1	Optical spectroscopy	50
5.1.2	HST ultraviolet spectroscopy	51
5.1.3	Optical photometry	53
5.2	Disc-rim absorption in EX-Dra	55
5.2.1	Evidence from the UV data	55

5.2.2	Evidence from the optical data	56
5.3	Radial velocities	58
5.3.1	Optical spectra: Balmer lines and CaI	59
5.3.2	UV spectra lines	60
5.4	Doppler tomography	62
5.4.1	Slingshot prominences	63
5.4.2	Spiral shocks	69
5.5	Summary	70
Chapter 6 Measuring the boundary layer of Z Cha		72
6.1	Introduction	72
6.2	Observations	73
6.2.1	XMM-Newton X-ray observations	73
6.2.2	XMM-Newton optical observations	76
6.2.3	Ultracam optical observations	76
6.2.4	Ultraspec optical observations	76
6.3	The boundary layer of Z Cha	78
6.3.1	Optical eclipse modelling	78
6.3.2	X-ray eclipse modelling	84
6.3.3	The X-ray spectrum	89
6.3.4	The area of the X-ray emission upon the white dwarf	89
6.3.5	The scale height above the white dwarf	92
6.3.6	The truncated inner disc	92
6.4	The second absorption site	94
6.5	The system parameters	95
6.6	The eclipse times	96
6.6.1	The sinusoidal ephemeris	96
6.6.2	The origin of the eclipse time variations	97
6.7	The colour of Z Cha's white dwarf	101
6.8	Summary	103
Chapter 7 X-Shooter observations of OY Car		106
7.1	Observations and data reduction	106
7.2	Testing the photometric method with X-Shooter	110
7.3	Iron lines in the spectrum of OY Car	115
7.4	Summary	118

Chapter 8	Improvements to the Ultracam system	120
8.1	Calibration of the Ultracam SDSS filters	120
8.1.1	The system response	121
8.1.2	Comparing the Ultracam filter set to SDSS	122
8.1.3	Theoretical extinction	125
8.2	The “peppering” phenomenon in Ultracam	128
8.3	The Ultracam archive software	134
8.3.1	The back-end software	135
8.3.2	The UAC SQL database	136
8.3.3	The “front-end” software	137
8.3.4	Improvements to the observer interface software	138
8.4	Summary	138
Chapter 9	Calibration and reduction of Ultraspec data	140
9.1	Observations and data reduction	140
9.2	Ultraspec calibration and data reduction	141
9.2.1	Calibration of CCD parameters	141
9.2.2	The low bias tail	147
9.2.3	Spurious charge	149
9.2.4	Electron traps	156
9.3	Results	168
9.3.1	Z Cha on the rise to superoutburst	168
9.3.2	The Rossiter-McLaughlin effect and QS Vir	171
9.4	Summary	172
Chapter 10	Conclusions	174

Declarations

The following contents of this thesis consist of work performed by myself over the course of 2007 to 2010 at the Astronomy and Astrophysics department at the University of Warwick. All data analysis was performed by myself, as well as data reduction with the notable assistance of S. Parsons regarding the X-Shooter data. Observations were performed by several others including the Ultracam/Ultaspec teams at the universities of Warwick and Sheffield, including T.R. Marsh, V. Dhillon, S. Littlefair, C. Copperwheat, and K. O'Brien, and the Paranal service mode observers for X-Shooter. WHT/ISIS observations were performed by T.R. Marsh and TWIN observations by D. Steeghs. The PI for the XMM-Newton data presented here is P.J. Wheatley. All the new software code described within the thesis was written by myself, with the original version of the 'lcurve' CV modelling code that was modified for use in the work presented here was made by T.R. Marsh. All figures shown have been produced for this thesis.

The contents of §5 is to be submitted to MNRAS with D. Steeghs and T.R. Marsh as co-authors. The contents of §6 will also be submitted to MNRAS with the coauthors P.J. Wheatley, T.R. Marsh, S. Littlefair, B. Gänsicke, and S.G. Parsons. My coauthors assisted in guiding the progression of the work and various technical aspects, as well as the presentation and style of the text herein.

Abstract

Observations of the cataclysmic variable systems EX Dra, Z Cha, and OY Car using multi-wavelength data ranging from near-infrared to X-ray bands are presented and analysed. Obscuration of the white dwarf by intervening material is discussed in context of all three systems, including the confirmation of iron lines in the spectrum of OY Car and the absorbing nature of the accretion disc upon the soft X-ray band in Z Cha. Evidence for spiral shocks and slingshot prominences from the secondary star are found in EX Dra, while the extent of Z Cha's boundary layer is probed with extensive modelling of simultaneous optical and X-ray data. Combined with a spectral analysis, this leads to the conclusion of a truncated inner disc present within the system. The first radial velocity measurement of OY Car's red secondary star is presented, and used to produce a fully spectroscopic determination of the system's mass ratio which is compared to previous attempts using the photometric method. The second part of this thesis deals with the improvement of the Ultracam and Ultraspec instruments. For the former, increased colour precision is achieved by careful calibration of the instrument's filter response with respect to the SDSS system; a charge modulation problem is investigated and quantified; and the development of a software system designed to improve observation and reduction workflow is detailed. For Ultraspec, a careful analysis of the CCD parameters is performed in order to prepare the instrument for scientific quality data publication, including an analysis of spurious charge found present. It is also shown that the AC coupling of the video circuit in the CCD can cause erroneous values in the readout but can be rectified, and presence of charge trailing caused by electron traps is investigated with preliminary corrections shown.

Abbreviations

- * ADC - analogue-to-digital converter
- * ADU - analogue-to-digital unit
- * ATG - array transfer gate
- * CCD - charge-coupled device
- * CTE - charge transfer efficiency
- * CDS - correlated double sampling
- * CPU - central processing unit
- * CV - cataclysmic variable
- * DMA - direct memory access
- * DN - digital number
- * DSP - digital signal processor
- * EMCCD - electron-mutliplying CCD
- * ENF - excess noise factor
- * FWHM - full-width at half-maximum
- * MOS - metal oxide semiconductor
- * MTF - modulation transfer function

- * MPP - multi-pinned phase
- * NIR - near infra-red
- * OSW - output summing well
- * OTG - output transfer gate
- * PCI - peripheral component interconnect
- * PMT - photo-multiplier tube
- * QE - quantum efficiency
- * SDSS - Sloan digital sky survey
- * SNR - signal-to-noise ratio
- * SQL - structured query language
- * UV - ultra-violet
- * VPH - volume phase holographic
- * XML - extended markup language

Chapter 1

Introduction

Cataclysmic variables are remarkable stars that display frequent and dramatic changes in their brightness by orders of magnitude, which has fueled an intense observation and study of such objects for what is now over 150 years. The origin for this observed variability was explained not by a process in a single star, but in a binary system [Walker, 1954], specifically the transfer of mass between the two objects. This mechanism presents many good reasons for studying such systems, as they have implications regarding stellar and Galactic evolution; accreting white dwarfs may be the source of Type Ia supernovae [King et al., 2003] and nova eruptions influence the interstellar medium and from that the galaxies which are constructed from such material [Amari et al., 1990]. Often, however, it is the process of accretion itself that is the matter of study. Accretion is the most efficient form of energy and is not unique to cataclysmic variables - black holes present at the centres of galaxies are another case, and jets of material are typically associated with accreting objects. The complete physical workings of such accreting discs that form around objects like the white dwarf remain a strong source of debate, however. The source of viscosity within accretion discs - while necessary for such a disc to operate - is under debate, and the disc instability model used frequently to describe the accretion and subsequent outburst within cataclysmic variables has been called into question several times due to inconsistencies with the observed evidence [Meyer and Meyer-Hofmeister, 1989; Pringle, 1988]. Due to their steadily increasing number of detections, cataclysmic variables provide a useful and readily available source to observe the phenomenon of accretion by professionals and amateurs.

As will be explained, a wide spectral range is required to successfully examine all the components within a cataclysmic variable system, but due to their variability and - in some cases - eclipses, efficient detectors that can run with a high cadence

are also of a huge benefit. The requirements of a wide spectral range and high time resolution are difficult to achieve in union; however several instruments have been developed that are well-suited to one regime or the other. Ultracam and Ultraspec are two instruments with very high time resolution; the former providing frame rates of 100Hz+ with triple-arm broadband photometry, and the latter with spectroscopy at up to 30Hz. The ESO X-Shooter instrument that has recently come online provides an unprecedented spectral range from 3000 Å to 25000 Å. All three of these instruments potentially offer great insights into the study of cataclysmic variables, and results from each will be shown during later chapters.

This thesis makes use of data from UV to X-ray wavelengths to study three cataclysmic variables (EX Dra, Z Cha, and OY Car), and also focus upon improvements to the well-established Ultracam instrument and the relatively new Ultraspec. Initially a more detailed background of the physics of cataclysmic variables will be given in §2, whereas the detectors that are used in all modern observations of these systems are detailed in §3, along with typical reduction methods for observational data. The hardware and software of Ultracam and Ultraspec are briefly described in §4.

In §5, optical and UV data of EX Dra are presented, examining the possibility of the white dwarf being obscured by its accretion disc along with signs that the disc may become elliptical in outburst. The thin regime above the surface of Z Cha's white dwarf, known as the boundary layer, is measured and discussed in §6 by using a combination of Ultracam optical data with X-ray observations using XMM-Newton; this chapter also investigates the remarkably variable ephemeris of the object. Following this, §7 details observations of OY Car with the new spectroscopic instrument "X-Shooter". Here the previous photometric results are tested against new independent spectroscopic measurements, and the presence of iron lines previously reported by Horne et al. [1994] is investigated.

The next two chapters focus upon instrumentation. In §8, this is Ultracam; where the accuracy of the filters with respect to the SDSS system are calculated, along with an investigation into a long-standing issue with the Ultracam CCDs that affects the accuracy of the data. The chapter also details the development of software modules that will automatically catalogue Ultracam data and provide improvements to the observing workflow. §9 details calibration of Ultraspec and the development of a software pipeline to account and correct for various technical issues that affect the accuracy of the data. Some preliminary science results of Ultraspec are presented. The results from each of the chapters will be summarised along with conclusions in §10.

Chapter 2

Cataclysmic variables: a brief overview

This chapter will discuss the nature of cataclysmic variable (CV) systems, including their evolution, geometry, important parameters and notable characteristics. While this will be a concise guide to the necessities of CV systems in context of research done, more detailed texts are readily available, such as Warner [1995a].

2.1 Stellar evolution

All stars begin their life in the form of collapsing, gravitationally-bound clouds of gas. These clouds produce young star clusters, and it is unsurprising that since they are formed in this manner, many of the produced stars are also gravitationally bound to one another. Binary systems, with two stars orbiting one another, are the most common result of such a scenario, though triple systems or pairs of binaries are also possibilities. A recent study shows that two-thirds of stars are single [Lada, 2006], perhaps ejected from an unstable gravitational configuration.

Most stars are in a hydrostatic equilibrium state; the constant force to collapse due to gravity offset by the thermonuclear reactions present in its core. This can only function as long as the core reactions have a fuel source (hydrogen), which eventually runs out. At this point, if the star has a mass greater than $0.2M_{\odot}$ the core begins to contract and any remnant hydrogen in the envelope of the star ignites resulting in a small, hot core and an expanded, cooler envelope. While many stars may go on to burn helium in their cores (resulting in a carbon/oxygen core), those that are under half a solar mass do not [Benvenuto and De Vito, 2005]. The remnant core continues to collapse until it becomes supported by degeneracy pres-

sure. At this point the Pauli exclusion principle comes into effect; the star can no longer contract because no two electrons can occupy the same state as dictated by Fermi-Dirac statistics. This is a white dwarf – a compact, degenerate object that as the cooling remains of a star that had insufficient mass to fuse any further material. It is important to note that stars more massive than ~ 8 solar masses will burn material in their core potentially up to iron and will not become white dwarfs [Heger et al., 2003], instead becoming neutron stars and eventually, perhaps, black holes. The lifetime of a star is approximately inversely proportional to the square of its mass (qualitatively, greater mass leads to increased core burning), and hence, it is expected that most observed white dwarfs will have carbon-oxygen cores since their helium counterparts will not yet have run out of fuel. Some are still observed however [Liebert et al., 2004], implying some external influence upon the star to allow it to proceed far enough along its evolution, which is attributed to binarity as discussed in §2.2. The ultimate state of a white dwarf will be a non-radiating object in thermal equilibrium with its local environment (the cosmic background radiation) but this will never in reality be achieved due to the long timescales involved.

The spectra of white dwarfs can be approximated by blackbody spectra, but in detail show hydrogen or helium features, despite of their expected carbon-oxygen cores [Sion, 1999]. The explanation is that the extremely high surface gravity separates elements, so heavier ones sink to the bottom whereas the lighter ones rise to the top, where they are observable; this is the white dwarf ‘envelope’ containing residue material from its asymptotic giant branch phase and potentially accreted material from the interstellar medium [Schatzman, 1945]. This envelope is a tiny fraction of the white dwarf’s mass but is extremely important, as it is highly opaque and hence slows the cooling rate of the core.

Low mass stars on the order of $0.1 - 0.3M_{\odot}$ are thought to be fully convective; they will exhaust all of their hydrogen fuel but never manage to become red giants. These have extremely long stellar lifetimes and are referred to as red dwarfs.

2.2 The common envelope phase and Roche potential

Cataclysmic variables begin their evolution as binaries with the heavier of the two stars evolving more rapidly, before expanding to become a red giant. Single stars or binaries with significant separation will always revert to a lowest-energy state, which in terms of gravitational potential is a spherical shape. When brought close enough together, this no longer holds true, and the less massive, less dense star will distort to match the equipotential surface which in a gravitational and centrifugal

potential is given by

$$\phi_R(\mathbf{r}) = -\frac{GM_1}{|\mathbf{r} - \mathbf{r}_1|} - \frac{GM_2}{|\mathbf{r} - \mathbf{r}_2|} - \frac{1}{2}(\boldsymbol{\omega} \times \mathbf{r})^2. \quad (2.1)$$

This equation is true under the assumptions of mass and angular momentum conservation, as well as a circular orbit, and is in a reference frame of the rotating binary. M_1 and M_2 are the masses of the stars, \mathbf{r}_1 and \mathbf{r}_2 are position vectors of the centres of the stars, and $\boldsymbol{\omega}$ is the orbital angular velocity which from Newtonian gravitation and Kepler's third law is

$$\boldsymbol{\omega} = \sqrt{\frac{G(M_1 + M_2)}{a^3}} \hat{\mathbf{z}}. \quad (2.2)$$

Here, a is the orbital separation and $\hat{\mathbf{z}}$ is a unit vector that is perpendicular to the orbital plane.

ϕ_R is known as the Roche potential, and a very important result is that the lines of Roche equipotential depend only upon the mass ratio, $q = \frac{M_2}{M_1}$, with the scale being a function of a alone [Warner, 1995a]. A double-lobe with an inner Lagrangian point L_1 is of particular interest; these are known as the Roche lobes of the two stars. If the surfaces of both stars are well within their Roche lobes, they will remain spherical, and the system is categorised as a detached binary. However, as a star grows larger in radius, it will begin to distort to match the equipotentials of its Roche lobe. If this trend continues it will eventually fill its lobe and then overflow; material passing through the L_1 point will then be caught in the potential of the other star, and mass transfer between the two will begin.

This is what happens in the binary case described above as the more massive star goes through its red giant phase. The envelope expands and overflows the Roche lobe, and mass transfer onto the less massive star begins. At this point in time, the star that will eventually be considered to be the primary in a CV is in fact the donor, or secondary, and the mass ratio $q > 1$. This is an unstable scenario, however; the material transferred from the heavier star means this mass is moving away from the centre of the binary, and since the material gains angular momentum, the stellar separation must decrease in order to conserve it. The consequence of this decrease is also a decrease in the Roche lobe size (as it scales with a). The L_1 point moves inwards towards the heavier star, and now even more mass is transferred. The runaway scenario continues until all material is effectively dumped onto the other star. The timescale for such a transfer may be on the order of years or less (limited by the rate at which material transfers), and the result is a cloud of material surrounding both stars which is known as the 'common envelope

phase'. This envelope of material causes a drag force, and the envelope is heated with energy being extracted from the orbit; meaning their separation will decrease along with the orbital period. At some point the thermal energy deposited into the surrounding material will exceed the binding energy of the common envelope and it will be expelled into what we term a planetary nebula on the order of $10^3 - 10^4$ yr. What remains in this standard treatment of binary evolution is now an unevolved secondary star and the exposed core of the former red giant, which is a white dwarf and becomes the primary in the newly formed system. The role of the two stars is reversed, and the resulting object after this phase is a binary that is either detached (if the Roche lobes do not permit mass transfer) or otherwise a cataclysmic variable binary.

2.3 The primary and the secondary

In terms of a CV, the primary refers to the compact, central object (the white dwarf), whereas the secondary refers to the distorted red star that material is transferring from. In semi-detached systems such as CVs, where the secondary is overflowing its Roche lobe, the star rotates at the same rate that it orbits. Material closest to the primary will exhibit a bulge on the surface of the star due to a gravitational potential, and should the orbital period not match the rotation of the star, a torque will be exerted effectively dragging this bulge towards the primary. This torque energy is supplied from the orbital rotation leading to the minimum energy scenario of the rotation matching the orbital period, which is known as tidal locking. This process also circularises the orbit [Tassoul, 1988].

The primary is expected to dominate at UV wavelengths because of its high effective temperature [Panek and Holm, 1984], and in systems with a high inclination angle with respect to our line of sight, can be eclipsed by the larger secondary star. In the optical band, the primary may be detected in Balmer lines or the Lyman- α series in near-UV, though in some systems the flux from the white dwarf may be negligible in comparison to the accretion disc (§2.4). The star is expected to have temperatures from 10,000K to 50,000K [Townesley and Bildsten, 2002]. While the cooling rate for an isolated white dwarf will depend upon its stored energy (and envelope transparency), in a CV, the primary is compressionally heated as material is accreted from the mass transfer.

The secondary (typically of M spectral type) also provides a unique feature to the observer because of its non-spherical geometry. When seen side-on to the observer, the star appears larger because of its distorted nature and hence more

flux is observed than when it is end on. This variation is known as ellipsoidal modulation and can help provide limits on parameters of a CV system, since the Roche lobe shape is dependent only on q , and the amplitude of ellipsoidal modulation is dependent upon the inclination angle of the system.

Because of its lower ($\sim 3,000\text{K}$) temperature, the secondary is expected to only be detected at red and infrared wavelengths. While the star may have a low flux density even in this band, its extremely large surface area compared to the white dwarf means that it is frequently visible at sufficiently long wavelengths. Calcium is a strong indicator of the secondary star, and since molecular features will be present in the secondary as temperatures are low enough that atoms can form and sustain bonds, Titanium Oxide (TiO) may be found as well. Due to the many modes of rotation and oscillation of a molecule, such features are more complex than the atomic lines one may observe from the white dwarf.

2.4 The disc and the bright spot

When mass transfer onto the primary begins, the material crosses the L_1 point at the speed of sound in gas (generally a few km/s, Hellier [2001]) but the orbital velocity can be 50 times this speed. For this reason, material cannot directly fall onto the primary, but instead swings around it (conserving angular momentum). Inside the rotating binary frame, this can be looked upon as the Coriolis force exerting itself upon the stream. The stream enters an orbit of lowest energy (a circle) that maintains its angular momentum; the result being a ring surrounding the primary star. This is known as the circularisation radius (see Frank et al. [1992]).

Considering Kepler's third law once more, material at a radius slightly closer to the white dwarf will be moving at a higher speed than those further out (differential rotation). These annuli of material will cause friction as they move past one another, and energy is dissipated in the form of turbulent heating. Because energy is lost, some material must move inwards to a small radius, but to also conserve angular momentum some material must move outwards, the result being a thin disc that is known as an accretion disc. An exhaustive review of accretion discs is given in Pringle [1991].

The inner boundary of the disc is set by the white dwarf, where material can eventually settle down onto the surface of the primary. The outer disc radius is set by tidal interactions of the outer annuli with the secondary star. Beyond the limiting radius, these interactions cause dissipation, draining the material's angular momentum and effectively halting further growth of the disc. Accretion discs vary in




Figure 2.1: A 3D representation of a cataclysmic variable system showing the white dwarf, secondary star, disc, bright spot (represented by the star) and the stream (solid line). Left: phase 0.7, showing the distortion of the secondary star due to the filling of its Roche lobe. Right: phase 0.96 just before the white dwarf is eclipsed.

temperature depending upon radius, and can be roughly approximated by annuli of blackbodies with temperatures ranging from 5,000K at the outer radius to 30,000K near the white dwarf. Balmer and helium emission lines are frequently seen in the spectra of accretion discs, but sometimes show absorption. The physical origin for these lines is the absorption and re-emission of photons at discrete energies of excited atomic states. Should the observer be seeing the continuum through a cloud of absorbing gas, the re-emitted photons are unlikely to be in the same direction as the absorbed ones which will result in a continuum spectra with dark, absorption lines. Alternatively, observing the scattered photons will result in emission lines. A disc surrounded by a thin corona hence may provide both absorption and emission lines.

Where the incoming stream of material from the secondary impacts the accretion disc, a supersonic, turbulent impact occurs, releasing large amounts of kinetic energy to heat the target site known as the bright spot. At high inclinations, the bright spot is an important contributor to the flux observed, and the luminosity upper limit is

$$L_{bs} \leq \frac{GM_1(\dot{M}_2)}{r_d}, \quad (2.3)$$

where \dot{M}_2 is the mass loss rate from the secondary and r_d is the radius of the disc. In most cases it is assumed the bright spot is a point on the rim of the disc, but it may in fact be extended as some material from the stream continues on its previous trajectory. Some simulations show the impact may create a hole in the accretion disc, or that some material may even swing around and create a second impact site [Lubow, 1989]. Later in §6 and §7, possible evidence of such a scenario are shown with the cataclysmic variables Z Cha and OY Car, respectively.

A diagram of a CV showing the white dwarf, secondary, disc, and bright spot can be seen in figure 2.4.

2.5 The boundary layer

Just above the surface of the white dwarf, the Keplerian velocity is approximately 3,000km/s. However, the white dwarf rotates at considerably lower velocities, and so the accreting material must be slowed down to accrete onto the primary. This regime is known as the boundary layer, which is a very thin layer ($\ll R_{WD}$) above the surface of the white dwarf. The Keplerian speed and gravitational potential of some mass m at distance r from the primary are given by

$$v = \sqrt{\frac{GM_1}{r}}, \quad (2.4)$$

$$U = -\frac{GM_1 m}{r}. \quad (2.5)$$

Since the kinetic energy is $\frac{1}{2}mv^2$, it can be seen that in this case it is $\frac{1}{2}$ of the potential energy, meaning that up to half of this gravitational potential energy can be released and radiated away with the other half converted into kinetic energy of the infalling matter.

In situations where the accretion rate is low, the material is optically thin. Cooling comes via collisions between particles, but due to the low density, this is generally insufficient to provide any notable cooling, and so the hot gas expands lowering the density further, resulting in a large, hot, and diffuse corona of material surrounding the white dwarf [Hellier, 2001]. In this case, the boundary layer will dominate the X-ray flux in the system. At higher accretion rates, the boundary layer will become optically thick, increasing the cooling rate at which point the X-ray flux will diminish and will instead be detected primarily in the UV.

The boundary layer may also be responsible for some of the wind outflow. Ions close to the boundary layer will be pulled closer by the gravitational pull of the primary, but also experience radiation pressure from the strong luminosity being emitted from the hot boundary layer which pushes in the opposite direction. At sufficient radiation pressure, the atom is driven out and hence becomes part of the wind outflowing from the binary. However, older models that used non-rotating radial outflow winds have been superseded by biconical rotating ones which in general provide much better agreement with line profiles and reasonable mass loss rates [Shlosman and Vitello, 1993]. Hence while the boundary layer may be of some influence to the wind, it is now quite often attributed to an inner disc wind [Mason et al., 1995].

The boundary layer temperatures can reach sufficient temperatures that its

flux will appear in the X-ray band. The spectral lines formed by outflowing winds are examined in §5 with UV data from the Hubble Space Telescope, and the extent of a CV's boundary layer will be probed in §6 with the use of X-ray observations.

2.6 Mass transfer and angular momentum

Assuming that angular momentum is conserved, when material is transferred from the lighter secondary star to the more massive compact primary, the secondary star will move further away (the orbital separation will increase). This can be shown by starting with J , the angular momentum

$$J = M_1 a_1^2 \omega + M_2 a_2^2 \omega, \quad (2.6)$$

where ω is the angular velocity and M_1 , M_2 , a_1 and a_2 are the masses and separations of the primary and secondary stars respectively. Given that $a = a_1 + a_2$ and that $M_1 a_1 = M_2 a_2$ we can write this as

$$J = \frac{M_1 M_2}{M_1 + M_2} a^2 \omega. \quad (2.7)$$

If differentiated with respect to time assuming total mass is conserved ($\dot{M} = 0$, where $M = M_1 + M_2$) and angular momentum is conserved ($\dot{J} = 0$) the result is

$$\frac{\dot{a}}{a} = 2 \left(\frac{M_2}{M_1} - 1 \right) \frac{\dot{M}_2}{M_2}. \quad (2.8)$$

Hence under the condition that $M_2 < M_1$, transfer from the secondary results in an increase in orbital separation (further details are found in King [1988]). There must be a source of angular momentum loss for the binary to stay in contact and thus match observations.

According to general relativity, the repetitive orbiting of two massive objects causes a warping of space we term gravitational radiation. This emission extracts energy and angular momentum from the orbit and causes a slow inwards spiral. However, for all but the shortest period binary systems, it is expected that this will be a weak effect.

The majority of the angular momentum loss is thought to come from magnetic braking [Mestel, 1968]. Ionised material will be captured and “frozen” on to magnetic field lines from the star, and begin to rotate along with it. If this material is being propelled away such as in a wind outflow, it will effectively be removing angular momentum from the system, hence, “braking” the rotation of the star. This

in turn causes the orbital separation to decrease.

2.7 Classification of CVs

While CVs are named for their rapid variability, there are in fact many variations upon the theme. Older classifications tend to have specific, agreed upon titles whereas some of the newer types are simply named after the first CV of its type. Outbursts are the reason for cataclysmic variables to have been titled so; an outburst is a semi-regular event that varies from one system to the next whereby the luminosity of the system rises by several magnitudes over the course of hours or days and then settles back to its previous level over the period of a week or two.

Novae show extremely large outbursts from 6 to 19 magnitudes which are the largest amplitude variations seen in CVs and the origin of the ‘cataclysmic’ name. With novae, the eruption occurs when a thermonuclear runaway is initiated upon the surface of the white dwarf primary due to sufficient material accumulating through accretion from the inner disc boundary layer; this increases the temperature and pressure until hydrogen fusion can begin [Shara, 1989]. This reaction will continue while degeneracy conditions still hold. When lifted, a nova shell of material expands and is then released. The recurrence time for such novae events are predicted to be between 10^4 yr and 10^6 yr, depending upon the accretion rate in the system [Fujimoto, 1982]. Classical novae are defined as systems which have had a single eruption of this kind, whereas recurrent novae are those which have been known to have had multiple eruptions.

Dwarf novae (DNe) are the class of CVs primarily dealt with in the research detailed here. Their outbursts are typically on the scale of 2-5 magnitudes but can be larger. These events are also semi-regular (weeks to months) and some also show another semi-regular but less common super-outburst of an even higher amplitude in luminosity. These systems are referred to as SU UMa stars, and also have photometric modulations called superhumps. Dwarf novae that show frequent bursts of activity followed by long periods in a “standstill” state (which are fainter than in outburst, but not quiescent) are termed Z Cam systems, whereas the more common dwarf novae that do not show either of these characteristics are referred to as U Gem stars.

Approximately 20% of CVs show a significant magnetic field [Ritter and Kolb, 2003] which can disrupt the accretion disc or prevent its formation at all. These are magnetic CVs and are further classified into polars or intermediate polars, depending on whether they show strong circular and linear polarisation which

are modulated by the orbital period or whether they show stable pulsations at frequencies less than the orbital period. Polars are expected to have stronger white dwarf magnetic field strengths and lower accretion rates than intermediate polars. Neither of these are dealt with in this thesis, but comprehensive reviews of these systems can be found in Cropper [1990] and Patterson [1994] respectively.

2.8 Dwarf novae outbursts

As dwarf novae are the main subject of this thesis, the physics of their outbursting behaviour will now be reviewed. Dwarf novae have known to have been observed and followed since 1855 with the CV system U Gem [Pogson, 1857].

Two competing models attempted to explain dwarf nova outbursts: the disc instability model by Osaki [1974] and the mass transfer burst model by Bath et al. [1974]. The former is based upon the idea that if the mass transfer rate is constant, but at a higher rate than can be transported through the disc then material will build up and cause the disc to jump to a hot, high viscosity state (often termed as the disc becoming “unstable”). The highly viscous disc now spreads this new material out and the higher accretion rate onto the white dwarf results in the observed luminosity increase. Eventually, the disc returns to its low-viscosity state and the build up begins again. The Bath et al. [1974] model suggests that the secondary occasionally floods the disc with bursts of material, raising its temperature and luminosity.

Observations have made the disc instability model the accepted method for outbursts in dwarf novae [Hellier, 2001] for several reasons. The bright spot has been observed to be stable during a rise to an outburst state (in direct opposition to the expected scenario for a mass transfer burst) and also the disc size expands during outburst and is then drained afterward as Osaki’s model predicted. This model was given a physical grounding later, described by Hōshi [1979]. The main parameter that drives the model is the ionisation of hydrogen (which occurs at a temperature of $\sim 10^4\text{K}$). The opacity of hydrogen is a very sensitive function of temperature ($\kappa \sim T^{10}$) and was shown to reveal two very clear states in which the hydrogen was either fully ionised or neutral. Between these two states is an unstable situation in which the system would tend to either one of the quiescent or outburst states depending upon a critical surface density. In a partially ionised disc, an increase in temperature results in an increase in opacity, which further enables the disc to trap heat and cause the opacity to rise further until it is entirely ionised. This process is represented by the well-known “S-curve” in figure 2.2, where the solid line indicates equilibrium states where the heating matches the cooling. Starting from position




Figure 2.2: A plot of the equilibrium states in an accretion disc showing how the multi-valued function arising from the opacity of hydrogen can lead to an unstable state, and the system jumping from quiescence to outburst, and back again (see text for details). Any mass transferring binary with an accretion rate between the two critical points is expected to have periodic outbursts.

A on the diagram, an increase in surface density causes heating within the disc, and hence the disc changes to the $\Sigma - T_{eff}$ represented by B. Beyond this point, the system heating exceeds the cooling leading to a thermally unstable scenario resulting in the disc changing to state C, whereby the disc is now in outburst. The disc drains and cools, changing to state D, whereby it drops back to the quiescent state at A, and the process begins again. Hence, for a mass-transferring binary system to have disc outbursts, it must have an \dot{M}_2 between \dot{M}_{crit1} and \dot{M}_{crit2} as shown on the diagram. Z Cam behaviour is qualified by such systems having an \dot{M}_2 very near the \dot{M}_{crit2} point, thus sustaining an outburst state.

2.9 Viscosity

In an accretion disc, the adjacent annuli attempt to force each other to corotate meaning that outer radii are given angular momentum from the inner. Thus while a majority of the material flows inwards, the minority flowing outwards carries the majority of the angular momentum [Lynden-Bell and Pringle, 1974]. It is the viscosity in the disc that opposes the motion between annuli and hence is responsible for angular momentum transfer within any accretion disc; without this mechanism

it cannot operate.

Molecular viscosity familiar from kinetic theory cannot account for viscosity within an accretion disc. These discs are diffuse, and thus molecular viscosity is extremely weak and cannot transport significant angular momentum. While the process giving rise to the viscosity may be unknown, it has become standard practice to quantify accretion disc viscosity in terms of a dimensionless α parameter as in Shakura and Sunyaev [1973]. They make the assumptions that turbulent eddies must be smaller than the scale height of the disc (H) and that the eddies transfer material at up to the sound speed in gas, c_s . The former assumption comes from the simple fact that if the eddies were larger than H , the disc scale height would increase, and violation of the latter would mean that the turbulent eddies are supersonic and would cause shocks that heat the disc. Hence the viscosity v is described by

$$v = \alpha c_s H, \quad (2.9)$$

where $\alpha \leq 1$. Discs using this model are thin ($H \ll r_d$), negligible in mass compared to the white dwarf, and slightly concave, flaring at the outer edges. While this has been used to successfully quantify viscosity in a disc, it gives no prescription as to the source of the turbulence.

One favoured answer to the turbulence problem is magnetic turbulence [Balbus and Hawley, 1998]. Charged particles readily flow along magnetic field lines but have difficulty crossing them. These field lines are then dragged along by ionised matter that is moving, and so in terms of magnetohydrodynamics, the field lines and the material are frozen together. There will be field lines between adjacent annuli of ionised gas, and due to the aforementioned differential rotation, these field lines will become stretched and generate a force opposing this motion and exchanging angular momentum. Since the inner annuli will lose angular momentum it will tend to fall inward and stretch the field line further; the small initial field has now grown considerably. At some point, orderly flow breaks up in the accretion disc due to the magnetic turbulence.

2.10 Elliptical discs

SU UMa stars show hump-shaped modulations of the light curve over a period slightly longer than the orbital period (termed superhumps) and undergo less frequent superoutbursts which are longer and more luminous than a standard outburst. Elliptical discs are the solution to this problem as shown by Vogt [1982]. Here, the disc precesses on a timescale that is considerably longer than the orbital period. The

interaction of these two periods creates a beat frequency which is the superhump period (P_{sh}) as defined by

$$\frac{1}{P_{sh}} = \frac{1}{P_{orb}} - \frac{1}{P_{prec}}, \quad (2.10)$$

where P_{orb} is the orbital period and P_{prec} is the precessional period. The cause of the ellipticity is due to the disc interaction with the secondary star. The outer regions of the disc bulge slightly as they experience the gravitational force of the secondary star, and usually this is simply the mechanism for setting the outer region of the disc as this process drains angular momentum. If the Keplerian orbit of the outer disc has a period that resonates with the orbital period, however, resonance occurs and this can exaggerate the radial component affecting the disc, pulling it into an elliptical state [Whitehurst, 1988]. While a 1:2 resonance would require the most extreme mass ratios of $q < 0.025$, the 1:3 resonance can occur for $q < 0.3$. It should be noted that an important limiting factor on whether a disc can become elliptical depends upon the white dwarf's Roche lobe. If the orbits of the outer disc material would extend outside of the tidal limit then the disc will not precess. This tidal limit, as per Warner [1995a], is

$$\frac{r_{tidal}}{a} = \frac{0.6}{1+q}. \quad (2.11)$$

Another potential side effect of the distortion of the edge of the disc by the secondary is shocks forming within the gas as it attempts to follow the non-circular orbits at speeds greater than the sound speed [Spruit, 1987]. These shocks slow down the gas, and the faster, inner orbits within the disc turn the band of shocked material into a spiral arm. Theoretical treatment showed no solutions for a single shock arm (two are always expected) and that the shock decreases if the disc is thinner, thus indicating that an outburst state is more likely to produce spiral shocks. Evidence for elliptical discs will be shown and discussed in §5.

2.11 Observing CVs

Each component within a cataclysmic variable system also has attributed to it a different temperature; the secondary star is red and cool at $\sim 3000\text{K}$, with the white dwarf temperature ranging from 10,000 - 30,000K. In practice, this means that a wide range of the spectrum is needed to detect each of the components clearly. For DNe, whereas optical observations allow us to easily see the accretion disc and sometimes the white dwarf, the red star is typically obscured by the flux from these brighter components at typical optical wavelengths. The observer has a

greater chance in the infra-red bandpass to find a clear detection of the secondary ($\sim 10,000\text{\AA}$). Shorter wavelengths in the ultra-violet and X-ray bands allow a detection of the boundary layer, though the presence of an accretion disc surrounding the central object means that often there is an obscured view of the white dwarf and boundary layer itself. The amount by which this affects observations depends upon the viewing angle (inclination) and the density of the column of material in the disc.

Arguably the most interesting cataclysmic variables (at least for the observer) are those at a high inclination angle with respect to our line of sight. These systems show eclipses that can last as short as minutes (with the ingress and egress of the white dwarfs themselves lasting seconds), and have been used extensively to constrain parameters of these systems [Smak, 1979; Bailey, 1979]. Precise measurements of the eclipse times also provide another avenue of research with cataclysmic variables, which is the evolution of their period. Increasingly, it is observed that more and more binary systems show a quadratic or sinusoidal nature in their orbital ephemeris [Pringle, 1975; van Amerongen et al., 1990; Baptista et al., 2002; Parsons et al., 2010b]. A third body has often been utilised in explaining this phenomenon, though other mechanisms have been proposed [Applegate, 1992] that would avoid the necessity of having a large quantity of these systems containing three objects.

2.12 Summary

Key physics involved in the description of cataclysmic variables have been outlined including their formation, classification, and the proposed models for their “cataclysmic” outbursts. In the context of this thesis, observations of three cataclysmic variables - specifically dwarf novae - will be presented: EX Dra, Z Cha, and OY Car. A detached system, QS Vir, will also be briefly examined.

In §2.11 above, the observational requirements of observing CVs were detailed: a wide spectral range varying from infra-red to observe the donor star to UV and X-rays to observe the boundary layer, or a high time resolution to determine system parameters from a photometric method or study the variation of the CV period. The next chapter will now proceed to detail the devices and methods we use to measure these variable systems; the detectors themselves, and the typical data reduction techniques.

Chapter 3

Detectors for modern astronomy

Astronomy is the study of astrophysical objects which are always at a great distance, and so our inability to go to our source material means that our targets are typically very small and faint. This has meant that astronomy has continuously been pushing the limits of detector technology in sensitivity. In addition to this, with the increasing study of rapid variability objects such as CVs in the past 30-40 years, speed has become a concern as well. Almost every observatory worldwide now uses charge coupled devices (CCDs) as their detectors whether it be for a photometer or a spectrograph camera, and in the past ten years or so, amateurs as well. While there will remain niche roles for particular devices, CCDs are now almost ubiquitous, and the reason for their rise to prominence must be explained, as well as the technology behind them. A thorough understanding of the important parameters and features of the devices is required so that any data taken from such a device can be processed correctly to give accurate scientific results.

Important measures of a detector technology consist of quantum efficiency (QE), linearity, spatial resolution and more practical concerns such as cost. The first measure is a direct indication of how much infalling light is actually recorded by the detector, where a QE of 100% indicates every single photon will be recorded. Linearity (sometimes known as reciprocity) is an indication of how well behaved the detector responds to a varying level of light; a perfectly linear detector indicates that twice the level of flux will generate twice the number of counts. Spatial resolution is a measure of how much detail a detector will record. In modern devices such as CCDs where each pixel is in effect its own detector, the spatial resolution is dictated by the 'pixel pitch' or the size of each MOS capacitor. Photographic plates, or any

	Plate	PMT	Digicon	CCD (80s)	CCD (90s)
QE	2%	25%	5%-20%	20% - 40%	> 80%
Non-linearity	5%	2%	10%	5%	1%
Spatial	High	None/Poor	Poor	Average	High
Cost	Cheap	Average	High	Very high	Average - High
Noise	Average	Low	High	High	Low
Space	No	No	Yes	Yes	Yes

Table 3.1: Summary of important properties of detectors used in astronomy, including photographic plates, photo-multiplier tubes (PMTs), Digicons, and charge coupled devices (CCDs). Listed parameters are the quantum efficiency (QE), the non-linearity or reciprocity, the spatial resolution, effective cost, noise, and whether it is suitable for deployment on spaceborne devices.

camera system with optical elements, is usually measured by a modulation transfer function (MTF), which specifies contrast for a range of spatial frequencies.

A summary of the important differences between detector technologies can be seen in table 3.1.

3.1 Detector technologies

3.1.1 Photographic plates

Originally put into use for astronomy in the late 19th century, photographic plates consisted of a light-sensitive emulsion of silver salts that were applied to a glass plate. Discovery of minor planets using these plates was pioneered by Max Wolf in 1891 [MacPherson, 1932]. While consumer use rapidly evolved towards the use of film, plates remained in use in the astronomical field due to their rigidity; they did not bend or distort like film, particularly in the extremely large formats used in such applications. While some photographic plate manufacturers claimed quantum efficiencies of up to 10%, the more typically used Kodak plates claimed 3%. Studies into the QE of such detectors revealed more conservative values between 1% and 2% [Argyle, 1955]. Many photons were merely reflected; others did not encounter halide crystals and thus did not trigger a response. Another problem with photographic plates is the characteristic ‘S’ shape reciprocity curve that comes from the chemical reactions involved. At low flux levels, the few photons that impact only cause a partial change in the halide grains which will revert unless another photon hits in a short time. To counter this, plates were often exposed to a low-level uniform light source which began triggering the chemical response in the emulsion to avoid this problem. This light offset is termed ‘fat-zero’ (also known as the D-min point on the reciprocity curve) and introduced a background level of flux as well as shot noise (see

§3.2.5). At high count levels the crystals can react no more, so the response trails off again. Photographic plates are heavily limited as far as time resolution is concerned, mostly due to the manual preparation and interchange of the plates, but also by the minimum exposure time mentioned previously. Finally it should be noted that such plates are awkward to use during observing because of the preparation required and their cumbersome size.

Despite these difficulties, photographic plates were used in many sky surveys, such as USNO [Monet et al., 2003] and POSS [Reid et al., 1991]. No other detector provides the low-cost and large field that photographic plates inherently have. Most emulsions were primarily photosensitive in the UV band.

3.1.2 Photomultipliers

Photomultipliers were primarily driven by the need for a practical television camera technology, and are based around the photoelectric effect. A vacuum sealed glass container holds a photocathode which the photons strike, producing an electron which is directed towards the dynodes. These produce further electrons as a result of secondary emission; multiple stages can be employed to increase the gain further. Electrical connectors at the far end then measure the signal.

Initial PMTs were single stage and provided a QE of a mere 0.4% at 800nm; this increased to 12% within a few years. Modern PMTs can have a peak QE of around 90% at specific wavelengths, which also allow good time resolution of a few seconds or less. Very large voltages (100-200V) are required for operation, and the devices must be shielded from ambient light when powered or the detector will be destroyed from the extremely high amplification. Multi-anode arrays in PMTs can provide some spatial resolution, and the linearity of these devices is better than 2% [Lush, 1965].

3.1.3 Digicons

Digicons were early solid-state detectors driven by the need for a technology that was feasible for a space-based telescope (primarily the Hubble Space Telescope, in the 1970s). The basic principle is that photoelectrons are accelerated and then impact on silicon diodes. These diodes consist of two regions (one of which is an n-type semiconductor, or donor, while the other is p-type semiconductor, or acceptor), which together form a pn-junction. These electrons and holes diffuse and recombine, leaving behind lattice ions near the interface area. This area is devoid of charge carriers, and is known as the depletion region. The ions set up an electric

field, which at some point becomes strong enough to prevent the diffusion of further charge carriers; this potential is known as the diode forward voltage drop (V_d). An interface voltage that exceeds V_d will then allow recombination to resume, and a large output voltage is recorded. Below this cut-off value, the diode provides no output. Unfortunately, for the 0.5-11 μ regime, the required V_d means that at 273K the detector would suffer an extreme amount of thermal noise, and hence needs to be cooled with liquid nitrogen. The spatial resolution is provided by an array of these diodes where the electrons are directed by electrostatics; hence, for example, with the FOS digicons on the HST, it was found that the geomagnetic field caused a positional shift with relation to the orbit. Digicons had a QE varying from 5% to 20% depending upon the wavelength of light and suffered non-linearity (up to 10%) at high count levels [Tull et al., 1975]. While the solid-state nature of the Digicon can theoretically allow very high time resolutions, one is typically limited by the QE, though FOS on HST can still operate on the order of seconds for bright targets).

3.2 Charge-coupled devices

CCDs were initially conceived as an electronic memory device in 1969 at Bell Laboratories, before it was realised that they had far more commercial capability as imaging devices, the first of which was released in 1974. Early devices revealed that towards red wavelengths, CCDs were 100 times more sensitive than photographic plates and 5 times more sensitive than other technologies at the time (such as vidicons). Early problems such as poor yields, residual image, and poor transfer efficiency were solved by the 1990s and, as a greater understanding of the technology became mainstream, yields increased as did the chip sizes, resulting in the massive proliferation of CCDs that we see today.

3.2.1 Principles of operation

A CCD is, in simplest terms, a collection of MOS (metal-oxide-semiconductor) capacitors with the necessary electronics to measure and output its array contents. CCD operation can be split up into four critical phases: charge generation, charge collection, charge transfer, and charge measurement. It is frequently visualised as a ‘bucket brigade’, where buckets sit upon conveyors filling up with water until they are moved along onto a horizontal conveyor to where they are measured; one bucket at a time, one row at a time (see figure 3.2.1). On a CCD, each bucket is a potential well generated inside a MOS capacitor, and the water is simply the charge which is




Figure 3.1: The ‘bucket brigade’ analogy of the CCD. Empty buckets fill up with rainwater (photons) before being moved down via parallel transfer one row at a time onto the serial register, where each ‘bucket’ (pixel) is measured at a time in the sense node.

a byproduct of the photoelectric effect. This is what is responsible for the ‘charge generation’, and it should be noted that since the band-gap of silicon is 1.14eV, the equivalent wavelength of photons at such an energy is $\lambda = \frac{hc}{E} = 10,868\text{\AA}$ [Janesick, 2001]. This implies that long wavelength (infra-red) photons will have a quantum efficiency of zero as they will pass straight through the chip. Infra-red CCDs use materials such as germanium instead, though this comes with its own set of difficulties [Schroder, 1974].

A p-type MOS capacitor consists of a boron-doped silicon substrate with a thin $\sim 100\text{nm}$ oxide layer grown above, then a conductive polysilicon gate deposited atop, where the gate capacitance is $C_{OX} = \frac{\epsilon_0 \epsilon_X}{d}$ as with a parallel plate capacitor (where ϵ is the permittivity and d is the distance between surfaces). When a bias voltage is applied to the gate, mobile carriers are driven away from the substrate/oxide interface leaving behind only the dopant ions, the MOS instance of a depletion region. Photoelectrons (known as signal electrons) are now collected at the interface, while the bias voltage is held. The well capacity of such a MOS capacitor is defined as the amount of charge required to bring the surface potential to 0V, and can be estimated by

$$Q = C_{OX} \Delta V_S, \quad (3.1)$$

where ΔV_S is the change in surface potential [Janesick, 2001]. Above the full well, more charge will “spill over” causing a phenomenon known as blooming, whereby charge collects in adjacent parallel potential wells (channel stops prevent it spilling horizontally).

By altering the bias voltages on each of the capacitors, the charge can now be shuffled down the chip, most commonly in a three-phase setup (that is, three different configurations of potential barriers during clocking that result in the movement of charge). During integration (when charge is being collected) potential barriers between subsequent pixels prevents charge from spreading. After this is done, each of the pixels is moved downwards eventually onto the array transfer gate (ATG). This downward movement is referred to as parallel, or vertical clocking. By altering the potential of the next gate to match, the potential barrier is removed and the charge balances itself between the two wells. Next, the original well is collapsed and the charge is now moved entirely into the next well; simultaneously with this all the charge going through the ATG is now moved along a serial (horizontal) register where each charge packet is sent to the output summing well (OSW). The voltage steppings between phases must be carefully calculated to to have sufficient swing to move the signal electrons while remaining fast enough for reasonable operation. Many newer instruments do not have a shutter mechanism, thus while the charge is being transferred out of the chip after integration, light is still falling on the CCD (leading to a signal trail as the CCD is read out, known as “smearing”). The OSW is where any pixel binning occurs; that is, the combination of several pixels into one, and the OSW itself is several times larger than the ordinary CCD pixels to accomodate the large amount of charge it may have to handle.

After charge has passed through the OSW, it is passed through the output transfer gate and reaches the sense node, which is where charge measurement occurs. The popular mechanism for this in sensor devices is correlated double sampling (CDS), which usually involves a preamplifier, postamplifier, clamp, sample and hold, and ADC circuits [White et al., 1974]. CDS works on two differential signals, the clamp circuit sets the ‘reset’ value (the zero level) before the video signal is dumped onto the sense node and sampled for a time t_s . Together with the time constant, τ_D , a CDS circuit acts as a band-pass filter where low-frequency noise is rejected by a low t_s and τ_D sets the bandwidth to reject high-frequency noise. It should be noted that a longer t_s will approach the absolute (ideal) signal value exponentially (a common setting is $t_s = 2\tau_D$, where the signal is 86% of its final value).

3.2.2 Buried channel potential well

Difficulties arose with the p-MOS CCD arrangement and the process of charge transfer. The efficiency of this mechanism is given as a percentage known as CTE (Charge Transfer Efficiency) and a high percentage is mandatory for accurate scientific work. A value less than 100% indicates charge is being left behind or otherwise lost, so the final value is either an under-representation or distortion of the actual number that would otherwise be obtained. The arrangement with p-MOS CCDs is known as a “surface channel potential well”, because signal electrons are collected near the surface of the capacitor. Unfortunately, many trapping states exist in the Si-SiO₂ interface, meaning that signal electrons here frequently remain stuck during the charge clocking [Kim, 1979]. These states can have very long lifetimes, and in early CCDs, saturated pixels would fill many of these states and slowly be released over the course of hours or even days, which was the problem of residual image. Better quality silicon will result in better CTE because of fewer impurities which produce potential trapping sites. Early CCDs demonstrated a CTE of $\sim 95\%$ and, through the best processing techniques, up to $\sim 99\%$. These numbers may seem satisfactory, but for the first Fairchild CCD arrays of 100x100 pixels, charge in corner pixels would lose $1 - 0.99^{200} = 87\%$ of its initial value before it was sampled.

The answer was to move the collection channel away from the surface interface to avoid the trapping states [Walden et al., 1970]. This is done by thermally driving n-dopant (typically phosphorous) into the epitaxial layer of silicon, moving the point of maximum potential down into the MOS by $\sim 0.5\mu\text{m}$. See figure 3.2.2 for a comparison of the surface and buried channel potential wells. Note that now there are two depletion regions; the surface interface region and the region formed at the n-p interface as with the silicon diode. Raising the bias voltage sufficiently will increase the size of the depletion zones until they merge at which point the depletion region extends all the way down into the p-layer. This technique allows modern CCDs to achieve CTEs of 99.999999%. Even for a 4096^2 square pixel array, less than a single electron is lost in a transfer from the chip corner with a charge packet of 3,000 electrons.

3.2.3 Front and back illumination

Many CCDs work as expected with light falling upon the designated ‘front’ of the processed silicon. However, for applications that require a high quantum efficiency, this has inherent disadvantages. Incoming photons have a limited site in which to generate electrons, because the electrodes on the front use a notable proportion of

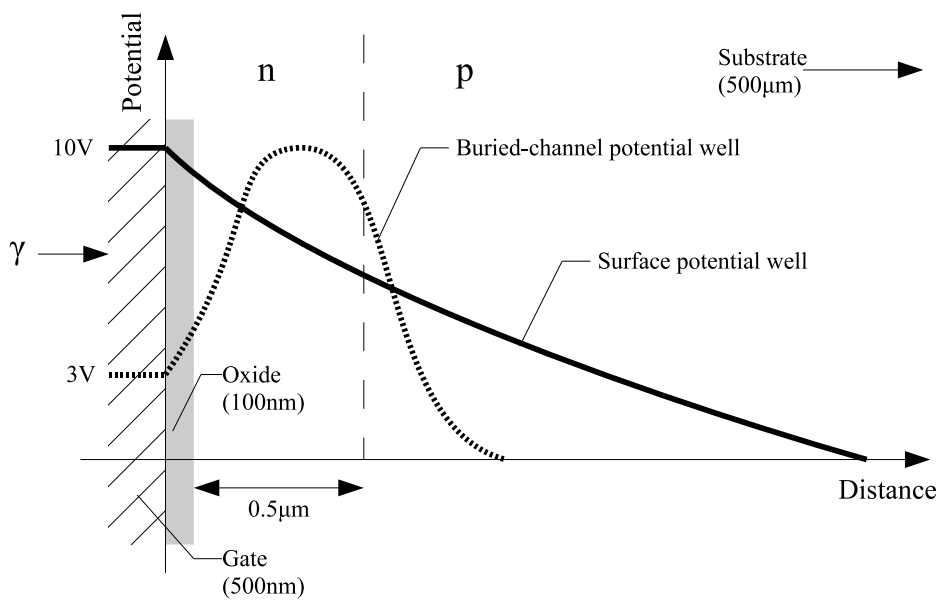


Figure 3.2: Demonstration of how potential changes in a CCD depending on the use of doping (producing a buried channel potential well) or not (a surface potential well). Note that the figure is not to scale, to fully represent the different elements present in the structure. Photons enter from the left (this is a front-illuminated CCD). The transition between n and p dopant in the case of the buried potential well is at $\sim 0.5\mu\text{m}$ (for the surface well, the entire epitaxial layer is considered to be p-doped).

potential pixel area. The solution in this case is to instead illuminate the CCD from the ‘back’. Simply turning a normal CCD around will not function; the thick bulk layer is too far away from the potential well (even buried) and thus the charge collection efficiency tends to zero. For this to be a viable option, the bulk layer must be ‘thinned’ away by a process involving complicated and careful acid etching. The result is a CCD with a very high QE, but lower yields consequently mean higher prices. There is a further difficulty with back-illuminated CCDs, which is etaloning. Because the CCD is now thin ($\sim 20\mu$), it becomes semi-transparent to near infra-red wavelengths. In the 700-800nm regime, the refractive index of silicon is approximately 4, meaning that the entire optical path is 120μ , equivalent to 160 wavelengths at 750nm, and thus a strong pattern of constructive and destructive interference can build up. These ‘fringes’ become increasingly strong towards the infra-red. While variations in the thickness of the silicon can affect the etaloning, the etching process is increasingly accurate and almost all etaloning effects are due to variations in wavelength, not thickness. This problem can partially be overcome by improved anti-reflection coating for near-infrared wavelengths (affecting QE in the red band) or making the CCD silicon thicker (potentially degrading the charge collection efficiency).

3.2.4 Dark current and inverted operation

Dark current physically is the thermal generation of minority carriers which then collect in the potential well. Practically, the amount of dark current a CCD generates defines the length of time a potential well can collect useful charge; as more thermal electrons collect, a background builds up, until eventually the well is full. The dominant source of dark current comes from the many impurities at the gate/oxide interface which act as “hopping conduction states”, where mid-band states offer an easy step to the conduction band for these thermo-electrons. The general dark current formula is [Sze, 1981]

$$D(e^-) = CT^{\frac{3}{2}} e^{-\frac{E_g}{2kT}}, \quad (3.2)$$

where D is the average dark current generated in electrons per second per pixel, T is the operating temperature, E_g is the silicon bandgap energy, and k is Boltzmann’s constant. The constant can be solved at room temperature for $C = 2.5 \times 10^{15} P_S D_{FM}$ where P_S is the pixel area and D_{FM} is the dark current in nA/cm² at 300K. It should be noted that the silicon bandgap energy varies with temperature and can be estimated empirically. For such a CCD, the only methods of reducing dark

current are higher quality silicon (resulting in fewer mid-band hopping states) or more practically reducing the CCD temperature via cooling.

One method used to drastically cut dark current is the use of multi-pinned phase (MPP) technology. If a CCD phase is said to be ‘inverted’, the gate bias is set to a negative voltage such that the surface potential is equal to the substrate potential. At this point holes from the p-layer migrate to the surface interface and ‘pin’ the potential here to zero. During CCD operation this also has the excellent property of using the holes to recombine with the surface state thermo-electrons, effectively eliminating surface dark current. Unfortunately the use of such a technique means that the potential well is now completely unaffected by the gate bias because the layer of holes is acting as a conducting shield; a further negative bias merely attracts more holes. This means that a normal CCD operating inverted will have no collecting capacity (that is, the CCD full well is effectively zero) as each phase is now at the same potential. To get a CCD to operate in inverted mode, phases have to be carefully doped with p-material (using various proprietary methods) to create a potential difference between the barrier and collecting phases for signal electrons to accumulate. When the CCD comes out of inversion the charge is then shuffled and read out as usual. While an increasingly well understood technology, the full well capacity still suffers from employing such techniques as typically the potential barrier formed by doping is not as effective as that generated by a normal CCD. Another issue can be that of spurious charge (clock induced charge) which occurs when a phase comes out of inversion. Holes that migrated to the Si-SiO₂ interface sometimes get trapped, and as the electric field increases when coming out of inversion, these can be accelerated with such force that they cause impact ionisation generating an electron which is then collected. These issues are typically negligible when compared to the benefits it offers in practice, and MPP CCDs have been shown to produce 400 times less dark current than a traditional CCD. For a full treatment of this phenomenon, see Saks [1980].

3.2.5 Noise sources

Even if the MOS capacitors and clocking process itself could generate no noise, the process of measuring the charge and amplifying it for output will always set a read noise floor for any CCD. At the lowest level there is Johnson noise, which is in effect the Brownian motion of charge carriers in the material, theoretically described by [Nyquist, 1928]

$$W = \sqrt{4kTBR}, \quad (3.3)$$

where W is the rms noise (in volts), T is the temperature, R is the resistance in the conducting path, k is Boltzmann's constant and B is the equivalent frequency bandwidth of the material (Hz). Thermal noise of the same type is also present in the reset level transistor, giving one of the two CDS input parameters as a slightly different level from pixel to pixel, which is known as reset noise. On top of this, amplifiers also show flicker noise which takes the form of $1/f$ noise (whereas Johnson noise, being white noise, is independent of frequency). Flicker noise is believed to be generated by surface interface states and tunneling of electrons into the oxide layer [Alers et al., 1996]. Buried channel devices still show flicker noise despite the minimal surface interface interactions, though it is measurably less. All three of these noise sources form the read noise floor, which is noise that will be present even if the CCD is not integrated for any time. It should be noted that the CCD output is configured in such a manner that there will always be a base offset from zero, known as the 'bias', to avoid negative values. A bias frame on a CCD will hence be this offset with a measurable standard deviation that comes from these read floor noise sources. This bias should not be confused with the "fat zero" level as used in photographic plates (§3.1.1); the bias does not contain shot noise (described below). Early CCDs had a read noise of several hundred electrons, whereas more recent chips have a few tens of electrons readout noise for mid-level CCDs and < 5 electrons for professional level devices.

Shot noise will be present in any image that has a positive integration time. This simply is the random arrival of signal photons at the sensor, meaning that some pixels will receive more photons than others. This distribution is well described by Poisson statistics and hence unlike the read noise is dependent upon the signal level (specifically, the square root of it). It should be noted that dark current noise also takes the form of shot noise (though it is dependent upon the dark current, not the signal) and hence will be indistinguishable in an exposed image.

Finally there is fixed pattern noise. This comes about almost entirely from manufacturing processes, most notably the slight misalignment of the poly-layers as the CCD is built up. This simply results in some pixels being more sensitive than others, and hence is directly proportional to signal. CCD manufacturing tolerances have achieved $\sim 1\%$ in fixed pattern noise of this type, but less effort has been expended in recent years on lowering this further as other types of noise became more significant as this noise source decreased. Fixed pattern noise is typically quantified as a dimensionless pixel non-uniformity, the mean difference in registered electrons between pixels given a flat uniform signal, $P_N \sim 0.01$.

3.2.6 Signal-to-noise

Summarising the quality of a data set in a single number without having to state every single CCD noise parameter is a useful facility provided by the signal-to-noise ratio (S/N). Adding the noise sources as errors in quadrature we quickly arrive at the equation

$$\frac{S}{N} = \frac{S}{\sqrt{\sigma_{SN}^2 + \sigma_{FP}^2 + \sigma_R^2}}, \quad (3.4)$$

where σ_{SN} is the shot noise, σ_{FP} is the fixed pattern noise, σ_R is the readout noise, and S is the signal, all in units of electrons. Since $\sigma_{SN} = \sqrt{S + D}$ where D is the dark current, and $\sigma_{FP} = P_N S$ where P_N is the fraction of pixel non-uniformity, we arrive at the more useful equation

$$\frac{S}{N} = \frac{S}{\sqrt{(1 + P_N^2 S)S + \sigma_R^2 + D}}. \quad (3.5)$$

Here it should be apparent that σ_{FP} is a systematic error (as opposed to a statistical error) that in theory should be completely corrected by correct usage of a flat-field (§3.3.4), meaning that $P_N = 0$ and this term will vanish from the equation.

3.2.7 CCD gain and dynamic range

A critical parameter of the CCD is its gain, which is in units of e^-/ADU where ADU is an analogue-to-digital unit, or sometimes known as a DN (digital number). These are units which are output by the OTG and must be interpreted by the end user. Formally, the CCD gain constant, K is

$$K = \frac{1}{S_V A_{CCD} A_1 A_2}, \quad (3.6)$$

where S_V is the sensitivity of the sense node (V/e^-), A_{CCD} is the output amplifier gain (V/V), A_1 is the gain of the signal processor (V/V) and A_2 is the gain of the ADC (ADU/V). A high gain value indicates a low granularity in charge measurement, and hence photometric accuracy can be lost. With CCD gains of $5 e^-/\text{ADU}$ and $200 e^-/\text{ADU}$ a signal packet of $26,700 e^-$ would be detected as 5340 ADU and 133 ADU respectively. Another aspect to consider in designing the camera gain is quantising noise which comes from the fact the ADC outputs discrete numbers and not a continuum. This quantising noise is determined by integrating the rms of a periodic function for a triangular wave with amplitude of $\frac{1}{2}$ [Janesick, 2001] and it

is found that

$$\sigma_{QN}(e^-) = 12^{-\frac{1}{2}}K. \quad (3.7)$$

Using a high gain value such that the read noise is completely negligible results in the output showing clear quantisation levels (steps) from one illumination level to the next. This quantisation will then be hidden when $K(e^-/\text{ADU}) = R(e^-)$.

The number of discrete levels the ADC is required to output is a function of the number of electrons each potential well can hold and the camera gain such that

$$N_{ADC} = \frac{S_{FW}}{K}. \quad (3.8)$$

It should be noted that the number of bits an ADC can output may limit operational choices; consumer devices may only output 8-12 bits, high level devices 16 bits, and for some advanced audio applications 24 bits or more are available. In the case that the camera gain is equal to the read noise, σ_R

$$N_{ADC} = \mathcal{D} = \frac{S_{FW}}{\sigma_R}, \quad (3.9)$$

which is termed the “dynamic range” of the CCD. Deeper potential wells clearly give the device a better scope of values to record, but a greater readout noise diminishes this.

3.3 Photometry

In general, photometry with relation to astronomy is the measurement of a flux source over time. This can simply be a bolometric flux of all light the detector receives, but in such a case values are difficult, if not impossible, to compare with similar measurements from another detector. Typically, a filter is inserted into the system to specify a bandpass of light that is known. When a similar filter is used in another setup, disregarding instrumental effects and calibration errors, it is now possible for these values to be examined together. The various processes of measuring and reducing photometry will now be examined.

3.3.1 Aperture photometry

One of the simplest, but most effective methods for extracting data is that of aperture photometry. Assuming a 2-dimensional image, one picks a small window of these pixels (called an aperture) that encloses the desired target. This aperture is usually circular to match the star profile. Normal aperture photometry now simply

adds up the counts within this window, but this will include an offset caused by the sky background. While it is not possible to measure the sky directly behind the object, it is likely to be a randomised, gaussian distribution, and so an estimation of the sky background is performed. Another aperture, usually an annulus around the original target aperture is then extracted and a mean of these pixels is calculated and then subtracted from each pixel in the target aperture. In most cases this would work without any further issue, and it would seem wise to increase the sky annulus further to perform a better measurement of the mean (as long as it did not encroach other stars). However, this increases the danger of the sky annulus encountering non-random noise such as cosmic rays (highly energetic particles hitting the sensor), bad pixels, or other stars. One can take care with observation and avoid the latter problem, but not the former, which is somewhat countered by performing a “clipped mean” of the annulus, whereby values that are several σ above the straight mean are thus rejected, and the mean is recalculated ignoring the pixel (the most accurate but slower algorithm is to reject a single pixel per cycle). A small target aperture would reduce the error in the sky calculation within reason, but practical ground based observing dictates that one should have a reasonable sized aperture; atmospheric seeing can degrade the quality of the image, blurring the point source out (potentially beyond the target aperture) and losing counts and thus accuracy.

An alternative approach is that of profile fitting, whereby the flux level of the star in two dimensions is assumed to be known to resemble a particular function, and a least-squares (or similar) fit is performed on the variables (such as peak flux and FWHM). This fit can then be used as a good estimator for the current atmospheric conditions, and then a multiplier of this fit width can be used to decide the size of the circular aperture. This then will almost certainly guarantee the inclusion of all the detectable target photons as well as growing (and shrinking) with the variations in the atmosphere to minimise sky error. Typical profiles used are the Gaussian or Moffat functions, the latter of which is based upon studies of intensity profiles of stellar images taking into account seeing distribution and diffraction profiles [Moffat, 1969].

A final method is that of optimal photometry, which uses weighted sums within an aperture rather than simply including or excluding pixels simply based upon their position relative to the aperture circumference. In the case of faint objects with a bright sky background, optimal photometry shows a 10% increase in signal to noise [Naylor, 1998], but offers no benefit in the limiting case of zero sky background and readout noise.

3.3.2 Bias frames

As noted in §3.2.5, any CCD will have some native offset that is a function of its output circuitry and, along with it, some inherent noise level (but note, as pointed out earlier, this is not “fat-zero” as there is no shot noise). A measure of this offset can be made by reading out the CCD after a zero (or near-zero) integration time, preferably with some manner of shutter closing any stray light out from the system. This bias frame then allows a subtraction of this base offset from the science images, resulting in pixel values ~ 0 now where no flux is present. Several bias frames are usually taken and then processed together with a clipped mean algorithm; this provides a statistically good measure of the bias and allows the user to also accurately estimate the readout noise of the instrument for that observation. Another good use of this is to counter any static pattern that is apparent in the readout, typically due to electromagnetic interference from the presence of nearby power and data cables. This should not be relied upon however: interference can change without warning. The bias offsets of a CCD tend to show a slow variation with time simply because of voltage values drifting within the electronics. Altering the CDS t_s time will also affect the bias offset and noise.

3.3.3 Dark frames

These are a non-zero integration of the CCD chip with again no light falling upon it. They are designed to measure the dark current signal and associated noise which builds up over time and depletes the potential wells of useful charge (§3.2.4). Because dark current has shot noise and is not simply a uniform value that can be removed by subtraction, long exposures may be significantly degraded in S/N due to this. The quantity of dark current itself can be estimated if there is a dark frame (or set of dark frames that are averaged) of the same integration length as the science data frames, but different times can be extrapolated from a master dark since it is expected (and found in practice) that dark current is almost completely linear with exposure time. However, some CCDs (typically non-MPP devices) show dark spikes which are generated in the surface, often by damage from energetic protons, which is particularly an issue for space based detectors. For this reason, it is also expected that dark current will increase with the age of the detector. Further study of dark current and dark spikes in CCDs is found in Tanikawa et al. [1976]. A dark frame should be suitable for use in any data assuming the temperature difference is negligible. The use of a dark frame of a time equal to the exposure time also means that science data frames do not require to be debiased as a dark frame includes this

offset. Typical behaviour however is to do a series of darks at increasing steps of exposure time and generate a ‘master’ dark frame, which extrapolates the amount of dark current per pixel for an arbitrary exposure time. Such a process will still require the use of a bias frame to remove the offset.

3.3.4 Flat fields

The last type of noise to take into account via calibration is the fixed-pattern noise, and this is accomplished by flat fielding. The frame should be illuminated by a uniform source (such as quartz lamp, or a twilight sky) so as to give a steady count rate to all pixels. Flat fields must be of a high count rate (though clearly less than saturation) to ensure the significant majority of the noise is in the fixed-pattern regime. Very low counts will be readout noise dominated, whereas the intermediate rates will be dominated by shot noise (dictated by the S^0 and $S^{\frac{1}{2}}$ dependence respectively, where S is the signal). Once a significant number of flat fields have been recorded, these are typically averaged by a median algorithm, and then divided by the average value. In many cases the lamp is not uniform, so the flat is collapsed to one dimension and a polynomial fit to this is then subtracted from the original calibration frame. The result should be a field with all the values ~ 1 with variations on the percent level. All science data is then divided by this frame to account for the inter-pixel variation; at this point, assuming the flat field is an accurate representation of the pixel-nonuniformity, σ_{FP} can be eliminated from the S/N equation since $P_N = 0$ (§3.2.6). Flat fields should not depend upon CDS time as the dominant noise should not be in the readout regime, but they will depend upon the filter used in the light path. One aspect of this is the simple fact that contaminants upon the surface of the filter will affect the pixel sensitivity, but other factors also tend to come into play. Longer wavelength photons tend to penetrate deeper into the silicon where manufacturing processes such as acid etching make small differences in thickness and sensitivity. Flat fields cannot correct for etaloning, however, due to the strong dependence upon wavelength and angle of incidence (implying that the target will generate a different etaloning pattern than the flat-field lamp or twilight sky). Note that flat field frames should be bias subtracted.

3.3.5 Flux calibration

While not necessary for all data, turning the ADU values into mJy or magnitudes gives a value that can be utilised and compared to external data sets independent

of detector. Standard stars, which are non-variable and have well measured magnitudes in several colour bands are regularly utilised for flux calibration (modern photometry has gravitated towards using the SDSS filter system). The standard must be observed with the same CCD gain and filters used for the science frames to give a useable calibration. For the magnitude system, the equation of importance is

$$m_\mu = -2.5 \log_{10}(C_\mu) + x_\mu, \quad (3.10)$$

where μ represents the filter used, m_μ is the literature magnitude of the standard star, C_μ represents the counts from the instrument, and x_μ is the instrumental zero-point in band μ to be computed. Once x has been solved for in each of the apparent bands, the actual magnitudes of the target star can be determined. However, for precise colour measurements (particularly away from zenith), extinction must be taken into account. In this case, magnitudes above the atmosphere should be quoted. Stable comparison stars can be used to give a fit to magnitudes per airmass assuming a run on a target lasts a sufficient length of time (typically hours), where airmass X in the simplest case is assumed to be

$$X = \sec(z), \quad (3.11)$$

where z is the zenith angle, which is the degrees from local zenith. Due to time constraints when observing, the typical workflow involves imaging a set of standards at a variety of airmass in the twilight hours in the necessary filters, which is usually accurate enough for all but the most demanding colour work.

3.4 Spectroscopy

Compared to photometry where the observer studies the broadband flux of an object, spectroscopy involves spreading the light out over the chip and looking at a range of it. This allows a level of analysis otherwise impossible with photometry, including the study of spectral emission and absorption lines present in the system. The former of these come from the fact that in the case of optically thick gas, photons at the wavelength of a spectral line have a shorter mean free path. When re-emitted at the cooler surface, the light appears fainter in comparison to the continuum from the hotter interior. In optically thin conditions, collisional excitation and photo-ionisation can produce extra flux at particular wavelengths, resulting in emission lines.

Calibration is similar to the photometry case, with a few notable additions.

For example, the flat fields now depend upon the grism (used to spread the light out) and potentially the beam slit in the spectrograph as well as the filter (detailed further in 4.4.1). Most spectrograph detectors have an extended region of the chip which is not exposed to light which can be used to measure the bias level for each and every frame.

3.4.1 Arc line calibration

To accurately pinpoint the features detected in a spectrum, the wavelength scale must be calibrated. This is not necessarily static; instrumental flexure as the telescope is moved and rotated can potentially cause features to move on the order of several pixels. Ideally, images of an emission line lamp (such as helium, argon, etc) will be taken before and after an observation. For long runs, the images should be alternated with target exposures while the telescope is kept tracking the target position. These atomic excitation lines are known with extremely good precision. When reduction of the data begins, these lines can therefore be identified and a fit is performed (using a polynomial typically of 3-5 orders) that matches pixels to a specific wavelength. An important feature of arc line images is that each line should essentially be a delta function, so the width of these arc lines is a measure of the spectral resolution of the instrument.

3.4.2 Flux calibration

While using a standard star, as with the case for photometry, a model of the standard that relates wavelength to (relative) flux must be matched against what is observed. This is to account for various spectrograph features including vignetting (natural or instrumental light falloff towards the edges) in the system, which can get extremely pronounced towards the borders of a CCD chip. Once this fit is performed, the counts per mJy is known for a specific wavelength and can be used to flux calibrate the target. The granularity of the model can affect the quality of results, since if model data points are spaced too far apart, various features (which may be instrumental) will not be corrected properly and will affect the final flux values. As expected, the flux calibration frames must be using the same filter and grism as the target data.

3.5 Summary

The detectors used for astronomy - ranging from the traditionally used wide-field but awkward photographic plates to the efficient, but complex CCDs of the modern era - have been detailed. The two major categories of observations (photometric and spectroscopic) were outlined along with the methods used to reduce such data into the final products used for science. Taking into consideration the many challenges and compromises involved in making an efficient scientific charge-coupled device, the next chapter discusses Ultracam and Ultraspec, two astronomical instruments designed specifically for observing cataclysmic variables.

Chapter 4

Ultracam & Ultraspec: detectors for the study of CVs

Initially begun in 1999 as a collaboration between the universities of Sheffield and Southampton (now Warwick), Ultracam [Dhillon et al., 2007b] is a photometer specifically designed with the study of cataclysmic variable systems in mind. The project was funded by PPARC (now STFC), and the instrument saw first light in 2002. Following this success, a similar instrument named Ultraspec [Dhillon et al., 2007a] was commissioned as a test of the practicality of new CCD technologies in 2006, and works as a camera for the EFOSC2 spectrograph. These instruments share many common hardware elements and a software pipeline, and have been used extensively for the research presented here. This chapter will briefly introduce these instruments, the design choices made to meet the necessary science goals, and the implementation of the hardware and software.

4.1 Science goals

As explained in §2, CVs show rapid variability compared to the majority of stellar phenomena. This is particularly true for eclipsing CVs in which the eclipses may be over in a few short minutes. If the observer wishes to accurately examine the geometry of the system by studying when particular parts of the system are eclipsed, however, the detector may require a time resolution on the order of seconds or less. The instrument should have colour bandpasses that are carefully defined, on a scientifically accepted system, so that any results may be readily compared. The scientific use of very short exposure times would also be lessened without the ability to accurately time such events, so an instrument of this kind must be able

to synchronise with an external and accurate clock and keep itself updated at all times. Finally, due to the manner of operation of a visitor instrument at a professional observatory, it would be beneficial that there be reasonable scope within the instrument to perform observations of non-CV objects, so as to make the instrument more appealing to others.

4.2 The common hardware stack

While Ultracam and Ultraspec are different detectors, they have similar roles and rely upon the same basic camera technology. For this reason, excluding the CCDs and any hardware that is specifically relevant to photometry or spectroscopy respectively, most of the layout and setup of the instruments is the same. The critical hardware elements of the two instruments will now be briefly detailed.

4.2.1 GPS timing

Both Ultracam and Ultraspec use a GPS device to precisely set the system timer. The serial connection to the GPS is polled periodically and the internal system timer is then set and computes times in between these polling events. Each frame is then timestamped before readout begins and this is sent along with the frame data to a buffer before being written to disk. It should be noted that in early 2010, the Ultracam GPS was changed. The newer Meinberg GPS achieved a precision of better than 1ms, whereas estimates of the older system indicated better than 10ms (V. Dhillon, private communication).

4.2.2 CCD controller & rack hardware

Other shared equipment between Ultraspec and Ultracam includes a San Diego State University CCD controller (referred to as SDSU) and a set of hardware installed on a rack with the instrument on the telescope. This includes an instrument control PC, disk array, temperature monitoring system, and networking switches. The instrument control PC has a bi-directional link to the SDSU controller and disk array, as well as taking input from the GPS via a serial link. Access to this PC is usually via copper or fibre ethernet on a private network, and the use of a server/client architecture allows not only any external PCs with the appropriate software to drive the cameras, but *multiple* PCs to interact with the instrument and data being collected at the same time.

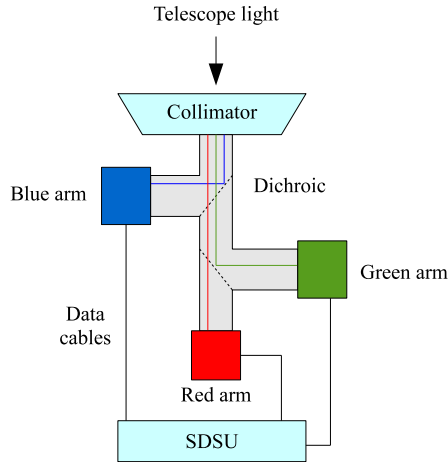


Figure 4.1: A schematic of the key layout of the Ultracam instrument. Light enters through a (changeable) collimator connected to the telescope, and is then split off using dichroics to each arm which consists of a filter slot and CCD detector. The SDSU (CCD) controller is mounted at the bottom of the instrument. The blue and green data cables are of considerably longer length than the red, the greater impedance of which can increase the effects of CCD problems such as self-clocking and substrate bounce, which will be discussed in §8.2.

4.3 Ultracam

4.3.1 Collimators & beam-splitters

The Ultracam photometer possesses the ability to record three bands simultaneously, potentially giving the observers information on the target colour. This feature is seen as important for moving Ultracam forward not only for the study of CVs, but also flare stars, and extrasolar planets amongst others. After the light leaves the telescope, it passes through the collimator, focussing the resulting beam into Ultracam’s light path. Since the collimator is separate, a different one can be attached for each telescope, giving the instrument a chance to participate at varying observatories. To date, Ultracam has been operated at the VLT UT3 8m, WHT 4.2m, and the NTT 3.5m (using the WHT collimator, which gives a wider field of view and some vignetting towards the edges of the field). The light path is then split into three by using two highly efficient beam-splitters, effectively giving three paths termed “red”, “green”, and “blue”, with the blue path diverging first at the top of the instrument, and the red light going all the way through to the bottom, which is nearest the controller. A schematic of this layout can be seen in figure 4.1.

4.3.2 Ultracam CCDs

With three light paths, Ultracam has the requirement of three independent CCDs, but also detectors that run at sufficiently high speeds and with a low noise level so that faint objects can be readily detected. For almost all of its targets, Ultracam will be receiving few photons (due to an inherently faint target or the high-time resolution) and so maximising the quantum efficiency was the primary goal in most decisions. Firstly, back-illuminated CCDs are used as detailed in §3.2.3. This is because the necessary electronics (gate electrodes, etc.) placed on the front obscure potential sites for photo-electrons to be generated, hence lowering the quantum efficiency. Secondly, each of the three Ultracam CCDs are coated to provide an optimal QE for the bandpass they are to cover. The red and green CCDs use a broadband anti-reflection coating whereas the blue CCD has a special UV coating. This latter coating works by reprocessing photons of these shorter wavelengths to a longer one that more readily interacts with the silicon.

For speed and efficiency concerns, each of the three chips also employ frame transfer technology. While the imaging area for each CCD is 1024x1024 pixels, there is an area of the same size on the other side of the array transfer gate which is light shielded. When the integration time is completed, the charge packets in the light-sensitive area are then rapidly clocked in the parallel (vertical) direction into the shielded “storage” area, where they are read out. The reasoning behind this is that while vertical clocking can happen extremely fast, the slowest element of the system is always the charge sampling; thus everything from that point thereon is bottlenecked by this delay (the horizontal *and* vertical clocks). For a full 1024x1024 pixels, this imaged area can be transferred to the storage area via vertical clocking in 24.8ms, whereas the full readout will take several seconds. Such a method also allows Ultracam to operate without a shutter which would otherwise be required to avoid smearing of the image as light fell upon it as it was read out. Hence with Ultracam, while the storage area is being processed, the image area is actively integrating once more, improving the efficiency of the system. This still leaves the fastest exposure time in Ultracam equal to the readout time of the CCD, however, so to decrease this the first step is to employ two amplifiers. This works as two sets of readout electronics: when clocked through the ATG, columns on the left hand side of the chip go to the left amplifier and vice versa. This effectively halves the readout time, which now in a standard full-frame Ultracam setup would equal 6.2s. Another method offered in Ultracam is an alternate “fast” mode readout, which uses a much lower t_s sample-and-hold time in the CDS circuit, but thus increases the

readout noise ¹. Pixel binning can be employed, of course, which can get readout times down to the order of ~ 1 s, but to break this barrier, windowing is required.

A common objective with astronomical observations is to observe point sources, though sometimes several, which take up a tiny fraction of the CCD. By only reading out particular sections of the chip (“windows”), the readout speed can be increased significantly since the unwindowed pixels are clocked but *not sampled*. The pixels in the serial register of a CCD are larger to support vertical binning, and this in turn allows a higher clock speed in the serial direction (full wells are harder to transfer and larger pixels have larger wells). Because of this, it can be seen that with a small, square window, increasing the height of the window will affect the readout speed more than increasing the width because each pixel in a row is clocked down onto the horizontal register simultaneously. Increasing the height by 1 adds a vertical clock and a width multiple of horizontal clocks whereas increasing the width by 1 simply adds an extra horizontal clock for each row. Other subtle ways can increase the speed with windowing, such as placing the windows nearer to the amplifiers on the chip.

Finally, Ultracam supports “drift mode”, for the highest speed operations (to date, the highest speed science run was ~ 400 Hz). In this case, after each pair of windows is integrated, it is shifted in the readout direction by the number of parallel clocks equal to the vertical extent of the window. One pair of windows is read out at a time as they reach the readout register. Because light would otherwise be falling on regions of the chips exposed to starlight, a mask must be inserted into the beam to block the rest of the chip. Another subtlety affects operation speed here; namely the fastest drift mode runs will occur when an integer number of windows fits on the vertical extent of the chip (if this is not the case, there must be a set of parallel clocks equal to the remainder which effectively count as wasted time).

Ultracam is cooled using Peltier devices with water to exchange the heat. A flow sensor inside the water tubes will cut power to the CCD controllers should water flow cease, preventing overheating. The absence of a cryostat makes the instrument more compact and easier to operate as no refilling will be required, but it does mean the CCDs run at a comparatively high temperature. The choice of a fully inverted (MPP) CCD in Ultracam ensures that the dark current is low despite of this design choice.

¹In the Ultracam software stack, this readout/CDS time is occasionally referred to (incorrectly) as “gain”. Modifying the CDS sample time will have no impact upon the detector e^-/ADU

4.4 Ultraspec

4.4.1 EFOSC2

Spectroscopy is inherently more complex than photometry, and requires much more hardware to implement. Rather than design a whole new spectrograph, Ultraspec acts as a replacement camera for the ESO Faint Object Spectrograph and Camera 2 (EFOSC2), which has been mounted upon the ESO 360 (3.6m) and NTT telescopes. The spectrograph consists of a slit (to select the object of interest and mask the rest of the field), a collimator, a filter wheel, and a grism wheel. A grism means a combination of a prism and a grating that disperses collimated light around a designed central wavelength. The Ultraspec project procured two highly efficient, high resolution volume phase holographic (VPH) grisms for use with EFOSC2 (see Barden et al. [1998] for further information on VPH grisms).

4.4.2 Ultraspec CCD

Because Ultraspec only employs one light path (from EFOSC2) rather than the triple-arm used in Ultracam, only one CCD is required. Due to the spreading of the light, count rates are expected to be markedly lower than with Ultracam, and hence frame rates equivalently less. Ultraspec employs different CDS sampling times, frame transfer, and windowing operation as with Ultracam, however Ultraspec only has a single amplifier and does not offer drift mode. This puts a practical upper limit on the speed of Ultraspec at $\sim 30\text{Hz}$. When the slit and grism are removed from EFOSC2, Ultraspec acts in a similar manner to a single-arm Ultracam.

The extremely low count rates expected offered a great possibility for testing the use of L3CCD/EMCCD technology; where signal electrons are amplified resulting in readout noise that is effectively negligible. The principle is not dissimilar to the photomultiplier as discussed in §3.1.2. An alternate serial register (termed the “avalanche” register) creates a huge instrumental gain by the process of secondary emission before it is passed to the sense node, and the voltage applied can be varied leading to a differing amount of secondary emission. In Ultraspec, this amount of secondary emission (or secondary gain) is modified by an “HV gain” (high-voltage gain) parameter between 0 and 9, which increases the voltage applied to the L3 register from 0V (no secondary to emission) to 50V. This extra register will also suffer from degradation should it be exposed to extremely large count rates, and so should be turned off except for when it is needed. The use of such technology absolutely requires Ultraspec to be cryogenically cooled to temperatures on the order of 160K. Any dark current thermo-electrons would be equivalently multiplied by the L3 reg-

ister and at temperatures similar to Ultracam this would defeat the purpose of such technology because, even when inverted, the small amount of dark current would be multiplied and become indistinguishable from signal. At Ultraspec temperatures, dark current is estimated to be on the order of $1e^-/\text{pixel}/\text{hour}$, which is negligible.

EMCCD multiplicative noise

While the HV gain in Ultraspec allows marketing of itself as a “zero read-noise” instrument, the very process of including a gain register adds a complication. In actual practice, each stage of the serial gain register only has a small probability of generating a secondary electron, and many stages are employed sequentially to make the gain considerable. At any stage with n electrons, each of them independently goes through the same multiplication process for the next stage. A critical parameter to determine is the variance of the gain introduced by the r sequential stages, which is a measure of the noise caused by this gain register.

The analysis of this noise source in EMCCDs was first undertaken by Matsuo et al. [1985], and was determined to be a special case of the Galston-Watson branching process as described by Harris [1963]. The first step is to investigate the source of this noise by using this branching process, and, in order to do so, a few definitions are first required. In probability theory, a moment generating function (MGF), μ of a random variable X is defined as

$$\mu_X(t) := E(e^{tX}) \quad t \in \mathbb{R}, \quad (4.1)$$

where $E(X)$ means the expectation of X . Hence $\mu_0(t) = e^t$. The first derivative of an MGF at $t = 0$ is thus described by

$$E(X^n) = \left. \frac{d^n \mu_X}{dt^n} \right|_{t=0}. \quad (4.2)$$

Similarly, the second derivative is used to obtain the variance.

Starting from a single electron at the stage $r = 0$, and associating the distribution with a random variable N , the mean gain g after r stages is

$$g_r = E(N_r) = \sum_{n=0}^{\infty} P(N_r = n)n. \quad (4.3)$$

The MGF for this case of the Galston-Watson branching process [Harris, 1963] is

$$\mu_{N_{r+1}}(t) = (1 - p)\mu_{N_r}(t) + p\mu_{N_r}^2(t), \quad (4.4)$$

where p is the probability of multiplication within a stage. Performing the differentiation at $t = 0$ gives

$$\frac{d\mu_{N_{r+1}}}{dt} = (1-p)g_r + 2pg_r g_{r+1} = (1+p)g_r, \quad (4.5)$$

using the fact that $g_r = E(N_r)$. Also, given that $\mu(t=0) = 1$ for any MGF, it is found that

$$g_r = (1+p)^r, \quad (4.6)$$

which is what may be anticipated. Continuing to the second derivative,

$$\frac{d^2\mu_{N_{r+1}}}{dt^2} = (1-p)\frac{d^2\mu_{N_r}}{dt^2} + 2p\mu\frac{d^2\mu_{N_r}}{dt^2} + 2p\left(\frac{d\mu_{N_r}}{dt}\right)^2, \quad (4.7)$$

which, after simplifying becomes

$$E(N_{r+1}^2) = (1+p)E(N_r^2) + 2pg_r^2. \quad (4.8)$$

Using the standard equation $\sigma^2 = E(X^2) - E(X)^2$ (where σ^2 is the variance) in conjunction with equation 4.6, it is found that

$$\begin{aligned} \sigma_{N_{r+1}}^2 &= (1+p)E(N_r^2) + 2pg_r^2 - [(1+p)(1+p)^r]^2 \\ &= (1+p)(\sigma_{N_r}^2 + (1+p)^{2r}) + 2p(1+p)^{2r} - (1+p)^2(1+p)^{2r} \\ &= (1+p)\sigma_{N_r}^2 + p(1-p)(1+p)^{2r}. \end{aligned} \quad (4.9)$$

This equation is described by the recurrence relation

$$\sigma_{N_r}^2 = \frac{1-p}{1+p}(g_r^2 - g_r). \quad (4.10)$$

As the number of stages, r , becomes large, $g_r \rightarrow \infty$. Also, since p is small, this equation tends to

$$\frac{\sigma_{N_r}}{g_r} \rightarrow \sqrt{\frac{1-p}{1+p}} \sim 1. \quad (4.11)$$

This implies that the spread in the gain is equivalent to its mean. Also as $p \rightarrow 1$, $\sigma_{N_r} \rightarrow 0$, implying that the gain would in such a simple case would be 2^r , as expected.

With the source of this extra noise apparent from the branching model presented above, it must be quantified. Consider the variance, given a mean input of

α counts per exposure with a probability distribution

$$P(n) = \sum_m P_P(m|\alpha)Q_m(n), \quad (4.12)$$

where $P_P(m|\alpha)$ is the probability of m counts given the mean rate α and $Q_m(n)$ is the probability of output n given input m . The mean is

$$E(n) = \sum_n nP(n) \quad (4.13)$$

$$= \sum_m P_P(m|\alpha) \sum_n nQ_m(n). \quad (4.14)$$

Now utilising the fact that $Q_m(n)$ is equivalent to the sum of m independent single electron input distributions (this is a special case of a binomial distribution called a Bernoulli trial) where g is the mean gain,

$$E(n) = \sum_m P_P(m|\alpha)gm \quad (4.15)$$

$$= g\alpha. \quad (4.16)$$

Equivalently it is also seen that

$$E(n^2) = \sum_m P_P(m|\alpha) \sum_n n^2 Q_m(n). \quad (4.17)$$

Since $E(X)^2 = E(X^2) + \sigma^2$, and that the variance of the Q_m distribution is equal to the sums of the individual events (again, due to the process being a Bernoulli trial),

$$E(n^2) = \sum_m P_P(m|\alpha)(m^2 g^2 + m\sigma^2) \quad (4.18)$$

$$= g^2(\alpha^2 + \alpha) + \sigma^2\alpha. \quad (4.19)$$

Using both the expressions for $E(n)$ and $E(n^2)$, the variance in the output after avalanche/L3 multiplication given the mean input rate α is

$$\text{Var}(n|\alpha) = (g^2 + \sigma^2)\alpha. \quad (4.20)$$

The output variance is thus increased by

$$\frac{g^2 + \sigma^2}{g^2} = 1 + \frac{\sigma^2}{g^2}, \quad (4.21)$$

which is termed the excess noise factor (ENF). Utilising equation 4.11, which assumed a small probability of multiplication and a large number of gain stages, $ENF \rightarrow 2$, meaning the process of using the secondary gain to make readout noise negligible inherently adds a factor of $\sqrt{2}$ to the noise in equation 3.5.

4.5 The common software stack

As both Ultracam and Ultraspec are designated ‘visitor’ mode instruments, they require their own full and independent software stack to operate without input or information from the observatory and telescope software. Since they are both intended to operate on multiple telescopes, the software must also be generalised so as to be adaptable to varying hardware configurations and architecture. Full details are given in Beard et al. [2002]. All the protocols used in the software communication in Ultracam are common, non-proprietary methods, including HTTP for the stream and XML for the content, which has the added bonus of being self-descriptive.

Taking the modular approach, Ultracam is split into four main pieces of software. There is a camera driver, whose main goal is to configure the SDSU and send it the appropriate “application”. This code is DSP and comes from a central repository, where each application gives the SDSU controller a set of instructions to perform a particular task; these include power on/off, full frame readout, windowed readout, and drift mode. The SDSU controller in turn deals with the CCD electronics and passes the raw data back to the shared software buffer, as well as a synchronisation pulse dictating the start of exposure. Software threads that deal with SDSU communication as well as the buffer are run with “real-time” scheduling; in theory this should allow all the necessary information and data to be stored in memory on demand from the SDSU even if the other user-space software elements are not ready yet. In addition to this the PCI card that receives the raw frame data from the SDSU controller uses Direct Memory Access (DMA), meaning that the majority of data does not in fact have to pass through the CPU at all. The data handler deals with the GPS and timestamping, as well as saving the finished data product to disk once all the appropriate information (raw data, status, timestamp) are present in the shared buffer. The optional file server will then serve out individual frames from the finished product to the end-users, allowing data reduction or real time monitoring by the observer. Finally, coordinating all these is the Ultracam/spec driver application, which is the observer interface program. Here, the user inputs information such as target name and filters, then picks a mode (full-frame, drift mode, etc.), sets up the windowing and binning if desired, then “posts” the

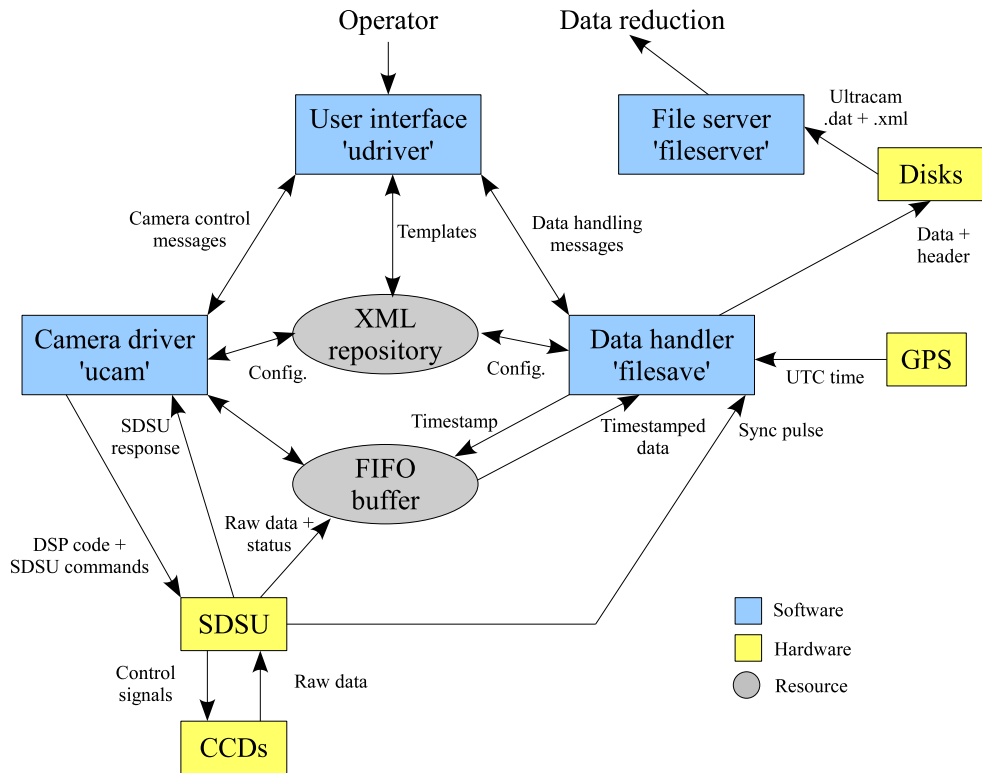


Figure 4.2: A schematic of the flow of information in the Ultracam software stack with endpoints meeting the hardware and end-user.

application to the camera driver and data handler before beginning the sequence of exposures. This final operation involves retrieving a template from the XML repository for the desired mode, using it to set the desired user configuration and send it to the other pieces of software to appropriately configure them. A schematic of this data flow can be seen in figure 4.2.

Ultraspex can harness the entirety of this software stack. The only real changes are in fact the DSP code for the SDSU controller, since Ultraspex has a single CCD without the split amplifier, and also an alternate output register. When at an observatory, the camera controller, data handler, and file server are all run on an SMP rack PC mounted near the instrument (so it can connect to the SDSU controller). The actual user interface software is on a remote PC in the observing building, connected by ethernet (preferably a fibre link). In theory, any disruption on the observer PC should not cause a malfunction with the rack software, and the instrument should keep running. The observer PC can also be exchanged for any other as long as it can run the Java interface client.

4.6 Summary

The layout and design of Ultracam and Ultraspec have been described along with their choices of CCD detector and other hardware to achieve their scientific goals of being ideal instruments for the observation of CVs while retaining some flexibility for other targets. Ultracam data will be presented in §6, whereas improvements to the Ultracam software and calibration of the filter system is detailed in §8. As a relatively new instrument, the technical details of data reduction for Ultraspec will be visited in §9, along with initial spectroscopy and spectro-photometry.

With the science behind CVs and detectors now described, the next chapter will proceed to examine EX Dra, an eclipsing CV system, using optical spectroscopy from CCDs and UV data from a space-based Digicon device, which was present on the Hubble Space Telescope.

Chapter 5

Disc-rim absorption in CVs: the case for EX Dra

EX Dra (also known as HS 1804+6753) is an eclipsing cataclysmic variable with recurrent accretion-driven variability resulting in outbursts triggering typically every few weeks. EX Dra was originally detected in the Hamburg Quasar Survey [Bade et al., 1989] with a V magnitude of 14.5 and was officially classified as a dwarf nova some years later [Barwig et al., 1994]. Outbursts were noted to be on the order of 1 - 2.3 magnitudes in V, with an orbital period of 5.04 hours, implying a system with an orbital separation such that magnetic braking is the primary reason for angular momentum loss. Spiral shocks have also been reported in the accretion disc in the Balmer and helium lines [Joergens et al., 2000]. The mass ratio and inclination respectively have been determined to be $q = 0.72 \pm 0.06$, and $i = 85 \pm 2^\circ$ [Baptista et al., 2000].

This chapter will utilise a mixture of optical and ultraviolet data to investigate properties of EX Dra. The combination of several years of optical spectroscopy gives an opportunity to investigate the differences present between its quiescent and outburst states, and the UV spectra offers the potential to study wind-driven lines, an area of continued study in the field of CVs.

Many studies of cataclysmic variables have shown that observations of cataclysmic variables are affected by obscuration, either by the accretion disc itself or from other material present in the system. It has been indicated that EX Dra itself has a white dwarf surrounded by extended cloud that is three times its radius [Baptista et al., 2000] - when the central object is treated as a white dwarf alone, mass-radius relations give unreasonably low values of $M_1 = 0.2M_\odot$. With OY Carinae (another dwarf nova), there is strong evidence that the boundary layer

Telescope	Instrument	Start UTC (Duration)	State	Comments
WHT	ISIS	June 10, 1995 00:12 (3.6h)	Low	Cloudy. Seeing 0.7"-1.2".
		June 11, 1995 00:01 (4.0h)	Low	Some cloud. Seeing 1.0"-1.3".
		June 11, 1995 22:30 (5.1h)	Low	Clear. Seeing 1.0"-1.3".
		July 27, 1996 22:56 (3.9h)	High	Clear. Seeing 1.3"-2.0".
		July 29, 1996 00:09 (4.9h)	High	Clear. Seeing 1.4"-2.0".
30"	CCD	Sept 05, 1996 20:52 (1.4h)	Low	
30"	CCD	Sept 13, 1996 20:25 (2.4h)	High	
HST	FOS	Sept 13, 1996 03:24 (12.2h)	High	Two full orbital periods.
JGT	CCD	Sept 14, 1996 20:24 (4.8h)	High	
Caha	TWIN	July 21, 2000 20:26 (5.6h)	High	Cloud. Seeing 2.0"-3.0".

Table 5.1: Journal of observations. A high state means the system was in outburst.

white dwarf is ‘veiled’ by an Fe-II curtain of absorbing material [Horne et al., 1994]. This material was anticipated to be in the upper atmosphere of the disc near the outer rim. Knigge et al. [1997] showed that absorbing gas of similar temperature and column density also could explain observations of another system, Z Cam. In both cases, the use of ultraviolet spectra was essential in detecting the absorption lines and attributing them to the curtain. Finally there is the remarkable case of DW UMa, where the disc completely obscures the line of sight to the inner disc and white dwarf, resulting in a reduction in UV flux during the high state at a time when the optical flux is 3 magnitudes brighter than in the low state [Knigge et al., 2000].

From the above cases mentioned, it is clear that partially obscured primary stars are not uncommon within CV systems. Presented here are several sets of spectra and photometry of the dwarf nova EX Dra, this includes continuous HST observations for over 12 hours, one of the longest observed dwarf novae with this telescope. As will be shown, these data sets reveal extensive evidence of obscuration.

5.1 Observations and data reduction

Our observations comprise of spectroscopy and photometry taken at UV and optical wavelengths, resulting in comprehensive coverage of EX Dra from 1100 to 9200Å across both quiescent and outburst states. For specific dates and details see table 5.1.

5.1.1 Optical spectroscopy

Initial spectroscopy was taken using ISIS, a twin-arm spectrograph mounted on the 4.2m William Herschel Telescope in La Palma [Clegg, 1992]. It offers simultaneous spectra in two bands known as the red and blue arm, the former typically used to cover the H- α regime, and the other H- β or H- γ . The spectra from 2000 were taken with the TWIN spectrograph based at the Calar Alto observatory using the 3.5m telescope [Graser, 1998].

The 1995 ISIS run used the R1200R grating in the red arm and the holographic H2400B grating in the blue arm, resulting in wavelength ranges of 4200-4500 Å in blue and 6400-6800 Å in red with resolutions of 1.23 Å and 0.57 Å respectively. The 1996 ISIS data used the R1200B grating (blaze, not holographic) in the blue arm and again R1200R in red. This resulted in a slightly longer wavelength in blue of 4600-5000 Å compared to the 1995 data (resolutions were measured to be 1.64 Å and 1.23 Å). Slit widths were 1.0" and 1.25" respectively. The 2000 TWIN run used the T05 grating in blue with 3700-4800 Å with resolution 1.85 Å and T06 in red resulting in 8000-9200 Å at 2.43 Å. The red arm was heavily polluted by telluric lines but was primarily targeting the Ca-II triplet. The observations down to H- γ with ISIS typically required exposure times of 3-5 minutes. The red H- α bands had significantly higher counts and hence had 1-2 minute exposure times. With TWIN, the exposure times were equal throughout most of the run, but the slit was set at 1.5" due to the highly variable seeing.

The ground-based spectra from ISIS and TWIN were reduced using the *STARLINK* package with the *PAMELA* routines developed and written by T.R. Marsh. The procedures involve debiasing, flat-fielding (3.3), curved spectra tracing, arc-line calibration (3.4), sky fitting, and an optimal extraction [Marsh, 1989]. A comparison star on the slit was used to correct for transparency variations, and a wide-slit exposure was used to correct for the slit loss factor. Some of the 1995 data, notably in the blue-arm, were slit-loss corrected using custom routines due to the extremely faint nature of the comparison which caused difficulty with the fitting procedure. Since the spectra were not notably curved, this custom routine initially took the top and bottom thirds of each spectrum and collapsed them, with a straight line being fitted between the two points where maximum counts were detected. A gaussian with a fixed position determined by the line but variable width and height was then fitted to the full frame, and the counts within summed (polynomial sky subtraction was used). Finally, all data sets had a flux standard star allowing the data to be put onto an absolute flux scale. It should be noted that there were no flat field calibrations taken for the 1996 data.

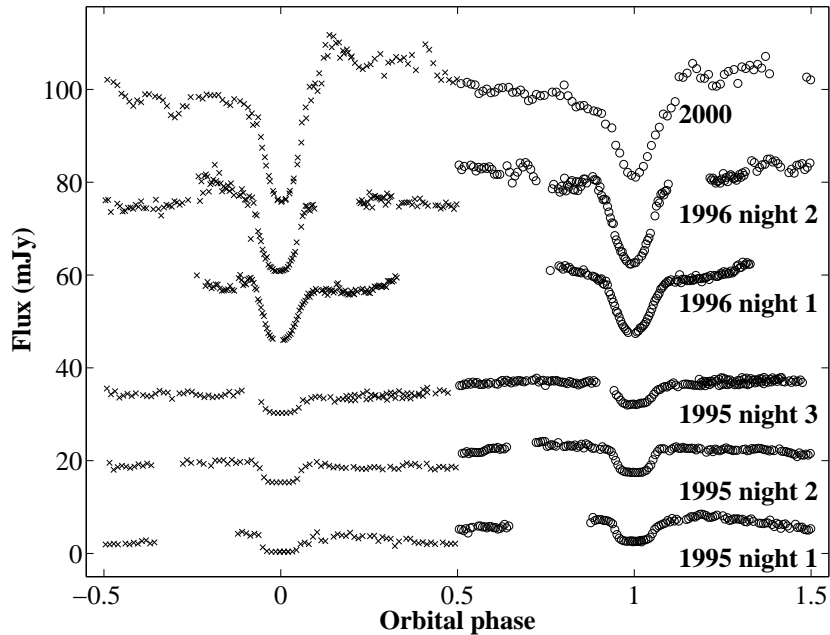


Figure 5.1: Stacked light curves derived from spectra of EX Dra, with the blue arm data on the left and red arm data on the right. The bottom three spectra are WHT/ISIS 1995, the next two are WHT/ISIS 1996, then at the top is Caha3.5/TWIN 2000. Each curve is offset by 15mJy.

The light curves constructed from all optical spectra are shown in figure 5.1, and their average spectrum in figure 5.2. From the continuum flux observed, it was found that the 1995 ISIS data were in a state of quiescence, whereas the 1996 ISIS and 2000 TWIN data were in outburst.

5.1.2 HST ultraviolet spectroscopy

For the UV observations, the Faint Object Spectrograph [Harms et al., 1979] on the Hubble Space Telescope (HST) was used in September 1996. This used photon-counting digicon detectors with a spectral resolution of $R \approx 1200$. As EX Dra lies in HST's continuous viewing zone, 2 eclipses with full orbital coverage were acquired. The HST FOS data runs between 1100 and 1800Å. The data came pre-calibrated from the automated HST archive pipeline and custom routines were written to reduce the data into a useful and workable format.

Radial velocities can be obtained from the HST FOS data despite the low resolution, but there is an important correction to be made. This is the 'geomagnetic image deflection problem' (sometimes termed 'GIMP'), whereby the impacting

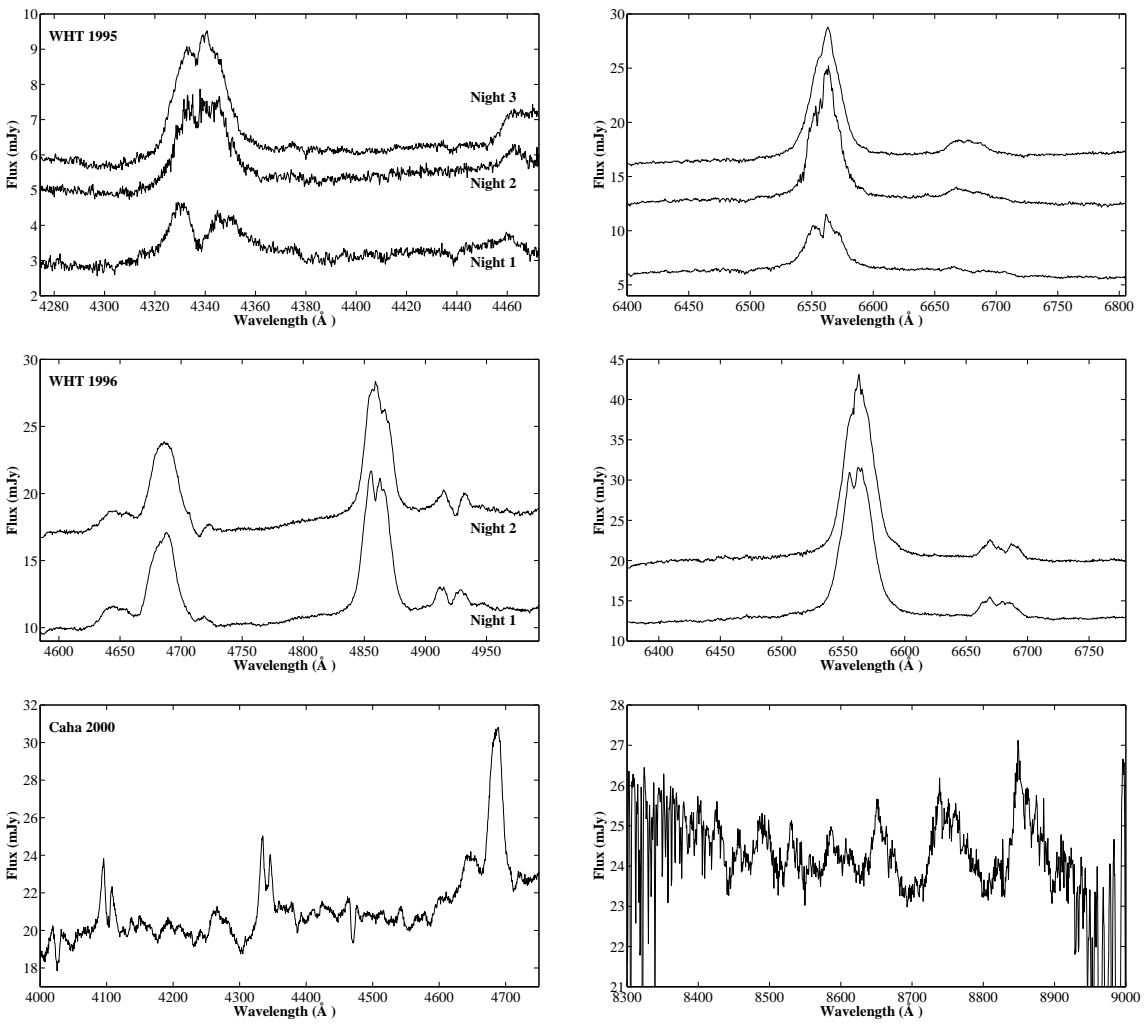


Figure 5.2: Average spectra from ground based optical observations with the blue arm on the left and red arm on the right. Top: WHT/ISIS 1995 (offsets: blue, 1mJy; red, 5mJy). Middle: WHT/ISIS 1996 (offset 2mJy). Bottom: Caha3.5/TWIN 2000.




Figure 5.3: The average spectrum generated from the HST FOS observations.

location of the charge in the digicon upon the silicon diodes varied with the orbit as the earth's magnetic field changed. In effect, there is a shift on the order of 60km/s every 48 minutes. This was corrected for by computing a power spectrum of the radial velocity curve of the strongest emission line (CIV) and taking the peak closest to the 48 minute mark to be the GIMP period. It should be noted there were no other strong peaks around this point (though it was not as strong as the orbital period itself) or clear aliases of this period. A sinusoid at this period was then fitted and subtracted from the data, and the power spectrum was shown to have no noticeable features at this period or multiples of it.

The average UV spectrum can be seen in figure 5.3, whereas continuum light curves of the entire data set increasing in wavelength, along with the zeroth order [Horne and Eracleous, 1993] light curve ($\sim 2000 - 5000\text{\AA}$, may be seen in figure 5.4.

5.1.3 Optical photometry

To determine the state of the system during the HST/FOS observations, photometry of the system was acquired during the lead-up to the observing date, which was augmented by observations with the UK-based 0.9m James-Gregory telescope in St. Andrews. Differential photometry with a nearby comparison provided V-band magnitudes in all cases. At the start of these observations, EX Dra was seen rising to an outburst state. This is confirmed by AAVSO data available for the period, which shows that EX Dra actually reached a peak visual magnitude of 12.8 on the day of the HST observation, before returning to its typical magnitude over the following ten days. The light curves of the photometry are shown in figure 5.5.




Figure 5.4: Split continuum light curves from differing wavelengths constructed from the HST FOS observations, excluding emission lines. From the top downwards: 1600-1800Å , 1400-1600Å , 1200-1400Å , zeroth order light curve.

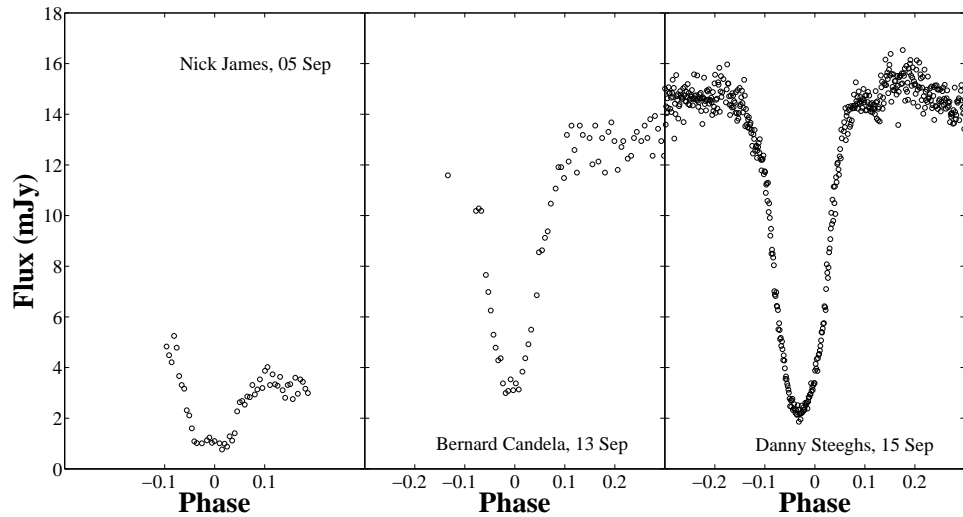


Figure 5.5: Optical V-band data taken with amateur telescopes and the 0.9m James-Gregory telescope (right) leading up to and after the HST run.

5.2 Disc-rim absorption in EX-Dra

5.2.1 Evidence from the UV data

The continuum light curves shown in figure 5.4 show significant changes between wavelength bands. A pronounced flickering behaviour is seen which is most emphasised prior to the first eclipse (at HJD 50339.18) and between the two eclipses at HJD 50339.34. While this can be seen at all wavelengths, if the ratio of the maximum to minimum out-of-eclipse flux is examined, it is found that while in the 1200-1400 Å range this is 3.2, this decreases to 2.5 at 1400-1600 Å, then 2.2 in the 1600-1800 Å band and is only 1.5 of the zeroth order light curve. These light curves were examined for periodic signals such as the white dwarf spin, but none were found.

To examine the possibility that there is obscuration along the line of sight leading to this rapid variability present in the light curve, we can utilise the wind-driven emission lines present in the UV spectrum. The large number of UV photons from the disc act to push the nearby material outwards since the absorbed photon is from a preferred direction, but the re-emitted photon is in a random direction. In effect, these photons from the disc are scattered into the observer's line of sight at some distance away from the primary (see figure 5.6), and so depending on the distance, should be affected less by any line of sight obscuration. Continuum subtracted light curves of these lines together with the continuum for comparison are shown in figure 5.7. The amount of flux in mid-eclipse immediately gives information with regards to the location of these lines in the system. HeII is seen as a flat-bottomed eclipse with near-zero flux indicating that it is fully eclipsed and hence it does not extend far from the disc. Lines from SiIV to CIV show an increasing amount of mid-eclipse flux and a U-shaped eclipse proving that these must be extended away from the disc and are not completely eclipsed. This is as anticipated; CIV is a resonance line and will hence be strongly affected by scattered photons from the wind whereas HeII is not. Most importantly, it is seen that CIV, NV, and SiIV show a steady flux at the same time the heavy flickering behaviour in the continuum is observed. HeII shows traces of flickering but it is clearly less pronounced, and does not show the continuum feature present at HJD 50339.18. This appears to give evidence towards the theory of obscuration along the line of sight, potentially in the disc rim, however this may not be conclusive due to the time scale that the wind can vary on ($\sim R_{wind}/v_{wind}$). At a typical wind velocity of a few km/s this results in a timescale of a few minutes, not a dissimilar period to many of the flickering features observed.

If there is a line of sight absorption, it is important to consider physical sources for such a phenomenon. HST FOS data of OY Carinae presented in Horne

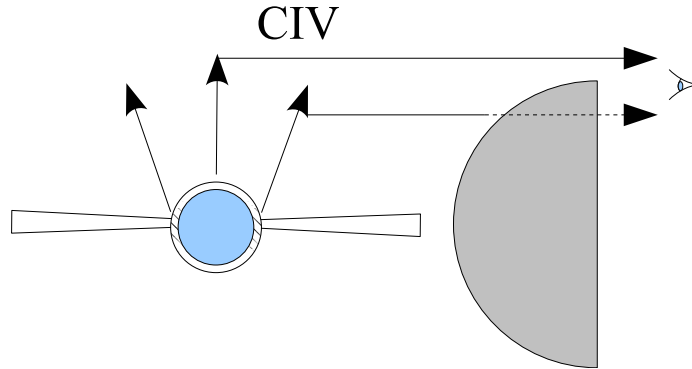


Figure 5.6: A diagram (not to scale) demonstrating the observation of a scattered CIV resonance line from EX Dra. UV photons push CIV material away from the inner disc and are hence scattered at such a distance that they are only partially eclipsed by the red dwarf to the observer's line of sight. Disc lines like HeII will be entirely eclipsed.

et al. [1994] showed a highly variable continuum in the average spectrum attributed to an 'FeII curtain', showing deepest absorption between 1600 and 2500Å. The FeII lines described are theorised to be in the outer disc, but require Mach ~ 6 velocity disturbances in such a scenario for the data to fit. The FOS spectrum (figure 5.3) for EX Dra shows a relatively smooth continuum in comparison to OY Car, however the spectrum does not extend to the higher wavelengths used by Horne et al. [1994]. No single absorption feature seen can be easily attributed to intervening iron lines, and this small number of features is compounded by a lack of synthetic spectra of the correct wavelengths. While past work has shown that the fitting of UV absorption features is of great use [Lanz et al., 1992; Shore and Aufdenberg, 1993; Knigge et al., 1997], EX Dra would require data at wavelengths more suited for the observation of the FeII lines to definitively reach a conclusion on the source of the obscuring material. The optical data is now examined for corroborating information.

5.2.2 Evidence from the optical data

The quiescent optical spectra taken in 1995 (see the top of figure 5.2) show a remarkable change in line profiles over the course of three nights. Initially, the first night shows $H\gamma$ and $H\alpha$ with weak emission (1.04 and 4.4 mJy respectively) and a heavily absorbed fraction of the line in the centre. These lines return to a more typical state (3.14 and 10.68 mJy) over the course of the next two nights. This central absorption component also appears to exaggerate the line widths, where it




Figure 5.7: Continuum eclipse and continuum subtracted emission line light curves from HST FOS data. The original $\sim 13,000$ data points have been binned to 1,300. Flux has been normalised to the line maximum and the error bar gives an indication of the typical error on each data point.

is found that the FWHM of H γ changes from 2990 km/s when the absorption is present to 1570 km/s afterwards. HeI (6678 Å) is also heavily affected by the same phenomenon, to the extent it is nearly undetectable on the first night. The measured width is 4750 km/s with a peak of 0.55 mJy above the continuum on night 1, returning to 1501 km/s and 1.29 mJy by the third night. It is also of note that the continuum is 0.5 mJy lower on average than the other quiescent observations. Considering that the system is not near an outburst state at this point, it seems likely that the continuum flux changing is the result of obscuration of the disc and white dwarf, and the absorption seen in the cores of the emission lines is along the line of sight. This can be adequately explained by material in the system's disc-rim, which absorbs flux from the white dwarf and inner disc and also changes on a sufficient timescale to match the observations.

5.3 Radial velocities

A traditional radial velocity analysis of several emission lines is now performed with the optical and UV spectra by cross-correlation. The white dwarf radial velocity, K_1 , is determined by measurement of lines associated with the inner part of the disc (eg. Balmer lines), whereas calcium lines are measured for the companion star. A measurement of both provides a spectroscopic method for determining the system's mass ratio.

Following the method of Schneider and Young [1980], the wavelength of a spectral line is solved for by computing

$$\int_{-\infty}^{\infty} S(\Lambda)K(\lambda - \Lambda)d\Lambda = 0, \quad (5.1)$$

where S is the emission line and K is a suitably defined function for the data (described below). This equation will equal zero at the position where K best matches S . While K can be a simple step function, a typical form of K is

$$K(x) = \exp\left[-\frac{(x - a)^2}{2\sigma^2}\right] - \exp\left[-\frac{(x + a)^2}{2\sigma^2}\right]. \quad (5.2)$$

This process convolves the spectrum with two gaussians with separation $2a$ and a width σ , and is a far better solution with digital data, particularly in a low signal-to-noise regime. Once the offset of each line has been determined in each spectra

they are fit to a sinusoid of the form

$$V(t) = \gamma + K \sin 2\pi \frac{(t - t_0)}{P} - \phi, \quad (5.3)$$

where γ is the radial velocity offset, K is the amplitude of the radial velocity in km/s, t is the mid-exposure time, ϕ is the phase offset and P and t_0 are defined by the ephemeris. For the purposes of fitting, the ephemeris of Fiedler et al. [1997] is adopted, where it was previously determined that $K_1 = 167 \pm 3$ km/s and $K_2 = 223 \pm 1$ km/s. In addition to this, Billington et al. [1996] adopted $K_2 = 210 \pm 15$ km/s by excluding data points from specific phases to reduce the effect of radiation-induced distortion. These two values for K_2 are consistent given the uncertainty quoted.

Analysis of the radial velocities was performed using the traditional method of a ‘diagnostic diagram’. By changing the width and separations of the gaussians fitted above, for each case a radial velocity amplitude, phase offset, and σ/K (where σ is the error) is found. Since the fitting method revolves around the unlikely scenario of a Keplerian disc (in reality, distortions in the outer disc from the bright spot are significant), the wings of the disc lines are expected to give a more realistic radial velocity. At some point however, the gaussian wings trail into noise, and the diagnostic diagram is a tool for attempting to determine at what point the gaussian parameters accurately reflect the true radial velocity of the line being measured. Distortions in the disc will give rise to an greater measurement of the amplitude K , so a region on the diagnostic diagram where K is at a minimum, with a stable σ/K and small phase offset is the most reliable estimate of the radial velocity using this method. In some scenarios, the diagnostic diagram converges towards a preferable K , whereas on some occasions there are no width/separation combinations that give a sufficiently low phase offset. This is likely indicating that the bright spot is significant in the disc, and that any K values quoted from such data should be regarded as questionable at best.

5.3.1 Optical spectra: Balmer lines and CaI

A common issue found when attempting to measure the Balmer lines of CVs is a large phase shift (ϕ) relative to the ephemeris [Shafter, 1983]. In addition to this, where ϕ was small, the value of σ_K/K grew large. Even when discarding the clearly absorbed first night of 1995, every Balmer line measured converged with a relative phase shift of greater than +0.1. The values at the point of lowest phase shift were determined to be 179.9 ± 10.4 km/s for H γ , and 155.9 ± 13.6 km/s for H α . HeII




Figure 5.8: Trail of the calcium lines present in the 1995 EX Dra data.

provided the best fit with 165 ± 4.6 km/s, but even this had $\phi = 0.08$.

Analysis of CaI (6439.1 Å, 6462.6 Å, and 6493.8 Å, figure 5.8) did allow for a reliable measurement of K_2 , though these absorption features were only strong enough to perform a fit on the first night of observations in 1995. There are clear signs of radiation induced distortion (see Hessman et al. [1984]; Billington et al. [1996]) which occur around phase 0.5 and cause a departure from a sinusoid, which can be seen in figure 5.9. Attempting to fit all data points results in $K_2 = 225.3 \pm 15.1$ km/s, whereas eliminating the distorted points either side of phase 0.5 instead gives $K_2 = 217.5 \pm 3.9$ km/s; the χ^2 falling from 3.8 to 1.4.

5.3.2 UV spectra lines

All the radial velocities from HST FOS showed a consistent offset in velocity of the data, which is not unexpected with this instrument, but was ~ 500 km/s in the data due to the low resolution. The typically recommended solution, realigning with the geocoronal Lyman- α , was not used due to the line being weak and clear signs that the system itself affected this line by the presence of an orbital modulation and partial eclipse. Given that the HeII line is fully eclipsed (as seen in figure 5.7) and hence must be coming from the disc, the solution was to treat the offset found in the HeII radial velocity as the actual instrumental offset.




Figure 5.9: The radial velocity curve of the Ca-I triplet from the donor star of EX Dra during the first night of 1995. Note the outlying points either side of phase 0.5, which distort the complete fit (dashed line). A more reliable estimate is hence made by not fitting these points (solid line). Points discarded by this fit are shown in blue.

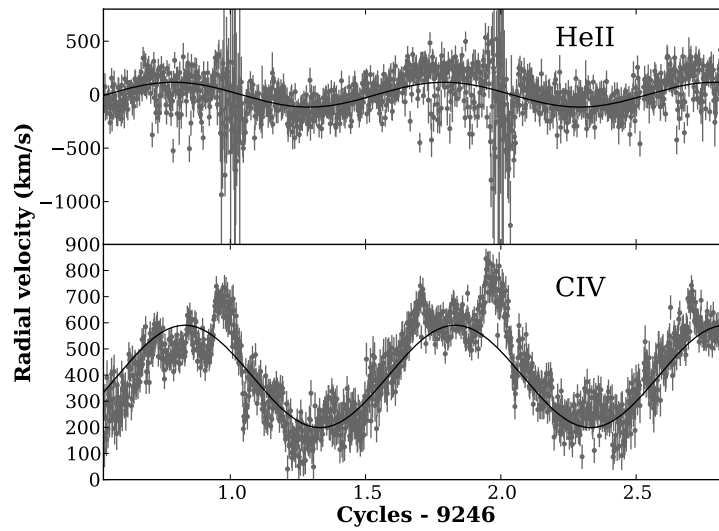


Figure 5.10: The GIMP-corrected radial velocity curve of HeII and C-IV (from HST FOS).

Line	K (km/s)	γ (km/s)	ϕ	Reduced χ^2
HeII	116.5 ± 8.0	0.0 ± 6.1	0.2 ± 0.08	2.0
SiIV	185.9 ± 7.4	66.2 ± 5.6	0.3 ± 0.05	1.3
NV	169.7 ± 8.9	611.5 ± 6.6	0.1 ± 0.06	11.4
CIV	196.0 ± 5.2	394.9 ± 3.8	0.0 ± 0.03	3.5

Table 5.2: Emission line radial velocities from HST FOS UV spectra. Note that γ for HeII was assumed to be the instrumental offset and was subtracted from the offset of each line.

Using this along with the GIMP correction, the computed radial velocities are detailed in table 5.2. In addition to this, departures from the typical radial velocity sinusoid in both HeII and CIV are found. The former shows a sudden blue-shift upon eclipse, as is associated with a rotating object being eclipsed. CIV alternatively shows a distinct red-shift on the order of 200km/s prior to eclipse, with little blue-shift afterward. In the case of CIV this may be a signature of a rotating disc wind [Knigge and Drew, 1997]. Such disc winds may be expected to have a ‘‘P Cygni’’ profile, a feature that consists of a combination of a scattered resonance line and the blue-shifted absorption that comes from gas along the line of sight to the UV source. The sum of the components gives a consistently red-shifted line, which could explain the large velocity offset observed in the NV and CIV lines. However, when the absorbed part of a P Cygni spectral line is eclipsed, this leads to an effective *increase* in flux from the blue side of the line, meaning a blue-shift is observed. With CIV, there is a counter-intuitive red-shift seen just before eclipse. This remains a mysterious feature that is currently unexplained.

5.4 Doppler tomography

Doppler tomograms of the spectra were produced to further analysis; this is the process of mapping a series of spectra into two-dimensional velocity space. This is possible since the measurement of a line in a spectrum gives a projection of velocity along the line of sight. In essence, each S-wave present in the spectra trail corresponds to a pixel (V_x, V_y) in the map, such that

$$V_R = \gamma - V_x \cos 2\pi\phi + V_y \sin 2\pi\phi. \quad (5.4)$$

For the process performed here, a maximum entropy method [Marsh and Horne, 1988] was used, which employs a ‘goodness of fit’ (χ^2) statistic to determine how far to optimise the image; these datasets achieved χ^2 values of between 1.2 and

1.5. The theoretical secondary lobe and stream velocities are overlaid on the maps. Each of the doppler map figures consists of a trail of the original spectrum data, followed by the computed doppler tomogram, then the trail that would be required to produce the previous map (excluding eclipse regions) and finally the residuals of the computed trail subtracted from the data trail as a measure of the quality of the doppler map. The tomogram itself in the second panel also has the secondary star and theoretical gas stream superimposed as a solid line.

Maps for the quiescent 1995 data are shown in figures 5.11, 5.12, and 5.13. For the 1996 outburst data, only the second night of data is shown in tomogram format in figure 5.14 due to insufficient phase coverage on the first night. The 2000 outburst map can be seen in figure 5.15.

It should be noted that the $H\gamma$ map on the first night of the 1995 run had a considerable absorption feature in the line centre which the doppler tomography optimisation method cannot reproduce. To improve the map, a 2-dimensional gaussian was added to the data before the map was optimised, and then it was subtracted afterward.

5.4.1 Slingshot prominences

On all three nights in 1995, significant emission in $H\alpha$ is seen from the secondary, with an extended region outside of the computed Roche lobe towards lower velocities. The same effect can be seen to a lesser extent in the 1996 $H\beta$ map and the 2000 $CaII$ map. This low velocity component does not model well and can be seen repeatedly in the residuals. Looking individually at the spectra, it is clearly present even during mid-eclipse with a flux density of 1.5mJy, and when subtracted from the $H\alpha$ line profile, the single peak becomes a much more clearly defined double peak typical of an accretion disc line. Producing a light curve from this region near the zero velocity offset (figure 5.16), it is found that the flux drops by 80% simultaneously with the eclipse of the white dwarf, but there is also a shallower dip at phase 0.5 by 35%.

This low velocity combined with its apparent relation to the secondary star makes it seem likely that this phenomenon is explained by a slingshot prominence on EX Dra's secondary, previously seen on another dwarf nova IP Peg [Steeeghs et al., 1996]. Active magnetic regions on the secondary star experience a centrifugal force associated with the stellar rotation and gravitational force from the white dwarf. The result is that magnetic field lines are pulled outward until the gravitational force and magnetic tension balance the centrifugal force. In a cataclysmic variable system, the shape of the binary potential means that material is less strongly bound on the

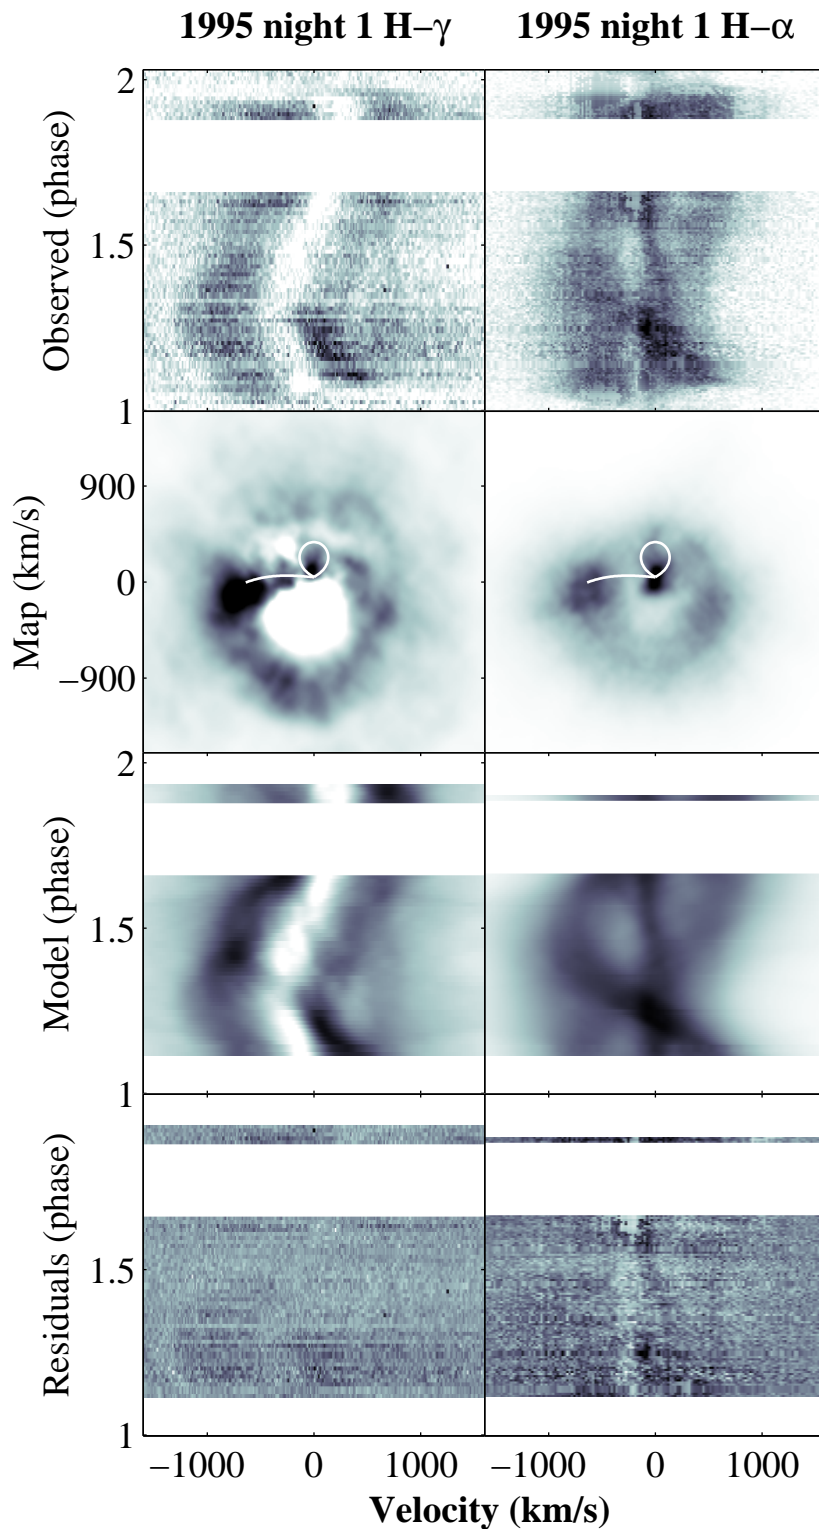


Figure 5.11: Trailed data, doppler maps, computed trail and residuals for the first night in 1995. See text for details.




Figure 5.12: Trailed data, doppler maps, computed trail and residuals for the second night in 1995. See text for details.

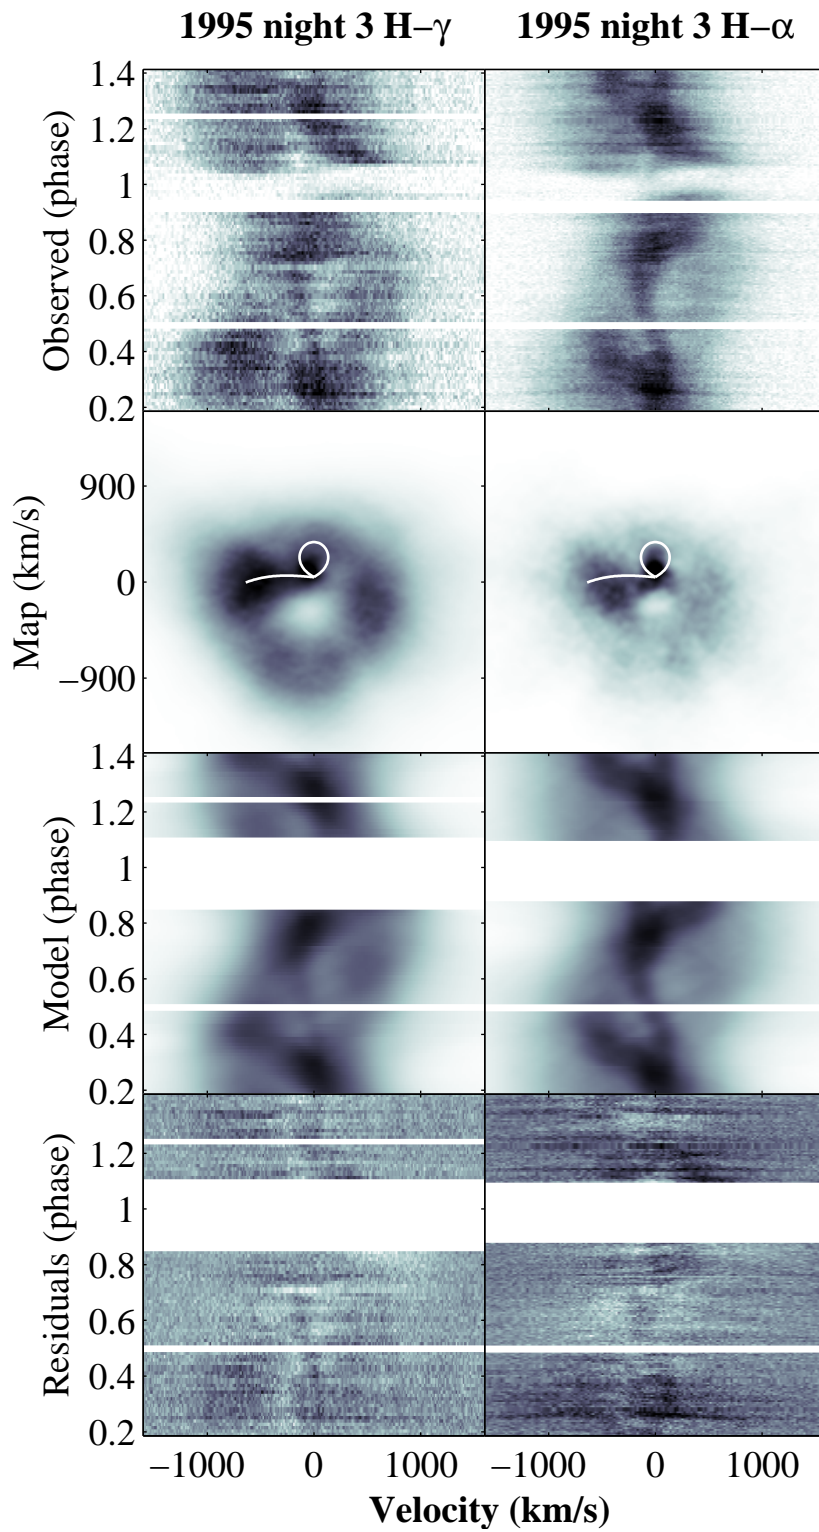


Figure 5.13: Trailed data, doppler maps, computed trail and residuals for the third night in 1995. See text for details.




Figure 5.14: Trailed data, doppler maps, computed trail and residuals for the second night in 1996. See text for details.




Figure 5.15: Trailed data, doppler maps, computed trail and residuals for 2000. See text for details.




Figure 5.16: Light curve of the centre of the $H\alpha$ emission line in the 1995 data.

inner face of the secondary, and is a likely place for stable loops to form. Plasma flows along the magnetic field lines and is illuminated not only by the secondary star but now by the disc which creates the emission source seen in the trails and doppler tomograms.

While the data are not of high spectral resolution, and the Balmer lines do tend to inherently provide less clearly resolved tomograms, it is of note that $H\gamma$ and HeI show a well-defined secondary that is clearly within the bounds of the Roche lobe, whereas $H\alpha$ and $CaII$ (which is expected to come from the secondary) show this extended region in a very pronounced manner. This confirms that the extended emission seen in the latter lines is real and not just an artefact of the Doppler reconstruction. The partial eclipse also suggests that this prominence is out of the binary plane.

5.4.2 Spiral shocks

The 1996 HeI and 2000 $HeII$ maps show non-axisymmetric features with arms of higher intensity visible in the disc. These can be attributed to spiral shocks, a phenomenon that has been seen before in EX Dra [Joergens et al., 2000] and had been predicted in accreting binary systems for some time [Savonije et al., 1994]. The first observational evidence for it is described in Steeghs et al. [1997] for IP Peg. Spiral shocks are excited as a result of the secondary exerting a tidal force on

the outer regions of the disc, creating a spiral pattern at the binary angular speed. However, this is less than the gas speed within the disc, and thus causes shocks to form. This has been looked upon as a method of transferring angular momentum outward through the disc, but is generally considered too weak to be of significance, particularly during quiescence. Large, hot discs are required to generate strong spiral waves which agrees with the observations since the phenomenon has only been seen during outburst. As in Joergens et al. [2000], it is found here that the shocks are most clearly present in HeI and less so in HeII, and also that the upper right spiral (as shown in the tomogram) is higher in intensity than the other.

Another possible source of these features is the distortion of the outer accretion disc by a tidal interaction, giving some of this region a greater vertical extent which is then irradiated by the inner components [Ogilvie, 2002].

5.5 Summary

Using a combination of optical and UV spectra, this chapter showed that strong variability in the out-of-eclipse flux is present in EX Dra. Heavy central absorption present in the optical emission lines and analysis of the wind-driven UV lines suggests that obscuring material along the line of sight could be an explanation for this effect. Previous analysis of obscuring material such as the iron lines in the atmosphere of the outer disc as found by Horne et al. [1994] offers potential physical reasons with a similar type of system, though the actual nature of the variability seen in EX Dra cannot yet be confirmed. Further UV data at slightly higher wavelengths could help in narrowing down whether FeII material is present in EX Dra or not, and while the current data on the wind-driven lines presented shows interesting features, these prove difficult to explain and indicates that the understanding of winds in CVs in general is not complete.

The concept of radial velocity analysis and Doppler tomography was also introduced and applied to EX Dra, noting the possible presence of “slingshot prominences” on the donor star. While prominences are well documented with single stars, the observations displayed in this chapter seem much clearer than those previously documented for binary systems by Steeghs et al. [1996]. There was also evidence for spiral shocks. The phenomenon of spiral shocks has now been reported in several systems only during outburst, a pattern which EX Dra continues, though it is at this stage difficult to determine whether the observations are better matched by an irradiated disc-edge with a vertical extent.

Further investigation into boundary layers and disc obscuration will now be

presented in the next chapter for the CV system Z Cha, this time with Ultracam photometry and X-ray data from the space-based telescope, XMM-Newton.

Chapter 6

Measuring the boundary layer of Z Cha

6.1 Introduction

Z Cha is a dwarf nova which was discovered by Mumford [1969], and the system's eclipses and nature were detailed some years later in Warner [1974]. The most comprehensive study of the system is detailed in Wood et al. [1986], giving parameters of the system obtained through the average of many eclipses. It was found that the mass ratio $q = \frac{M_2}{M_1} = 0.15 \pm 0.003$, the inclination $i = 81.75^\circ \pm 0.15^\circ$, and the disc radius $\frac{R_D}{a} = 0.0195 \pm 0.0004$, where a is the orbital separation. The nature of the central object has been debated several times, with an analysis of the contact points of optical data implying that the lower white dwarf may be obscured in the visual band by the disc [Smak, 1986]. It was shown, however, that these observations could be explained by an extended, optically thick toroidal boundary layer forming during outburst. Modelled white dwarf light curves did seem to confirm the presence of an entire white dwarf in the optical band from data in its low state [Wood, 1987]. During quiescence, the first X-ray observations Z Cha have been shown to agree with the theory of the optically thin boundary layer around the white dwarf [van Teeseling, 1997].

Eclipse timings have revealed that a quadratic ephemeris was needed to predict accurate mid-eclipses for the system [Cook and Warner, 1981], and the possibility of a third body was initially discussed by Wood et al. [1986]. Further work on the ephemeris progressed from a cubic form [van Amerongen et al., 1990] to a sinusoidal one [Baptista et al., 2002] that was attributed to the Applegate effect Applegate [1992], which will be discussed in §6.6.2.

In 2003, another dwarf nova, OY Car, was observed with 51 kiloseconds of XMM-Newton time. Several discrepancies in the X-ray eclipse in comparison to the optical were noted, including an apparently shorter ingress/egress duration and that the eclipse appeared narrower [Wheatley and West, 2003]. Because the ingress and egress lengths depend upon the size of the emitting region, it was concluded that the unusual difference in the eclipse ingress/egress duration was caused by a vertically displaced emitting region. Therefore it was concluded that the X-ray emitting region was in fact from the polar regions of the white dwarf (specifically, the upper region).

In this chapter, the nature of the boundary layer in Z Cha will be probed by use of simultaneous optical and X-ray data provided by XMM-Newton, with the goal of furthering the understanding of the boundary layer and how well it compares to existing models of CVs. Another critical factor which will be investigated is whether the white dwarf is obscured by the accretion disc in observations, an aspect often overlooked in analysis of optical light curves which can have significant impact upon determining system parameters such as the radius of the white dwarf. The comparison of the optical and X-ray eclipses will also be used to determine whether the polar emission proposed for OY Car may be present in this system also. Finally, the new eclipses will be used to compute a new ephemeris for Z Cha, and the remarkable variations in the system's period will be discussed. Many systems are beginning to show similar sinusoidal variations after extensive monitoring (see Richman et al. [1994] for a few classic cases) and it is still a matter of debate as to the underlying physical source for this period altering mechanism.

6.2 Observations

A complete listing of the observations is given in table 6.1.

6.2.1 XMM-Newton X-ray observations

The X-ray data from XMM-Newton consists of observations using the EPIC MOS [Turner et al., 2001] and PN [Strüder et al., 2001] photon counting CCD cameras with a total of 101 kiloseconds time. The observations utilised the thin optical filter to avoid loss of soft X-rays for the object while still blocking the optical and UV flux that the cameras are sensitive to. The data event files were processed using the XMM SAS pipeline using the latest calibration files. Background subtraction was performed for all spectral data, though the extracted light-curves had minimal background flux and thus no correction was performed. In total, 11 full eclipses of Z Cha were recorded. The average X-ray eclipse in soft, medium, and hard X-rays

Table 6.1: Journal of observations.

Instrument (Telescope)	Date	Start (UTC)	End (UTC)	Comments
EPIC (XMM)	2003-12-19	20:45:52		
	2003-12-21		00:56:09	
OM (XMM)	2003-12-19	21:36:05	21:54:25	
		23:22:16	23:45:35	
	2003-12-20	02:23:08	03:20:07	Corrupt telemetry
		08:22:52	08:44:31	
		10:06:57	10:33:36	
		13:38:12	14:06:31	
		15:26:01	15:57:40	
		19:03:59	19:28:58	
		20:40:25	21:22:04	
		22:34:30	23:01:09	
Ultracam (VLT)	2005-05-07	23:58:50	00:34:38	Seeing 0.6", variable transmission
	2005-05-10	23:05:43	00:09:16	Seeing 0.6-1.5", photometric
	2005-05-15	22:56:47	23:55:11	Seeing 0.6-1.2", photometric
	2005-05-17	22:47:11	00:08:36	Poor weather. Seeing 1-2", 40-80% transparency
	2005-05-21	23:45:57	00:41:34	Seeing 0.6-1.5", 80% transparency
Ultraspec (ESO 3.6m)	2008-02-10	05:58:20	07:21:14	Seeing 1.5", average transparency.
Ultracam (NTT)	2010-04-27	00:45:28	01:04:10	Seeing 1.2", photometric.




Figure 6.1: Mean eclipses obtained from XMM-Newton, binned on a local ephemeris. From the top downwards: OM optical B band; soft X-rays (0.1 - 1.0 keV); medium X-rays (1 - 2.5 keV); hard X-rays (2.5 - 10 keV).

can be seen in the bottom three panels of figure 6.1.

6.2.2 XMM-Newton optical observations

XMM also contains the Optical Monitor [Mason et al., 2001] which uses a 30cm telescope to provide photometry simultaneous with the X-ray data. For these observations, the OM telescope was used in ‘fast mode’ to obtain timing data for a small region around the target. The B filter was used, and while 11 eclipses were again recorded, one suffered a timing glitch due a telemetry error and thus was discarded from further analysis. The average optical light-curve can be seen in the top panel of figure 6.1.

6.2.3 Ultracam optical observations

Ultracam (§4.3.1) was used in 2005 on the UT3 VLT telescope at Paranal, Chile, with 0.5 second exposures. Five eclipses were recorded in varying conditions. A further eclipse was recorded in 2010 on the NTT at La Silla, Chile with 1.5 second exposures. To correct for transparency variations a comparison star one arcminute away was used that appears to be a K star with measured magnitudes of $r' = 13.65$, $g' = 14.72$, and $u' = 16.88$. The red star that is present approximately 5 arcseconds from Z Cha [Rayne and Whelan, 1981] was masked from the photometric apertures linked to the position of Z Cha. In all cases, the r' g' u' filters were utilised, and the usual bias and flat field corrections were applied. Colour values were corrected for extinction using a long sequence of exposures on the standard star G 93-48 on a photometric night, and also for the (minor) bandpass differences between the Ultracam SDSS filters and the official SDSS equivalents (detailed in §8). It was noted that the flickering from the bright spot was much higher in the 2005 eclipses when compared to the 2010 data. The Ultracam eclipses can be seen in figure 6.2.

6.2.4 Ultraspec optical observations

A single eclipse of Z Cha in 2008 was acquired using Ultraspec (§4.4) at the ESO 3.6m telescope in La Silla, Chile. The exposure time was 3.6 seconds; seeing conditions varied between 0.8“ and 1.2” but were photometric.

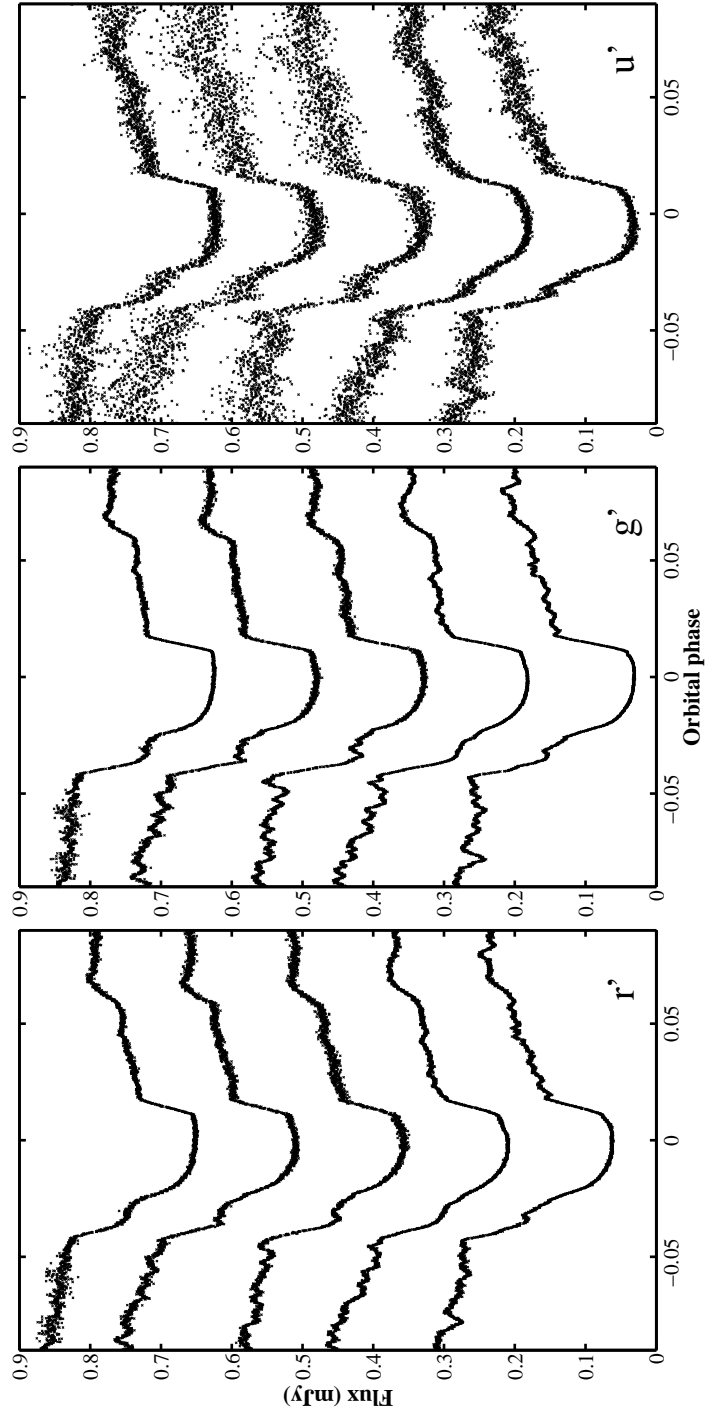


Figure 6.2: Light curves taken with Ultracam at VLT UT3 8.2m in 2005.. From left to right, the filters bands are r' , g' , and u' . Each curve is offset by 0.15mJy from the bottom.

6.3 The boundary layer of Z Cha

6.3.1 Optical eclipse modelling

To determine the nature of the emitting regions in Z Cha, a modelling approach is used to determine the system parameters by fitting the optical light curves, and then using these results for analysis of the X-ray data. Photometric methods described in detail in many places such as Warner [1995a], begin with determining the eclipse width which is a function of q and i . These two parameters are disentangled via the use of the bright spot contact points. Since the stream should only be a function of q , a unique mass ratio and hence inclination can be determined for the light curve via the use of the theoretical calculations of the ballistic gas stream [Lubow and Shu, 1975].

The modelling method used here utilises computer code which treats each of the four main components in a CV system (white dwarf, secondary, disc, and bright spot) as separate collections of flat elements, each one having the properties of area, position, orientation and luminosity. The computed light curve is therefore the summation of all light contributions towards the direction of the observer, taking into account necessary eclipses. Each element, with a vector angle θ towards Earth, has its light contribution altered by a limb-darkened Lambert's law such that

$$I \propto \mu(1 - \eta + \eta\mu), \quad (6.1)$$

where $\mu = \cos \theta$, and η is the limb darkening coefficient. Finite length exposures are accounted for by computing multiple light contributions during the exposure time, and are averaged together using a trapezoid weighting scheme.

The white dwarf is modelled as a linearly limb darkened sphere, whereas the secondary is Roche distorted. The disc is symmetric, with a height specified by

$$h(R) = h_0 R^\alpha, \quad (6.2)$$

where h_0 and α are fit parameters, as well as the disc radii. However, for this case $h_0 = 0.02$ and $\alpha = 1.5$ since these parameters were difficult to constrain in the data given it was quiescent observations and the disc is always expected to be thin. The surface brightness of the disc is determined by a power law in radius, and the overall level defined by an outer disc temperature. This brightness is translated to a light contribution by assuming a blackbody spectrum at a central wavelength of a filter specified. The bright spot is modelled as a series of elements along a straight line

in the orbital plane, with the brightness specified by

$$S \propto \left(\frac{x}{l}\right)^\beta \exp\left[-\left(\frac{x}{l}\right)^\gamma\right], \quad (6.3)$$

where x is the distance along the line defining the bright spot, β and γ are power law exponents, and l is the scale length. The maximum brightness is at the point

$$x = l \left(\frac{\beta}{\gamma}\right)^{\frac{1}{\gamma}}, \quad (6.4)$$

which is positioned on the ballistic gas stream at R_{BS} , the distance away from the white dwarf, which is also a fit parameter. The gas stream calculations are as per Lubow and Shu [1975], and since γ was causing degeneracies in the fitting procedures and was not required for a good bright spot fit in the data, it was fixed at 1. Limb darkening is accounted for by using fixed values expected for a white dwarf of its temperature [Robinson et al., 1995]. All components were initially fitted with a low resolution grid and a Nelder-Mead “simplex” algorithm before moving on to a Levenberg-Marquardt minimisation with a high resolution grid. The number of grid elements was increased until the chi-squared stabilised. Other parameters and discussion about the code is supplied in the appendix of Copperwheat et al. [2010].

Providing the model is sufficient to give no major discrepancies in the fit, it can be argued that such a technique is inherently more accurate and less biased than attempting to determine the contact points of the white dwarf and bright spot via numerical derivatives (as in Wood et al. [1986]). However, since the XMM-Newton OM instrument has an aperture of only 30cm, the set of eclipses obtained with Ultracam (using the 8.2m VLT) that have high signal-to-noise were used for the fitting method. Subsequently, the mass ratio q , system inclination i , and white dwarf radius R_{WD} were fixed for the later XMM-Newton OM models.

With q and i as free parameters, the fit results in values that are inconsistent with those determined by Wood et al. [1986] by a factor of ~ 10 the quoted uncertainty (± 0.0035). This is explained further in §6.5, where it is argued that this was not likely a real change, but is caused by the differing analytical methods and flickering perturbing q and i which are extremely correlated in the modelling methods. Over a limited range, a decrease in q with a respective increase in i results in a near identical fit. As the several eclipses obtained showed significant night-to-night differences with variable weather conditions, it was opted to choose q and i to be those determined by Wood et al. [1986], as it was felt that the greater number of eclipses used in the paper gave a more accurate depiction of mean state of Z Cha.

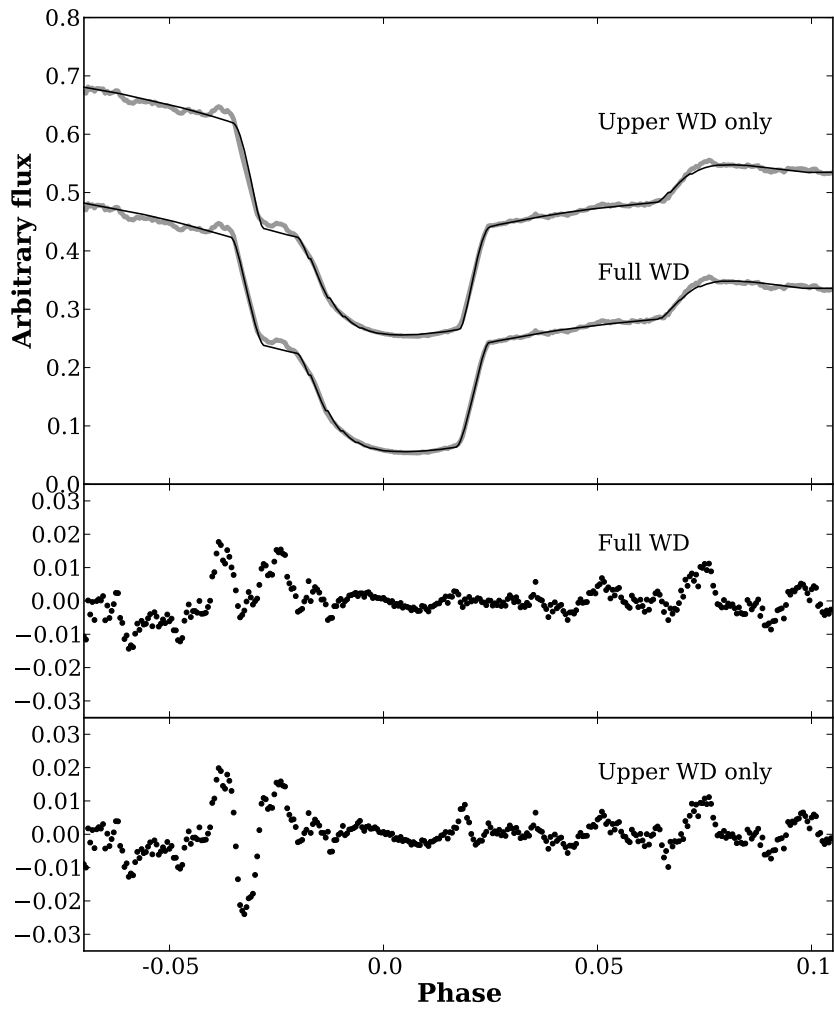


Figure 6.3: The average Ultracam g' band eclipse with fit to entire white dwarf, and a white dwarf with an obscured lower hemisphere, with residuals beneath.

The fitting proceeded by fixing these values and then refitting the other parameters free as before. This gives us a master set of parameters to continue with using a combination of the mass ratio and inclination from Wood et al. [1986] and the white dwarf radius, disc, and bright spot parameters from the g' band Ultracam data. This band was chosen because it had the highest signal-to-noise and also seemed to be least affected by any flickering effects. The r' and g' band fits produced values of q , i , and disc radius to within their margin of error. The u' band data values showed some divergence from the other two bands, likely due to the significantly worse signal-to-noise, but also potentially due to the fact the integration times were three times longer in this channel. The white dwarf radius is found to be $0.0164a$ (where a is the orbital separation), and the bright spot is found to be at $0.32a$ on a disc extending to $0.38a$. These parameters are summarised in table 6.2.

Theoretical errors in the fitted components from the Levenberg-Marquadt algorithm were noted to be unrealistically low, so the final fit uncertainties were computed using a Markov-chain Monte Carlo (MCMC) process. A Markov chain is a system that undergoes transitions between a finite series of states in a chain-like manner, where each transition depends only upon the current state and not past ones. In this context, each state is a set of fit parameters that generates a light curve, and an ensemble of 'random walkers' is used to explore the multi-dimensional space. Whereas a traditional Monte Carlo process will evaluate the space uniformly, steps in an MCMC process will be correlated. After each trial, if the χ^2 is found to be lower, then it is accepted and it becomes the current step. If this is not the case, there is still some probability of the new trial becoming an accepted state. Parameters of the MCMC are tuned to achieve an acceptance rate of $\sim 25\%$, with adequate mixing so that there are no fit parameters that end up stuck in local minima. Ultimately there will be a convergence of the fit parameters, though how long this will take is difficult to determine, particularly with dozens of parameters. In the process used here, a million iterations were performed in order to sample the probability distribution to achieve a more realistic range of the error.

Subsequently this model was used to fit the average OM light curve, keeping the mass ratio, inclination, white dwarf radius and limb darkening fixed. The disc and bright spot parameters were started from the Ultracam results but allowed to be free; we expect there to be some difference here as there was a 2 year gap between the two sets of data. This final optical fit from the average OM data then allows for direct comparison with the X-ray data, and was also used as a basis to fit the central eclipse times in each individual OM eclipse, which is discussed in §6.6.1. This fit achieved a reduced chi-squared of 1.11.

There remains one important issue to be addressed if we are to consider the optical parameters reliable enough to provide a model for the X-ray data; that is, whether the central object we see is truly an entire white dwarf or if it is partially obscured by the accretion disc. The latter scenario was suggested by Smak [1986] from an analysis of the white dwarf contact points in optical data. However, it was suggested in the same paper that detailed analysis of the light curves would give a more satisfactory result, which is now investigated with the Ultracam optical model. In the case of a white dwarf emitting from the upper hemisphere only, it is expected that the radius of the white dwarf would have to increase in an attempt to match the ingress and egress, which is indeed what is found. Our previous white dwarf radius of $0.0164a$ increases to $0.018a$. Importantly, we also see an increased χ^2 from 16 to nearly 40 (see figure 6.3), caused by significant residuals mostly in the ingress (it should be noted the difference between the two white dwarf models are completely negligible in the OM data due to the lack of S/N with the 30cm telescope). The difficulty with this scenario is presented in Wood [1987] which shows synthetic white dwarf eclipses for various emitting regimes, including a distinctly asymmetric eclipse in the case of an obscured lower hemisphere (see Fig. 1 (b)). A further test is performed by estimating the secondary star radial velocity K_2 from the data using

$$\frac{PK_2}{2\pi G} = \frac{M_1 \sin^3 i}{(1+q)^2}, \quad (6.5)$$

where M_1 is inferred from the Eggleton mass-radius relation [Verbunt and Rappaport, 1988], including a small oversize factor of $\sim 3\%$ due to the offset from zero temperature which is then checked against synthetic white dwarf models described in Holberg and Bergeron [2006], adjusting to match. In Wade and Horne [1988], a detailed study is performed to determine $K_2 = 430 \pm 15$ km/s. Given the white dwarf radius found for a uniformly emitting white dwarf ($M_1 = 0.90 \pm 0.03M_\odot$), $K_2 = 415 \pm 15$ km/s is consistent with this. With only the upper hemisphere emitting however ($M_1 = 0.72 \pm 0.03M_\odot$), the radial velocity amplitude decreases resulting in $K_2 = 380 \pm 15$ km/s. The underlying reason for this difference is the fitted white dwarf radius, which is used in turn to compute M_1 . Hence, given that the fit residuals indicate that the χ^2 more than doubles when using only an upper white dwarf, and also that the computed K_2 becomes inconsistent with previous studies when using the white dwarf radius determined from such a scenario, the X-ray fitting proceeded with confidence that the entire white dwarf is seen in the optical data.

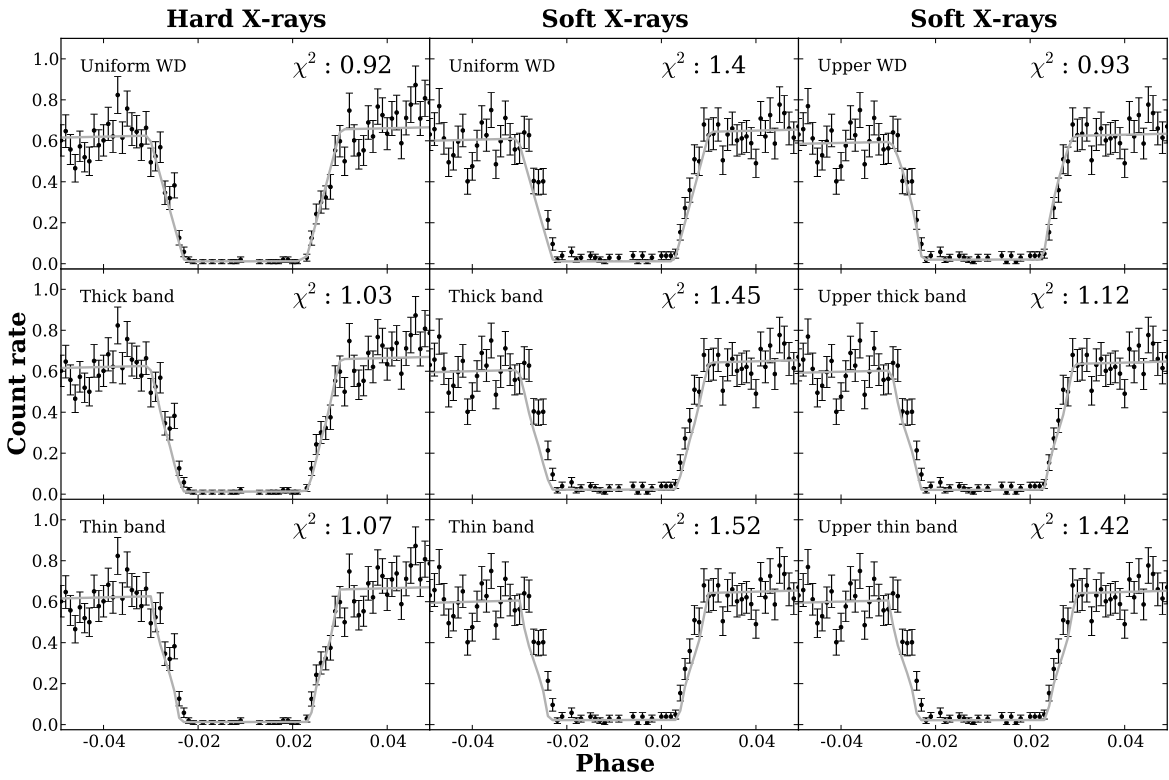


Figure 6.4: Best fit of different white dwarf X-ray models to data. The left panels show hard X-ray data, the middle and right are soft X-rays, with the right also obscuring the lower hemisphere of the white dwarf. From top to bottom for each column are: the entire white dwarf emitting; a thick equatorial band of $0.5R_{WD}$ emitting; and a thin equatorial band of $0.1R_{WD}$ emitting. Note that the phase range -0.01 to 0.01 has been omitted for clarity.

Parameter	Value	Source
q	0.149 ± 0.003	Wood et al. [1986]
i ($^\circ$)	81.3 ± 0.3	Wood et al. [1986]
R_{WD}/a	0.0164 ± 0.0004	Ultracam
R_D/a	0.359 ± 0.010	Ultracam
R_{BS}/a	0.317 ± 0.002	Ultracam

Table 6.2: Table of parameters used in final models for eclipse comparison purposes.

6.3.2 X-ray eclipse modelling

The average light-curve from the PN and MOS data was extracted by mapping each event directly to a phase and then extracting in bins of phase. This avoids any potential issues introduced by the process of re-binning data. Due to the continually changing nature of Z Cha’s ephemeris, a local ephemeris was computed from the OM optical data for the phase mapping.

The final OM fit parameters were then used to produce a model consisting only of a central object (ie. the disc and bright spot were disabled). Also, since it is expected that the X-rays come from an optically thin region near the white dwarf, limb darkening as well as the modelling of Lambert’s law were removed from the central object. This law is a computation of the projection effect (such that area elements nearer the limb of the star show less flux towards the line of sight) and when disabled the central object effectively becomes a limb-brightened disc with a radius matching that of the white dwarf in the optical data. The only parameter allowed to vary is a slope in the out-of-eclipse flux, which is artificial, but approximates the local gradient in the X-ray flux. This ensures that the majority of the chi-squared will originate from the ingress/egress and their associated contact points.

From the observations, there appears to be significantly different amounts of absorption between the hard and soft X-rays as seen from the XMM light curves (figure 6.1) and also from the spectral analysis performed in §6.3.3. For this reason the 0.1-1.0keV (soft) regime was extracted separately from the 1.0-10keV (hard) regime. This split achieves roughly equal count rates in each band.

The hard X-rays

The initial model with the limb-brightened disc and invariant radius resulted in a good fit to the hard X-ray data with $\chi^2 = 0.92$. This indicates that the entire white dwarf is seen in hard X-rays as with the optical case. To investigate whether the emission area of these hard X-rays was restricted to specific areas upon the




Figure 6.5: Diagram depicting the “vertical emission extent” parameter modelled, which dictates the amount of the object, in latitude, above and below the equator that will be emitting flux.

white dwarf surface, the model was modified to allow different emitting regimes and refitted. The ‘vertical emission extent’ parameter defines how far out (in latitude) from the equator of the white dwarf the emission region is modelled in terms of the radius, which is explained visually in figure 6.5.

In figure 6.4, the emission extent shown for the entire white dwarf, a thick equatorial band that is half of the white dwarf (from the middle of the lower hemisphere to the middle of the top hemisphere; meaning an extent of 0.5) and a thin equatorial band (an extent parameter of 0.1). A steady increase in χ^2 is found as the model moves from the whole white dwarf emitting to a thinner band (0.92 to 1.07). Stepping these models in fractions of 0.05 of the vertical emission extent, the χ^2 rises steadily with a smaller emitting fraction, as plotted in figure 6.6. It is found that, to 90% confidence, the hard X-ray emission comes from an extent of greater than 95% of the white dwarf radius.

The soft X-rays

Starting again with the entire white dwarf emitting (the initial assumption as with the hard X-rays), it is very clear that the limb-brightened disc does not fit the data; achieving a χ^2 of 1.4 (the middle panels of figure 6.4). The light curve appears to fall less rapidly in ingress and rise more rapidly in egress than the model performs. This is similar as to what was found by Wheatley and West [2003] with OY Car, and was attributed to emission from the polar caps of the white dwarf. Again performing a test with a thin and thick equatorial band finds that neither fits in a satisfactory manner. However, by restricting the emission region to the upper white dwarf only (right-hand series of panels, figure 6.4), it is found that the χ^2 reduces to 0.93. Finally, the test of thin and thick equatorial *upper* bands is performed, once again finding that the χ^2 grows worse with a decreasing vertical emission extent.

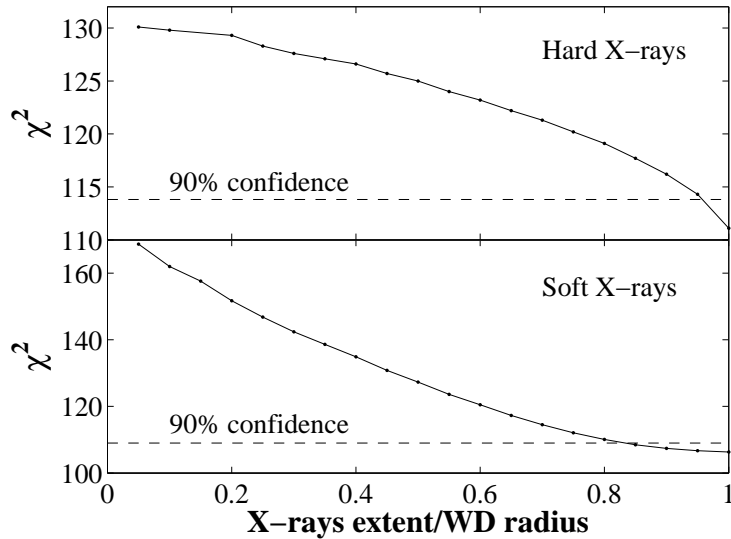


Figure 6.6: Plot of white dwarf emission height (as from figure 6.5) against χ^2 for hard and soft X-rays. The dashed line is the 90% confidence margin.

Stepping the size of this upper band in a similar fashion to the hard model, a plot of χ^2 against upper emission extent is presented in the bottom panel of figure 6.6. It is concluded that the emission extent in soft X-rays (on the upper hemisphere) is greater than 85% of the white dwarf radius with 90% confidence.

The scale height of the emitting material

The hard X-ray data shows a thin emission shell coating the white dwarf. Next, it was examined whether the scale of this emitting material could be determined. Since the evidence suggests that the whole white dwarf region is emitting in the X-ray band, a scale height of an emitting atmosphere around the white dwarf allows exploration of whether this region is actually larger than just the white dwarf itself. Given the assumptions of an optically thin boundary layer, the model is adjusted to better approximate an optically thin shell with a finite thickness. To this point, the model simply approximates the effect of a thin shell by removing Lambert's law, which results in an effective limb-brightening. While such a configuration could be altered with the limb-darkening parameters, the previous model has no shell thickness; it is merely approximating the desired appearance. To determine an actual scale height of such an object, a more physical model is required. The

impact parameter, \mathbf{p} , of a point in the shell is defined as

$$\mathbf{p} = \sqrt{r_1^2 - (\mathbf{r} \cdot \mathbf{E})^2}, \quad (6.6)$$

where r_1 is the radius of the white dwarf, \mathbf{E} is the vector towards Earth, and \mathbf{r} is the current position vector. By using figure 6.7, and the chord rule $(AO)(OB) = (CO)(OD)$, each grid element on the central object has its flux F_0 altered by one of the following equations to give the final element flux F

$$F(p) = F_0 \left\{ \begin{array}{ll} \frac{\sqrt{(r_1 + H)^2 - p^2} - \sqrt{r_1^2 - p^2}}{2\sqrt{(r_1 + H)^2 - p^2}} & p \leq r_1 \\ & r_1 < p < r_1 + H \end{array} \right\}, \quad (6.7)$$

where H is the material scale height. In this extended shell, $(AO)(OB) = (R - P)(R + P)$, where $R = R_1 + H$, which results in $(AO)^2 = (R_1 + H)^2 - p^2$. Since when beyond the edge of the white dwarf we are actually seeing $(AO)(OB)$, a factor of 2 must be included in this case. When $R_1 < H$, we do not see this, and the extra $\sqrt{R_1^2 - P^2}$ comes from the consideration that we are only looking at an area element P away from the white dwarf centre. A quick verification can be performed, because in the case of $H = 0$, the $P \leq R_1$ factor becomes zero. This gives a function that produces least flux at the centre and then increases rapidly as it reaches the outer limb before increasing by a factor of 2 just outside the edge of the white dwarf (where one sees twice the shell material). This then drops off sharply in accordance with the H parameter.

Due to the sharp decrease in flux at the edge of the boundary layer and the discontinuity between the two regimes, the number of vertices in the model was increased substantially. It was determined that $\sim 800,000$ grid points representing the white dwarf were sufficient for the chi-squared to stabilise and hence reach a reliable representation of the model. The scale height has a very shallow gradient in chi-squared space, so instead of fitting, the models were manually stepped in terms of H , while fixing the radius of the white dwarf to be the value from the optical data as before. The result for both regimes can be seen in figure 6.8. The best scale heights are found to be $H_{hard} = 0.0018a = 1.76 \times 10^8$ cm and $H_{soft} = 0.0004a = 3.92 \times 10^7$ cm. The scale height in the soft band is determined to be smaller than $0.00016a = 1.57 \times 10^8$ cm with 90% confidence and the hard band scale height is between $0.0003a - 0.003a = 2.94 \times 10^7 - 2.94 \times 10^8$ cm, also with 90% confidence.

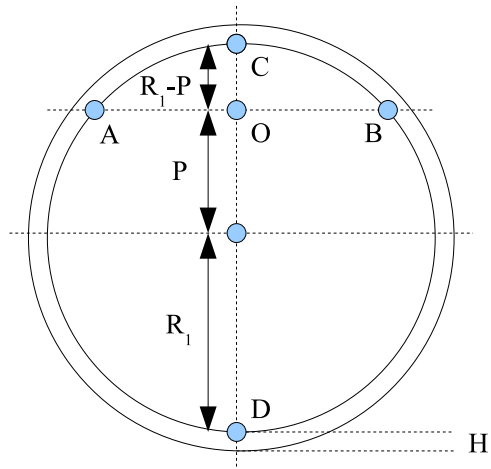


Figure 6.7: Example of the chord rule in terms of parameters used to determine the scale height of emitting material in a white dwarf. In this figure, $(AO)(OB) = (CO)(OD)$.

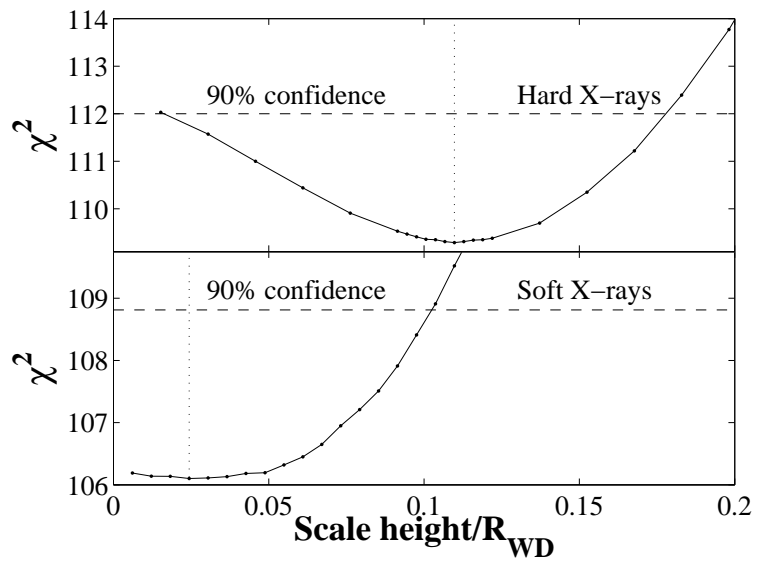


Figure 6.8: Plot of boundary layer scale height against χ^2 for hard and soft X-rays. The dashed line is the 90% confidence margin. The dotted line is lowest χ^2 found by stepping.

6.3.3 The X-ray spectrum

To determine the origin of the orbital modulations seen in figure 6.1, ten phase bins of the XMM-Newton data were extracted to model using the XSPEC software. Since there were clear signs of absorption in the light-curve, a partial covering fraction absorber was included (accounting for intervening disc material), and then a *mekal* model [Liedahl et al., 1995] at two separate temperatures, which also incorporates plasma density and abundance (relative to solar quantities). While it is clear that the system will in fact be found at a range of temperatures, this is approximated by two temperatures in an attempt to model the hot and cold material together. These *mekal* parameters were free, but forced to fit over all phase bins simultaneously as it was not expected that there would be phase dependence with temperature and abundance. The absorber parameters were fitted with each bin, representing a changing level of obscuration along the line of sight.

The resulting fits are seen in figure 6.9, with a final reduced chi-squared of 0.95 for the 3594 total spectral bins. There is a clear distinction between the soft and hard X-rays, where there is evident absorption present for energies lower than 1keV; this matches with the discovery in the previous section that there is a considerable amount of absorption of the soft X-rays. The *mekal* temperatures were determined to be $kT_1 = 8.05 \pm 0.275\text{keV}$ and $kT_2 = 1.52 \pm 0.12\text{keV}$, whereas the abundance was $A = 0.696 \pm 0.03$, where all uncertainties are for a 68% confidence margin. The variation of the fit parameters with orbital phase can be seen in figure 6.10, where the hydrogen column effectively remains fixed at $0.1 \times 10^{22}\text{cm}^{-2}$ except for a peak at phase 0.8 to $0.7 \times 10^{22}\text{cm}^{-2}$. The covering fraction shows a double-hump profile with peaks at 0.95 at phases 0.4 and 0.8. The normalisation shows a slightly lower value around phase 0 as expected, but otherwise maintains a flat and featureless profile.

6.3.4 The area of the X-ray emission upon the white dwarf

For constraining the emission region of the white dwarf in CV systems, the brighter sources such as HT Cas were originally targeted, with Wood et al. [1995] finding that there was indeed an eclipse of the X-rays, and Mukai et al. [1997] providing a constraint upon the emitting area of up to a maximum of 15% larger than the white dwarf radius. These measurements, along with the determination that hard X-rays were being emitted during the low state gave significant support to the general view of the boundary layer, though the observations of OY Car during superoutburst raised questions due to the complete lack of an X-ray eclipse [Naylor




Figure 6.9: Fits to each phase bin of X-ray spectral data.




Figure 6.10: Fitted parameters from the X-ray model against phase. The range in phase is repeated for clarity.

et al., 1988], which was attributed to the state of the system. The optical data presented here clearly indicates Z Cha was in a low state which is confirmed by AAVSO observations, hence it is expected that the object observed will be quiescent with an optically thin boundary layer consisting of bremsstrahlung emission in the hard X-ray wavelengths.

The modelling process clearly shows that the best fit to observations consists of the entire surface of the white dwarf emitting, not an equatorial band, but that the lower hemisphere of the white dwarf is obscured at soft X-ray energies. This agrees with the spectral modelling (§6.3.3), which concludes that the covering fraction of the observed X-rays is always above 0.6; any lower than 0.5 and this would be hard to explain given the assumption that nearly half of the boundary layer is obscured. The column density of 10^{21}cm^{-2} has been shown to be sufficient to obscure soft X-rays [Naylor and La Dous, 1997], and so all the evidence is consistent with the accretion disc absorbing the soft X-rays from the lower half of the primary. The conclusions put forth in Wheatley and West [2003], which called for X-ray emission from the polar caps, could in fact be mis-interpreted obscuration of the lower hemisphere of the white dwarf. This configuration will produce a similar effect in the timing of the contact points, but a different shape of the ingress and egress. Given current understanding of the boundary layer, obscuration of the lower white dwarf seems a more logical and physical explanation.

6.3.5 The scale height above the white dwarf

The eclipse of Z Cha itself in X-rays was reported by van Teeseling [1997] providing a constraint of $R_{BL} < 1.4R_{WD}$. While it has already discussed that the data presented shows convincing evidence that the entire surface of the white dwarf is emitting, this still leaves the possibility that it may come from an area larger than the white dwarf. The modelling of a scale height above the primary was performed, which constrained $R_{BL} \leq 1.18R_{WD}$ at a 90% confidence margin for hard X-rays. The result from the soft X-ray emission also agrees with the general picture of the boundary layer; it should come from a cooler temperature region and hence have a lower scale height. A rough approximation of the atmospheric scale height, H , is given by

$$H = \frac{kT}{mg} \quad (6.8)$$

Where $T = 8.05\text{keV}$ from the spectral fitting, $g = 10^6\text{m/s}^2$ (from GM_{WD}/R_{WD}) and m is the average particle mass which is taken as half the mass of a proton (from ionisation). The boundary layer atmosphere is expected to have an exponential falloff with bremsstrahlung emission, in which the intensity is proportional to the density squared. This means that

$$I \propto \left(e^{-\frac{kT}{mg}} \right)^2 \equiv e^{-2\frac{kT}{mg}} \quad (6.9)$$

Since the model used in the X-ray fitting uses the ‘emission scale height’ rather than the ‘density scale height’, all compared values must be done to $0.5H = kT/2mg$. Taking this into account, the scale height is found to be $1.54 \times 10^8\text{cm}$ ($0.094R_{WD}$). The best fit for the hard X-rays indicates $0.11R_{WD}$, which shows good agreement. Taking the upper limit of the temperature from the spectral fit with the 68% confidence margin results in $0.10R_{WD}$. It is also important to investigate the effect that changing R_{WD} has upon these results. When using $R_{WD} = 0.0151a$ (from the alternate Ultracam model discussed in §6.3.1), a smaller scale height of $0.076R_{WD}$ is found for hard X-rays, which in that case means the fit is slightly smaller than the calculated value (though both results are consistent).

6.3.6 The truncated inner disc

The fitted line-of-sight column density (cm^{-3}) allows the calculation of the surface density of the inner disc of Z Cha. We find this to be very low 10^{-3}g cm^{-3} (using solar abundances and masses), and may even be an overestimate given the model used. There are good reasons for believing this fit; while the covering fraction in

the model dictates how strongly the soft X-rays are absorbed, the column density has the property of dictating at how high an energy the ‘turn-off’ point of this absorption occurs, as higher energy photons will require a greater column density to be affected. It is found that the energy at which the detected photons become clearly absorbed is at the 1keV level; the chosen boundary between the hard and soft X-ray regimes used. Increasing the column density would push this value to higher energies, and the X-ray modelling would then indicate an obscured white dwarf in the hard X-ray band, which is inconsistent with the rest of the observations. Since the whole lower half of the white dwarf has disc material of this column density in front of it, an estimate of the minimum radial extent of the disc can be made from

$$R_{in} = \frac{R_{WD}}{\tan\left(\frac{\pi}{2} - i\right)} = 6.9R_{WD}. \quad (6.10)$$

See figure 6.11 for a diagram showing this. Typical values from disc instability models have surface densities, Σ , on the order of 10-100 g/cm⁻³ [Warner, 1995a; Lasota, 2001]. Since our Σ is 4-5 magnitudes lower than the standard disc models out to $\sim 7R_{WD}$, it is reasonable to conclude that this part inner disc may be truncated, though the disc outside this region may be of a more typical surface density.

Many arguments have been put forward for the latter case to explain a variety of issues, such as an observed delay between the optical and UV outburst light-curves [Meyer and Meyer-Hofmeister, 1989] and longer than expected recurrence times (in systems such as WZ Sge, see Hameury et al. [1997]). Livio and Pringle [1992] suggest that during quiescence, the inner accretion disc is evacuated by the presence of a magnetic field on the white dwarf of the order $\sim 10^4$ G, which can push material outwards if the white dwarf is rapidly spinning, except during outburst where the effect can be overcome. King [1997] argues that the irradiation of the disc by central component ionises the inner disc, shifting into a high state and thus accounting for the low surface density. Finally there has also been the model of a coronal siphon flow [Meyer and Meyer-Hofmeister, 1994] which calls for the evaporation of the inner part of the disc due to the interaction between the cool disc and the hot corona.

One other reason many of these arguments are brought into play is that standard disc instability models have predicted quiescent accretion rates far below what has been observed from the inner disc with X-ray instruments. For example, with SS Cyg [Wheatley et al., 2003], and also Z Cha itself by van Teeseling [1997]. A quiescent accretion rate of $\dot{M} > 10^{-12}M_{\odot}/\text{yr}$ was found, with a similar number determined by Wood et al. [1986] from optical data. Integrating over the observed X-ray wavelengths, a peak (absorbed) flux of 6.6×10^{-13} ergs/cm²/s is found. This




Figure 6.11: The estimation of the radius of disc seen through that is obscuring the lower white dwarf at an observed inclination angle.

value is extrapolated to an unabsorbed bolometric luminosity at a distance of 97pc [Wood et al., 1986], and hence an estimate of the inner disc quiescent accretion is found from

$$\dot{M} = \frac{2L_{BL}R_{WD}}{GM_{WD}}. \quad (6.11)$$

This results in $\dot{M} = 3.40 \times 10^{13} \text{g/s} = 5.37 \times 10^{-13} M_{\odot}/\text{yr}$. This is an order of magnitude less than the result previously determined from van Teeseling [1997], and indicates that this measurement of the inner disc accretion rate is not anomalous. Since the disc instability model requires the surface density of the quiescent disc to be less than the critical surface density at all points, it must satisfy the following [Lasota et al., 1995]

$$\dot{M}(r) \leq 4.1 \times 10^{14} t_6^{-1} r_9^{3.11} M_1^{-0.37} \left(\frac{\alpha}{0.01} \right)^{-0.79}, \text{g s}^{-1} \quad (6.12)$$

with t_6 being the outburst recurrence time in 10^6 s, r_9 being the radius in 10^9 cm, M_1 being the white dwarf mass in solar masses, and α being the viscosity parameter. Using a recurrence time of 51d [Warner, 1995b] and the model radius with its related mass (§6.3.1), the accretion rate is found to be $\dot{M} \leq 1.14 \times 10^{14} \text{g s}^{-1} = 2.21 \times 10^{-12} M_{\odot}/\text{yr}$. This confirms that these observations of Z Cha suggest an accretion rate in line with the standard disc instability model. However, it still leaves Σ that is orders of magnitude less than such models predict, particularly at $7R_{WD}$.

6.4 The second absorption site

Figure 6.1 shows two notable absorption dips that weaken towards higher energies. Given their location in phase, the one prior to eclipse is almost certainly the disc-

stream impact as this can be seen to be most prominent where the bright spot flux is rising in the optical band, and is at the expected phase 0.8. The second dip seen is weaker and with less width in phase. It has been shown by Lubow [1989] that some of the stream could pass over the disc and impact on the other side at a phase of ~ 0.5 , though in this case it appears to be at a phase of ~ 0.3 . The alternative view is that the disc possesses an extended vertical height at phase 0.3 which is partially absorbing the X-rays. This is a particularly interesting solution as it has already been shown that the hard X-rays pass through the disc material but the soft X-rays do not. Another result to contribute evidence to this theory is that while the impact point shows an increase in covering fraction and column density, the smaller impact point consists only of an increase in covering fraction by the model fit. If both were attributable to disc-stream impacts, then a similar rise in column density at the second point would be anticipated. The source of such phase-dependent vertical thickening has been attributed to tidal effect [Ogilvie, 2002].

6.5 The system parameters

Performing the fit as described in §6.3.1 on the average Ultracam eclipse with q and i as free parameters resulted in $q = 0.18$ and $i = 80.4^\circ$, which are inconsistent with the previous analysis performed [Wood et al., 1986]. Despite the apparent good fit with very high signal-to-noise data, such a change must be examined carefully since these values should not be variable. There are two important differences between our optical light curve analysis and that performed by Wood et al. [1986]; firstly there is the number of eclipses. We have 5 compared to 29, although ours are from much larger telescopes. A large number of eclipses is still preferable when faced with such a deviation however, since there is a clear difference between night-to-night eclipses, particularly at these levels of S/N. There is a notable quantity of flickering present in the light curves that would be made negligible by a larger stack of eclipses to average, and also the actual shape of the eclipse is varying a significant amount from night-to-night. This is potentially caused by the disc, as some nights show a sharper ingress/egress, whereas others show a smoother transition typically associated with a stronger disc. These factors present significant difficulties when modelling the averaged light curve. The bright spot and sharpness of the contact points, which prove challenging to fit due to the factors mentioned, provide strong constraints upon q in the model. In effect, it is not believed that there is a sufficiently accurate light-curve of the actual average of Z Cha in quiescence. The second difference is the method itself; the technique presented here utilises a model

based upon physical parameters and then simply tries and optimise the chi-squared. Wood et al. [1986] use a numerical derivative method to determine the contact points. While a limb darkening parameter is added in the model code used here (based upon the white dwarf temperature), Wood et al. [1986] neglect this entirely, but this is unlikely to explain the differences found. If the individual eclipses are fitted alone with the resulting values of q and i plotted against one another, a clear straight line can be fitted through them, indicating strong correlation (individually fitted mass ratios are found between 0.162 and 0.185). This is not surprising, as both parameters are effectively working on the contact points in an opposing manner. For this reason, it is concluded that the uncertainty on q and i is more than that quoted in Wood et al. [1986]. More importantly, the q and i parameters of Wood et al. [1986] lie clearly on the same q - i line as the data presented here. Hence, the distribution of parameters in the individual fits comes from the variation in the eclipse shape (likely the disc) and flickering from night to night, and the general offset towards higher q may be from a bias in one method or the other; of which the only clear conclusion can come from analysing the other set of data with the same method. Despite this offset, the greater number of eclipses taken by Wood et al. [1986] presents a better measure of the average eclipse light curve of Z Cha (particularly given the night-to-night variability observed) and was the reason that these q/i parameters were chosen and fixed in the modelling performed here. Finally, if either set of q and i is picked, the difference on the white dwarf radius is $0.0013a$ (resulting in $R_{WD} = 0.0151a$ for the data presented here with all parameters free). This produces a slightly smaller scale height parameter than with the q/i parameters from Wood et al. [1986] but it remains consistent, and makes no difference upon the conclusion regarding the vertical emission scale height on the white dwarf.

6.6 The eclipse times

6.6.1 The sinusoidal ephemeris

For each of the optical eclipses obtained, the same fit was performed as detailed in §6.3.1, using the model from the average optical light curve and allowing the T_0 to vary until the chi-squared was lowest. Given a number of central eclipse times, an ephemeris can be calculated, consisting of an HJD equating to an arbitrary ‘cycle zero’ and a period. This can then be extended to quadratic, cubic, or even sinusoidal form. Such an ephemeris is typically analysed in the form of an O-C plot, which plots the difference in time between the observed eclipse time and the calculated eclipse time against the cycle number. Utilising the previous timings

listed in van Amerongen et al. [1990] along with Robinson et al. [1995] and Baptista et al. [2002], the 7 Ultracam optical eclipses, 10 XMM OM eclipses, plus a further two from Ultraspex, a new ephemeris was fitted with a least squares method. A table of the new measured white dwarf mid-eclipse times can be seen in table 6.3. The difference between the observed and calculated T_0 points was fitted using a sinusoidal ephemeris. This choice is motivated by the fact it was determined by Baptista et al. [2002] that the sinusoidal component was relevant in an F-test to 99.9% confidence whereas a quadratic ephemeris was not. The starting parameters for the fit were taken from the same paper. Each epoch was weighted equally (that is, sequential eclipses were downweighted by their quantity); this is because the residuals far outweigh the uncertainty on each data point. The new O-C plot can be seen in figure 6.12. The best fit ephemeris is

$$\begin{aligned}
 JD &= T_0 + PE + \alpha \cos\left(\frac{2\pi}{\gamma}(E - \beta)\right) & (6.13) \\
 T_0 &= 2440264.68323(11) \text{ days} \\
 P &= 0.0744992829(13) \text{ days} \\
 \alpha &= 0.001046(65) \text{ days} \\
 \beta &= 12.79(20) \times 10^3 \text{ cycles} \\
 \gamma &= 18.36(78) \times 10^3 \text{ cycles}
 \end{aligned}$$

This form allows for a cyclical variation with an effective amplitude α , with β controlling the phase offset and γ the frequency. An error of 5×10^{-5} d was assigned to previously quoted values without errors. The mid-eclipse times from the new fits have an error of the order 2s. While it remains to be seen with certainty whether this sinusoidal motion will be truly periodic, following on from previous suggestions regarding a third body in the system [Wood et al., 1986; Dai et al., 2009], the best fit to such a case results in mass $m_3 \sin i_3 = 19.9 \pm 2.1 M_J$, period $P_3 = 56.35 \pm 1.37$ years, and eccentricity $e = 0.17 \pm 0.03$. Uncertainties were obtained with a Monte Carlo method.

6.6.2 The origin of the eclipse time variations

The computed O-C plot shows the considerable variation in Z Cha's period, which, from peak to peak is currently in excess of three minutes. A third body seems a feasible reason for this observed variation, however caution must be advised on any presumptions made from current data as not a whole sinusoid has yet been recorded, which means that this fit is also not yet stable. Even the addition of the single

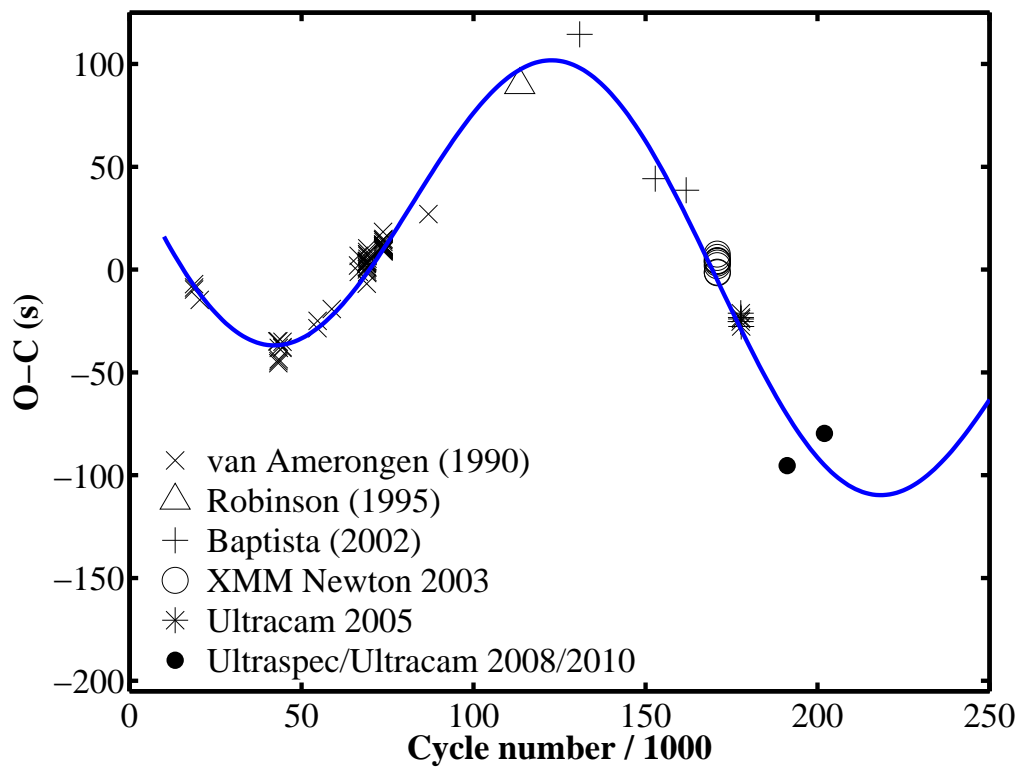


Figure 6.12: The new ephemeris O-C, generated from previous data, ten XMM OM eclipses, six Ultracam eclipses, and an Ultraspes eclipse.

Eclipse time (BJMD)	Source	Telescope
52992.907348(29)	XMM	OM 30cm
52992.981867(29)		
52993.205282(27)		
52993.354318(33)		
52993.428817(26)		
52993.577881(25)		
52993.652274(40)		
52993.801328(24)		
52993.875841(34)		
52993.950336(32)		
53498.0121821(22)	Ultracam	VLT UT3 8.2m
53500.9922010(64)		
53505.9836285(26)		
53507.9951018(19)		
53512.0180207(20)		
54506.284757(32)	Ultraspec	ESO 3.6m
55313.037681(14)	Ultracam	NTT 3.58m

Table 6.3: Listing of new white dwarf mid-eclipse times, barycentric corrected and in the TT time system.

2010 eclipse altered parameters significantly, and the newly computed ephemeris is outside the uncertainties quoted in the previously published values. The error values on the third body parameters are also likely to be largely underestimated for this same reason. While the Monte Carlo method may give reasonable errors for this fit, realistic uncertainties will only be achieved by studying the effect that the number of eclipse timings has on the fit, which was not performed at this time.

Given the 56.35 ± 1.37 year period, this behaviour will not likely be confirmed for some time, assuming the system continues upon its current predicted trend of variation. Residuals also indicate large discrepancies even with this sinusoidal fit, with some eclipses still approximately 20 seconds away from their predicted T_0 times. It should be noted that these residuals are not sinusoidal in nature, and while the points after 1995 seem to indicate there is further behaviour here to be modeled, an attempt to fit another body to account for these does not result in a notably improved fit.

Another potential candidate for this variation is the Applegate effect [Applegate, 1992]. In principle, the secondary star undergoes deformations caused by the changing distribution of angular momentum generated by a magnetic torque within the star. This change in shape couples to the orbit and drives a change in

the system’s period. From Applegate [1992],

$$\frac{\Delta P}{P} = 2\pi \frac{O - C}{P_{mod}} = 2\pi \frac{\alpha}{\gamma} \quad (6.14)$$

Where α and γ are from our computed sinusoidal ephemeris as per §6.6.1.

It has been shown that using such a model to explain behaviour in objects such as Algol stars and other CVs indicates such a phenomenon can drive a change of $\frac{\Delta P}{P} \sim 10^{-5}$. Using the new ephemeris parameters, we find for Z Cha that $\frac{\Delta P}{P} = 3.6 \times 10^{-7}$, which is two orders of magnitude smaller, indicating this would seem to be a feasible mechanism. A key element in separating the Applegate effect from a third body would be determining whether the sinusoidal O-C function is strictly periodic (favouring the brown dwarf) or shows regular, but variable behaviour where many cycles may exist. This variation would also manifest synchronously with a change in luminosity of the secondary star, on the order of $\frac{\Delta L}{L} = 0.1$ [Applegate, 1992]. It has been shown that the ΔP in many binary systems that show an O-C behaviour similar to Z Cha is not strictly truly periodic [Richman et al., 1994], leading to the proposal that the Applegate mechanism could indeed drive this varying nature of Z Cha’s eclipse times [Baptista et al., 2002]. The notable residuals (up to 30s) present in the O-C fit may give weight to this.

A conclusive method of determining whether the Applegate effect can account for the observations is to consider whether the required energy from the secondary star can drive the variation in period. The method provided by Applegate [1992] considered the transfer of angular momentum from within the star to a thin outer shell, hence increasing the oblateness of this outer component. This paper utilised a fixed outer shell of $0.1M_{\odot}$, however the total secondary mass of Z Cha is $0.12M_{\odot}$ [Wade and Horne, 1988], meaning that such an assumption would place almost the entire mass of the companion star in its thin outer shell; clearly an unacceptable scenario. Instead, the calculations proceed using the method laid out by Brinkworth et al. [2006], which generalises the approach by Applegate [1992] by splitting a star of arbitrary mass into an inner core and an outer shell by integrating over finite ranges of radii. To do this, the density of the star is required as a function of its radius, which is provided by the Lane-Emden equation with an $n = 1.5$ polytrope (a fully convective secondary star) and the expansions in Clayton [1968] for around the zero point where the function diverges. Angular momentum is transferred from the core to the shell, leading to a change in their angular frequencies, which affects their oblateness and therefore leading to a change in the quadrupole moment. Using this, it is found that the minimum energy required for the computed $\frac{\Delta P}{P}$ with the

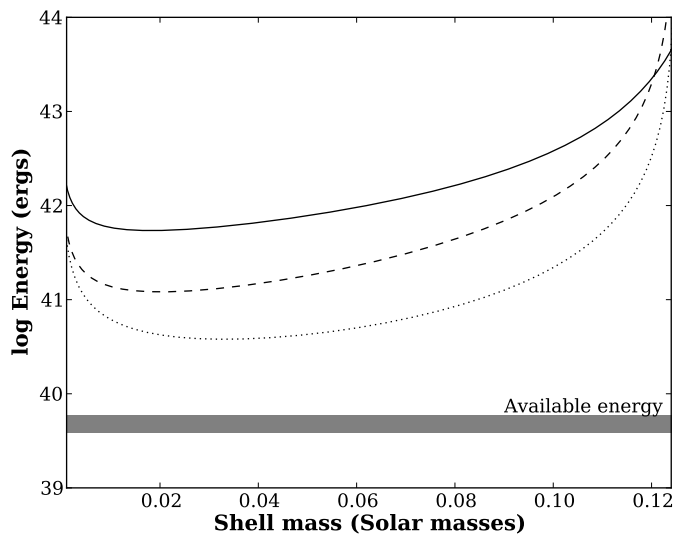


Figure 6.13: Plot of shell mass in the secondary star against minimum energy required for the Applegate effect to produce the observed $\frac{\Delta P}{P}$. The bar indicates the estimated range of energy available from the secondary star. The solid line is for the full computation as per Brinkworth et al. [2006], the dashed line is the same but ignoring the star’s quadrupole moment. The dotted line is the original Applegate calculation.

Applegate effect is 5.42×10^{41} ergs. Using $L = 4\pi R^2 \sigma T^4$ for a star of $2850 \pm 150\text{K}$, the energy output of the secondary is $4.86 \times 10^{39} \pm 2.02 \times 10^{39}$ ergs. In figure 6.13, the shaded area indicates the energy output from the secondary star that could be used to drive the Applegate mechanism, whereas the solid line indicates the energy *required* from the calculation performed. Even when the simplified model as per the original Applegate [1992] paper is used (dashed and dotted lines), there is still insufficient energy available.

These energy considerations imply that the Applegate effect is entirely unfeasible given the energy output of the secondary star, by 2 orders of magnitude. Though the residuals are not particularly sinusoidal after the third body fit, even here the available energy can only just account for this $\sim 20\text{s}$ variation.

6.7 The colour of Z Cha’s white dwarf

The colour of Z Cha’s white dwarf was determined using the multi-band Ultracam optical data, with the corrections between the Ultracam and SDSS filters presented

in §8 taken into account. To begin with, the extinction at the site was examined by using a long photometric observation of a non-variable star. Since this object went through an airmass range of 1.3-2.2, this was long enough to quite clearly see the effects the atmosphere was having upon the data. When converted to a magnitudes versus airmass plot, a straight line fit was performed on each arm resulting in $k_r = 0.059$, $k_g = 0.136$ and $k_u = 0.46$. Using a previous estimate for the zero point of Ultracam (T.R. Marsh, private communication), for this object it was found that $u' - g' = 1.009$ and $g' - r' = 0.494$. The theoretical extinction plots presented in §8.1.3 suggest that $k_r = 0.071$, $k_g = 0.150$ and $k_u = 0.490$. Here it was seen that the theoretical extinction overcompensated by 20% in r' , 10% in g' , and 6% in u' , almost certainly due to the fact that the atmosphere conditions in Paranal are better than those in La Palma.

Because the extinction depends upon a measurement that itself already uses the assumed zero point, this procedure was iterated upon to find a new zero point. The SDSS standard star G93-48 was observed on a nearby night with excellent conditions and was used for this calculation. The (actual) extinction correction found on the previous target was applied, and the zero point for each colour band was found with

$$m_x = -2.5 \log(f_x) + p_x, \quad (6.15)$$

where m_x is the observed magnitude (in bandpass x), f is the observed flux (in counts) of the target and p is the zero point. These values were found to be $p_r = 27.495$, $p_g = 28.131$ and $p_u = 26.137$. The process can now begin again to ensure the colours of the first target (and hence the extinction) are correct, which in turn leads to another zero point. In practice this is found to converge very rapidly. An alternate set of zero points were calculated for the standard star SA 113-260 on a different night and found to be $p_r = 27.583$, $p_g = 28.165$, and $p_u = 26.127$. The g' and u' zero points are within the standard uncertainty of the SDSS system. The red band is slightly less fixed with a shift of 0.07 mags.

Next, the colour of Z Cha's comparison star was determined. This was found to be $g' - r' = 1.05$ and $u' - g' = 2.14$. The theoretical extinction for this colour was read from the appropriate figures in §8.1.3 and then corrected by the percentages listed earlier which account for the site differences. The extinction coefficients were determined to be $k_r = 0.057$, $k_g = 0.136$ and $k_u = 0.462$ mags/airmass. From the optical modelling performed, the contribution of the the white dwarf to the total flux (out of eclipse) is used to find the colour which then must be corrected. Since the

Z Cha flux is divided by that of nearby companion star to correct for transparency

$$m_t - m_c = -2.5 \log \left(C \frac{f_t}{f_c} \right), \quad (6.16)$$

where m_t, f_t are the target magnitude and flux, m_c, f_c is the comparison magnitude and flux, and C is a correction factor to find. Using the extinction correction in magnitudes per airmass (k_t for target and k_c for comparison)

$$C \frac{f_t}{f_c} = 10^{\frac{1}{2.5} m_c - m_t} 10^{\frac{1}{2.5} (k_t - k_c) X}. \quad (6.17)$$

So the flux ratio of the target to comparison star must be corrected by the factor

$$C = 10^{\frac{1}{2.5} (k_t - k_c) X}, \quad (6.18)$$

where X is the airmass. The result of the colours obtained and plotted against white dwarf models is shown in figure 6.14. It can clearly be seen that the scatter of the points for the determined white dwarf colours is greater than their uncertainty. One particular point, located at $g' - r' = -0.45$, $u' - g' = 0.18$ was from a night of particularly variable transparency, indicating that even when care is taken in the reduction and correction of the target, there is still some residual effect upon the data that is being transferred to the fitting process. Previous studies of Z Cha have indicated that the temperature is of the order of 15,700K [Robinson et al., 1995]. While the average eclipse does cross into this regime at the 90% confidence level, the fact remains that this only happens because of the suspect data point from the night with variable conditions. When eliminated, the average moves towards the red and away from the models. If the data were to lie in the regime predicted by the models, it would indicate a consistent lack of g' band flux (the points are too red in $g' - r'$ and too blue in $u' - g'$). IP Peg has been shown to have a particularly variable u' band flux which also makes the white dwarf colour discrepant with the expected models, and was explained by absorption due to disc material [Copperwheat et al., 2010]. Such an argument cannot be made in this case for Z Cha however, and the consistent lack of g' flux is not easily explained.

6.8 Summary

This chapter showed that after constructing a CV model from Ultracam photometry, such a model can not only be used to conclusively prove that the white dwarf is unobscured in the optical band but can also be utilised in the comparison of X-ray




Figure 6.14: The colour-colour plot of Z Cha's white dwarf against white dwarf models calculated by Holberg and Bergeron [2006]. The solid lines are those of constant temperature (indicated), whereas the dashed lines indicate white dwarf masses from $0.6M_{\odot}$ (top) to $1.2M_{\odot}$ (bottom). Each triangle is a single optical eclipse whereas the filled circle indicates the average eclipse. The error bars are for 68% confidence. It is expected that Z Cha should be at 15,700K [Robinson et al., 1995] at $\sim 0.8M_{\odot}$, implying that the observed average point is too faint in g' .

data to constrain the emitting region of the boundary layer. The disc was shown to obscure soft X-rays, but somewhat surprisingly, not be dense enough to obscure the hard X-rays, in apparent disagreement with expected surface densities from the standard disc-instability models. This discrepancy is attributed to a “truncated” inner disc, for which several possible models are suggested, but at this time cannot be convincingly narrowed down. It is also of note that the inner-disc accretion rate was in line with such models in contrast to previous studies of the same system [van Teeseling, 1997]; whether this implies that the accretion rate is non-anomalous because of, or in spite of the truncated disc would require further study and X-ray observations.

The precise timings from Ultracam and XMM were also utilised to investigate the ephemeris of Z Cha, which has been showing a strong modulation over the past 30+ years of observations. Due to energy constraints, the reason for this variation is most likely the presence of a third body - a brown dwarf - present in the system.

The next chapter will utilise data from a new instrument, “X-Shooter”, which operates on a wide spectral range. The subject will be the cataclysmic variable OY Car, and as will be shown, X-Shooter’s range allows a simultaneous measurement of both the white dwarf and the secondary star, and could provide an essential source of information for CV observers in the near future.

Chapter 7

X-Shooter observations of OY Car

In the north of Chile at the Paranal observatory, the ESO instrument “X-Shooter” has recently come online and has begun processing requested targets. This instrument operates in an entirely different regime compared to Ultracam and Ultraspec: while the exposure times are long (on the order of minutes), X-Shooter has extremely wide wavelength coverage, from 3000 Å in the near-UV to 25000 Å in the infra-red, achieved by the use of an echelle grating in three arms. Presented here is an initial analysis of OY Car data, giving a conclusive and unambiguous measurement of K_2 for the first time in this system, along with a K_1 . Together these provide a measurement of the system’s mass ratio q , independent of previous photometric methods. In addition to this, detections of FeI and FeII lines as reported by Horne et al. [1994] are presented together with their radial velocities and line strengths.

7.1 Observations and data reduction

The service mode X-Shooter observations of OY Car were taken during February 2010. The X-Shooter data was processed with the latest pipeline as of November 2010. Results from this show a greatly improved throughput compared to previous versions of the pipeline due to better tracking of the slit along the curved paths the echelle grating produces. Signal-to-noise is generally good in the UV and visual arms, but it decreases sharply in the infra-red until the tracking is lost completely, hence only half of the infra-red range is actually useable. The infra-red is also heavily affected by tellurics. Extracted average spectra of the IR, visual, and UV arms can be seen in figures 7.1, 7.2, and 7.3 respectively.




Figure 7.1: Extraction of the red/IR arm from the OY Car X-Shooter data.




Figure 7.2: Extraction of the green/VIS arm from the OY Car X-Shooter data.




Figure 7.3: Extraction of the blue/UV arm from the OY Car X-Shooter data.

7.2 Testing the photometric method with X-Shooter

The system parameters for OY Car have been a source of discussion in many papers [Vogt et al., 1981; Schoembs and Hartmann, 1983; Cook, 1985; Wood et al., 1989]. Typically, the mass ratio q , and white dwarf radius R_{WD} , have been determined through photometric means. This generally involves examining contact points in eclipses, and using various combinations of assumptions such as white dwarf mass-radius relations or that the secondary star is of a specific type. Ideally, the photometric q is backed up with a spectroscopic q . This useful relation can be seen by starting from

$$K_x = \frac{2\pi a_x}{P} \sin i, \quad (7.1)$$

where a_x is the the separation for star x (where $a = a_1 + a_2$) and P is the orbital period. Then, using Kepler's law,

$$a^3 = \frac{G(M_1 + M_2)P^2}{4\pi^2}. \quad (7.2)$$

Utilising $M_1 a_1 = M_2 a_2$,

$$a = a_1 \frac{M_1 + M_2}{M_2}, \quad (7.3)$$

which can be rearranged to find that

$$\frac{(M_x \sin i)^3}{(M_1 + M_2)^2} = \frac{PK_x^3}{2\pi G}. \quad (7.4)$$

Hence it is seen that

$$\frac{K_1}{K_2} = \frac{M_2}{M_1} = q. \quad (7.5)$$

K_1 is usually a very difficult parameter to determine by measurement, as found in §5. The disc lines dominate over those of the white dwarf, and the simple assumption that the disc is circular proves invalid, because of the stream-disc interactions. The secondary lines are not affected by such problems, and hence a measurement of K_2 can be valuable in constraining system parameters. However, because the secondary star is the least luminous part of the system, any signature of it is usually obscured in the optical band making K_2 quite an elusive measurement in most CVs. Fortunately, X-Shooter extends out to the infrared, where the secondary's luminosity will peak due to its $\sim 3000\text{K}$ temperature. Here, a clear measurement of K_2 was found in the form of KI at 12539.34\AA (figure 7.4). Using a gaussian fit to the profile for each spectrum, the radial velocity of this line is found to be $480.5 \pm 1.5\text{km/s}$, and a plot of this from fit can be seen in figure 7.5.




Figure 7.4: Trail of the KI line in X-Shooter OY Car data.




Figure 7.5: The radial velocity curve of the KI 12539.34 Å line in OY Car found with X-Shooter.

In addition to this, the considerably fainter MgII 4481.327Å line was found in the visual arm of X-Shooter. This line is associated with the white dwarf atmosphere [Holberg et al., 1997; Gänsicke et al., 2007] and since it is not expected to be present in the accretion disc, should provide an unambiguous measurement of K_1 without the usual difficulties associated with the bright spot or a non-circular disc. There is, however, a nearby helium line (4471Å) which, while clearly separate, could affect the measurement. The trail of the MgII line is presented in figure 7.6. The radial velocity measurement can be seen in figure 7.7 with a fit to all data points and a fit excluding the data points near eclipse which are quite clearly non-sinusoidal. In the former case, $K_1 = 78.5 \pm 1.5\text{km/s}$, and in the latter $K_1 = 103.7 \pm 1.19\text{km/s}$. It is notable, however, that the lower velocity result is at phase $\phi = 0.015 \pm 0.02$ (which is at the white dwarf phase to within the uncertainty) whereas the higher velocity result is at $\phi = 0.040 \pm 0.02$. This seems to indicate that the result is at least partially affected by the HeI line which would be expected to be at a higher velocity coming from the disc. The lower K_1 is also more in line with the anticipated values that have been inferred in literature over the years, of which a summary and comparison is detailed in table 7.1. As anticipated, with the spectroscopic method, K_1 is the greatest source of uncertainty when determining q . The measured K_1 and K_2 values are both higher than those derived in Wood and Horne [1990], as is the mass ratio. Although mass ratios near this value (or higher) have been predicted in older papers, most recent estimates from photometric methods predict $\sim 0.102 - 0.11$. Assuming that the K_2 measurement with X-Shooter is correct, which seems reasonable given its fit, the K_1 required to match such a q is 49 km/s. The closest values to those measured here are found in Schoembs and Hartmann [1983], who measure a K_1 from the Balmer lines and apply a correction using a model by Hensler [1982], resulting in $K_1 = 75 \pm 15\text{km/s}$ (the K_2 value is then predicted assuming the correct K_1).

It is of note that the results by Wood et al. [1989], Schoembs et al. [1987] and Cook [1985] who all predict a $q \sim 0.10$ use the photometric method from separate measurements of the bright spot contact points. This would imply that these bright spot points are measured sufficiently well for this method, and that if this mass ratio is incorrect, the method may be at fault. To achieve a higher value of q , the point at which the stream impacts the disc must be moved further upstream, as opposed to the typical photometric method assumption of the stream impacting at the brightest point (the ‘light centre’). However, even at its most extreme position, this still only leads to $q = 0.135$ [Wood et al., 1989]. The contrary set of results such as Vogt et al. [1981] and Schoembs and Hartmann [1983] which predict $q \sim 0.15$, is more in line with the result achieved here. These use measurements of K_1 to determine q ,

q	K_1	K_2	Reference	Source
0.163-0.216	78.5-103.7	480.5	This thesis	X-Shooter
0.102	46-48	455-470	Wood and Horne [1990]	1,2
0.1	45	440	Wood et al. [1989]	1,2
0.11	54	475	Schoembs et al. [1987]	1,2,3
0.11	59	535	Cook [1985]	1,3
0.2	75	375	Berriman [1984]	2,4
0.15	70	481	Schoembs and Hartmann [1983]	4,5
0.45	125	275	Bailey and Ward [1981]	3,4
0.15	69	472	Vogt et al. [1981]	3,4

Table 7.1: Comparison of q , K_1 , and K_2 from various sources. The assumptions used are as follows; 1: stream trajectory, 2: white dwarf mass-radius relation, 3: main sequence secondary star, 4: K_1 , 5: bright spot position.

and involved attempts at measuring the white dwarf using the high-velocity wings of Balmer lines. These however, are expected somewhat to be distorted by the asymmetric disc (as noted in §5). While it would be a surprise that even the MgII line measured here (which should only be present in the white dwarf) is affected, the result is similar enough to previous attempts at measuring K_1 to raise suspicion, particularly with the effective nature of the photometric method for other systems (such as Z Cha as discussed in §6). The other option is that K_2 may be incorrect, potentially caused by a radiation induced distortion. However, given the very good fit to K_2 produced by X-Shooter, and also that such a distortion would have to be extremely large ($\sim 780\text{km/s}$ with the detected K_1) to bring q in line with the photometric method, this seems somewhat unlikely. Since there is plenty of signal-to-noise in the X-Shooter data, further attempts with a higher time resolution may be of assistance in determining any non-sinusoidal distortion present and give a more accurate K_2 .

The two options to reconcile the views of the spectroscopic and photometric methods are, firstly, that the same distortion of the disc lines is affecting the MgII line in some manner and moving it to a higher velocity. Secondly, there is a systematic error in the positioning of the bright spot which is affecting the photometric measurement of q . Given the strikingly different fits that can be made to the MgII line depending on the elimination of a few points (varying from 78.5 – 103.7 km/s), it would be premature and unlikely to consider the photometric method is wrong, and that in all probability K_1 is still the biggest problem in measuring q spectroscopically for CVs.




Figure 7.6: Trail of the MgII line in X-Shooter OY Car data.




Figure 7.7: The radial velocity curve of the MgII 4481Å line in OY Car found with X-Shooter. The solid line is a fit to all data points whereas the dashed line excludes the non-sinusoidal points near the eclipse.




Figure 7.8: A portion in wavelength of the trailed OY Car data in the UV arm from X-Shooter. All the absorption lines seen can be attributed to iron.

7.3 Iron lines in the spectrum of OY Car

The high spectral range and resolution of X-Shooter also lead to another observation. A cursory glance at the spectrum and trails reveal a wealth of observed lines, most of them attributed to FeI with some FeII (see figure 7.8). These iron lines were shown to be veiling the white dwarf UV spectrum by Horne et al. [1994], though the actual location of this material was uncertain. The theory put forward was that the material may occur in the upper atmosphere of the outer accretion disc since the other option, that the material is actually located at $\sim 5R_{WD}$, would mean the view of the upper white dwarf hemisphere was unobscured, which was incompatible with their observations.

A rough estimate of the radial velocity of the observed lines was used to “straighten” the trailed spectra, at which point they were averaged together and an identification of the lines in the UV and VIS arms of X-Shooter data was performed, listed in table 7.2. Avoiding the blends, the remaining iron lines were fitted with a shared offset from their rest wavelengths and individual peak heights with a shared scaling factor. The radial velocity curve and varying peak height with phase is shown in figure 7.9.

The uncorrected fit results in $K_{Fe} = 83.7 \pm 1.7$, but since this fit is clearly distorted by the non-sinusoidal points near the eclipse, a fit excluding these points drops the χ^2 (with 28 points) from 361 to 60 and results in $K_{Fe} = 116.7 \pm 0.86$. The

Line	Å	Line	Å	Line	Å
FeII	3089.38	FeI	3120.44	FeI	3125.65
FeII	3133.05	FeII	3154.20	FeII	3162.80
FeI	3168.85	FeII*	3180.15	FeII*	3187.30
FeII*	3193.86	FeI	3196.93	FeII	3210.44
FeII	3213.31	FeI	3228.25	FeI*	3234.61
FeI*	3237.22	FeI	3239.43	FeI	3242.26
FeI*	3249.19	FeII	3259.05	FeI	3262.01
TiII ?	3278.29	FeII*	3281.26	FeII*	3323.06
FeI*	3329.52	FeI*	3336.26	FeI*	3342.29
FeI	3349.72	FeI*	3361.95	FeI*	3368.50
TiII* ?	3372.80	FeI	3380.11	FeI	3383.98
FeII	3388.13	FeI*	3394.80	FeI*	3403.40
FeI	3409.59	FeI	3422.66	FeI	3433.57
FeI*	3442.36	??	3460.90	CaI	3474.76
FeI	3483.01	FeI	3505.06	FeI*	3565.58
FeI*	3570.25	FeI*	3581.65	FeI	3614.56
FeI*	3619.77	FeI	3625.14	FeI*	3631.46
TiII ?	3641.33	KI ?	3648.50	FeI*	3686.26
FeI	3760.05	FeI	3762.20	CaII	3933.67
CaII	3967.47	FeI	4046.06	MgII	4481.13
FeI*	4584.72	FeII	4923.92	FeII	5018.43
FeII*	5169.03				

Table 7.2: Table of measured lines in the UV and VIS arms of X-Shooter in OY Car, excluding the usual Balmer and helium lines. An asterisk indicates that the line is a doublet or blend, in which case the line is identified as the one with the highest transition strength.




Figure 7.9: The fit of several FeI lines present in OY Car data with X-Shooter. Top: the scaling of each of the line peaks. Bottom: the radial velocity amplitude, where the solid line is a fit to all points and the dashed line excludes points near the eclipse.




Figure 7.10: A diagram (not to scale) showing an example of a stream trajectory and impact points predicted by Lubow [1989] and found in OY Car and Z Cha. Material that flows over the disc edge preferentially impacts half an orbit later.

offset $\gamma = 78.6 \pm 0.63$. This is notably similar to the measured K_1 , and hence seems to agree with the view of the absorbing material present in the disc as the radial velocity appears to be tracking the white dwarf. The scaling of the line strengths shows a double-peaked behaviour, with low points at phases 0.45 and 0.85, in a similar fashion to the X-ray absorption found in Z Cha in §6. Given that the bright spot becomes a significant proportion of the flux at phase 0.8 due to foreshortening, and that the extent of the material from the impact is several times the disc scale height [Lubow and Shu, 1975], it may be plausible that the disc/stream impact is the cause for the absorption. The second absorption point can be explained by material passing over the disc edge and “swinging around” to impact the disc near the stream periastron, shown by Lubow [1989], as shown in figure 7.10. The points of absorption match those predicted by [Lubow and Shu, 1975] much better than the observations of Z Cha in §6, but it should once again be mentioned that the alternative is a disc with a tidally thickened outer rim [Ogilvie, 2002]. Distinguishing between these two options would be an interesting point of research in the future.

7.4 Summary

This chapter showed the enormous advantages of using an instrument with a wide spectral range for the study of cataclysmic variables. A measurement of K_2 of OY Car was shown for the first time, along with K_1 which provided a completely independent measurement of q by which to compare the photometric method previously applied to determine system parameters in OY Car. The presence of numerous FeI and FeII lines was also detailed, with a radial velocity suggesting an outer disc location and a “double-humped” modulation in line strength which could either be attributed to a second stream-disc interaction or a tidally thickened disc rim.

The attention of the thesis will now turn to instrumentation. During the

course of this work it was found to be necessary to perform a careful examination of the Ultracam filters in an attempt to determine the colour of the white dwarf in Z Cha (§6). These calibrations will be shown in the next chapter, along with an investigation into Ultracam's CCDs and software designed to greatly improve the workflow of operating with the now extensive archive of Ultracam data.

Chapter 8

Improvements to the Ultracam system

Ultracam, which was introduced in §4.3.1, has now been operating for eight years and has produced a wealth of refereed papers on subjects ranging from neutron stars and flare stars to its more native realm of cataclysmic variables. As with all instruments, Ultracam is constantly evolving and being improved in many areas, of which three will be detailed in this chapter. Initially the photometric accuracy of the filters will be examined, followed by a CCD phenomenon known in Ultracam as “peppering”, which has so far limited the range of target brightness/exposure length that the instrument can operate at. Finally, a separate and automated software pipeline for indexing and cataloguing the now 4TB+ of Ultracam (and Ultraspec) data will be presented, along with the relevant updates to the user interface and backend software to accommodate these features.

8.1 Calibration of the Ultracam SDSS filters

From its initial concept, Ultracam was designed to use the Sloan Digital Sky Survey [York et al., 2000] filters. This is comprised of five filters (u’g’r’i’z’) encompassing 3000-11000 Å, effectively covering the entire useful response of a modern CCD. While the filter bandpasses do not overlap due to a strong cutoff (caused by the shortpass glass design and longpass interference film), they feature a wide plateau rather than a single individual peak wavelength and this assists the study of faint objects when compared to older filter systems like the Johnson filter system [Johnson and Morgan, 1953]. The optical path of each filter is made to be the same length by using neutral glass to give equal thicknesses. The only filter that does not have a cutoff at the

long wavelength end is the z' filter, due to a lack of materials that act upon such wavelengths. However, this turns out to be irrelevant in modern astronomy as the CCD response rapidly tails to zero at these wavelengths (§3). In fact, the detector response dictates the system response in the case of z' more than the filter itself. The SDSS filter bands are also designed to avoid common night sky lines such as OI 5577 Å and HgI 5460 Å, which was a problem sometimes associated with the Thuan and Gunn griz filter system [Thuan and Gunn, 1976]. Since the SDSS system uses a monochromatic “AB” magnitude system, it is not prone to calibration uncertainties with the Johnson system, which was based on comparing each object to α Lyr (Vega). An object in the AB system will have an equal magnitude in each passband if f_ν is constant. The actual calibration of the system is done by measuring several F subdwarfs which have flat spectra with low metallicity. Every filter is defined to give zero magnitudes when a flux density f_ν of 3631 Jy is detected. This allows compatibility (or at least similarity) with other AB systems that preceded it.

The Ultracam filters are copies of the SDSS filters not only for the benefits of this system, but to also easily offer comparison with the vast databases produced by the Sloan survey. Typically this filter set is referred to as u'g'r'i'z' within Ultracam, though the “prime” has been queried. The origin of the prime comes from the fact SDSS unintentionally ended up with one set of filters on USNO-40 that was slightly altered due to dehydration of the interference film in a vacuum decreasing the effective refractive index. There is no particular reason for this notation within Ultracam, but it has been kept for consistency since it has now been marked as such for many years. In sections below, where both the SDSS filters and the Ultracam clone filters are discussed in comparison, the official SDSS filters are described using the prime notation and the Ultracam filters with a capital, ie. u' vs. U.

8.1.1 The system response

The AB magnitude of an object is defined by [Fukugita et al., 1996]

$$m = -2.5 \log \frac{\int d(\log \nu) f_\nu S_\nu}{\int d(\log \nu) S_\nu} - 48.6, \quad (8.1)$$

where f_ν is the flux density (at frequency ν), and S_ν is the system response. To initially determine S , lab measurements of each of the filters were performed (with results supplied by the filter manufacturer), determining the fraction of light transmitted in steps of wavelength. In addition to this, the CCD quantum efficiency must be taken into account, as well as the coatings upon the filters and CCD as well (obtained from their respective data sheets). Other materials in the light path in-

Arm	Surfaces
Blue	Aluminium (x2), CaF2, LLF1 (x2), CCD broadband coating, SK16 (x2)
Green	Aluminium (x2), LLF1, CaF2, Dichroic 1, CCD broadband coating, LAK10 (x2), Fused silica
Red	Aluminium (x2), CaF2, LLF1, Dichroic 1, Dichroic 2, CCD midband coating, LAK22 (x2), SF1
Arm	Filters
Blue	u', NBF3500Å
Green	g', HeII, Blue continuum, NBF4170Å
Red	r', i', z', Red continuum, Sodium, H α broad, H α narrow

Table 8.1: List of surfaces necessary in computing the quantum efficiency of each light path in Ultracam, and the separate filters available for each arm.

cluding the dichroics, collimator, and aluminium mirrors of the telescope (assumed clean) were also considered. Software was written to compute a transmission (in terms of quantum efficiency) for given combinations of surfaces. Those used in the computations of the light paths are shown in table 8.1. Linear interpolation was used to ensure each of the surfaces was on the same wavelength scale and there is the assumption that every surface has a quantum efficiency of exactly zero outside their specified ranges to avoid anomalies in the numerical integration. The resultant plots showing the transmission regime and quantum efficiency of each Ultracam filter in the red, green, and blue arms can be seen in figure 8.1.

8.1.2 Comparing the Ultracam filter set to SDSS

Now S_λ has been computed, a magnitude m can be determined from a flux f_λ input. Starting from equation 8.1, and converting from ν to λ (where $\frac{d\nu}{d\lambda} = \frac{c}{\lambda^2}$), the two integrals to perform are $\int S_\lambda f_\lambda \lambda d\lambda$ and $\int S_\lambda / \lambda d\lambda$, where the factors of λ come from the appropriate Jacobian transformations. Template spectra of standard star types from Pickles [1998] were used, as well as white dwarf models from Holberg and Bergeron [2006] to serve as f_λ . Computing these values for each colour band, the colour-colour plot of the output through the Ultracam filters can be seen in figure 8.2.

A similar task can now be performed but instead using the SDSS r', g', and u' filter throughput data which is available on the SDSS website (a direct comparison of the system responses can be seen in figure 8.3). Using the same set of models, the set of values r' - R, g' - G, and u' - U for g' - r' and u' - g' can be determined. This gives an estimate of the difference between the two sets of filters for a given

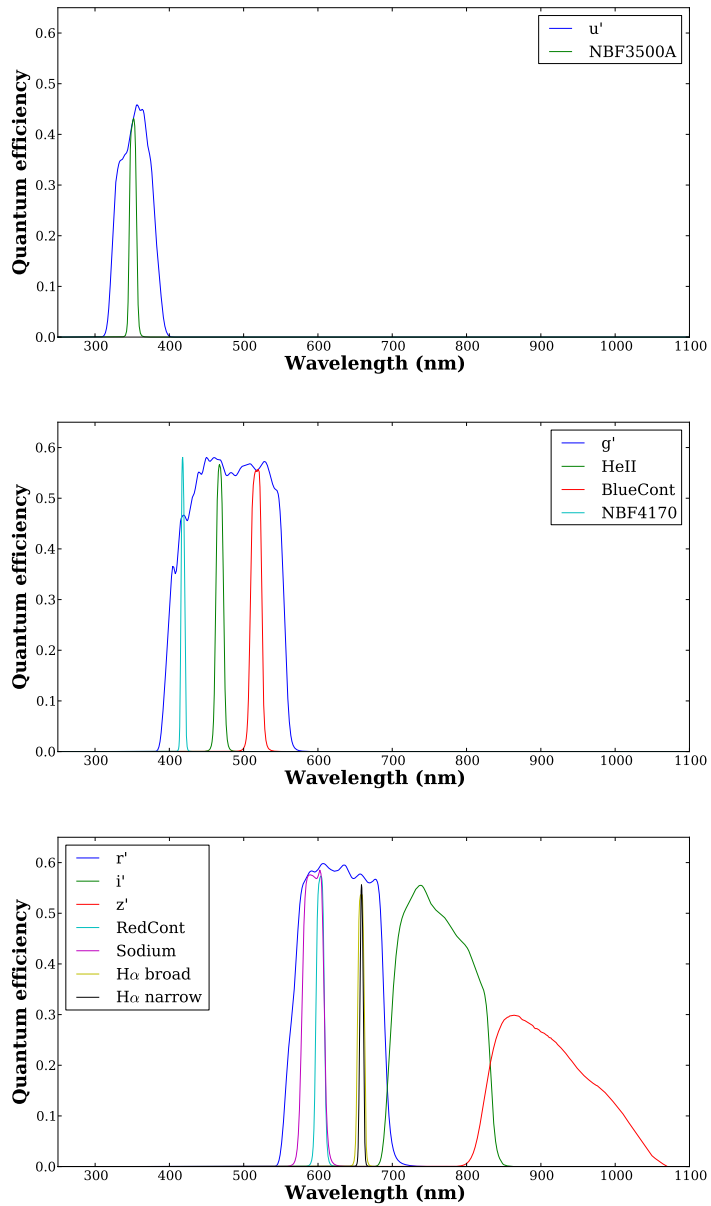


Figure 8.1: Ultracam system response for all available filters. Top: blue arm. Middle: green arm. Bottom: red arm.




Figure 8.2: Colour-colour plot of spectral type templates from Pickles [1998] convolved through the Ultracam light paths using the r' , g' , and u' filters.

colour, and can be seen in figure 8.4. For the red filter sets, the difference is minimal (≤ 0.05 mag) for all but the reddest star types, with the m8iii spectral type (of $g' - r' = 2.5$) resulting in a 0.22 mag colour shift. For $g' - G$, a steadier trend is seen as $g' - r'$ increases. This is almost linear in the sense that a $g' - r'$ of 1 gives a $g' - G$ of ~ 0.1 and a $g' - r'$ of 2 gives a $g' - G$ of ~ 0.2 . The m6, m7, and m8 spectral types once again show the greatest difference. The $u' - U$ plot is very hard to quantify though the greatest colour shift found is 0.1. The majority of the star types show a shift of ≤ 0.04 , with all the white dwarf models plotted yielding $u' - U$ of ≤ 0.01 . Since the absolute error in an SDSS magnitude is estimated to be on the order of 0.03 mag [Fukugita et al., 1996], the shift between the filters in the u' bandpass is effectively negligible, though the g' and r' bandpasses should be considered carefully if the object of interest has a $g' - r'$ of 1 or greater.

This process does depend upon the assumption that the SDSS (and Ultracam) filter bandpass measurements supplied are accurate. Should greater accuracy be required, the analysis could be performed in count space, or another possible method would be to see whether the filter bandpasses (along with the other necessary light path transmission curves) reliably reproduce the colours for various targets within the SDSS database. Since the Ultracam and SDSS filter difference was on




Figure 8.3: The u'g'r'i'z' filters from Ultracam above their SDSS equivalents. Note this is the complete system response taking into account the telescope and CCDs, explaining the significant difference in QE, particularly at the extreme ends of the wavelength scale where Ultracam has CCDs with specifically optimised coatings.

the order of ≤ 0.1 mags for all but the most extreme colour targets, further analysis was not pursued.

8.1.3 Theoretical extinction

The program is now extended to include extinction curves that have been compiled for the La Palma observatory [King, 1985], where the flux is now modified to be

$$f_x(\lambda) = f_0(\lambda) \times 10^{\frac{k_x(\lambda)Z}{2.5}}, \quad (8.2)$$

where f_x is the flux in bandpass x , f_0 is the flux above atmosphere, k_x is the extinction coefficient, and Z is the airmass. This process is performed for a range of Z for each spectral template, and the estimated extinction in magnitudes per airmass is computed and shown against colour for the same star models for the Ultracam filters. These are presented in figure 8.5.

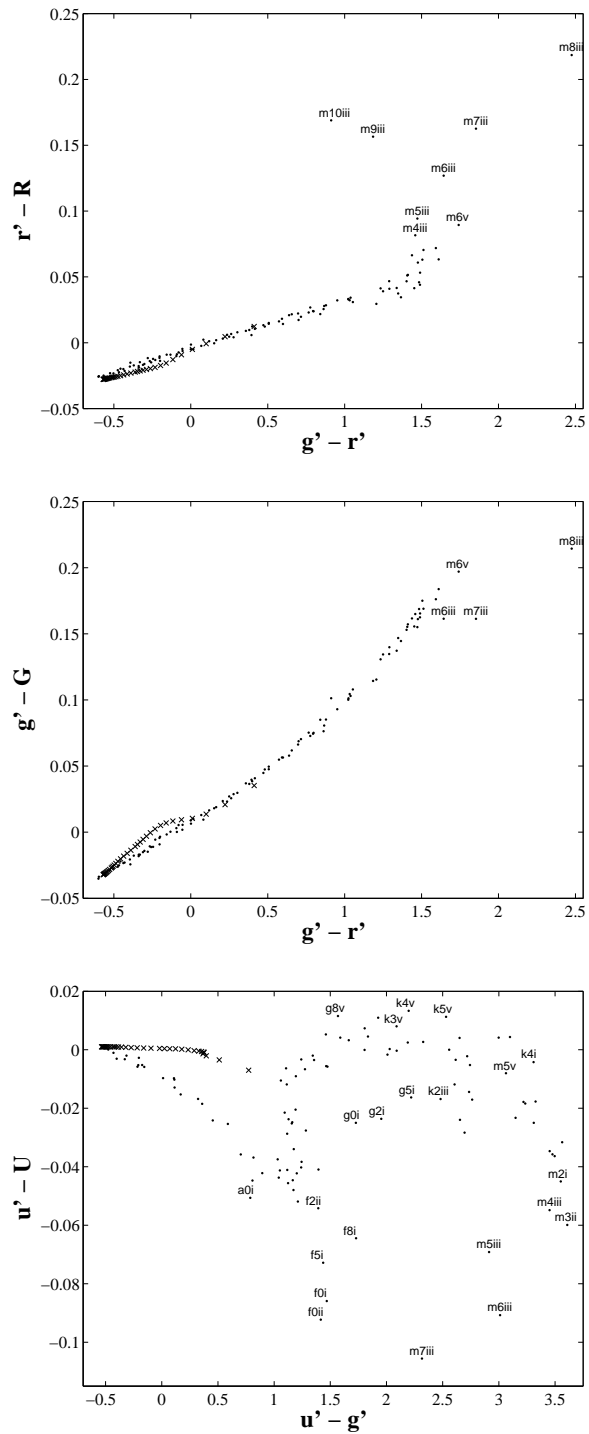


Figure 8.4: A plot of the difference between the SDSS filter set (primed) and the Ultracam filter set (capitalised) for a $g' - r'$ or $u' - g'$ colour. Points are spectral templates provided by Pickles [1998] and the crosses are log $g = 8$ white dwarfs of various temperatures from Holberg and Bergeron [2006].

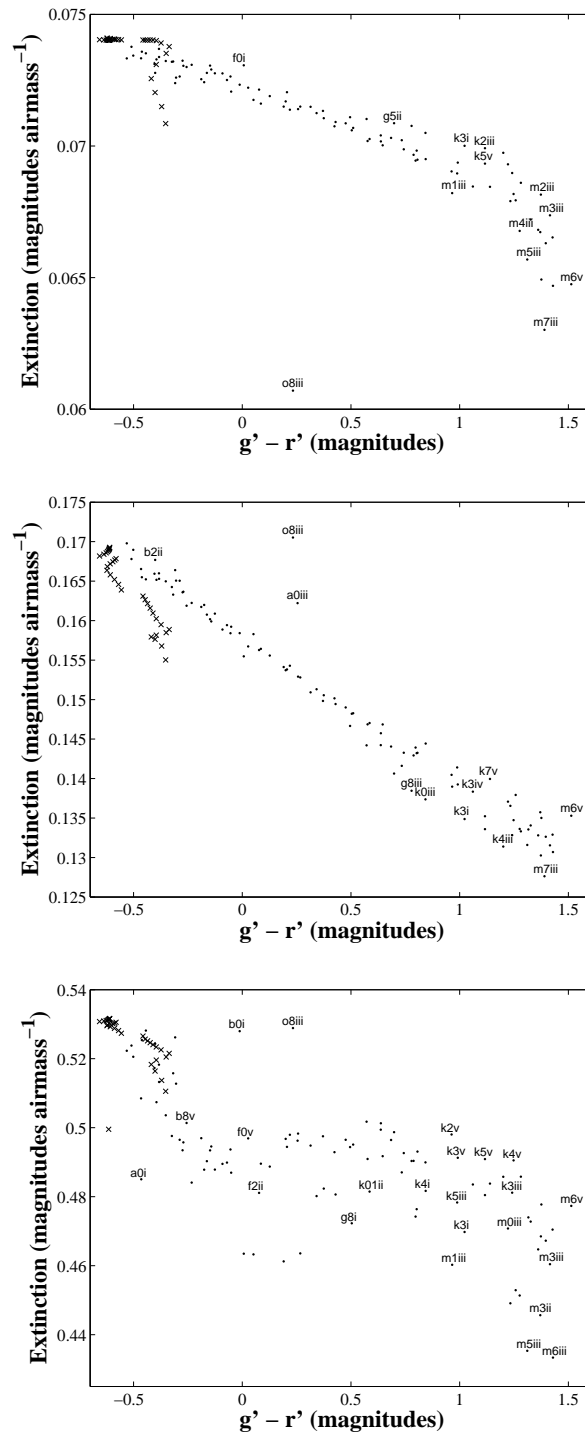


Figure 8.5: Theoretical extinction plots for Ultracam filters. Top: r' band. Middle: g' band. Bottom: u' band.




Figure 8.6: The peppering phenomenon in Ultracam CCD1 (red arm) in both left and right amplifiers just before the onset of saturation.

8.2 The “peppering” phenomenon in Ultracam

The name “peppering” is given to an effect that is present in all the Ultracam CCDs at some level, observed to be at a well fraction just before the onset of saturation. It is typically characterised by a “noisy” circle near the centre of the chip, and has been alarming enough in appearance to instruct all observers to avoid the peppering level; figure 8.6 shows an example frame from the red CCD showing the problem. The diagonal lines show a typical coating pattern and the blurred streak near the middle of the chip indicates a local saturation effect whereby charge is spilling in the vertical direction as it is clocked. Below that in the circular shape is a speckled region which is the problem area.

The phenomenon has been seen as early as the first telescope run, usually on flat field operations as one has to reach a signal level beyond the shot-noise limiting region to get a good measure of the pixel non-uniformity. It was originally suspected that peppering was a function of charge level alone, as it appeared close to saturation and there is some fraction of vignetting meaning that the centre of the chip will saturate first. From study of flat field runs that go completely through




Figure 8.7: The division of two sequential flat field frames in the peppering regime. The pattern divides out almost completely.

the saturation level though, it is clear that this is not true, and there is a spatial component involved. While the peppering regime will increase in radius (and in fact, eventually approximate a teardrop shape towards the bottom of the chip), some areas of the chip will never go into peppering, and instead go directly to saturation.

The pattern varies from run to run, but is not random. The frame-to-frame pattern over a period of time will in fact be equal and divides out perfectly. An example of this is shown in figure 8.7, which are sequential flat-field frames. Here the centre regime of the chip shows no peppering at all. The thin ring surrounding it comes from the fact the second frame has a brighter sky level and so the peppering regime has extended outward. The pattern from a different night, however, is not the same.

Taking a long flat field run from 2009, 50 by 50 pixel areas were chosen near the centre of the chip for each CCD with the mean counts and variance measured for each exposure. The result of this is seen in figure 8.8. The mean counts within a window rise smoothly until the onset of saturation at 65,535 ADU - hence *no counts are lost within the sampling window*. Despite the dark/light speckled appearance,

this does not seem to be a normal charge transfer issue (also suggested by the fact this occurs in a circular regime, whereas a transfer issue would cause the same effect to all rows above the onset). This means that peppering is not a disaster for aperture photometry, as almost certainly all the counts are going to be collected unless the aperture is very small. The onset of peppering becomes very clear in the plot of variance, marked by a sudden discontinuity in what should otherwise be a smooth curve with a log slope of 1/2 (from shot noise). The numerical derivative of the log variance is in fact a reliable and accurate method for determining the onset of peppering. While it is typically accepted that the red channel exhibits peppering at $\sim 50,000$ ADU upward, a single run in 2009 was observed to have peppering in this channel occurring at as low as 29,000 ADU. This indicates that should the peppering regime wish to be avoided, flat fields that cross the peppering regime (which are not usually taken as they are effectively wasted data) are needed to definitively determine the trigger level at the time. The last thing the plot reveals is the notable fluctuation in the variance in the blue and green channels which is apparently missing in the red channel. This double peaked behaviour is quite visible in the CCD images very close to the saturation point (figure 8.9), manifesting as two rings that move outward.

During live testing, the only things that have been observed to have had an effect upon the peppering have been the parallel vertical clock times after inversion during frame-transfer (D. Atkinson, private communication), which were settled upon to be 24 microseconds per clock. This reduced the phenomenon but maintained a high transfer speed. A few possible options exist to explain the difference between the red and the blue/green peppering effects. Firstly, the red chip has a slightly different manufacturing process to enhance it for that end of the spectrum. Secondly, the red chip has a slightly higher clock swing (clock voltage amplitude), which was implemented as a side effect of this chip being somewhat different in practice. The higher swing brought the dark current level down to the same as the other two chips, as the higher amplitude drives it further into inversion. Thirdly, the red chip is simply much closer to the SDSU controller, and thus has a short cable. The blue and green cables are longer, but of the same length. Given the dependence upon parallel clock rates, the reason the peppering is preferentially seen towards the center of the chip is adequately explained by clock propagation. The resistance to ground each pixel sees is greatest at the centre of the array. This is because the ground connections are usually at the edges of the chip, and so in the centre, the path to ground is longer. At a sufficiently high clock frequency, the time constant ($\tau = RC$, where R is resistance and C is capacitance) caused by self wave shaping




Figure 8.8: The mean and variance of a sampled window near the centre of each Ultracam CCD as the count level progresses from a saturated sky through the peeping region and back into normal operation. Left: red CCD. Middle: green CCD. Right: blue CCD.




Figure 8.9: The green Ultracam CCD with the central wells beyond the saturation point. Two rings of peppering are seen which shrink and eventually merge into a single ring before disappearing altogether as the count level is lowered.

begins to affect what should be the dominant time constant set by the external clock [Janesick, 2001]. In practice, this extra resistance due to the long ground path shapes (distorts) the clock swing nearer the centre of the array, and is clearly present at some level because the saturation/blooming in Ultracam always begins in this region. A weaker clock swing means that the potential between the collecting and barrier phases is lessened which in turn means less well capacity in this region of the chip. Charge is most efficiently transferred between pixels on a CCD via fringing fields, with the ultimate limit being thermal diffusion. Unfortunately as a well nears its capacity, the effect of the fringing field becomes negligible. Potentially, a clock swing that was too fast with a pixel near full-well condition would cause a “charge back-up” problem, whereby the small fringing field at this signal level means there is insufficient time for all the charge to diffuse to the next phase in the given clock cycle. This will result in the remnant electrons collecting in the next nearest potential well, accounting for the modulation loss. Examining neighbouring pixels in the vertical direction does reveal a mild correlation (figure 8.10). If this is an accurate depiction of what is occurring, the remaining questions are what is in fact causing this problem when it comes out of inversion, and how may it be fixed.

While most solutions for this appear to revolve around slowing the parallel clock times, particularly the delay after integration (which would decrease the fram-

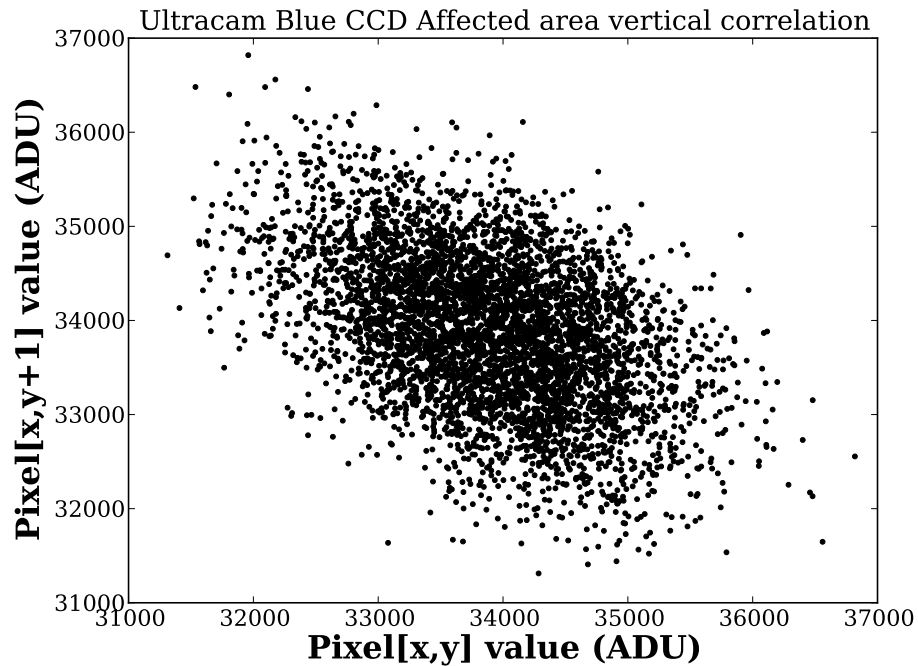


Figure 8.10: The vertical pixel correlation in a peppered area of the blue CCD.

erate of Ultracam and also increase the smearing of the image), several suggestions have been made (P. Jorden, private communication) that may progress investigation into the effect. Firstly, drive the CCDs less into inversion. This will have the side effect of increased dark current. Secondly, modulating one of the CCD phases during integration which will cycle the inversion periodically. An investigation into whether temperature has an effect upon the phenomenon could also help narrow down the physical cause. Finally, there is the issue of the peppering pattern apparently being fixed for a length of time before changing. Ultracam, being a visitor instrument, is quite prone to electromagnetic interference, especially when data cables are placed close to power cables. If these cables had a varying resistance to ground, this could potentially affect the clock propagation effect, with the resistance changing with the interference and thus telescope position. An exchange of cables between the blue/green and red arms could prove revealing as to whether impedance in the data cables is a factor with this problem.

8.3 The Ultracam archive software

Ultracam has been run at least once per year (sometimes three times) since 2002, with up to 80GB of raw data generated per night. More recent runs with the instrument at the NTT in La Silla, Chile have involved operating Ultracam for nearly 2 months during winter, meaning that the actual quantity of data being handled is now growing exceptionally large. Neglecting issues such as transport and backup of the data sets, there is still the more practical difficulty of finding one data set in over 8,000 or even remembering what targets have been observed with the instrument. Initially this was done from a generated HTML page for each night, using scripts. While this has sufficed until now, as it offers a method for browsing the data, there is still no convenient mechanism for searching for a target.

Other issues have become increasingly prevalent during the actual observation phase. Most of these originate with the fact that Ultracam is operated as a visitor instrument, and so has no integration with the telescope software. The observer must enter a target name and the filters used every time there is a change. Inevitably, errors creep in as target names are spelled incorrectly or are described by an alias, and each observer ended up describing the filters used in their own style (with or without primes, with or without spaces between arms, etc.). The only way to tell apart the type of the data taken is to examine it or hope the observer has written accurate comments for the appropriate set of exposures.

The Ultracam archiver software introduced here rectifies most of these problems. The core of the entire code base is built around the construction and utilisation of an SQL (Structured Query Language) database, which uses a simple language that is ideally suited to performing the types of searches needed in large data sets. The “back-end” software is created with the function of building this database and interacting with it and the Ultracam data at a low level. The “front-end” software is a web-based interface that transparently performs queries and returns information on searches and requested data for the end-user.

All of the software is designed to be simple to maintain, modular, cross-platform and open source. The front-end, back-end, SQL database and file server can potentially each be on a completely independent computer system, and the flow of information is concisely shown in figure 8.11. The first and last goals limit the choice of languages and servers considerably. The front-end was written in the PHP 5.2 language, whereas the back-end is in Python 2.5. Both of these languages are open source and have particular strengths. For example, PHP is not typically used (or often available) as a command-line language, whereas Python is still maturing

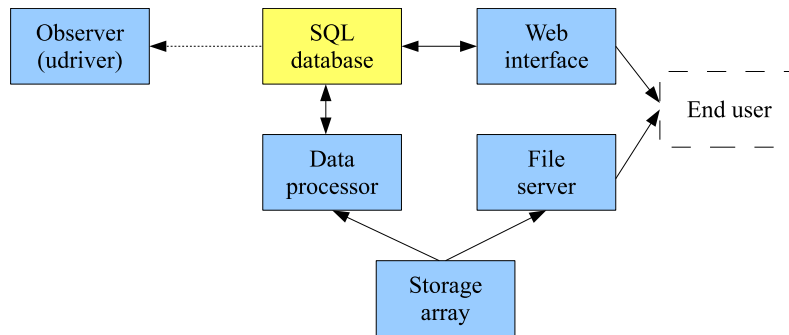


Figure 8.11: Flow diagram of the interacting components of the Ultracam archive software.

for web-development. The HTTP server was chosen to be Apache 2.2 which is a well-known, robust server that can equally host on Linux as Windows. The last difficulty was that while the back-end code could potentially harness a lot of the algorithms made for the Ultracam software pipeline, this software is not cross platform or open source (due to the inclusion of SLALIB), plus it requires a large dependency stack.

8.3.1 The back-end software

The code responsible for database generation and low-level interaction is the “UAC” module for Python. This is split into various namespaces that include `uac.meta`, `uac.raw`, `uac.astro`, `uac.db`, and `uac.simbad`. The reason for this is that only necessary functions related to a script can be loaded, as the import overhead in Python can be large. In addition to this, `uac.meta` can be imported with a “plain” Python installation, whereas `uac.raw` requires the Numpy module and `uac.db` requires MySQLdb (again, both are freely available and open source), so the actual dependency stack can be very low.

The disassociation with the Ultracam pipeline does require that the back-end software be able to directly read the raw Ultracam data, however. Most information can be extracted from the accompanying XML file, but timestamps are only in the raw frame headers, and it was decided that it would be useful to store preview images of each run as well. The `uac.meta` code deals with the XML (and user comments), as well as timing correction and estimation of the actual cycle time (to predict the duty cycle). The `uac.raw` code can access individual frames and de-multiplex them into individual arrays, plus perform cosmetic operations that produce useful preview images. Each of these pieces of code either adds to or modifies a Python dictionary that describes all the useful information obtained (this data is not hidden away in a

user-defined class). The `uac.simbad` code can perform object lookup queries to the online simbad database with an arbitrary query format, but by default returns the “accepted” object name, location, and magnitudes. Finally, `uac.db` acts as a buffer between the user and MySQLdb directly, generating queries from the dictionary and other parameters supplied.

A subtle problem is identifying each Ultracam run individually. Internally, the Ultracam server is only aware of the files for a particular observing night (starting from “run001”) which are then archived and cleared out at the beginning of the next night. Using a GPS start time should in theory give a unique timestamp for every run, but potentially the GPS may be disconnected or satellites may not be visible to generate the timestamp. The file timestamp should again not be considered a robust measure as the data that ends up on the Ultracam storage array which is imported is a copy, and the timestamp may not be preserved. There is also the path of the individual data set, but since Warwick and Sheffield maintain their own copies of the Ultracam data, enforcing a particular structure would make the software inflexible to different network architectures. It was decided that the SHA-1 checksum would be used to identify an Ultracam raw data file, which is a 160 bit one-way cipher based upon the file contents. While slow to generate (the whole data file must be read), it will generate a unique code that not only identifies the run, but can be used to verify the data has not been changed. Even a “power on” run of 24 bytes (which does not contain any frame data) can typically be uniquely identified due to the timestamp. There is a small probability that in the case of a GPS failure/disconnect two of these power on data files may be identical, but these are not usually of any interest to the observer. Once at least one frame is present in an Ultracam data file, the chances of a SHA-1 collision become vanishingly small, even without a timestamp.

8.3.2 The UAC SQL database

An SQL database consists of a series of tables with a certain structure of fields that contain data (termed entries). For the Ultracam archive, a semi-normalised structure is adopted that optimises query speed and accuracy while maintaining a manageable set of tables. The largest table (in terms of fields) is the one containing an entry for each data set. A new data entry is only made by the import code if the SHA-1 hash is unique. This new entry then gets a new number termed a “UAC ID” which is a unique (and human-friendly) integer that identifies a specific run within the database software. The images table is the largest in size, where each image is parented to a UAC ID (so more than one image can belong to a single

run). Other tables include filters, grisms, telescopes, and for the purposes of the front-end software, a users table. The uac.meta code attempts to disassemble the string of filters by matching long filter names (such as “HAlpha”, “Sodium”, etc.) first and then deals with the remaining characters. Once a filter in each arm is positively identified, the filter fields in a data run entry becomes an integer for the ID in the filters table. A similar process is performed for the grism and telescope fields. Another important table is the one containing a list of targets. Each target will have a unique ID together with a SIMBAD name (which the code will prefer), a possible list of aliases (which are automatically updated every time a new one is found), and a right ascension and declination. One aspect of the data table that is not normalised is the inclusion of a “flags” field. This is in fact just a 64-bit number, where each bit can be attributed to a specific flag name, and is used for declaring the data type (automatically inferred by the back-end software from the user comments) and other useful information. A new addition here was the inclusion of flags that specify caution (such as indicating that a data run is actually an acquisition, or there is a timing glitch) or that the data is completely useless (as marked by an observer or a CCD failure in a calibration set).

8.3.3 The “front-end” software

The front-end written in PHP has no dependencies other than a modern language version and a compatible web server. Most of the code is written in an object-oriented style with the HTML completely separated from the actual code via use of templates. This means that the front-end layout can be changed merely by bringing up the template files in an HTML editor, and the style by editing the CSS (Cascading Style Sheet) file. Should extra functionality be required, the object-oriented style results in a concise code format that is not entirely dissimilar to Java, with most of the objects automatically performing their necessary functionality simply by their instancing (for example, query dialogs will retain their entries between submissions, so that they can be edited and tweaked). A large class will generate UAC compatible queries from an array passed to it, and includes a layer of security so that, optionally, only runs with a specific ID or belonging to a certain PI will return in searches for specific user accounts.

Accounts with a specified access level are granted the ability to modify the meta-data in the database. As mentioned, errors will undoubtedly occur when entering information during an observing run; perhaps the filters are wrong or what was previously thought to be junk data is actually useful. The front-end code allows users to add extra comments to specific runs and update the database accordingly

(these changes are logged in a messages table).

Finally there is the capability of a user (with sufficient access level) to request a data set. The request process is in fact simply the posting of a new entry to a jobs table where a job server (which polls the database) will execute a specific set of code depending on the job type. The data set will be copied (and compressed) from the storage archive to the file server with a randomised path (to prevent anonymous access). Once prepared, the job is marked complete and optionally an email can be sent to the user with a download link. The job server includes dependency in the sense that one job can only execute once another is done, which allows for a second task to be submitted that will clear the data from the file server after a set amount of time.

8.3.4 Improvements to the observer interface software

The presence of a master database of targets and an accepted “flag” system to mark the data type allows for several improvements to the user interface presented to the observer during a run. The filters text box has been replaced with three drop-down boxes with filters appropriate to each arm, preventing conflicting styles and errors. Check-boxes are supplied that flag the run type (data, flat-field, bias, technical, etc.). Finally the target name can now be verified before starting a run. The application will connect directly to the UAC database and query the target names to see if the supplied name is an accepted one or an accepted alias. If this fails, it will then query the SIMBAD database similarly. Should both return a negative result, the user will be warned. In some cases, the target will have no name (such as for GRBs, optical transients), in which case the observer is encouraged to mark the coordinates in the observer comments log and the target can be classified later when the data is imported.

8.4 Summary

This chapter detailed three separate investigations that led to the improvement of Ultracam as an overall system. For observers concerned about the photometric accuracy of the filters, the difference between the SDSS filters and the Ultracam equivalents have been examined, and for most cases found to be under 0.1 mag. Should very accurate absolute photometry be required, the code produced here could predict this, given a template spectrum that approximated the desired target. Secondly, another concern to observers, “peppering” was shown to be simply a modulation problem and not an actual loss of charge, downgrading this issue from

severe (indicating unuseable data) to quite low (an increase in noise). Possible courses of action are suggested for Ultracam to further track down the source of this issue and perhaps offer a correction. Finally, a newly developed archive system was described that not only allows for increased productivity during data reduction, but also observing due to the inclusion of a normalised database that indexes targets, filters, and grisms for Ultracam (and Ultraspec).

The next chapter will examine early commissioning data of Ultraspec. Unlike Ultracam, it is still a relatively new instrument and requires careful attention to extract accurate data. The chapter will detail the calibration of critical CCD parameters, and investigate and offer corrections for several issues that affect the quality of the data. This will be followed up with preliminary scientific data to show the effect and capability of the reduction.

Chapter 9

Calibration and reduction of Ultraspec data

The virtues of Ultraspec when used for observing cataclysmic variables and similar phenomena have been described in detail in §4. Essentially, it allows spectra to be obtained at comparatively high frame rates, so that they can be studied at an otherwise impossible time resolution. Ambitious projects already investigated with Ultraspec include phase-binned spectra from the Vela pulsar (which has a period of 89ms). Ultraspec saw first light in December of 2006 with a short 3-night run at the ESO 3.6m telescope in La Silla, Chile, and returned in February of 2008 for further testing and the first science products. Using commissioning data from Ultraspec, preliminary results will be shown as an example of the output that the data calibration and correction pipeline detailed here can do.

9.1 Observations and data reduction

The data used here are from three separate targets; OY Car, Z Cha, and QS Vir. The first two targets have been introduced in §6. QS Vir is a detached system where no mass transfer is taking place and consists of a DA white dwarf and a dMe white dwarf. Its status as a possible “hibernating” CV has been proposed [O’Donoghue et al., 2003], but brought into question by Parsons et al. [2010b], who conclude it is a pre-cataclysmic binary.

Observations of all three targets were taken with Ultraspec during February 2008 on the ESO 3.6m telescope, using the EFOSC2 G3 grism for OY Car and QS Vir (3750Å-5150Å), and the VPH656 holographic grating for Z Cha (6350Å-6830Å). Conditions for OY Car were variable, but corrected by a comparison target. Z

Cha and QS Vir had photometric conditions with 1"-2" seeing. Reduction of the Ultraspec data is detailed in the next section. All the observations are listed with times of observation in table 9.1.

9.2 Ultraspec calibration and data reduction

While a basic reduction utility is present in the Ultracam/spec pipeline, it was primarily designed to monitor live observations and as such only offers basic features. No accurate numbers for key quantities such as gain and readout noise have been previously collected for Ultraspec. Also there are a number of features present in the data which must be taken into consideration for an accurate reduction; these include a low-bias tail, spurious charge, and electron trapping, which will now be dealt with in turn.

9.2.1 Calibration of CCD parameters

As discussed in §3, readout noise and gain are critical parameters for a CCD, since they define tangible values such as dynamic range and signal-to-noise. Firstly, the normal gain (e^-/ADU) is determined, separate to the “avalanche” or “HV gain” which defines how many electrons are produced in the secondary-emission register (e^-/e^-). This can simply be done by reading out from the normal output with Ultraspec.

The formal definition of gain in terms of the physical CCD electronics is shown in equation 3.6. While, in theory, each of these parameters could be measured, it is challenging to do so in a precise manner and even more difficult to do so once the instrument is assembled and the CCD is hidden away inside. Instead, parameters are measured via a useful practical method instead. Starting from

$$S = \frac{P}{K}, \quad (9.1)$$

where S is the signal (ADU), K is the camera gain constant (e^-/ADU) and P is the number of detected photons per pixel, the error propagation is considered and it is found that the variance is

$$\sigma_S^2 = \left(\frac{\delta S}{\delta P}\right)^2 \sigma_P^2 + \left(\frac{\delta S}{\delta K}\right)^2 \sigma_K^2 + \sigma_R^2. \quad (9.2)$$

The readout noise has been added in the last term which is effectively the output amplifier noise floor. If it is assumed that K has negligible variance ($\sigma_K^2 = 0$) then

OY Car (Ultraspec, G3 grism)		
Date	Start (UTC)	End (UTC)
4th February 2008	06:17	08:19
5th February 2008	02:22	02:49
	04:30	06:35
6th February 2008	03:39	07:35
7th February 2008	00:58	01:40
	04:27	05:18
	06:53	07:49
8th February 2008	04:07	06:42
	08:50	09:29
10th February 2008	04:10	05:56
Z Cha (Ultraspec, VPH656 grism)		
Date	Start (UTC)	End (UTC)
5th February 2008	01:43	02:01
	03:32	04:14
	08:18	09:24
8th February 2008	00:49	01:03
	07:53	08:40
9th February 2008	03:44	04:17
	05:21	05:59
	07:06	07:42
10th February 2008	01:15	02:00
	05:17	05:55
QS Vir (Ultraspec, G3 grism)		
Date	Start (UTC)	End (UTC)
6th February 2008	07:49	08:48
7th February 2008	05:45	06:41
8th February 2008	06:55	07:40
9th February 2008	08:19	05:59

Table 9.1: Journal of observations.

the second term vanishes and $\frac{\delta S}{\delta P} = \frac{1}{K}$ from differentiation, resulting in

$$\sigma_S^2 = \left(\frac{\sigma_P}{K}\right)^2 + \sigma_R^2. \quad (9.3)$$

Making the assumption that the output is shot-noise limited, hence $\sigma_P^2 = P$ this results in

$$K = \frac{S}{\sigma_S^2 - \sigma_R^2}. \quad (9.4)$$

This is a remarkable result, because the camera gain can now be determined through measurements of data from the CCD alone. Measurements of σ_S can be performed by using a flat illumination source and then increasing the exposure time in steps to cover the full range of the CCD. Due to the three separate noise sources (§3.2.5), if a plot of $\log S$ against $\log \sigma_S$ is made, three different regimes are observed that are determined by independent variables. The read noise floor will not depend upon the signal and hence have no gradient. The shot noise is \sqrt{S} and will thus have a slope of 1/2. Finally the fixed pattern noise is directly proportional to the signal and will have a slope of 1. It should be noted that if logarithms of equation 9.4 are taken,

$$\log K = \log S - \log(\sigma_S^2 - \sigma_R^2), \quad (9.5)$$

then, assuming that the data is shot-noise limited, $\sigma_R = 0$ and since $\sigma_S = 1$, this means that $\log \sigma_S = 0$. Thus it is found that, where the shot-noise slope intersects the x axis, the result is K, the camera gain constant. This analysis was performed for Ultraspec data, with the resulting photon transfer curve being seen in figure 9.1. The gain of Ultraspec was found to be $15e^-/\text{ADU}$ with a pixel non-uniformity of 0.3%. These parameters are expected to be fixed, as the former is determined by the CCD electronics and the latter by the CCD manufacturing process.

Calibration continues next with CCD parameters that are expected to change with time, such as readout noise and the secondary ‘‘avalanche’’ gain. These can be determined from a sufficient quantity of bias frames in the correct mode. When plotted as a histogram, the peak will be the bias level, with the width of the feature representing the uncertainty in this level, and hence the read noise. In a standard CCD there should be no other features, but in an EMCCD, a long tail towards the higher count level should be visible due to the nature of some electrons passing through the gain stage. Using a model for an EMCCD (T.R. Marsh, unpublished), parameters determined previously are fixed, such as K , the bias level, and the number of amplification stages which are known. The other parameters are fitted, including the readout noise, secondary gain, and also the probability of clock induced




Figure 9.1: The photon transfer curve generated from Ultraspec data of a flat-field with increasing exposure time.

Year	Readout mode	HV gain	g (e^-/e^-)	σ_R (ADU)	P(CIC)
2006	Medium	3	691	1.89	9.65×10^{-5}
2008	Fast	9	2689	4.1	4.92×10^{-4}
2009	Fast	9	2750	3.28	2.81×10^{-4}

Table 9.2: The mean avalanche gain, readout noise, and probability of clock induced charge for Ultraspex in 2006, 2008, and 2009. Since the gain and readout noise will change depending upon the readout mode and HV gain parameter, the most common values used during the year are given.

charge (dealt with in §9.2.3). The resultant plots can be seen in figure 9.2 for 2006, 2008, and 2009, with the relevant numbers detailed in table 9.2. The effect of clock induced charge on the best fit can be seen in figure 9.3. The abnormal long tail that is not predicted by the model is dealt with in §9.2.2.

The effective gain of an EMCCD

Most reduction software will only supply a single gain variable which is used for computing the variance of the input data and hence determine reliable uncertainties. The *effective* gain must be determined, which consists of both the CCD and avalanche gain to supply as input. In §4.4.2 the variance of an EMCCD was computed in electrons; if this is converted to ADU using $G(e^-/ADU)$

$$\sigma^2 = \left(\frac{\sigma_R}{G}\right)^2 + \frac{2g^2 S_e}{G^2}, \quad (9.6)$$

where S_e is the signal in electrons. By definition, the signal after the avalanche register is $S'_e = gS_e$, and the signal in ADU will be $S = \frac{S'_e}{G}$ which results in

$$\sigma^2 = \left(\frac{\sigma_R}{G}\right)^2 + 2\left(\frac{g}{G}\right) S. \quad (9.7)$$

Taking R to be the readout noise in ADU, the final result is:

$$\sigma^2 = R^2 + \frac{S}{\frac{G}{2g}}, \quad (9.8)$$

where $\frac{G}{2g}$ is the effective gain. For Ultraspex this is ~ 0.01 .

The VPH656 grism

Since no arc line map was available for EFOSC2/VPH656, one is presented here in figure 9.4 for reference, using the standard argon lamp available with EFOSC2.

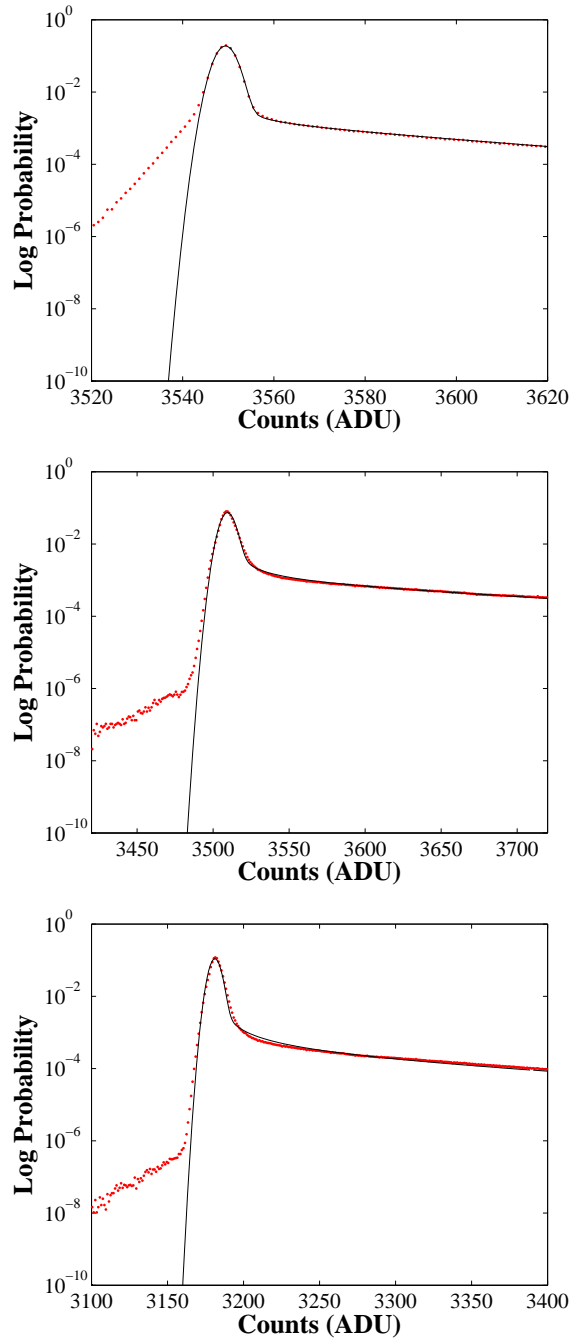


Figure 9.2: Ultraspec gain, spurious charge, and readout noise model fitting for 2006 (top), 2008 (middle), 2009 (bottom). The extended low bias tail is particularly notable in 2006, starting at $\log P \sim 10^{-2}$, but is well controlled by 2009, where it starts at $\log P \sim 10^{-6}$ with the lowest values a further two orders of magnitude down.

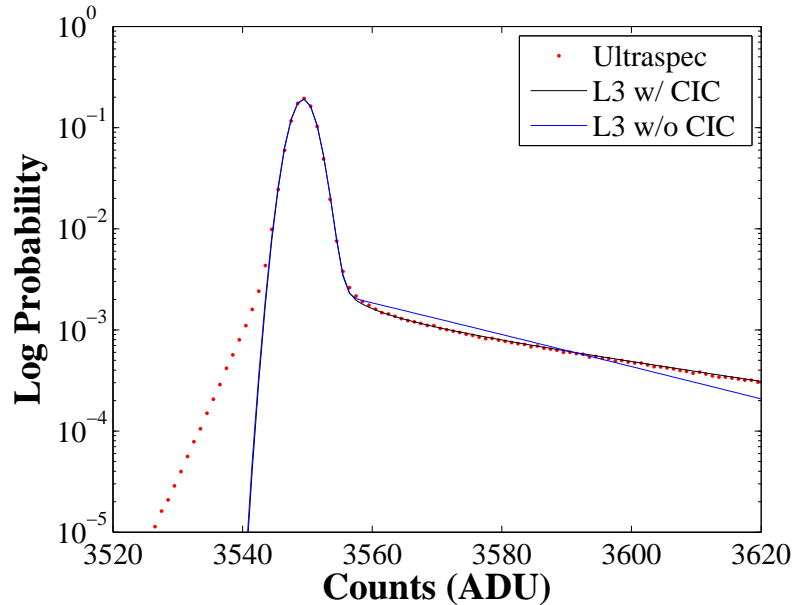


Figure 9.3: Best fit for an EMCCD model with and without including a probability of expected unstructured clock induced charge to Ultraspec data.

The residuals when using a 4-term polynomial for wavelength calibration resulted in 0.0256\AA .

9.2.2 The low bias tail

Figure 9.2 showed a notable discrepancy with the model of an EMCCD as seen by a long tail at low ADU; this is indicating there is a greater than predicted quantity of pixels with a low count. This seems in complete contrast to the basic configuration of the CCD - the bias level is an artificial offset to ensure that all possible numbers detected by the device can be output as an unsigned integer. There is no physical way to detect negative counts, and since the bias level should not be changing by any easily detectable fraction, this feature was investigated. Performing lowpass cuts on the bias frames with the observed low tails did not give any insight into the problem. Occasionally there was a gradient of affected pixels, but sometimes only lines of affected pixels. When a similar histogram was computed for the OY Car data, it was noted the tail became even more extended and with an increased probability, indicating that infalling light was proportional with the low bias tail. This was confirmed by performing lowpass cuts on the OY Car data and then adding the frames, as shown in figure 9.5. Here it is seen, somewhat

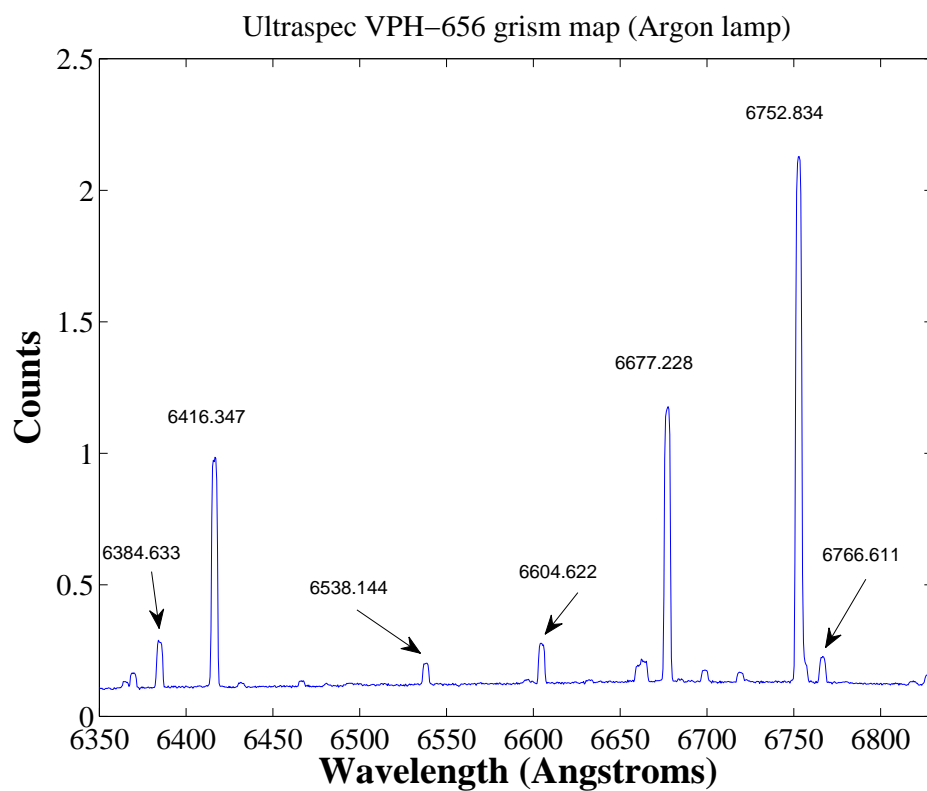


Figure 9.4: Arc line map for argon with the VPH656 grism on Ultraspec with EFOSC2.

surprisingly, that the pixels that are abnormally low directly represent the data; the probability of an abnormal count is directly proportional to the flux. When each of the lowpass cut frames was examined individually against their equivalent highpass cuts, it immediately became apparent that an abnormally low pixel directly trailed a high flux pixel in the serial readout direction (ie. horizontally). This correlation of pixel value against adjacent pixel value can be seen in figure 9.6.

A fit to this correlation can be performed but is typically made awkward simply due to the statistics involved; a very large number of frames must be processed to get a good display of the correlation. Proceeding in a crude manner, the fitting method performed was to find the minimum adjacent pixel value for each pixel value, and only accept those that were sequentially decreasing (shown in figure 9.6). For the 2006 data, this was found to be $3758 \text{ ADU} - 0.0461376 \text{ ADU/ADU}$. This is a measure of the effect a pixel count has upon an adjacent pixel count for these abnormal pixels, and can be used as a partial correction for this effect.

Since there is no way for charge to be distributed in such a manner leading to values less than the bias level before the sampling stage, the CDS circuit was marked as the likely culprit. It has since been confirmed that the video processor is responsible (S. Tulloch, private communication), due to its AC coupling. As the circuit was acting as a highpass filter, if the lower 3dB point was incorrectly set too high, it caused a negative overshoot on the following pixel when a signal resembling a delta function was passed through. In effect, some proportion of the first pixel's signal is being added to the zero point for the next one sampled, so a greater value than the real zero point is subtracted resulting in the effect observed. Starting from the right-hand side of the chip (ie. the direction of readout for the avalanche mode), taking the fraction of the pixel's count *above* the bias level and adding this to the following pixel undoes this artificial effect. See figure 9.7 to compare the original histogram with the corrected data for 2006.

Modifications to the video circuit for the 2008 data resulted in a correlation that can be seen in figure 9.8. It is significantly weakened but now in fact undershoots slightly rather than overshoots, meaning that low bias pixels are inversely proportional to flux, albeit at a lower proportion to 2006. This problem appears to have been rectified to a negligible level from 2009 onwards, and appears to be common in this type of CCD, as testing has shown it is present in QUCAM as well.

9.2.3 Spurious charge

Spurious charge manifests in all CCDs that use inversion, the process whereby the gate voltage is driven negative and holes recombine with midband conducting elec-




Figure 9.5: Lowpass cuts of the Ultraspec OY Car data at 3535 ADU (top), 3530 ADU (middle), and 3525 ADU (bottom).

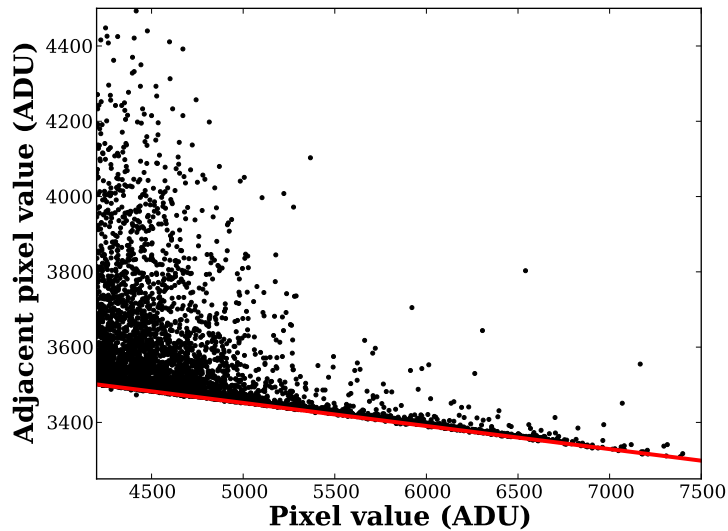


Figure 9.6: Pixel counts against their adjacent pixel counts for 2006 OY Car data. A strong correlation is seen indicating that bright pixels are having an abnormal affect on their neighbours. The red line is a fit to the minimum adjacent pixel value for each pixel value, effectively measuring this correlation.

trons at the surface, reducing dark current (§3.2.4). This is mandatory in the case of Ultraspec, because dark current will be amplified when using “avalanche” gain, thus negating the benefit of such a the electron-multiplying process as the dark signal will be indistinguishable from the actual signal. Typically, spurious charge manifests as random spikes that appear as signal charge (and the noise it carries is similarly characterised by shot noise). During inversion, when holes are driven to the interface layer and collect beneath the gates, some of these may become trapped. When the gate voltage is switched back to its non-inverted state, some of these holes are accelerated back towards the main silicon layer with sufficient energy to cause impact ionisation; this creates electron-hole pairs, and these electrons are collected in the nearest potential well. Occasionally spurious charge can be mistaken for dark current (as they both possess shot noise), but the former will depend upon clocking, not exposure time as with the latter. Also, due to carrier mobility, spurious charge will become a greater problem with a lower temperature, in contrast to dark current.

The process of transferring charge down a CCD is termed ‘clocking’, and is performed by a clock signal which effectively changes the potential barrier between phases and thus shuffles the charge packets along. This clock signal has an associated swing (amplitude) and rise time (the rate of change of potential), which can have

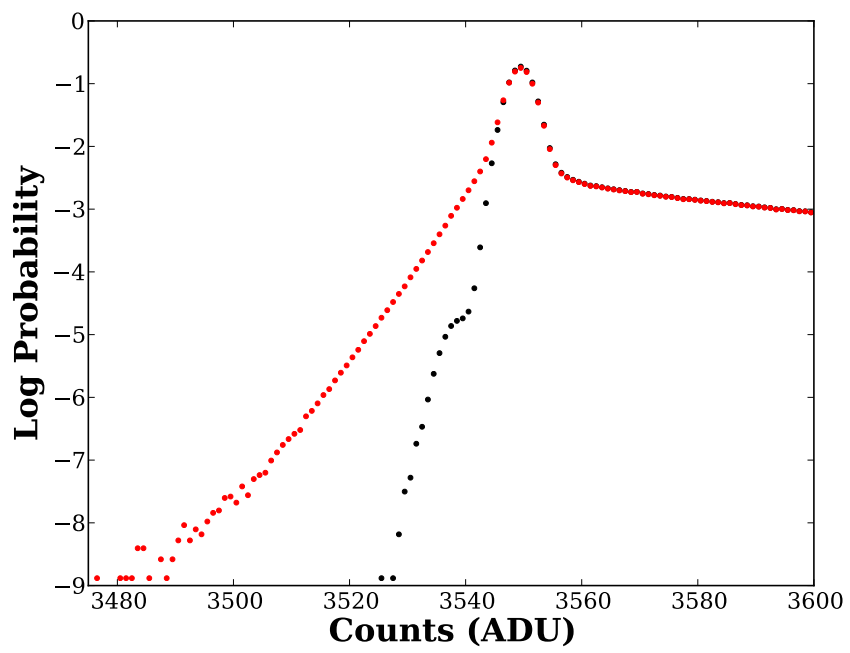


Figure 9.7: Counts histogram for 2006 OY Car data before and after the removal of the correlation between neighbouring pixels.

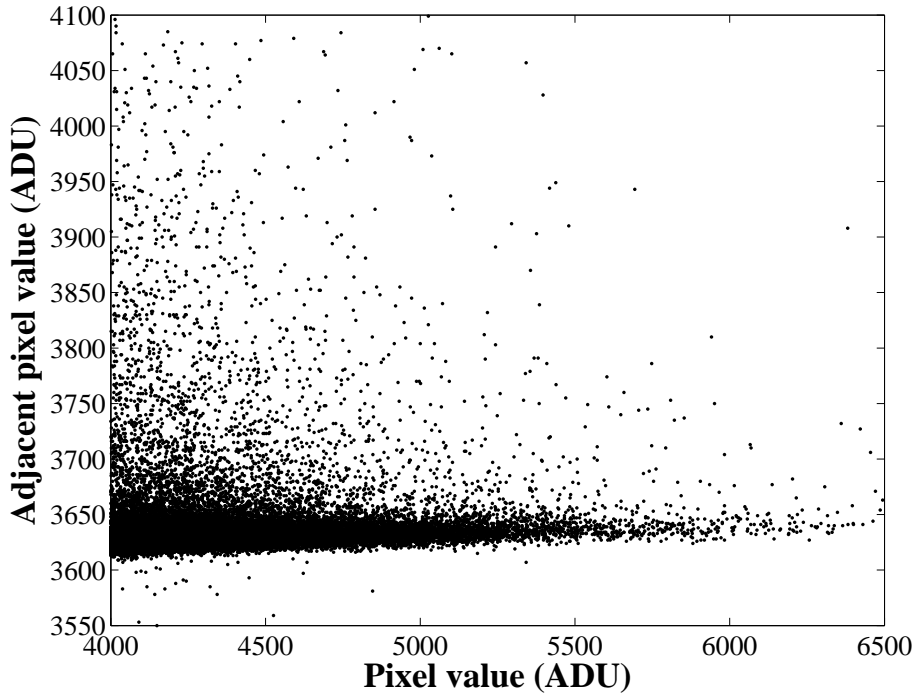


Figure 9.8: Neighbouring pixel count correlation plot for 2008 data.

great effect upon the performance of a CCD. While a stronger clock swing and faster clock rise time will worsen the problem of spurious charge as the holes will be accelerated further (which increases linearly with the number of transfers), the effect is still usually on the order of a few electrons. Thus, in a normal CCD, such an effect is masked by the read noise. This changes when using an EMCCD as they are collected as signal electrons and are thus subject to the same gain as any photo-electron.

It was noted that Ultraspec appears to suffer a higher than usual quantity of spurious charge when powered on, and that this charge appears to have structure. Figure 9.9 shows a frame from Ultraspec with gain enabled after power on, with a straight line through the summed rows achieving a good fit. This quite clearly shows the linear relationship of spurious charge with clocking. Performing this same fit on a sequential series of exposures, the gradient of these bias frames is plotted against time in figure 9.10. It can be seen that after power on, a vertical gradient with an amplitude of ~ 20 ADU from bottom to top will be seen, but that this is effectively halved after 10 minutes. This time dependency can be well-modelled by

an exponential with a linear term, in this case

$$\frac{\Delta Y}{\Delta \Sigma X} = 15.2 \exp^{-0.229t} - 0.0294t + 3.93, \quad (9.9)$$

where t is the time in minutes since power on¹. Since the exponential term becomes negligible at large time, it is found that the structure in the spurious charge should be eliminated after 133 minutes of run time. Random spurious charge events will still be present, but they should lack the vertical gradient structure. Further sets of bias exposures with the HV gain enabled would allow further investigation of the effect, particularly with the new 2009/2010 configuration of the instrument which is likely to have changed.

The source of this spurious charge could be imagined to be related to the operation of the electronics during power on. A large electric field could potentially drive a large number of holes into the interface. Their positioning must be reasonably uniform across the interface so as to achieve the structure seen (the linear proportionality with rising clock edge means that the top row will have undergone the most clocks and thus possess the most spurious charge in this case). Since figure 9.10 shows a decrease in the parallel gradient during time between exposures, this indicates that a significant number of these holes are decaying over time. The remaining *trapped* holes do not decay so rapidly, and it is here that clocking will help to remove them.

While potentially a minor issue to photometric accuracy, Ultraspec is mostly powered on in the mid-afternoon for calibrations, meaning that any vertical gradient should have vanished by the time of actual target observations. Any calibrations taken with the normal amplifier will make the spurious charge gradient invisible due to the readout noise, and in fact there should be little reason for the secondary gain to be on for calibrations. The only potential issue this could cause is if there was a system crash during the middle of the night when the HV gain was required for observing. Since with Ultraspec the slit is placed along the serial direction, any significant mean gradient would be removed by linear sky background subtraction, but would still add a small quantity of shot noise.

¹These values are for 2008 data, since long runs of HV gain biases are not done frequently. Since the operating temperature was changed in 2009 it is likely that Ultraspec will now reach a stable spurious charge state more quickly.

CCD 1

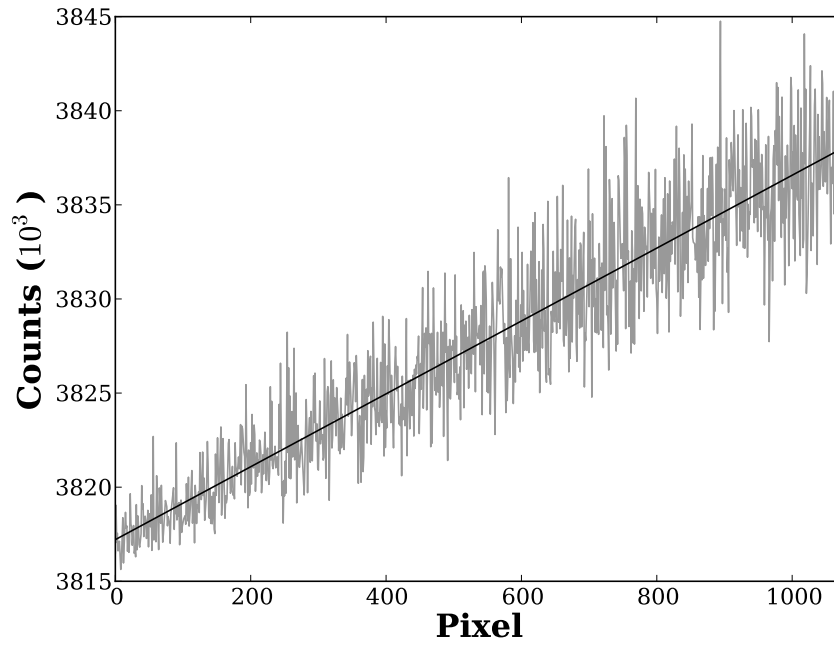
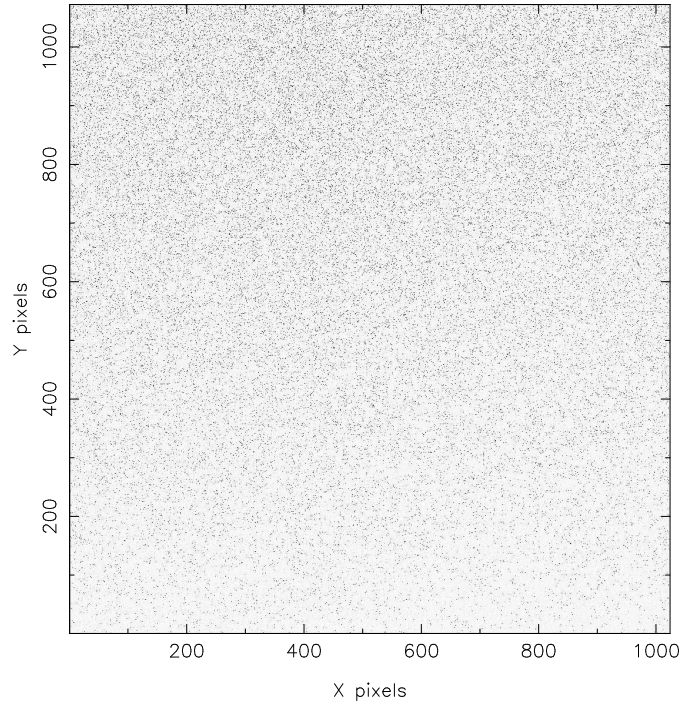


Figure 9.9: Top: a bias frame from Ultraspec straight after power on with avalanche gain enabled. Bottom: the sum of each row from the same frame with a straight line fit.




Figure 9.10: The parallel gradient of Ultraspec in relation to time since power on. The fit is modelled as an exponential plus a linear term. The instrument went through several stop/start cycles during the taking of this data.

9.2.4 Electron traps

While the CDS circuit overshoot and spurious charge gradient detailed in the previous sections are for the most part, subtle issues that only partially affect the quality of the reduced data, there remains one final problem that is much more significant and readily noticeable, which is charge trapping. An example of this can be seen in figure 9.11, which shows 2008 data of Z Cha. In the upper frame (the comparison star) nearly 1 in 10 columns shows an extended trail in the parallel direction, with the more extensively affected columns also showing a ‘gap’ of charge on the side of the slit opposite the tail. The net effect of this is the signal charge (the desired signal electrons from the target) on these columns appears shifted to a higher position and extended. A typical extraction will sum along the slit and hence will show single column dropouts of varying magnitude. This problem becomes two-fold when using utilities such as STARLINK as the sky area will now be polluted by this trailing charge and an excess will be subtracted from the already diminished columns.

This is generally considered a local CTE problem in regards to the CCD, as the charge is not being completely transferred with each clock cycle. Due to the speed Ultraspec is operated at, the effect can not usually be seen in single frames due to the very low count rate, but, as shown, is immediately apparent when summed.




Figure 9.11: 3000 summed frames from raw Ultraspec data of Z Cha in 2008. The data was read out in two windows which spanned the entire width of the chip (of which only a section is shown) but the height of the windows is as seen. The slit is in the serial (horizontal) direction, with the lower object being Z Cha itself and the upper showing its comparison star. The level cut on both images is identical.

Further study has shown that it only appears to affect data taken via the electron-multiplying mode and that fainter targets (such as Z Cha) are more heavily affected than brighter ones (OY Car).

Electron trap theory

The culprit for this is electron traps within the CCD in the parallel direction. Traps are broken down into four categories; design traps, process traps, bulk traps, and radiation-induced traps [Janesick, 2001]. The last option is almost entirely found in space-based detectors for obvious reasons, and bulk traps are caused by impurities in the silicon and are typically confined to single electron traps (which hence become the ultimate limiting parameter of global CTE). These two types of traps are usually found as proportional traps (extracting a quantity of charge from a packet that is proportional to its quantity) whereas for Ultraspec it is seen that a *fixed* quantity of charge is extracted since the effect was seen to be worse in fainter targets.

Local charge traps formed due to a process in the CCD’s manufacture are not particularly uncommon, whereas a design trap is an error in the actual layout or doping of the CCD, which creates a fixed (though unintentional) potential barrier which acts as a trap. The former option is considered to be the most likely here (as again, QUCAM does not exhibit exactly the same issue, implying it is not a design error), although the actual cause for most process traps is speculative. One possible model for the phenomenon is shown in figure 9.12. If there is a small hole in the oxide layer, the nitride layer is now close to the n-channel and electrons will collect at this interface and cause a potential “bump”. Like any potential well, sufficient electrons will fill the trap and subsequent ones will fall over and hence be transferred correctly along the parallel direction as expected, but those that remain require energy to escape. This energy mostly comes from thermal excitation, which is used in defining a trap’s emission time constant, which in a bulk trap is [Kim, 1979]

$$\tau_E = \frac{\exp(\frac{E_T}{kT})}{\sigma_e v_{th} N_C}, \quad (9.10)$$

where E_T is the trap energy level beneath the conduction band, v_{th} is the electron thermal velocity, σ_e is the electron cross section and N_C is the density of states in the conduction band. Using typical silicon and electron values, $\sigma_e v_{th} N_C \sim 2.8 \times 10^{11} \text{s}$.

Methods for correcting affected Ultraspec data

During the 2009 run, the temperature of Ultraspec was increased by 10K, which significantly reduced this problem, as now the electrons are re-emitted on much




Figure 9.12: A schematic for a possible cause of a local trap in a CCD. A “pinhole” in the oxide layer allows electrons to build up at the interface creating a small potential barrier.

smaller timescales. This has the drawback of generating an increased dark current, but this was considered a reasonable trade-off, due to this being quite a pronounced issue, and one that occurs frequently in the regime where Ultraspec is designed to be operated. Electron trapping will always be present on some level, however, and since data from 2006 and 2008 had already been taken, methods to rectify the data were investigated.

No correction needs to be taken on normal readout data. In this case, there are a large number of photoelectrons falling onto the CCD, and the traps will almost instantly be filled by the background sky level photons, plus there is a read noise floor. When observing faint targets (or targets with a very high time resolution), fewer photons mean that the traps are filled less quickly, and so the effect is more pronounced. Knowing the exact location of these charge traps requires a very faint but uniform light source that illuminates the CCD, and then a flat-field (but in gain mode) can be performed which should show sudden drops where the charge traps are present. Unfortunately this does require an extremely weak uniform light to see the traps and also not damage the avalanche register, and, to date, such a map has not been created for Ultraspec. It is obvious, however, that there are not many traps per column; examining the two slits in figure 9.11 does not show any sudden dips in the middle of the slits, indicating that the charge traps must be ahead of the data. However, there must be more traps between the target and comparison slit data as there are clearly more trails in the comparison slit which is higher on the CCD.

There are two potential objectives with such a problem. The first is to simply attempt to find how far the charge extends out from the slit and include it in the sum. While this will add in extra sky shot noise, the readout noise is still negligible, as the secondary gain will be on, and thus the degradation of the spectrum should not be significant. The second is to in fact correct the original image, and then any manner of tools that the end user typically prefers can be used instead. The former method benefits from modelling to predict how far the charge trail extends (instead of crudely just masking the entire column for inclusion in the mask); in the second modelling is a necessity so the process can be reversed. Modelling also has another advantage, in that it can predict charge that may be missing from the original data. The extremely deep trap that is present on column 82 (again, see figure 9.11) has a trail that clearly extends out of the windowing region. Thus, actual photo-electrons have been skipped out of the sampling process and even summing the entire column will not restore this charge. Modelling of electron trapping in CCDs and the reversal of this process has been performed for the heavily degraded Hubble detectors and is detailed in Massey et al. [2010].

The convolution method

Since the effect of a trap on the data is to apparently “smear” it out across the column with an exponential profile that originates from its emission constant, one simpler method is to attempt to fit a convolution of a gaussian with an exponential (of a specific height and falloff, representing the trap’s depth and probability of emission). The initial gaussian was found from a fitted gaussian to unaffected columns, and an offset to match the sky background was added in. This resulted in a reasonable but clearly not perfect fit, with most of the difficulty arising from an ambiguity in determining the peak of the gaussian. A moffat profile did not help here ($\beta \sim 5$ indicated that the shape was indeed very gaussian already) because it appeared to have quite a flat peak that the function could not match.

This is, in fact, quite clearly due to the fit being performed on the sum of the data. The ~ 3000 spectra were of varying flux, and the act of a local trap on a gaussian is to “eat into” the side of the gaussian (as electrons are dropped into the trap) which effectively moves its peak to a lower pixel in the parallel direction. From examining sums of high-count and low-count frames, it can be seen that the peak of the gaussian moves depending upon the incoming flux from the target during the given exposure. By simply splitting the spectra into two sets that are above and below the mean, significantly better fits are achieved with much less uncertainty in the gaussian parameters, which is shown in figure 9.13.




Figure 9.13: A single column from Ultraspec data of Z Cha, split into a low count regime below the mean of the spectra (top), and a high count regime above the mean (bottom). In each case, a convolved moffat profile with an offset is fitted to the data and overlaid.

While this method was not expected to offer a particularly good fit since (as discussed in the next section) the process of charge-trailing is not a linear process, the results are reasonable for such an easy method, not computationally expensive (one gaussian, one exponential, and a convolution per column) and provides a possible way to at least measure the extent of the tail in the generation of an extraction mask.

Modelling the electron trapping

Unsatisfied with the deficiencies in the simple gaussian convolution, the investigation moved to a model that approximates the actual physics of the issue. Here, a trap object that has a specific depth (number of electrons it can hold) and probability of emission (its emission constant) is created. Diagrammatically, this is shown in figure 9.14.

The initial, single-trapped gaussian case provided a rough fit for a trailed column (see the bottom panel of figure 9.15) but was clearly worse than the convolution model case, which was initially discouraging despite the fact all parameters were free. Since the position of the trap (or traps) was unknown but seemingly not within the slit area, the trap was placed far before the data so it would be in a steady-state condition before the data reached it. The two strange “jumps” in the model in this case are worth discussing and are perfectly reasonable in terms of the model used. The left side of the gaussian rises more slowly than typical before the first apparent discontinuity when it begins rising much more sharply. This slow rise is caused by the charge trap filling, and only a certain amount, P of what is captured is released per clock. When the trap is completely filled, the charge packets being clocked past are summed in addition to what is being emitted, causing this first jump in the model. The latter jump is for the opposite reason, where the trap is no longer in a full state and what remains of the gaussian goes completely into the trap to try and fill it. These discontinuities are not immediately apparent in the data, however, and are a source for the high χ^2 along with a gaussian profile that appears too broad. However, since it can be envisioned that the point the trap fills and empties will depend strongly upon the flux in any one exposure, it can be understood that with many stacked images to what is being fitted, these two points will be smoothed out as they vary from frame to frame.

Another alternative is that there are indeed more than a single trap in a column, of which varying combinations are shown in the other panels in figure 9.15. These further traps work to smoothe out these features, resulting in a believable fit with 3 pairs of electron traps (6 total, but only 3 different depths and probabilities

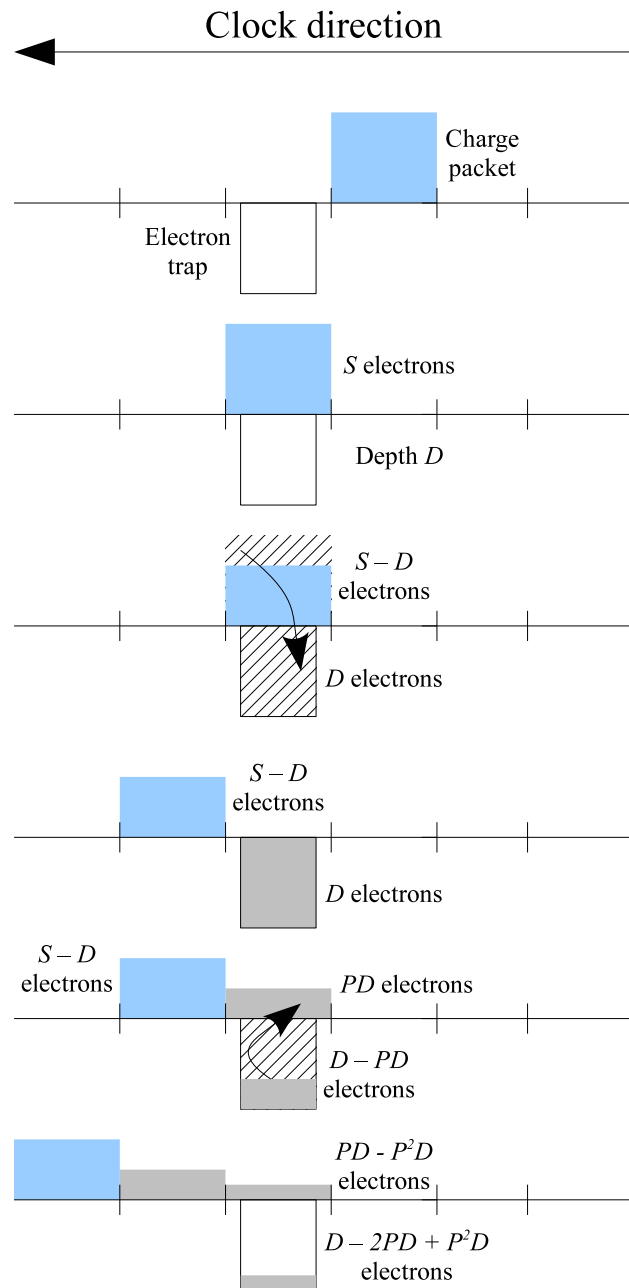


Figure 9.14: Schematic of the electron trapping object model in the simple scenario of a single charge packet, which is shuffled towards the left after each clock cycle. The trap has a depth D and emission probability P , and the charge packet has a quantity of electrons S .

Ultraspec image	$A = I + \delta$
Extra simulated readout	$B = I + 2\delta + \delta^2$
$A + A - B$	$C = I - \delta^2$
Simulated readout of C	$D = I + \delta - \delta^2 - \delta^3$
$A + C - D$	$E = I + \delta^3$
Simulated readout of E	$F = I + \delta + \delta^3 + \delta^4$
$A + E - F$	$G = I - \delta^4$

Table 9.3: Step by step procedure of the forward iterative algorithm used to remove trailing in images, where I is the original image (desired, but not available) and δ represents the trailing process in the readout.

and a χ^2 of 0.79. In this case, the fit is to the faintest third of the spectra (~ 1000) to try and narrow down the gaussian ambiguity further, but one could anticipate splitting the data set further to try and get more accurate fits to each regime. The ideal case, a fit to each exposure, is completely impossible due to the handful of photons present.

Reversal of charge trapping via a forward algorithm

Using the six electron trap model discussed in the previous section, the potential to reverse such a process to retrieve the original data is examined. A successful method detailed in Bristow [2003] uses a clever forward algorithm that works in iteration to reduce the effect of trailing. Considering the ideal image is I and δ represents the electron trailing, the procedure is detailed in table 9.3. In essence, if the model simulates the readout trailing, then the difference between the ‘simulated’ image ($I + 2\delta + \delta^2$) and twice the original image ($I + \delta$) results in an image closer to I ($I - \delta^2$). This process can be repeated as many times as deemed necessary. While the iteration will never analytically reach I , δ becomes insignificant very quickly in practice. The single iteration result can be seen in figure 9.16.

While the iteration process is very quick, it is still a time-consuming procedure to fit every single column with the model described. Whereas the CCD modelled in Massey et al. [2010] had a trap density greater than 1 per pixel, it is likely that the significant traps in Ultraspec are few in number (as it is a new CCD and is not subject to high-radiation conditions), and some columns are likely clean to detectable levels. Attempting to fit all six traps to every column results in fitting problems with some, usually flagged by unphysical parameters resulting from the fit (such as negative depth). For the data reduction presented here, the simpler convolution method is used to quickly gauge the length of each column trail and

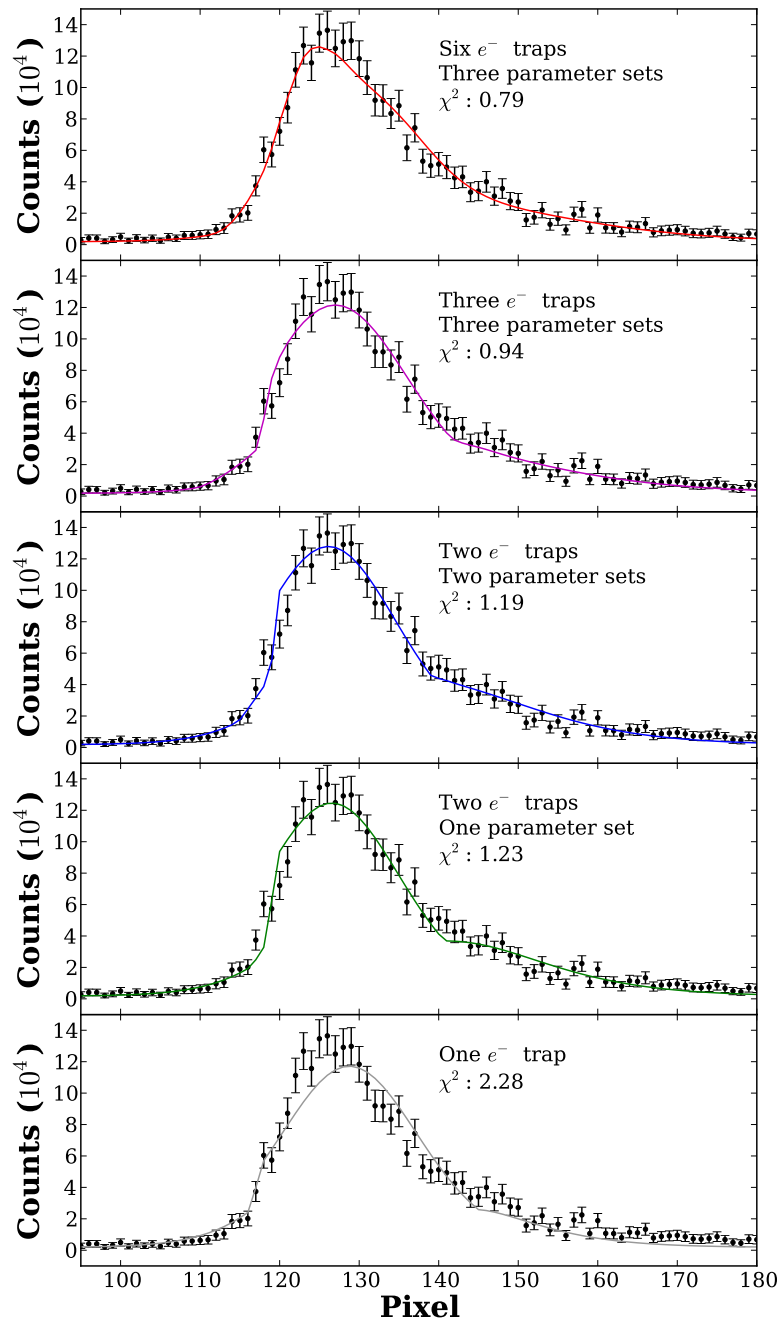


Figure 9.15: Fits to a column of Ultraspec data with varying numbers of electron traps and parameters.

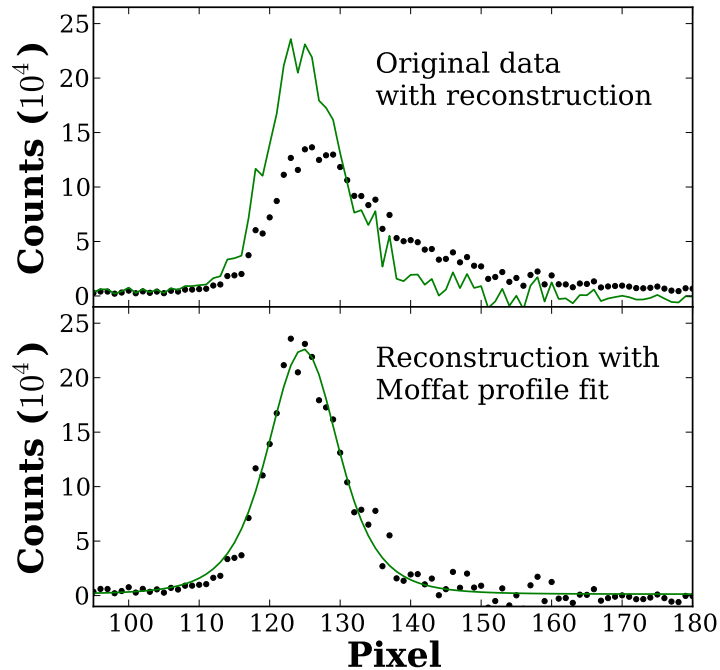


Figure 9.16: Top: The original column from the summed spectrum $I + \delta$ shown as data points with a line representing the reconstructed data, $I - \delta^2$. Bottom: The reconstructed data as points with a moffat fit to it.

perform a masked extraction, eliminating the trailed area from the sky estimation. The resulting comparison from a simple extraction and a masked extraction can be seen in figure 9.17.

Possible improvements for the modelling process

While it is mentioned that individual frame modelling is impossible at the count levels Ultraspec tends to produce at these frame rates, one possible method would be to try and fit the gaussian width, position and the traps to all of the exposures simultaneously with a varying gaussian amplitude for each frame. An important result from this analysis that conforms with theory is that the long smeared tail (which strongly defines the emission probability P in the model) does not appear to change its nature with target flux; hence as the equations suggest, as long as the instrument is at a stable temperature, P can be regarded as fixed. Unfortunately this method becomes intractable very rapidly at the frame quantities Ultraspec produces (for Z Cha, 2GB of RAM for the raw data and then fitting $\sim 3,000$ parameters to

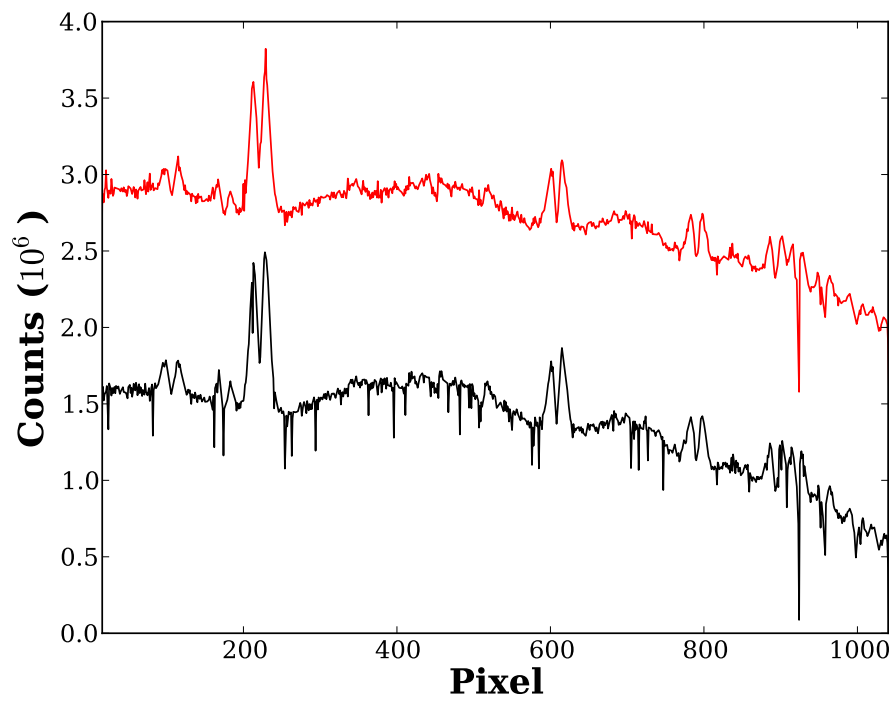


Figure 9.17: The simple normal extraction of the OY Car spectrum (bottom) against the masked extraction (above) to take into account electron trailing. The column at ~ 920 is unrecoverable.

240,000 data points *per column*). Simply scaling the trap depth by the number of stacked exposures is not feasible without knowing the quantity of traps in a column with reasonable certainty; the smoothing out of the discontinuous points formed by a trap can be created either by further numbers of traps or by the stacking process itself. Should further problems with charge trapping arise with Ultraspec, creation of such a map would be key to furthering models by removing this uncertainty.

The model currently also neglects the possibility of secondary capture. Since the emission of electrons from the traps is typically much less than the parallel clock times in a CCD at normal operating temperature, emitted electrons have a finite possibility of being re-captured and then trailed again. While a small effect, it has been noted as a measured improvement to the model used for the Hubble CCDs in Massey et al. [2010]. The more advanced models used to correct damage to CCDs in space missions typically model clouds of electrons with traps in a volume at different energy levels. This can produce an effect that is observed with some bulk traps whereby filling one trap actually creates energy states that can open up another trap. It is for this reason that applying fat-zero to such bulk trap affected CCDs does not significantly help. It has been reported that these traps can sometimes be eliminated or reduced by annealing; heating up a CCD and then cooling it back down to temperature again [Bautz et al., 2005]. While poorly understood, it is hypothesised that the annealing process allows these unstable carbon defects to form metastable complexes when moved towards room temperature. It should be stressed that these types of bulk traps (effectively damaged silicon) are almost certainly *not* the types of trap seen in Ultraspec; thus such a process would not be expected to change the Ultraspec CCD, and extending to a volumetric modelling code - while more physically accurate - is most likely unnecessary.

9.3 Results

9.3.1 Z Cha on the rise to superoutburst

During the observation period dating 10th February 2008 to 12th February 2008, spectra of Z Cha taken happened to coincide with one of its rarer superoutbursts. Using the standard star LTT 3864 to provide a flux calibration of the data sets, the continuum flux is tracked over the several sets of exposures and then the Ultraspec measurements are overlaid with AAVSO data, which can be seen in figure 9.18.

It is somewhat unfortunate that the VPH656 grism was being used for this set of observations, since it means only $H\alpha$ and the nearby helium line could be monitored as opposed to a wide spectral range. This limits the amount of analysis

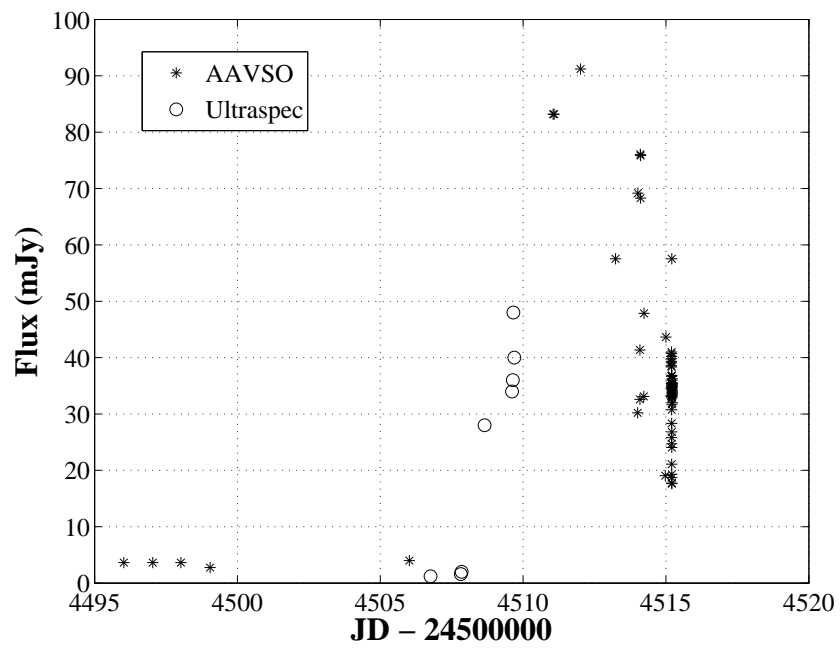


Figure 9.18: Continuum flux measurements of Z Cha from Ultraspec with AAVSO visual band observations.

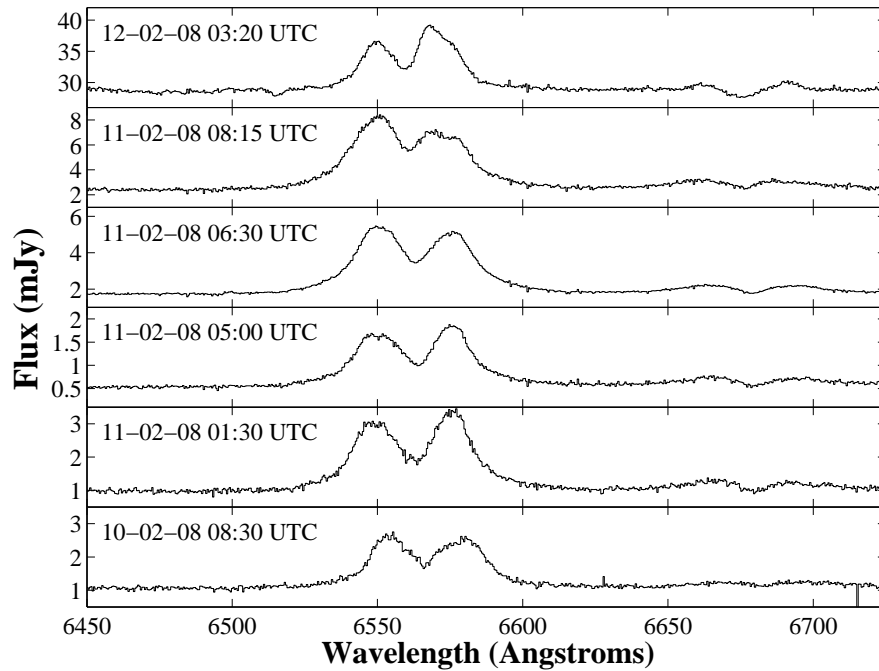


Figure 9.19: Average spectra from observations of Z Cha in successive runs. The initial sets of data (near the bottom) show a typical quiescent behaviour at $\sim 1\text{mJy}$ before a sudden dip in flux followed by a sharp rise.

that can be done with the data. Presented in figure 9.19 is the average spectrum from each observing epoch, showing it in a typical quiescent state with a flux density of $\sim 1\text{mJy}$, before actually dropping by 0.5mJy within 4 hours, rising to $\sim 2\text{mJy}$ only 90 minutes later and then being found at 28mJy the following day. AAVSO indicates that Z Cha reached in excess of 90mJy two days later (not observed) and then proceeded to settle again over the next 3-4 days. The $\text{H}\alpha$ emission becomes a much smaller fraction of the continuum light as expected, starting from nearly twice the continuum in quiescence and dropping to only one third of the continuum in the last observed epoch. It is also seen that the HeI (6678\AA) line becomes more apparent during outburst, including a central absorption component that dips beneath the continuum.

Finally, as proof that Ultraspec can indeed provide good spectrophotometry due to its low readout noise and high frame-rate, figure 9.20 shows the light-curve of Z Cha extracted from a data run in 2008 with a model fit to it using system parameters of Z Cha from Wood et al. [1986].




Figure 9.20: A light curve of Z Cha extracted from a set of Ultraspec spectra. A fit using parameters from Wood et al. [1986] is shown as a solid line, with residuals beneath.

9.3.2 The Rossiter-McLaughlin effect and QS Vir

Being a detached system, QS Vir has strong hydrogen absorption lines from the white dwarf atmosphere which are usually polluted or entirely hidden in CV systems by disc emission lines. In a similar fashion to the way disc lines are eclipsed, creating a subsequent red and blue shift depending on which part of the disc is obscured, the partial eclipse of the white dwarf should create a shift in the wavelength of these absorption lines, because the central object is spinning. This spin is considerably less than the rotation of the inner disc however, and so would produce a much weaker effect. In addition to this, the white dwarf is also much smaller than the disc and so the length of time over which the effect would be observed is a fraction of that otherwise. The phenomenon was first detailed almost simultaneously by Rossiter [1924] and McLaughlin [1924].

Satisfying both the requirements for high time resolution and good signal-to-noise that are required observe the effect, Ultraspec and a few other instruments are likely the only choices to accomplish such a feat. Using the four eclipses of QS Vir from 2008, these were phase-binned into 2s bins, and the hydrogen absorption lines were fitted with gaussians (with a shared offset from their rest wavelength). This




Figure 9.21: Radial velocity against phase of the $H\gamma$ and $H\delta$ absorption lines in QS Vir, modelled using a combination of gaussians. The red line is a model of the Rossiter-McLaughlin effect as an approximate guide of the anticipated phase and magnitude of the radial velocity deviation. The data is four eclipses phase-folded with ~ 2 s time resolution.

offset, plotted against phase, can be seen in figure 9.21, with a computed estimate of the expected effect overlaid. As can clearly be seen, this data unfortunately produces no visible result. While the fitting errors are small, the error that originates from the data is high. The instrument is clearly resolving the ingress and egress of QS Vir, as a dozen individual points can be seen on the light curve as the white dwarf is eclipsed and appears again, but the reason for the non-result appears to be a lack of signal-to-noise, so any rotation effect is completely masked in the noise. The modelling could be improved by further eclipses being averaged, and it should also be noted that a gaussian does not perfectly model the absorption (a much better fit is a gaussian plus a lorentzian function). However, adding further terms to the modelling process with this S/N quickly creates degeneracy in the fitting.

9.4 Summary

This chapter provided calibration results for the Ultraspec EMCCD before investigating various phenomena that are present in the instrument including spurious charge, AC coupling of the video circuit, and electron trapping. Comparisons of the

effect of the reduction code were presented, along with possible avenues for improvement of the reduction in the future. Z Cha was shown on the rise to superoutburst, and the difficulty of measuring the rotation of the white dwarf in QS Vir via the Rossiter-McLaughlin effect was discussed, showing that even with the low noise and high speed of Ultraspec, a considerable number of eclipses will be required to reach the necessary S/N for a detection.

The final chapter will now compile all the results obtained during the work presented in this thesis.

Chapter 10

Conclusions

Observations of the cataclysmic variable systems EX Dra, Z Cha, and OY Car have been presented, and were analysed using a variety of data from different wavelengths and instruments. Firstly, several data sets show that the white dwarf is indeed obscured by the disc in some manner. While a seemingly likely scenario due to our viewing angle with respect to the system, the effect can be subtle and requires careful analysis of the data to provide a definitive conclusion. In §5, this was shown by the use of the wind emission lines driven by UV photons, whereby instead of the flickering light-curve from the continuum, a steady and unvarying flux was observed. In §6, it was shown that the lower hemisphere of the boundary layer (and hence white dwarf) was obscured in the soft X-ray band. This was achieved by using fitted models, and was found to agree with the spectral analysis of the absorbing component. The obscuration of the white dwarf in OY Car originally reported by Horne et al. [1994] is again found in §7. The observation of dozens of FeI and FeII lines allowed for detailed measurements to be performed. These same iron lines could potentially be the source of veiling material for EX Dra as well, though lack of data at sufficient wavelengths prevents an absolute conclusion on this point. Again, the importance of data at a range of wavelengths should be stressed, particularly when examining this phenomenon. While optical spectra for EX Dra showed unusual behaviour during quiescence that helped build a picture for the disc-rim obscuration in the system, use of optical data alone with Z Cha would simply conclude that the disc does not obscure the flux from the white dwarf. Future observations of EX Dra with UV data in an extended range should easily be able to conclude whether iron lines are responsible for the veiling of the inner disc and white dwarf. A greater number of CV observations at the necessary UV wavelengths would not only help determine how common the presence of veiling iron material is in these objects,

but also drive forward the development of LTE (local thermodynamic equilibrium) models of such gas, perhaps conclusively determining where this material lies.

A second important conclusion from the careful modelling of the boundary layer of Z Cha in §6 is that the inner disc is truncated. While no other data presented here had X-ray wavelengths to perform a similar analysis on, it is believed that the presence of this extremely low surface density inner disc in a system that otherwise appears to comply with standard disc instability models may suggest truncated discs are more common than previously thought. A route for further study would be a careful reanalysis of the OY Car XMM-Newton data previously presented by Wheatley and West [2003]. The vertically offset emission that was argued for in the paper can alternatively be explained by the obscuration of the lower pole, and if this is found to be the case for soft X-rays as with Z Cha, the likely conclusion of a similarly low-density inner disc may follow. With the increasing amount of X-ray data available to those studying CVs and constant improvement to detectors in this region, the next logical step in this field after the confirmation of truncated discs is determining the correct model for the cause of this effect; of which there are currently many competing entries.

Thirdly, using eclipse timings from high speed Ultracam (and Ultraspec) data combined with XMM-Newton, it is found that the Applegate effect [Applegate, 1992] fails to account for the observed period variations in Z Cha (§6) by two orders of magnitude when using a generalised calculation originally performed by Brinkworth et al. [2006]. Studies of other binaries showing continued change in their periods has also failed to attribute any of the observations to the Applegate effect [Parsons et al., 2010a], once again concluding as with Z Cha, that a third body provides the only viable explanation. If this observed trend continues, it would result in a significant fraction of eclipsing binary systems containing a third body, and almost all of them having to be brown dwarfs. The counter-argument to this point is that a third body can almost always be used to match these period variations in some manner, meaning that conclusive evidence of a third body will require monitoring of the period variations for several cycles. In the case of Z Cha, this would take over a century should the fitted period of 54 years remain unchanged. At this current point in time, the increasing number of systems showing this sinusoidal period change is intriguing, but little beyond careful long-term study will reveal for certain the exact source of these variations.

Finally, two sets of data (Z Cha in §6 and OY Car in §7) showed the possibility of a second absorption point in the emission light-curves, theorised to be a second disc-stream interaction as discussed by Lubow [1989]. X-ray light curves of

Z Cha show two significant points of absorption, one at approximately the typically expected bright spot impact phase, and the other around half an orbit later. A similar phenomenon was found in the line strengths of the measured iron lines in OY Car. The alternative model put forth by Ogilvie [2002] is a tidally thickened disc rim, which seems logical from the results of §6 since it was determined the second point is caused merely by an increase in the covering fraction of the model, not an increased column density as what may be expected if the stream was continuing around the disc and impacting again. It is also true that the continuing stream should preferentially impact at a certain point which does not seem to match well with what is observed. Conclusive proof is difficult to obtain from the current data, but the fact that such a phenomenon seems to be present in so many other X-ray observations of CVs should in time reach a conclusion on this matter.

In addition to the study of cataclysmic variable systems, improvements to both the Ultracam and Ultraspec instruments have been detailed, providing increased accuracy and productivity as they continue to grow as premier astronomical instruments for the study of CVs and other objects that require high time resolution.

Theoretical extinction coefficients and a detailed comparison of the Ultracam filters to their SDSS equivalents were determined in §8, with the result that the difference is almost always ≤ 0.1 mag for all but the most extreme colour objects. The “peppering” phenomenon in Ultracam was investigated and found not to be as problematic as originally thought, particularly for aperture photometry. Since counts are preserved, photometric accuracy is not compromised, though there will be a marked decrease in signal-to-noise due to the modulation loss. Correction of this problem will increase the useable dynamic range of Ultracam and allow a greater margin for error in utilisation of Ultracam for brighter targets; several experiments were proposed in the chapter to continue the investigation. These suggestions should be acted upon the next time Ultracam is returned to the laboratory, and may help to narrow down the cause for this effect. A similar study of CCD features related to Ultraspec was performed in §9, showing that each of the features can be accounted for or even corrected, with a significant improvement to the resulting data. It was also shown that reconstruction of the original image unaffected by electron trapping is entirely feasible, and calibration of a “charge trap map” for Ultraspec will allow many parameters to be constrained, removing the ambiguity currently present. This should result in an accurately reconstructed image without degeneracy problems. With no significant issues remaining with Ultraspec, any future observations should be of a scientific quality that matches the instrument’s original goals, and it is hoped that the next run will result in several published journal papers as a testament to

this.

Finally, a new, open-source, modular software package with a small dependency stack was detailed, that provides a low-level interface to Ultracam (and Ultraspec) raw data and automatic indexing of data using a central database. Other software can then directly (or via higher-level written routines) leverage the contents of this database for their own purposes. This has led to an improvement in the observing workflow of Ultracam by enabling a target database and easier and more robust classification of data. The software should be of great benefit now that the instruments are operating for months at a time and generating even greater quantities of data. Future developments of this will involve a public/private access system with a file server, and also the continual evolution of the software as it continues to match the updates and improvements to Ultracam, such as their new sub-millisecond GPS which has recently been installed.

In summary, using a combination of multi-wavelength observations, it is found that each of the three cataclysmic variable systems reported on here show signs that the white dwarf is partially obscured by material in some fashion that is present in the disc. Instruments that observe at short wavelengths such as X-Shooter and FOS have positively identified intervening disc material in OY Car and EX Dra, whereas the high time resolution of Ultracam combined with XMM-Newton allowed constraints of the white dwarf emission region for Z Cha, as well as continued monitoring of its period variations. While these phenomena remain far from being well understood, continued observations of cataclysmic variables and continued evolution and improvement of the instruments used to observe them should provide an increasing amount of data to aid astronomers with. With X-Shooter now providing real scientific data, and most of the details of Ultraspec dealt with and documented, it is hoped that the near future may shed further light on these and other aspects of cataclysmic variable systems.

Bibliography

- G. B. Alers, K. S. Krisch, D. Monroe, B. E. Weir, and A. M. Chang. Tunneling current noise in thin gate oxides. *Applied Physics Letters*, 69:2885–2887, November 1996. doi: 10.1063/1.117351.
- S. Amari, A. Anders, A. Virag, and E. Zinner. Interstellar graphite in meteorites. *Nature*, 345:238–240, May 1990. doi: 10.1038/345238a0.
- J. H. Applegate. A mechanism for orbital period modulation in close binaries. *ApJ*, 385:621–629, February 1992. doi: 10.1086/170967.
- P. E. Argyle. The Effective Quantum Efficiency in Astronomical Spectroscopy. *JRASC*, 49:19–+, February 1955.
- N. Bade, H.-J. Hagen, and D. Reimers. Fast classification of ROSAT sources on objective prism plates. In J. Hunt & B. Battrick, editor, *Two Topics in X-Ray Astronomy, Volume 1: X Ray Binaries. Volume 2: AGN and the X Ray Background*, volume 296 of *ESA Special Publication*, pages 883–884, November 1989.
- J. Bailey. The dwarf nova Z Chamaeleontis. I - Photometry. *MNRAS*, 187:645–653, June 1979.
- J. Bailey and M. Ward. Spectrophotometry of the eclipsing dwarf nova OY Carinae. *MNRAS*, 194:17P–+, January 1981.
- S. A. Balbus and J. F. Hawley. Instability, turbulence, and enhanced transport in accretion disks. *Reviews of Modern Physics*, 70:1–53, January 1998. doi: 10.1103/RevModPhys.70.1.
- R. Baptista, M. S. Catalán, and L. Costa. Eclipse studies of the dwarf nova EX Draconis. *MNRAS*, 316:529–539, August 2000. doi: 10.1046/j.1365-8711.2000.03557.x.

- R. Baptista, F. Jablonski, E. Oliveira, S. Vrielmann, P. A. Woudt, and M. S. Catalán. Cyclical period changes in Z Chamaeleontis. *MNRAS*, 335:L75–L78, September 2002. doi: 10.1046/j.1365-8711.2002.05880.x.
- S. C. Barden, J. A. Arns, and W. S. Colburn. Volume-phase holographic gratings and their potential for astronomical applications. In S. D’Odorico, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 3355 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 866–876, July 1998.
- H. Barwig, H. Fiedler, and D. Reimers. Compact stars in binary systems. In H. van Woerden, editor, *Abstracts*, volume 165 of *IAU Symposium*, page 89, 1994.
- G. T. Bath, W. D. Evans, J. Papaloizou, and J. E. Pringle. The accretion model of dwarf novae with application to Z Chamaeleontis. *MNRAS*, 169:447–470, December 1974.
- M. Bautz, G. Prigozhin, Kissel S., B. LaMarr, C. Grant, and S. Brown. Anomalous annealing of a high-resistivity ccd irradiated at low temperature. *IEEE Transactions on Nuclear Science*, 52:519–526, April 2005.
- S. M. Beard, A. J. A. Vick, D. Atkinson, V. S. Dhillon, T. Marsh, S. McLay, M. Stevenson, and C. Tierney. Ultracam camera control and data acquisition system. In H. Lewis, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4848 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, pages 218–229, December 2002. doi: 10.1117/12.461354.
- O. G. Benvenuto and M. A. De Vito. The formation of helium white dwarfs in close binary systems - II. *MNRAS*, 362:891–905, September 2005. doi: 10.1111/j.1365-2966.2005.09315.x.
- G. Berriman. A visible and infrared study of the eclipsing dwarf nova OY Carinae. I. The visible eclipses of the central object. *MNRAS*, 207:783–799, April 1984.
- I. Billington, T. R. Marsh, and V. S. Dhillon. The eclipsing dwarf nova HS 1804+6753. *MNRAS*, 278:673–682, February 1996.
- C. S. Brinkworth, T. R. Marsh, V. S. Dhillon, and C. Knigge. Detection of a period decrease in NN Ser with ULTRACAM: evidence for strong magnetic braking or an unseen companion. *MNRAS*, 365:287–295, January 2006. doi: 10.1111/j.1365-2966.2005.09718.x.

- P. Bristow. CCD Charge Transfer Inefficiency. *Space Telescope European Coordinating Facility Newsletter*, 34:9–+, September 2003.
- D. D. Clayton. *Principles of stellar evolution and nucleosynthesis*. University Of Chicago Press, 1968.
- R. E. S. Clegg. *ISIS Astronomers' Guide. Version 1.0*. Isaac Newton Group, 1992.
- M. C. Cook. High-speed photometry of the eclipsing dwarf nova OY Carinae. *MNRAS*, 215:211–232, July 1985.
- M. C. Cook and B. Warner. Period Changes in Z-Chamaeleontis. *MNRAS*, 196:55P–+, August 1981.
- C. M. Copperwheat, T. R. Marsh, V. S. Dhillon, S. P. Littlefair, R. Hickman, B. T. Gänsicke, and J. Southworth. Physical properties of IP Pegasi: an eclipsing dwarf nova with an unusually cool white dwarf. *MNRAS*, 402:1824–1840, March 2010. doi: 10.1111/j.1365-2966.2009.16010.x.
- M. Cropper. The Polars. *Space Sci. Rev.*, 54:195–295, December 1990. doi: 10.1007/BF00177799.
- Z. Dai, S. Qian, and E. Fernández Lajús. Evidence of a Brown Dwarf in the Eclipsing Dwarf Nova Z Chamaeleonis. *ApJ*, 703:109–113, September 2009. doi: 10.1088/0004-637X/703/1/109.
- V. Dhillon, T. Marsh, C. Copperwheat, N. Bezawada, D. Ives, A. Vick, and K. O'Brien. ULTRASPEC: High-speed Spectroscopy with Zero Read-out Noise. *The Messenger*, 127:41–+, March 2007a.
- V. S. Dhillon, 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. ULTRACAM: an ultrafast, triple-beam CCD camera for high-speed astrophysics. *MNRAS*, 378:825–840, July 2007b. doi: 10.1111/j.1365-2966.2007.11881.x.
- H. Fiedler, H. Barwig, and K. H. Mantel. HS 1804+6753: a new eclipsing CV above the period gap. *A&A*, 327:173–182, November 1997.
- J. Frank, A. King, and D. Raine. *Accretion power in astrophysics*. Cambridge University Press, 1992.

- M. Y. Fujimoto. A Theory of Hydrogen Shell Flashes on Accreting White Dwarfs - Part Two - the Stable Shell Burning and the Recurrence Period of Shell Flashes. *ApJ*, 257:767–+, June 1982. doi: 10.1086/160030.
- M. Fukugita, T. Ichikawa, J. E. Gunn, M. Doi, K. Shimasaku, and D. P. Schneider. The Sloan Digital Sky Survey Photometric System. *AJ*, 111:1748–+, April 1996. doi: 10.1086/117915.
- B. T. Gänsicke, T. R. Marsh, and J. Southworth. SDSSJ104341.53+085558.2: a second white dwarf with a gaseous debris disc. *MNRAS*, 380:L35–L39, September 2007. doi: 10.1111/j.1745-3933.2007.00343.x.
- U. Graser. *Cassegrain TWIN Spectrograph*. Calar Alto Observatory, 1998.
- R. Hōshi. Accretion Model for Outbursts of Dwarf Nova. *Progress of Theoretical Physics*, 61:1307–1319, May 1979. doi: 10.1143/PTP.61.1307.
- J.-M. Hameury, J.-P. Lasota, and J.-M. Hure. A model for WZ SGE with ‘standard’ values of alpha. *MNRAS*, 287:937–940, June 1997.
- R. J. Harms, E. Beaver, E. M. Burbidge, R. Angel, F. Bartko, W. Bloomquist, J. C. Flemming, R. Bohlin, A. F. Davidsen, and H. Ford. Faint-object spectrograph for Space Telescope. In C. L. Wyman, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 183 of *Presented at the Society of Photo-Optical Instrumentation Engineers (SPIE) Conference*, pages 74–87, January 1979.
- T. E. Harris. *The Theory of Branching Processes*. New York Springer, 1963.
- A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann. How Massive Single Stars End Their Life. *ApJ*, 591:288–300, July 2003. doi: 10.1086/375341.
- C. Hellier. *Cataclysmic Variable Stars*. Springer-Praxis, January 2001.
- G. Hensler. Hydrodynamical Calculations of Accretion Disks in Close Binary Systems - Part Two - Models. *A&A*, 114:319–+, October 1982.
- F. V. Hessman, E. L. Robinson, R. E. Nather, and E.-H. Zhang. Time-resolved spectroscopy of SS Cygni at minimum and maximum light. *ApJ*, 286:747–759, November 1984. doi: 10.1086/162651.
- J. B. Holberg and P. Bergeron. Calibration of Synthetic Photometry Using DA White Dwarfs. *AJ*, 132:1221–1233, September 2006. doi: 10.1086/505938.

- J. B. Holberg, M. A. Barstow, and E. M. Green. The Discovery of Mg II 4481 in the White Dwarf EG 102: Evidence for Ongoing Accretion. *ApJ*, 474:L127+, January 1997. doi: 10.1086/310446.
- K. Horne and M. Eracleous. A rough photometric calibration for fos,blue,g160l,order0. *FOS Instrument Science Report*, August 1993.
- K. Horne, T. R. Marsh, F. H. Cheng, I. Hubeny, and T. Lanz. HST eclipse mapping of dwarf nova OY Carinae in quiescence: an 'Fe II curtain' with Mach approx. = 6 velocity dispersion veils the white dwarf. *ApJ*, 426:294–307, May 1994. doi: 10.1086/174064.
- J. R. Janesick. *Scientific charge-coupled devices*. SPIE Press, 2001.
- V. Joergens, H. C. Spruit, and R. G. M. Rutten. Spirals and the size of the disk in EX Dra. *A&A*, 356:L33–L36, April 2000.
- H. L. Johnson and W. W. Morgan. Fundamental stellar photometry for standards of spectral type on the revised system of the Yerkes spectral atlas. *ApJ*, 117:313–+, May 1953. doi: 10.1086/145697.
- C.-K. Kim. *The physics of charge-coupled devices*, pages 1–80. Wiley-Interscience, 1979.
- A. R. King. The evolution of compact binaries. *QJRAS*, 29:1–25, March 1988.
- A. R. King. The UV delay in dwarf novae. *MNRAS*, 288:L16–L18, June 1997.
- A. R. King, D. J. Rolfe, and K. Schenker. A new evolutionary channel for Type Ia supernovae. *MNRAS*, 341:L35–L38, June 2003. doi: 10.1046/j.1365-8711.2003.06639.x.
- D. L. King. Atmospheric extinction at the roque de los muchachos observatory, la palma. *La Palma technical notes*, 31, September 1985.
- C. Knigge and J. E. Drew. Eclipse Mapping of the Accretion Disk Wind in the Cataclysmic Variable UX Ursae Majoris. *ApJ*, 486:445–+, September 1997. doi: 10.1086/304519.
- C. Knigge, K. S. Long, W. P. Blair, and R. A. Wade. Disks, Winds, and Veiling Curtains: Dissecting the Ultraviolet Spectrum of the Dwarf Nova Z Camelopardalis in Outburst. *ApJ*, 476:291–+, February 1997. doi: 10.1086/303607.

- C. Knigge, K. S. Long, D. W. Hoard, P. Szkody, and V. S. Dhillon. A Self-occluding Accretion Disk in the SW Sextantis Star DW Ursae Majoris. *ApJ*, 539:L49–L53, August 2000. doi: 10.1086/312825.
- C. J. Lada. Stellar Multiplicity and the Initial Mass Function: Most Stars Are Single. *ApJ*, 640:L63–L66, March 2006. doi: 10.1086/503158.
- T. Lanz, I. Hubeny, F.-H. Cheng, and K. Horne. Line Blanketed Model Atmospheres for Metal-rich White Dwarfs. In *Bulletin of the American Astronomical Society*, volume 24 of *Bulletin of the American Astronomical Society*, pages 1203–+, December 1992.
- J.-P. Lasota. The disc instability model of dwarf novae and low-mass X-ray binary transients. *NAR*, 45:449–508, June 2001. doi: 10.1016/S1387-6473(01)00112-9.
- J. P. Lasota, J. M. Hameury, and J. M. Hure. Dwarf novae at low mass transfer rates. *A&A*, 302:L29+, October 1995.
- J. Liebert, P. Bergeron, D. Eisenstein, H. C. Harris, S. J. Kleinman, A. Nitta, and J. Krzesinski. A Helium White Dwarf of Extremely Low Mass. *ApJ*, 606:L147–L149, May 2004. doi: 10.1086/421462.
- D. A. Liedahl, A. L. Osterheld, and W. H. Goldstein. New calculations of Fe L-shell X-ray spectra in high-temperature plasmas. *ApJ*, 438:L115–L118, January 1995. doi: 10.1086/187729.
- M. Livio and J. E. Pringle. Dwarf nova outbursts - The ultraviolet delay and the effect of a weakly magnetized white dwarf. *MNRAS*, 259:23P–26P, December 1992.
- S. H. Lubow. On the dynamics of mass transfer over an accretion disk. *ApJ*, 340:1064–1069, May 1989. doi: 10.1086/167458.
- S. H. Lubow and F. H. Shu. Gas dynamics of semidetached binaries. *ApJ*, 198:383–405, June 1975. doi: 10.1086/153614.
- H. J. Lush. Photomultiplier linearity. *Journal of Scientific Instruments*, 42:597–602, August 1965. doi: 10.1088/0950-7671/42/8/328.
- D. Lynden-Bell and J. E. Pringle. The evolution of viscous discs and the origin of the nebular variables. *MNRAS*, 168:603–637, September 1974.

- H. MacPherson. Obituary: Max Wolf. *The Observatory*, 55:355–359, December 1932.
- T. R. Marsh. The extraction of highly distorted spectra. *PASP*, 101:1032–1037, November 1989. doi: 10.1086/132570.
- T. R. Marsh and K. Horne. Images of accretion discs. II - Doppler tomography. *MNRAS*, 235:269–286, November 1988.
- K. O. Mason, J. E. Drew, F. A. Cordova, K. Horne, R. Hilditch, C. Knigge, T. Lanz, and T. Meylan. Eclipse observations of an accretion disc wind. *MNRAS*, 274:271–286, May 1995.
- K. O. Mason, A. Breeveld, R. Much, M. Carter, F. A. Cordova, M. S. Cropper, J. Fordham, H. Huckle, C. Ho, H. Kawakami, J. Kennea, T. Kennedy, J. Mittaz, D. Pandel, W. C. Priedhorsky, T. Sasseen, R. Shirey, P. Smith, and J.-M. Vreux. The XMM-Newton optical/UV monitor telescope. *A&A*, 365:L36–L44, January 2001. doi: 10.1051/0004-6361:20000044.
- R. Massey, C. Stoughton, A. Leauthaud, J. Rhodes, A. Koekemoer, R. Ellis, and E. Shaghoulain. Pixel-based correction for Charge Transfer Inefficiency in the Hubble Space Telescope Advanced Camera for Surveys. *MNRAS*, 401:371–384, January 2010. doi: 10.1111/j.1365-2966.2009.15638.x.
- K. Matsuo, M. C. Teich, and B. E. A. Saleh. Noise properties and time response of the staircase avalanche photodiode. *Journal of Lightwave Technology*, 3:1223–1231, December 1985. doi: 10.1109/JLT.1985.1074334.
- D. B. McLaughlin. Some results of a spectrographic study of the Algol system. *ApJ*, 60:22–31, July 1924. doi: 10.1086/142826.
- L. Mestel. Magnetic braking by a stellar wind-I. *MNRAS*, 138:359–+, 1968.
- F. Meyer and E. Meyer-Hofmeister. Constraints from the UV delay in dwarf nova outbursts. *A&A*, 221:36–40, August 1989.
- F. Meyer and E. Meyer-Hofmeister. Accretion disk evaporation by a coronal siphon flow. *A&A*, 288:175–182, August 1994.
- A. F. J. Moffat. A Theoretical Investigation of Focal Stellar Images in the Photographic Emulsion and Application to Photographic Photometry. *A&A*, 3:455–+, December 1969.

- D. G. Monet, S. E. Levine, B. Canzian, H. D. Ables, A. R. Bird, C. C. Dahn, H. H. Guetter, H. C. Harris, A. A. Henden, S. K. Leggett, H. F. Levison, C. B. Luginbuhl, J. Martini, A. K. B. Monet, J. A. Munn, J. R. Pier, A. R. Rhodes, B. Riepe, S. Sell, R. C. Stone, F. J. Vrba, R. L. Walker, G. Westerhout, R. J. Brucato, I. N. Reid, W. Schoening, M. Hartley, M. A. Read, and S. B. Tritton. The USNO-B Catalog. *AJ*, 125:984–993, February 2003. doi: 10.1086/345888.
- K. Mukai, J. H. Wood, T. Naylor, E. M. Schlegel, and J. H. Swank. The X-Ray Eclipse of the Dwarf Nova HT Cassiopeiae: Results from ASCA and ROSAT HRI Observations. *ApJ*, 475:812–+, February 1997. doi: 10.1086/303571.
- G. S. Mumford. Z Chameleontis. *Information Bulletin on Variable Stars*, 337:1–+, March 1969.
- T. Naylor. An optimal extraction algorithm for imaging photometry. *MNRAS*, 296:339–346, May 1998. doi: 10.1046/j.1365-8711.1998.01314.x.
- T. Naylor and C. La Dous. The vertical disc structure of U Geminorum in outburst. *MNRAS*, 290:160–164, September 1997.
- T. Naylor, G. T. Bath, P. A. Charles, B. J. M. Hassall, G. Sonneborn, H. van der Woerd, and J. van Paradijs. The 1985 May superoutburst of the dwarf nova OY Carinae. II - IUE and EXOSAT observations. *MNRAS*, 231:237–255, March 1988.
- H. Nyquist. Thermal Agitation of Electric Charge in Conductors. *Physical Review*, 32:110–113, July 1928. doi: 10.1103/PhysRev.32.110.
- D. O’Donoghue, C. Koen, D. Kilkeny, R. S. Stobie, D. Koester, M. S. Bessell, N. Hambly, and H. MacGillivray. The DA+dMe eclipsing binary EC13471-1258: its cup runneth over ... just. *MNRAS*, 345:506–528, October 2003. doi: 10.1046/j.1365-8711.2003.06973.x.
- G. I. Ogilvie. Tidally distorted accretion discs in binary stars. *MNRAS*, 330:937–949, March 2002. doi: 10.1046/j.1365-8711.2002.05140.x.
- Y. Osaki. An accretion model for the outbursts of U Geminorum stars. *PASJ*, 26:429–436, 1974.
- R. J. Panek and A. V. Holm. Ultraviolet spectroscopy of the dwarf nova U Geminorum. *ApJ*, 277:700–709, February 1984. doi: 10.1086/161741.
- S. G. Parsons, T. R. Marsh, C. M. Copperwheat, V. S. Dhillon, S. P. Littlefair, R. D. G. Hickman, P. F. L. Maxted, B. T. Gänsicke, E. Unda-Sanzana, J. P.

- Colque, N. Barraza, N. Sánchez, and L. A. G. Monard. Orbital period variations in eclipsing post-common-envelope binaries. *MNRAS*, 407:2362–2382, October 2010a. doi: 10.1111/j.1365-2966.2010.17063.x.
- S. G. Parsons, T. R. Marsh, B. T. Gänsicke, and C. Tappert. A stellar prominence in the white dwarf/red dwarf binary QS Vir: evidence for a detached system. *ArXiv e-prints*, November 2010b.
- J. Patterson. The DQ Herculis stars. *PASP*, 106:209–238, March 1994. doi: 10.1086/133375.
- A. J. Pickles. A Stellar Spectral Flux Library: 1150-25000 Å. *PASP*, 110:863–878, July 1998. doi: 10.1086/316197.
- N. Pogson. On the Variable Star U Geminorum. *MNRAS*, 17:200–202, May 1857.
- J. E. Pringle. Period changes in eruptive binaries. *MNRAS*, 170:633–642, March 1975.
- J. E. Pringle. White dwarf heating and the ultraviolet flux in dwarf novae. *MNRAS*, 230:587–595, February 1988.
- J. E. Pringle. The properties of external accretion discs. *MNRAS*, 248:754–759, February 1991.
- M. W. Rayne and J. A. J. Whelan. The dwarf nova Z Chamaeleontis. II - Spectroscopy. *MNRAS*, 196:73–94, July 1981.
- I. N. Reid, C. Brewer, R. J. Brucato, W. R. McKinley, A. Maury, D. Mendenhall, J. R. Mould, J. Mueller, G. Neugebauer, J. Phinney, W. L. W. Sargent, J. Schombert, and R. Thicksten. The second Palomar Sky Survey. *PASP*, 103:661–674, July 1991. doi: 10.1086/132866.
- H. R. Richman, J. H. Applegate, and J. Patterson. Long-term periods in cataclysmic variables. *PASP*, 106:1075–1084, October 1994. doi: 10.1086/133481.
- H. Ritter and U. Kolb. Catalogue of cataclysmic binaries, low-mass X-ray binaries and related objects (Seventh edition). *A&A*, 404:301–303, June 2003. doi: 10.1051/0004-6361:20030330.
- E. L. Robinson, J. H. Wood, R. C. Bless, J. C. Clemens, J. F. Dolan, J. L. Elliot, M. J. Nelson, J. W. Percival, M. J. Taylor, G. W. van Citters, and E. Zhang. Hubble Space Telescope observations of the dwarf Nova Z Chamaeleontis through two eruption cycles. *ApJ*, 443:295–318, April 1995. doi: 10.1086/175526.

- R. A. Rossiter. On the detection of an effect of rotation during eclipse in the velocity of the brighter component of beta Lyrae, and on the constancy of velocity of this system. *ApJ*, 60:15–21, July 1924. doi: 10.1086/142825.
- N. Saks. A technique for suppressing dark current generated by interface states in buried channel ccd imagers. *IEEE Electron Device Letters*, 1:131–133, July 1980.
- G. J. Savonije, J. C. B. Papaloizou, and D. N. C. Lin. On Tidally Induced Shocks in Accretion Discs in Close Binary Systems. *MNRAS*, 268:13–+, May 1994.
- E. Schatzman. Théorie du débit d'énergie des naines blanches. *Annales d'Astrophysique*, 8:143–+, January 1945.
- D. P. Schneider and P. Young. The magnetic maw of 2A 0311-227. *ApJ*, 238:946–954, June 1980. doi: 10.1086/158059.
- R. Schoembs and K. Hartmann. The eclipsing dwarf nova OY Carinae. II - Spectroscopy and photometry during quiescence. *A&A*, 128:37–52, November 1983.
- R. Schoembs, H. Dreier, and H. Barwig. Simultaneous multicolour photometry of OY Carinae during quiescence. *A&A*, 181:50–56, July 1987.
- D. K. Schroder. A two-phase germanium charge-coupled device. *Applied Physics Letters*, 25:747–+, December 1974. doi: 10.1063/1.1655386.
- A. W. Shafter. Radial velocity studies of cataclysmic binaries. I - KR Aurigae. *ApJ*, 267:222–231, April 1983. doi: 10.1086/160861.
- N. I. Shakura and R. A. Sunyaev. Black holes in binary systems. Observational appearance. *A&A*, 24:337–355, 1973.
- M. M. Shara. Recent progress in understanding the eruptions of classical novae. *PASP*, 101:5–31, January 1989. doi: 10.1086/132400.
- I. Shlosman and P. Vitello. Winds from accretion disks - Ultraviolet line formation in cataclysmic variables. *ApJ*, 409:372–386, May 1993. doi: 10.1086/172670.
- S. N. Shore and J. P. Aufdenberg. On the Interpretation of the Ultraviolet Spectra of Symbiotic Stars and Recurrent Novae. I. *ApJ*, 416:355–+, October 1993. doi: 10.1086/173239.
- E. M. Sion. White Dwarfs in Cataclysmic Variables. *PASP*, 111:532–555, May 1999. doi: 10.1086/316361.

- J. Smak. Eruptive binaries. VII - Z Cha. *ActA*, 29:309–324, 1979.
- J. Smak. On the White Dwarf in Z-Chamaeleontis. *Acta Astronomica*, 36:211–+, 1986.
- H. C. Spruit. Stationary shocks in accretion disks. *A&A*, 184:173–184, October 1987.
- D. Steeghs, K. Horne, T. R. Marsh, and J. F. Donati. Slingshot prominences during dwarf nova outbursts? *MNRAS*, 281:626–636, July 1996.
- D. Steeghs, E. T. Harlaftis, and K. Horne. Spiral structure in the accretion disc of the binary IP Pegasi. *MNRAS*, 290:L28–L32, September 1997.
- L. Strüder, U. Briel, K. Dennerl, R. Hartmann, E. Kendziorra, N. Meidinger, E. Pfeffermann, C. Reppin, B. Aschenbach, W. Bornemann, H. Bräuninger, W. Burkert, M. Elender, M. Freyberg, F. Haberl, G. Hartner, F. Heuschmann, H. Hippmann, E. Kastelic, S. Kemmer, G. Kettenring, W. Kink, N. Krause, S. Müller, A. Oppitz, W. Pietsch, M. Popp, P. Predehl, A. Read, K. H. Stephan, D. Stötter, J. Trümper, P. Holl, J. Kemmer, H. Soltau, R. Stötter, U. Weber, U. Weichert, C. von Zanthier, D. Carathanassis, G. Lutz, R. H. Richter, P. Solc, H. Böttcher, M. Kuster, R. Staubert, A. Abbey, A. Holland, M. Turner, M. Balasini, G. F. Bignami, N. La Palombara, G. Villa, W. Buttler, F. Gianini, R. Lainé, D. Lumb, and P. Dhez. The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera. *A&A*, 365:L18–L26, January 2001. doi: 10.1051/0004-6361:20000066.
- S. M. Sze. *Physics of semiconductor devices*. Wiley-Interscience, 1981.
- K. Tanikawa, Y. Ito, and H. Sei. Evaluation of dark-current nonuniformity in a charge-coupled device. *Applied Physics Letters*, 28:285–+, March 1976. doi: 10.1063/1.88737.
- J.-L. Tassoul. On orbital circularization in detached close binaries. *ApJ*, 324:L71–L73, January 1988. doi: 10.1086/185094.
- T. X. Thuan and J. E. Gunn. A new four-color intermediate-band photometric system. *PASP*, 88:543–547, August 1976. doi: 10.1086/129982.
- D. M. Townsley and L. Bildsten. Faint Cataclysmic Variables in Quiescence: Globular Cluster and Field Surveys. *ApJ*, 565:L35–L38, January 2002. doi: 10.1086/339052.

- R. G. Tull, J. P. Choisser, and E. H. Snow. Self-scanned Digicon - A digital image tube for astronomical spectroscopy. *Appl. Opt.*, 14:1182–1189, May 1975.
- M. J. L. Turner, A. Abbey, M. Arnaud, M. Balasini, M. Barbera, E. Belsole, P. J. Bennie, J. P. Bernard, G. F. Bignami, M. Boer, U. Briel, I. Butler, C. Cara, C. Chabaud, R. Cole, A. Collura, M. Conte, A. Cros, M. Denby, P. Dhez, G. Di Coco, J. Dowson, P. Ferrando, S. Ghizzardi, F. Gianotti, C. V. Goodall, L. Gretton, R. G. Griffiths, O. Hainaut, J. F. Hochedez, A. D. Holland, E. Jourdain, E. Kendziorra, A. Lagostina, R. Laine, N. La Palombara, M. Lortholary, D. Lumb, P. Marty, S. Molendi, C. Pigot, E. Poindron, K. A. Pounds, J. N. Reeves, C. Reppin, R. Rothenflug, P. Salvétat, J. L. Sauvageot, D. Schmitt, S. Sembay, A. D. T. Short, J. Spragg, J. Stephen, L. Strüder, A. Tiengo, M. Trifoglio, J. Trümper, S. Vercellone, L. Vigroux, G. Villa, M. J. Ward, S. Whitehead, and E. Zonca. The European Photon Imaging Camera on XMM-Newton: The MOS cameras : The MOS cameras. *A&A*, 365:L27–L35, January 2001. doi: 10.1051/0004-6361:20000087.
- S. van Amerongen, E. Kuulkers, and J. van Paradijs. Secular brightness variation of the dwarf nova Z Cha during quiescence. *MNRAS*, 242:522–528, February 1990.
- A. van Teeseling. An eclipse in the X-ray flux from the dwarf nova Z Cha. *A&A*, 319:L25–L28, March 1997.
- F. Verbunt and S. Rappaport. Mass transfer instabilities due to angular momentum flows in close binaries. *ApJ*, 332:193–198, September 1988. doi: 10.1086/166645.
- N. Vogt. Z Chamaeleontis - Evidence for an eccentric disk during supermaximum. *ApJ*, 252:653–667, January 1982. doi: 10.1086/159592.
- N. Vogt, R. Schoembs, W. Krzeminski, and H. Pedersen. The Eclipsing Dwarf Nova Oy-Carinae - Ephemeris and Physical Parameters. *A&A*, 94:L29+, February 1981.
- R. A. Wade and K. Horne. The radial velocity curve and peculiar TiO distribution of the red secondary star in Z Chamaeleontis. *ApJ*, 324:411–430, January 1988. doi: 10.1086/165905.
- R. Walden, R. Krambeck, R. Strain, J. Mckenna, N. Schryer, and Smith G. The buried channel charge coupled device. *Bell System Technical Journal*, 51:1635–1640, 1970.
- M. F. Walker. Nova DQ Herculis (1934): an Eclipsing Binary with Very Short Period. *PASP*, 66:230–+, October 1954. doi: 10.1086/126703.

- B. Warner. Observations of rapid blue variables - XIV. Z Chamaeleontis. *MNRAS*, 168:235–247, July 1974.
- B. Warner. Cataclysmic variable stars. *Cambridge Astrophysics Series*, 28, 1995a.
- B. Warner. Systematics of Superoutbursts in Dwarf Novae. *Ap&SS*, 226:187–211, April 1995b. doi: 10.1007/BF00627371.
- P. J. Wheatley and R. G. West. The X-ray eclipse of OY Car resolved with XMM-Newton: X-ray emission from the polar regions of the white dwarf. *MNRAS*, 345:1009–1014, November 2003. doi: 10.1046/j.1365-8711.2003.07030.x.
- P. J. Wheatley, C. W. Mauche, and J. A. Mattei. The X-ray and extreme-ultraviolet flux evolution of SS Cygni throughout outburst. *MNRAS*, 345:49–61, October 2003. doi: 10.1046/j.1365-8711.2003.06936.x.
- M. White, D. Lampe, and Blaha F. Characterization of surface channel ccd image arrays at low light levels. *IEEE Journal of Solid-State Circuits*, 9:1–12, February 1974.
- R. Whitehurst. Numerical simulations of accretion disks. I - Superhumps - A tidal phenomenon of accretion disks. *MNRAS*, 232:35–51, May 1988.
- J. Wood, K. Horne, G. Berriman, R. Wade, D. O’Donoghue, and B. Warner. High-speed photometry of the dwarf nova Z Cha in quiescence. *MNRAS*, 219:629–655, April 1986.
- J. H. Wood. On the nature of the central object in Z Cha. *MNRAS*, 228:797–803, October 1987.
- J. H. Wood and K. Horne. White dwarf radii and boundary-layer constraints in three dwarf novae. *MNRAS*, 242:606–615, February 1990.
- J. H. Wood, K. Horne, G. Berriman, and R. A. Wade. Eclipse studies of the dwarf nova OY Carinae in quiescence. *ApJ*, 341:974–996, June 1989. doi: 10.1086/167557.
- J. H. Wood, T. Naylor, B. J. M. Hassall, and T. F. Ramseyer. The discovery of an X-ray eclipse during a low state of the dwarf nova HT Cassiopeiae. *MNRAS*, 273:772–784, April 1995.
- D. G. York, J. Adelman, J. E. Anderson, Jr., S. F. Anderson, J. Annis, N. A. Bahcall, J. A. Bakken, R. Barkhouser, S. Bastian, E. Berman, W. N. Boroski, S. Bracker,

C. Briegel, J. W. Briggs, J. Brinkmann, R. Brunner, S. Burles, L. Carey, M. A. Carr, F. J. Castander, B. Chen, P. L. Colestock, A. J. Connolly, J. H. Crocker, I. Csabai, P. C. Czarapata, J. E. Davis, M. Doi, T. Dombeck, D. Eisenstein, N. Ellman, B. R. Elms, M. L. Evans, X. Fan, G. R. Federwitz, L. Fiscelli, S. Friedman, J. A. Frieman, M. Fukugita, B. Gillespie, J. E. Gunn, V. K. Gurbani, E. de Haas, M. Haldeman, F. H. Harris, J. Hayes, T. M. Heckman, G. S. Hennessy, R. B. Hindsley, S. Holm, D. J. Holmgren, C.-h. Huang, C. Hull, D. Husby, S.-I. Ichikawa, T. Ichikawa, Ž. Ivezić, S. Kent, R. S. J. Kim, E. Kinney, M. Klaene, A. N. Kleinman, S. Kleinman, G. R. Knapp, J. Korienek, R. G. Kron, P. Z. Kunzst, D. Q. Lamb, B. Lee, R. F. Leger, S. Limmongkol, C. Lindenmeyer, D. C. Long, C. Loomis, J. Loveday, R. Lucinio, R. H. Lupton, B. MacKinnon, E. J. Mannery, P. M. Mantsch, B. Margon, P. McGehee, T. A. McKay, A. Meiksin, A. Merelli, D. G. Monet, J. A. Munn, V. K. Narayanan, T. Nash, E. Neilsen, R. Neswold, H. J. Newberg, R. C. Nichol, T. Nicinski, M. Nonino, N. Okada, S. Okamura, J. P. Ostriker, R. Owen, A. G. Pauls, J. Peoples, R. L. Peterson, D. Petravick, J. R. Pier, A. Pope, R. Pordes, A. Prosapio, R. Rechenmacher, T. R. Quinn, G. T. Richards, M. W. Richmond, C. H. Rivetta, C. M. Rockosi, K. Ruthmansdorfer, D. Sandford, D. J. Schlegel, D. P. Schneider, M. Sekiguchi, G. Sergey, K. Shimasaku, W. A. Siegmund, S. Smee, J. A. Smith, S. Snedden, R. Stone, C. Stoughton, M. A. Strauss, C. Stubbs, M. SubbaRao, A. S. Szalay, I. Szapudi, G. P. Szokoly, A. R. Thakar, C. Tremonti, D. L. Tucker, A. Uomoto, D. Vanden Berk, M. S. Vogeley, P. Waddell, S.-i. Wang, M. Watanabe, D. H. Weinberg, B. Yanny, and N. Yasuda. The Sloan Digital Sky Survey: Technical Summary. *AJ*, 120:1579–1587, September 2000. doi: 10.1086/301513.