Probing the time-resolved Universe

Liam K. Hardy

Department of Physics & Astronomy
The University of Sheffield

A dissertation submitted in candidature for the degree of Doctor of Philosophy at the University of Sheffield

June 2016

"Life can only be understood backwards but it must be lived forwards."
Johann von Goethe

Declaration

I declare that no part of this thesis has been accepted, or is currently being submitted, for any degree or diploma or certificate or any other qualification at this University or elsewhere, except for Section 3.6 which briefly describes the results of my MPhys project report as part of the relevant publication (Hardy et al., 2015b).

This thesis is the result of my own work unless otherwise stated at the beginning of each chapter.

The following chapters have been based wholly or in part on my publications:

- Chapter 3 - Hardy et al. (2015b)
- Chapter 4 - Hardy et al. (2016b)

The other chapters (2 and 5) include work that has contributed to publications (e.g. Dhillon et al., 2014), or will contribute to future publications.

Summary

Studying the Universe in a time-resolved manner, including both high-speed variability, and transient events, allows us to explore the most extreme astrophysical environments. The study of transient astronomy and high-speed variability requires specialist instrumentation to explore the fastest timescales, and follow-up rapidly fading events. For this reason I have helped to develop two facilities for time-resolved astronomy, and used them to conduct two research programmes.

In this thesis I present the lab-based and on-sky commissioning of the high-speed imager ULTRASPEC. I include a theoretical analysis of the instrument throughput, and compare this with the observed throughput, finding that ULTRASPEC is performing to within $\sim 20\%$ of expectations in all filters. I also present a method for flux calibrating the non-standard *KG5* filter, using bolometric corrections and model spectra of stellar atmospheres.

I describe the design and operation of the 0.5m robotic telescope *pt5m*. I have developed a collection of software programmes which enable *pt5m* to function as an automated transient follow-up facility, which listens for new events and acts on them immediately, collecting and reducing the data. I present the results of performance tests, and scientific results from follow-up observations of transients such as GRBs, supernovae, and an outbursting X-ray binary.

I then used *pt5m* to conduct a search for eclipsing cataclysmic variables, and further investigated the structure of their eclipses with ULTRASPEC. I found or confirmed 13 new eclipsing systems, of which two are definitely suitable and six are possibly suitable for detailed modelling of their eclipse structure.

I also used ULTRASPEC, and later *pt5m*, to search for possible optical counterparts to recently discovered fast radio bursts (FRBs). FRBs are bright, short pulses of radio emission from unknown sources. I followed-up two FRB detections in detail, finding no evidence for counterparts in one field, but 5 variable sources in the other, all of which require further study.

Acknowledgements

It is not easy to express how grateful I am to all the people who have helped me get through my PhD. From academic colleagues to family and friends, they have all guided me and helped me through good and bad times.

Thank you to Vik Dhillon and Stuart Littlefair, the best supervisory team anyone could ask for. For keeping me motivated, offering expert academic (and non-academic) guidance, and supplying sufficient entertainment on long nights at the telescope, I am enormously grateful. Thank you also for putting up with my never-ending interruptions, and 'one-minute' questions that often turned into an hour of discussion. Thanks also for the doughnuts.

Thank you to Tom Marsh, Richard Wilson and Tim Butterley, for always answering my annoying questions about ULTRACAM software or *pt5m* problems. Your contributions to both of these projects far outshine mine, and I'm very grateful for all the work you put in before I came along.

Thank you to Paul Kerry, for helping me solve every IT and Linux problem imaginable, even when they perplexed you as much as me. Your Yorkshiremans enthusiasm and depth of computing knowledge is unbeatable.

Thank you to Clive, Paul, Simon, Susan, James and Justyn, for first teaching me as an undergraduate, and then being supportive in your own ways. I am grateful that you would always make time for me even though I'm not your student.

Thank you to Martin McAllister, for being such a hard-working and easy-going teammate. Your enthusiasm for data reduction has helped me through tedious tasks on several occasions, and I'm grateful for all the work we have shared.

Thank you to Dave Sahman, for sharing your nova knowledge and your solidarity regarding having a busy supervisor.

Thank you to Marvin, Saida, Andy, Claire, Patricia, Pav, Emile, Darren, Manu, Rob, Dan and the newbies for providing a welcoming working environment and an excellent range of experience and ideas, and for always being willing to answer my questions.

Special thanks to Chris Rosslowe for being the best PhD and campaigning buddy. Your enthusiasm for academic rigour, partying and a better world are inspirational.

To Emily, thank you for everything. Your support has been unwavering and I wouldn't have made it through without you.

Thank you to my parents, sister and extended family for supporting me in your own ways. You've left me to my own devices and given me your attention at just the right times.

Thank you to everyone at Sheffield People and Planet, for inspiring me to keep environmental and social protections at the top of the agenda, and for being a supportive network of caring people.

Thank you to everyone at New Roots, for giving me something to look forward to every week, and for raising my spirits even on the darkest days.

More formally, I am grateful to the Harry Worthington Scholarship for funding my PhD, and to the Sheffield Alumni Fund and Leverhulme Trust for funding aspects of our work with *pt5m* and ULTRASPEC. Thank you also to NARIT for hosting me for a short time in Thailand and providing a warm welcome.

Contents

Acronyms

2MASS	2 Micron All Sky Survey
AAVSO	American Association of Variable Star Observers
AGN	Active Galactic Nucleus
AMI	Arcminute Microkelvin Imager
ASAS-SN	All Sky Automated Survey for Supernovae
BH	Black Hole
CCD	Charge-Coupled Device
CSS	Catalina Sky Survey
CRTS	Catalina Real-time Transient Survey
CV	Cataclysmic Variable
DEC	Declination
DM	Dispersion Measure
DN	Dwarf Nova
DSS	Digitized Sky Survey
ESO	European Southern Observatory
FoV	Field of View
FRB	Fast Radio Burst
FWHM	Full Width at Half Maximum
GRB	Gamma-Ray Burst
GSA	Gaia Science Alerts
GW	Gravitational Wave
IGM	Inter-galactic Medium
ING	Isaac Newton Group of Telescopes
IR	Infrared
IRAF	Image Reduction and Analysis Facility
ISM	Inter-stellar Medium
LRN	Luminous Red Nova
MASTER	Mobile Astronomical System of the TElescope-Robots
MLS	Mount Lemmon Survey
NARIT	National Astronomical Research Institute of Thailand
NEO	Near-Earth Object
NL	Nova-like
NS	Neutron Star
NTT	New Technology Telescope

PSF	Point Spread Function
PTC	Photon Transfer Curve
RA	Right Ascension
RFI	Radio Frequency Interference
RGB	Red Giant Branch
RMS	Root Mean Square
SDSS	Sloan Digital Sky Survey
SE	SourceExtractor
SN	SuperNova
SSS	Siding Springs Survey
SUPERB	SUrvery for Pulsars and Extragalactic Radio Bursts
TDE	Tidal Disruption Event
TNT	Thai National Telescope
UV	Ultraviolet
VLT	Very Large Telescope
WD	White Dwarf
WHT	William Herschel Telescope
WMAP	Wilkinson Microwave Anisotropy Probe

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Time-Domain Astrophysics

Many studies conducted in the field of astronomy, both observational and theoretical, are concerned with rather long timescales. Most measurable aspects of these studies vary only on timescales of thousands, millions or billions of years, which cannot be observed by any single human. Such measurements and interpretations are mere snapshots in the course of astronomical history. Of course, looking deeper into the Universe, at objects further away and at higher redshifts, allows us to look back in time and understand how conditions throughout the Universe may have changed since the Big Bang. However, most of these changes cannot be seen to be happening 'live'.

Time-resolved astrophysics is different, and in a way much more exciting, because the timescales involved (from seconds up to months or years) are comprehensible, and the changes can be observed by a human. Part of this thesis is devoted to the study of eclipsing cataclysmic variables, and I would argue that there is little more exciting to observe in astronomy than the real-time total eclipse of a star as it passes behind its close companion in a matter of seconds. Another aspect of time-domain astrophysics is transient astronomy, and I believe the observation of the sudden appearance of a

previously unknown astronomical source is equally as exciting. High-speed variability and transient astronomy are the two key science themes in this thesis.

1.1.1 High-speed variability

The observed variability of astrophysical sources can be due to a number of physical processes, including pulsations, rotation, flares, outbursts, transits, eclipses and occultations. This variability can have valuable information encoded within it. For example, the timing and duration of the ingress and egress of an eclipse or transit are related to the relative sizes of the two bodies (Sackett, 1999; Savoury et al., 2011). Studies of pulsations can reveal information about the densities and composition of certain objects (Kurtz et al., 2013; Maxted et al., 2013). Characterising the flaring activity in relativistic jets can constrain the size and emission mechanisms of the emitting region (Gandhi et al., 2016).

The physical size and mass of an object effectively dictates how fast it can pulsate or rotate. This is due to the limit placed on the speed at which the bulk properties of an object can vary. The limit is the dynamical timescale, t_D, which is the timescale on which a star might expand or collapse if either the gravity force or the outward pressure force were to suddenly be removed. By definition it is the same as the free-fall timescale, and therefore depends directly on the mass, M, and radius, r, of the object, as shown in Equation 1.1:

$$t_D = (2r^3/GM)^{0.5}, \tag{1.1}$$

where $G = 6.67 \times 10^{-11}$ m^3 kg^{-1} s^{-2} is the gravitational constant. Since the dynamical timescale of an object is proportional to its radius (to the power $\frac{3}{2}$), compact massive objects such as white dwarfs, neutron stars and black holes show the fastest variability, on timescales of seconds to milliseconds (t_D for a typical white dwarf is 10s, for a neutron star it is 0.1ms, and for a stellar-mass black hole it is < 0.1ms). Exploring the highest

speeds of astrophysical variability therefore allows us to probe the most gravitationally extreme environments.

High-speed observations can also provide access to greater spatial resolution when observing eclipses, transits and occultations. In eclipsing binaries, sampling with a time resolution of a few seconds can reveal the relative radii of the components. For example, Parsons et al. (2010a) measured the radius of the white dwarf in NN Ser to a precision of 140km. To achieve the same precision with direct imaging would require a spatial resolution of 2×10^{-9} arcseconds - an impossible task. For occultations of stars by Solar System objects, or by the Moon, millisecond sampling is required to enable measurement of the angular diameter of the stars (a few milli-arcseconds), and determine whether they are single or binary systems (Richichi et al., 2016).

Most existing astronomical instrumentation is designed to study the unchanging Universe, and to extract the maximum number of photons from non-variable sources. Because of this, conventional instruments have long detector read-out times (up to a minute) making them very inefficient for studying fast variability. If we are to explore the fastest time-domains of astronomy, specialist instrumentation is required. This is discussed further in Section 1.2.

1.1.2 Transient astronomy

In astronomy, a transient object is considered to be any source/event/detection which is unpredicted and astronomically short-lived, lasting from a few years down to a few milliseconds. Transients can sometimes be highly explosive events, occurring in environments of extreme gravity, temperature, density and energy. They allow us to probe theories of fundamental physics at the limits of their predictive powers. Because of their short-lived nature, transients can also be elusive and challenging objects to study, and there are still many unanswered questions regarding their nature.

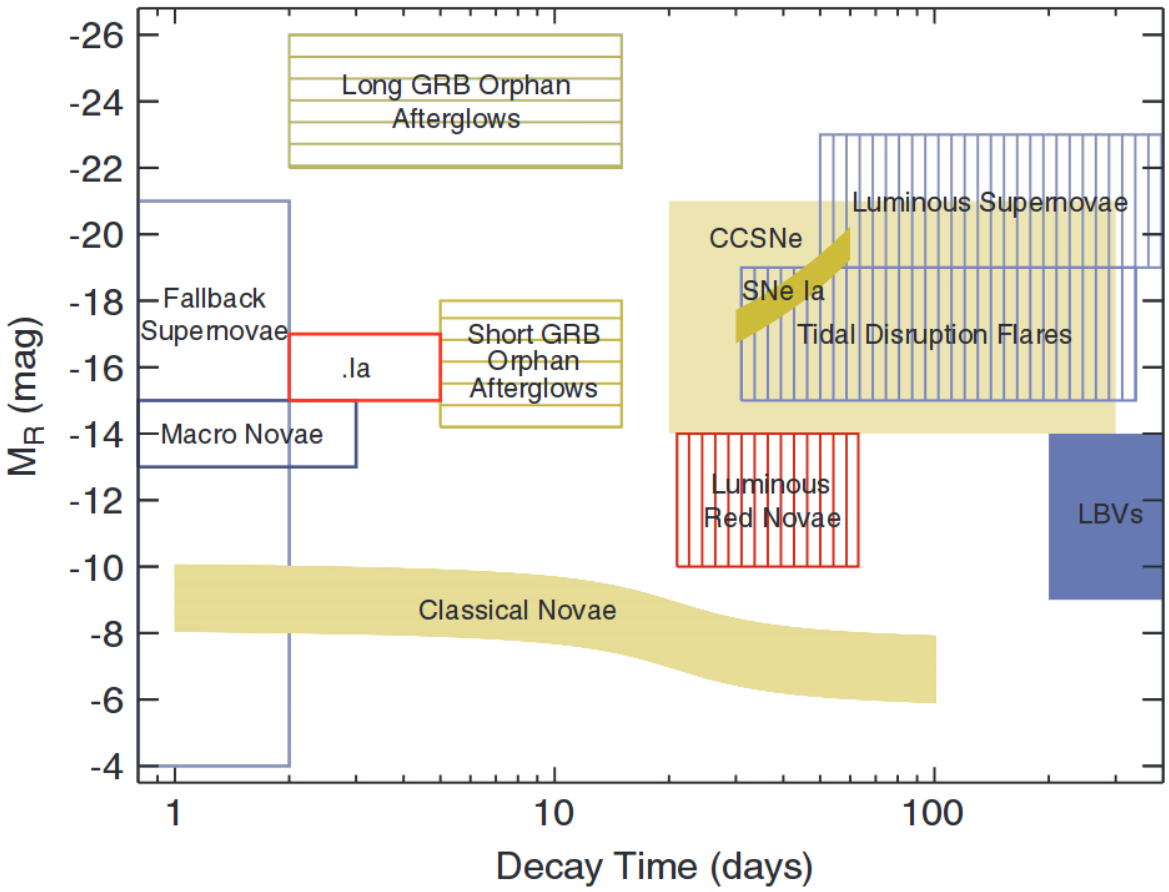

Figure 1.1: Plot of known or predicted optical transients, with peak R-band absolute magnitude plotted against the characteristic decay time (time to fade by 2 magnitudes from peak), reproduced directly from Rau et al. (2009). Filled boxes are well-studied classes with large numbers of known objects. Vertically hatched boxes are classes with very few candidates known. Horizontally hatched boxes are classes which are believed to exist but are yet to be detected. Empty boxes describe theoretically predicted events. Very bright transients (e.g. GRB optical afterglows) and Galactic transients (e.g. dwarf novae) are omitted.

Figure 1.1 shows a luminosity-phase space diagram marked with the regions occupied by known and predicted optical transients. It includes well known classes such as supernovae and gamma-ray bursts, as well as the less common and poorly understood classes such as luminous red novae and tidal disruption events. Supernovae, being the final energetic events associated with the deaths of stars, are of interest to studies of stellar formation and evolution, but they can also be used as distance indicators (della Valle & Panagia, 1992; Phillips, 1993). Type Ia SNe have of course been famously used to detect

and characterise the acceleration in the expansion of the Universe (Frieman, Turner & Huterer, 2008). Nevertheless, there are still many unknowns in supernovae research, and still some unknown systematics which prevent more precise distance measurements (Howell et al., 2009).

Gamma-ray bursts on the other hand represent some of the most energetic events in the Universe, but we still know little about their energy generation mechanisms (van der Horst et al., 2015). This is in part because they are so short-lived, fading at all wavelengths over the course of minutes to hours. We know even less about other rapid transients such as fast radio bursts (see Section 1.3.2), which have radio lifetimes of a few milliseconds, and may not show any optical counterparts at all.

One class of transient that has been well-studied for many decades are the classical and dwarf novae eruptions from cataclysmic variables (CVs). CVs may be the progenitors of Type Ia SNe, and are highly relevant to studies of stellar evolution and the enrichment of the interstellar medium (Starrfield, Iliadis & Hix, 2016). Despite being the focus of much research and observation for the best part of a century, there are still outstanding questions. Can the white dwarfs in CVs grow in mass with time? Will all CVs eventually undergo novae eruption? What causes the additional angular momentum loss required to explain CV evolution with orbital period? (Knigge, Baraffe & Patterson, 2011). Section 1.3.1 below discusses my contribution to answering these questions, which involves searching for and characterising eclipsing CVs, most of which were initially discovered via their transient outbursts.

There are several classes of transient events which have only been recently discovered, and which have emission mechanisms and progenitors which are still unknown. In addition, it is perfectly conceivable that completely unknown and unpredicted classes exist and are yet to be discovered. Transient astronomy has therefore become a large and busy field. This is also partly due to the advent of large-scale surveys which can detect large numbers of transient events, at ever faster cadences. In the last decade or

so we have seen major results from a number of surveys, including the Palomar-Quest survey (Djorgovski et al., 2008b), the Catalina Real-time Transient Survey (Drake et al., 2009), the Palomar Transient Factory (Law et al., 2009), the All-Sky Automated Survey for Supernovae (ASAS-SN, Shappee et al., 2014), MASTER (Lipunov et al., 2010), OGLE (Udalski et al., 2002), Pan-STARRS (Kaiser et al., 2002), SkyMapper (Keller et al., 2007), Gaia (Wyrzykowski et al., 2012) and *Swift* (Gehrels et al., 2004).

Alongside this boom in transient surveys and detections, there has also been great progress in the field of automatic detection and classification software, which allows us to investigate new transients more quickly and easily (Law et al., 2009; Bloom et al., 2012; Wright et al., 2015). It is now possible to automatically follow up gamma-ray bursts (GRBs), and other transient detections, without any input, or subsequent delay, from a human user (see e.g. Hardy et al., 2015b,c,d). As the follow-up community becomes more adept at rapid follow-up, we will probe deeper into the early-time behaviour of these events.

In the near future, the Large Synoptic Survey Telescope (LSST) will undertake the deepest, fastest survey of approximately half the entire sky (Ivezic et al., 2008). The survey will use an 8.4m aperture telescope with a 10 deg^2 field-of-view to detect 20 billion sources down to a magnitude of $r \sim 24.5$, and re-image them all every few nights, creating an unprecedented capability for detection of transients. Due to begin operations in 2021, the LSST will create 15 TB of data every night, and huge efforts are underway to process and analyse that data into usable products. LSST is expected to discover up to 1000 supernovae every night, and the volume of transient detections from the LSST will be larger than all previous surveys put together. Any transient event brighter than r=21 magnitudes and lasting longer than a few days will not be missed.

This enormous number of transient detections will need a substantial array of dedicated follow-up facilities to study these new objects, since any new detections will need confirmation and further classification, and many will need continued follow-up for many

months (e.g. supernovae or tidal disruption events). There has been some growth in the number of transient follow-up facilities available, e.g. Holmes et al. (2011); Topinka et al. (2014); Street et al. (2014), but not enough to fully exploit the number of transients now being detected, and which are expected to be found by future surveys like LSST. The bottleneck to gathering the most useful and relevant information remains at the follow-up stage. Section 1.2 discusses the requirements for such follow-up facilities.

For a more detailed overview of some of the most interesting and important classes of transients, see Appendix A.

1.2 Development of Instrumentation

Studying the Universe in a time-resolved manner often requires specialist instrumentation and analysis procedures. High-speed observations require instrumentation designed to achieve high cadences with good efficiency. Transient follow-up observations require rapid-response telescopes with schedule override options, and the ability to provide long-term monitoring of interesting targets. I have therefore focused approximately half of my efforts on developing the required instrumentation for high-speed observations and transient follow-up, and the other half of my efforts on using these new facilities I have developed to study transient and variable phenomena. The instrumentation/astrophysics split is reflected in this thesis, with half describing the development and commissioning of two major instruments/facilities for time-resolved astronomy, and half describing the use of this instrumentation.

1.2.1 High-speed photometry with ULTRASPEC

To probe the high-speed variability of compact objects and close binary systems, high-speed instrumentation is needed. Equation 1.1 shows that the more massive and compact an object is, the shorter its dynamical timescale, and therefore the more rapidly it can

show, for example, pulsational or rotational variability. If we want to investigate the physics of white dwarfs, neutron stars and black holes, or spatially resolve the components of close binary systems, high-speed observations will be required.

The requirements of such high-speed instrumentation include: the ability to take and process images at high frame rates in order to explore the fastest types of variability; good instrumental throughput and detector efficiency to ensure as many photons are collected as possible within the short exposure times; and low dead-times to make efficient use of observing time.

To this end we have developed and commissioned ULTRASPEC for high-speed photometry on the 2.4m Thai National Telescope. In Chapter 2 I describe the instrument and its operation, and present the laboratory and on-sky testing of ULTRASPEC which I have performed. I also describe a flux calibration procedure which I carried out for a non-standard filter, using bolometric corrections and model stellar spectra.

1.2.2 Transient follow-up with *pt5m*

To study transients, we need dedicated follow-up facilities to classify new events, and study their evolution over short and long timescales.

The requirements of such a transient follow-up facility include: full access to observing time, so that observations can be conducted immediately in the event of an exciting new detection; the ability to respond and move quickly, pointing to new targets within seconds so as not to miss the early evolution of fast transients; the ability to study transient sources in multiple colours to help with classification; and the ability to listen for and analyse transient event alerts, and act on them accordingly by scheduling the appropriate observations, autonomously and without the interaction of a human user. For this purpose we have fully roboticised the 0.5m *pt5m* telescope on La Palma (Hardy et al., 2015b), which now automatically listens for transient events and immediately

attempts the necessary follow-up observations, as well as automatically reducing the images as soon as they are collected.

A description of the facility and its performance is presented in Chapter 3. I also describe the transient follow-up software I have developed, and discuss some example studies which I have conducted.

1.3 Observational Campaigns with *pt5m* and ULTRASPEC

Having developed ULTRASPEC to be able to explore high-speed variability, and *pt5m* to be able to provide immediate and rapid follow-up of transient events, I then went on to use these tools to conduct two observational campaigns. One study has been the systematic search for eclipses in newly discovered cataclysmic variables (CVs), which are usually discovered by their transient dwarf nova outbursts. The other has been to conduct optical follow-up of a new class of radio transients called Fast Radio Bursts.

1.3.1 Eclipsing CVs

CVs are close binary systems in which a white dwarf accretes matter from its (usually) main-sequence companion via an accretion disk or magnetic stream. CVs undergo outbursts which make them detectable as transient events (Hellier, 2001). They are considered to be possible progenitors of Type Ia supernovae (Starrfield, Iliadis & Hix, 2016), which are themselves vital tools for studying cosmology and the expansion of the Universe (Frieman, Turner & Huterer, 2008; Betoule et al., 2014).

Some CVs show eclipses, and modelling of the eclipse structure of these systems can be used to evaluate the masses and radii of the binary components, as well as other system parameters (see e.g. McAllister et al., 2015 and references therein). Measuring

these parameters precisely through light curve modelling provides important data to test models of CV evolution (Knigge, Baraffe & Patterson, 2011; Schreiber, Zorotovic & Wijnen, 2016). Currently, only around 35 systems have had their parameters measured precisely, and more are needed to verify models of evolution and determine whether or not CVs can become SN Ia. Many more systems are known to be eclipsing, but not all eclipses show the necessary structure required for modelling to be conducted. For precise system parameter derivation, the eclipses must be studied at sufficient time resolution to be able to resolve the ingresses and egresses of the components, which requires high-speed imagers.

To contribute to the ongoing measurement of CV system parameters, I conducted a search for new eclipsing systems using *pt5m*. To evaluate their potential for light curve modelling I then used the newly-commissioned ULTRASPEC (as well as its predecessor ULTRACAM, see Chapter 2), to obtain high-speed light curves. This work has been submitted for publication (Hardy et al., submitted), and is discussed in full in Chapter 4.

1.3.2 Fast Radio Bursts

Fast radio bursts (FRBs) are a new class of radio transients, first discovered by Lorimer et al. (2007). These short (few ms), bright (peak flux $\sim$1 Jy) radio pulses are proposed to be extra-galactic. This requires them to have high energies, and thus many exotic progenitor models have been proposed (Katz, 2016a). If FRBs are extra-galactic they could impart valuable information about the density and content of the intergalactic medium (Macquart & Koay, 2013), and may help us understand the origins of dark matter (McQuinn, 2014), and even dark energy (Zhou et al., 2014).

To understand the emission mechanisms and identify the possible progenitors of FRBs, detections or constraints at other wavelengths are very important. The locations

of FRBs are currently poorly localised, with error circles of $10'$ to $30'$. However, if an optical counterpart to an FRB was to be discovered, it could allow us to measure a redshift and hence verify that they are extra-galactic, and also allow us to gain some understanding of the host conditions of the burst progenitor. A distance measurement would also allow us to investigate the baryon density of the intergalactic medium (Keane et al., 2016).

We have therefore used ULTRASPEC to attempt immediate optical follow-up of new FRBs discovered in real time. In Chapter 5 I present our attempted follow-up observations of two real-time FRB events, and the subsequent data analysis and results. For one burst in particular, I have also utilised our access to *pt5m*, using it to conduct further follow-up observations of several variable sources found in the FRB field. In this instance, as with the eclipsing CVs project, I have used both *pt5m* and ULTRASPEC to study FRBs, making the most of the facilities I have developed in this thesis.

Chapter 2

Commissioning ULTRASPEC on the 2.4m Thai National Telescope

This chapter discusses the lab-based and on-sky commissioning of ULTRASPEC, which I led. My work on this has been included in the published paper describing the instrument by Dhillon et al. (2014), on which I am a co-author. The following data collection, reduction, analysis and text are all my own.

2.1 Introduction

As discussed in Chapter 1, the study of compact objects requires high-speed observations. Such observations usually require purpose-built instrumentation to be able to efficiently probe the highest cadences. Some instrumentation of this type already exists. For example, our group developed ULTRACAM, a high-speed camera which can observe in three colours (e.g. u', g' and r') simultaneously, at frame rates up to 300 Hz in the most extreme cases (Dhillon et al., 2007). Table 2.1 shows an up-to-date list of operational or proposed high-speed instruments. In addition to this list there are also tests underway to prepare CMOS-type sensors, which are commonly used in commercial cam-

era technology, for astronomical instrumentation (Wang et al., 2014). Compared with the more established charge-coupled devices (CCDs), complementary metaloxidesemiconductor (CMOS) sensors can achieve much higher frame rates when reading out a full frame, because each pixel has its own amplifier and is read out simultaneously. Whilst the median noise levels on CMOS sensors tends to be lower than for CCDs, the noise pattern is less uniform. The main disadvantage of CMOS sensors is that they have worse quantum efficiencies than CCDs due to the presence of electronic components mounted directly onto pixels. This can be partially mitigated using microlenses. Further discussion on CMOS sensors is outside the scope of this thesis.

The ULTRASPEC instrument is an imaging camera which makes use of an electron-multiplying CCD (EMCCD) and a purpose-built data acquisition system (similar to that of ULTRACAM, see Dhillon et al., 2007). It began its life as a detector and cryostat only, which was originally designed to be used for high-speed spectroscopy in conjunction with the European Southern Observatory's EFOSC2 spectrograph on the 3.6m telescope at La Silla. Now ULTRASPEC is a fully functioning imager, with 12 re-imaging lenses and a 6-position filter wheel, mounted on the 2.4m Thai National Telescope (TNT) at the Thai National Observatory in Doi Inthanon, Thailand.

This chapter details the extensive testing completed during the commissioning phase of the new ULTRASPEC instrument. I start by describing and explaining the instrument design and operation in Section 2.2. Section 2.3 discusses the tests carried out in the lab in Sheffield, which explore the properties of the EMCCD itself, and the different options for reading out frames. The on-sky testing, including a full study of the telescope and instrument throughput, is presented in Section 2.4. In Section 2.5 I present a method for flux calibrating the non-standard KG5 filter. I discuss a short study of superluminous supernovae undertaken with ULTRASPEC in Section 2.6. Finally, for completeness I mention my experience in training the local observers in Thailand in Section 2.7, before concluding the chapter in Section 2.8.

Table 2.1: List of operational or proposed* high-speed instruments, including their maximum frame rates and detector types. FT: frame-transfer. SPADs: Single Photon Avalanche Diodes. STJ: Superconducting Tunnel Junction. MKID: microwave kinetic inductance detector.

Name	**Detector Type**	**Max Frame Rate (Hz)**	**Reference**
ULTRACAM	(FT)CCD	300	Dhillon et al. (2007)
ULTRASPEC	(FT)EMCCD	150	Dhillon et al. (2014)
HiPERCAM*	(FT)CCD	$> 10^3$	Dhillon et al. (2016)
CHIMERA	(FT)EMCCD	10^3	Harding et al. (2016)
ADVS	CCD	60	Barry et al. (2015)
Argos	(FT)CCD	1	Nather & Mukadam (2004)
Puoko-Nui North	(FT)EMCCD	100	Chote et al. (2014)
Xinglong 3-Channel Detector	CCD	10	Mao et al. (2013)
Sparc4*	EMCCD	>1	Rodrigues et al. (2012)
OCTOCAM*	(FT)CCD,EMCCD & IR-array	100	de Ugarte Postigo et al. (2010)
HESS 7-pixel camera	Photo-multipliers	10^6	Deil, Domainko & Hermann (2008)
AquEYE & IquEYE	SPADs	10^7	Barbieri et al. (2012); Zampieri et al. (2015a)
GASP	CMOS & EMCCD	10^5	Collins et al. (2009)
OPTIMA	SPADs	10^6	Kanbach et al. (2008)
S-CAM	STJ	10^5	Verhoeve et al. (2006)
SALTICAM	(FT)CCD	10	O'Donoghue et al (2006)
ARCONS	MKIDs	10^5	Mazin et al. (2013)

2.2 Instrument Design

This section describes the hardware and software design of ULTRASPEC. The reader is referred to the published paper for finer details (Dhillon et al., 2014).

2.2.1 Optics

The two optical design requirements in building a high-speed imager are that the field of view is sufficiently large to ensure suitable comparison stars are available (to enable differential photometry to be performed), and that the resolution provided by the pixel scale is sufficient to properly sample the expected image quality. The median seeing at Doi Inthanon is 0.9" (Dhillon et al., 2014), thus Nyquist sampling of the median image quality requires a pixel scale of approximately 0.45"/pixel. With a detector size of 1024×1024 pixels, this equates to a field of view of $7.7' \times 7.7'$, which is large enough to have an 80% chance of finding a bright ($R < 11$) star in the field at average galactic latitudes (Dhillon et al., 2014). This would provide sufficient comparisons stars for differential photometry. The optical design of ULTRASPEC therefore requires a 4× focal reducer in order to reach the required pixel scale, since the TNT plate scale is 8.6 arcsec/mm and each ULTRASPEC pixel is 13μm.

The ULTRASPEC optical design is shown in Figure 2.1, and is comprised of 12 lenses in 3 separate barrels. The first barrel (along the light path from telescope to CCD, called barrel 3), holds 5 lenses and acts as a collimator, producing a collimated beam which then passes through the plane-parallel filter. The next two barrels (barrel 1 and barrel 2) host the other 7 lenses and form the camera optics. Figure 2.2 shows the design of the barrels which hold the lenses in place. Included in the barrel design is a 'U-block' component which connects barrels 2 and 3, and provides sufficient space for the filter wheel to be positioned in the light path without being connected to the lenses or barrels at all. This isolation of the filter wheel minimises flexure by reducing the weight of

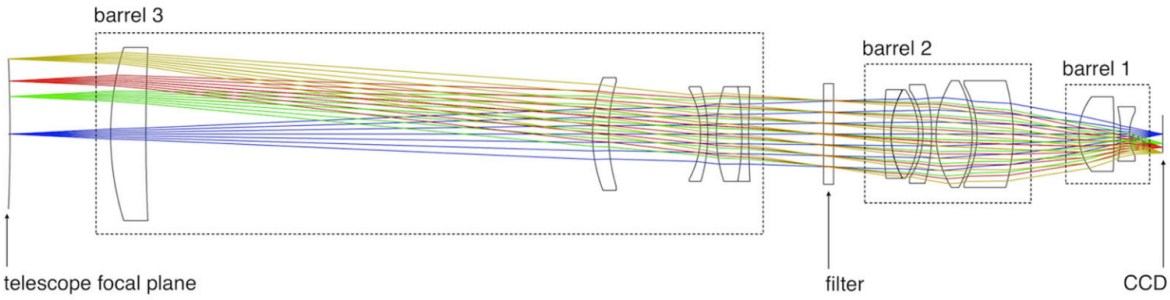

Figure 2.1: Ray tracing through the 12 lenses of ULTRASPEC. The three barrels are labelled, as well as the location of the telescope focal plane, the CCD, and the filter in the collimated beam. Reproduced from Dhillon et al. (2014).

components held by the barrels.

The design includes two lens doublets, which require high-transmission adhesives to hold them together. The glass materials in the second doublet (lenses 6 and 7 in Figure 2.1) have different thermal expansion coefficients, meaning a flexible silicon adhesive was required for these components.

The 6-position filter wheel is usually loaded with the normal Sloan Digital Sky Survey (SDSS) filters (u', g', r', i', z'), although other broad and narrow band options are available (see Table 2 in Dhillon et al., 2014).

2.2.2 Mechanics

The mechanical requirements for ULTRASPEC were driven primarily by the specifications of the TNT, and its rotating Nasmyth instrument cube (see Figures 2.3 and 2.4). The TNT requires that instruments are not heavier than 100kg, and do not extend more than 1m away from the cube. These tight requirements meant that carbon fibre struts and aluminium mounting plates were the ideal choice for constructing the instrument chassis. The carbon fibre has a low thermal expansion coefficient, and is strong and inflexible, which is vital for an instrument that rotates with the tracking of the sky. The design also allows the instrument to be open, allowing access to all components for maintenance.

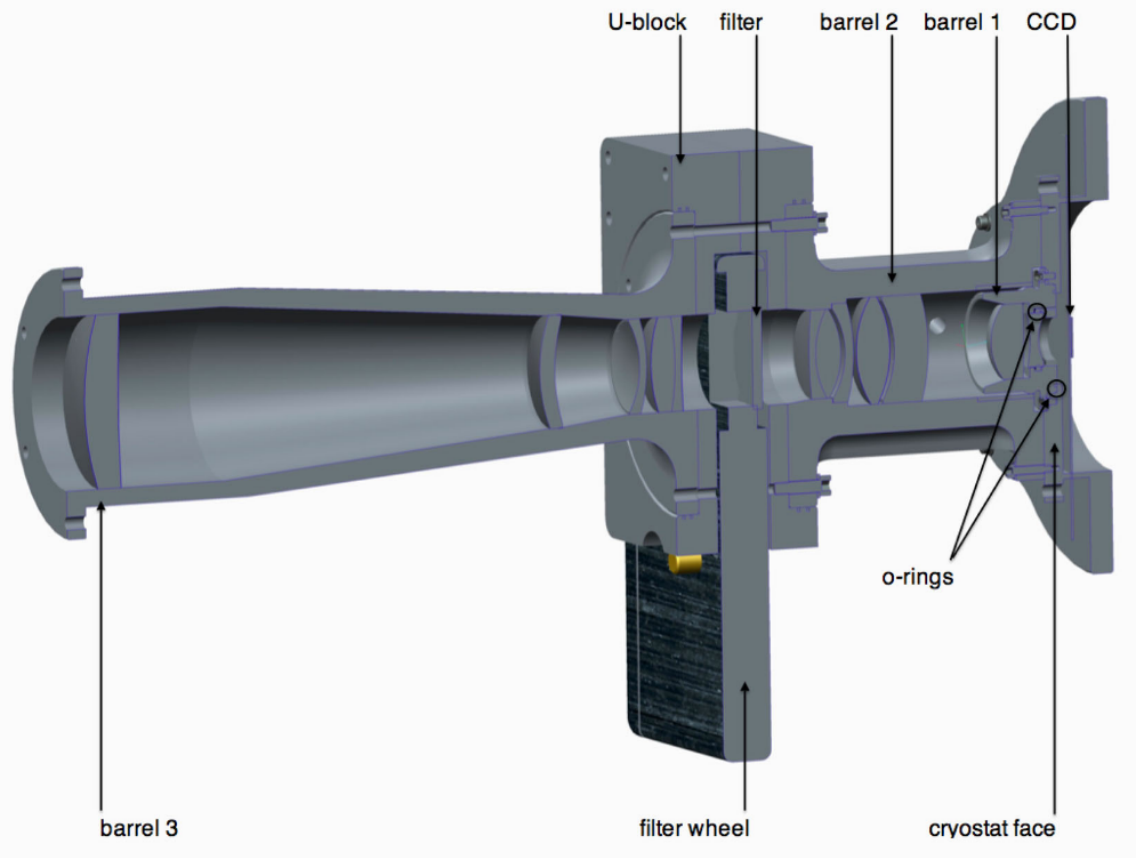

Figure 2.2: Cut-through of the optics barrels, including the U-block which holds barrel 3 to barrel 2, and in which the filter wheel is located but mechanically isolated from the U-block itself. Reproduced from Dhillon et al. (2014).

The optics barrels themselves are directly attached to the cryostat, and not to the chassis. Otherwise, any flexure in the chassis would also be enforced in the optics barrels, which would introduce further image degradation.

Instead of having a clear window between the cryostat and the optics barrels acting as a vacuum seal, the final lens itself acts as the seal, requiring polishing of the lens rim and an o-ring to maintain the vacuum. A cable twister arm is included in the design to allow the cables to be guided away from the instrument, to reduce electrical interference and tangling due to rotation. A diagram showing the instrument components is shown in Figure 2.3, and a photo of ULTRASPEC is shown in Figure 2.4.

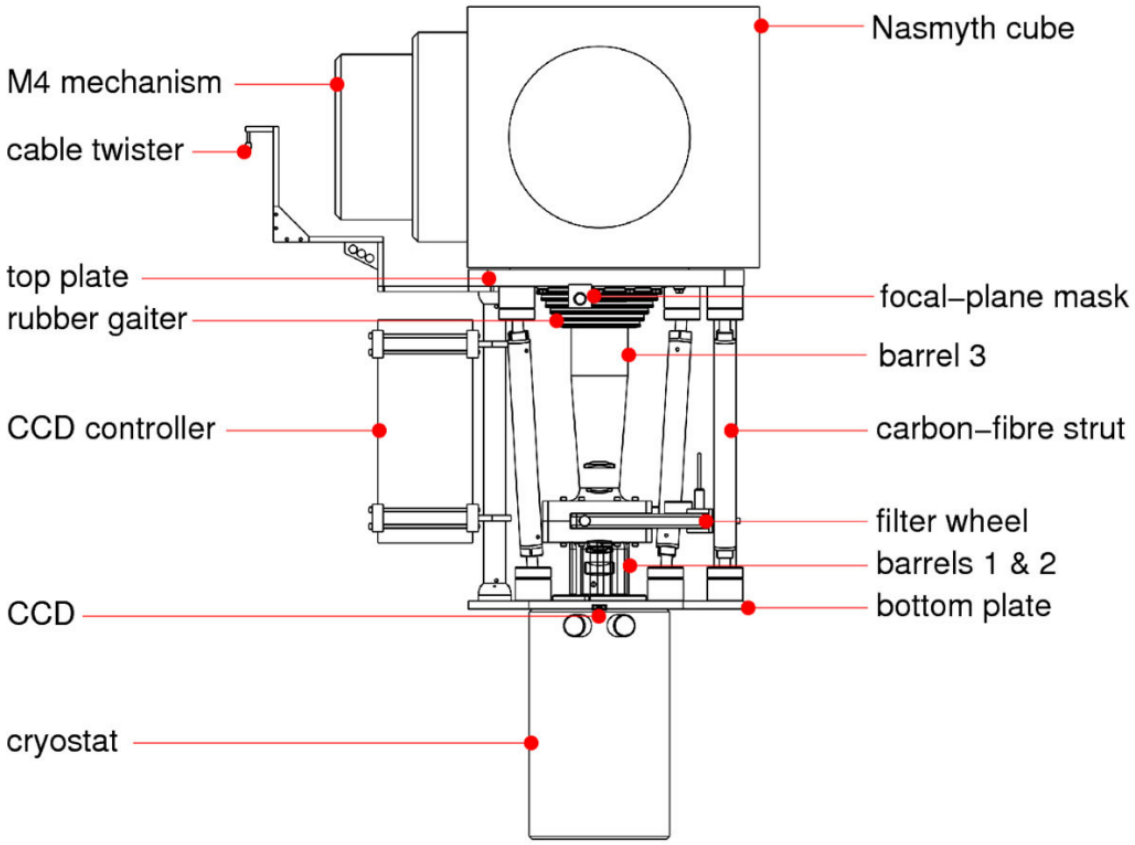

Figure 2.3: Diagram showing the mechanical components of ULTRASPEC including the chassis, cable twister, optics barrels, cryostat and CCD controller. The Nasmyth cube is used for up to 4 instruments which can be mounted simultaneously on the TNT, with the M4 mirror being physically moved to point to any given instrument. The labelled components can also be seen in the matching photo in Figure 2.4. Reproduced from Dhillon et al. (2014).

2.2.3 Detector

ULTRASPEC employs one of the best cosmetic grade, frame transfer electron-multiplying CCDs built by e2v. The imaging area is 1024×1024 pixels, each of which is 13μm in size. The CCD is thinned, back-illuminated, and coated with the standard e2v mid-band anti-reflection coating. A detailed explanation of the electron-multiplying functionality follows in Section 2.2.4.

The CCD has a full well capacity of 80,000 electrons, and thus we set the gain to be around 0.8 electrons/ADU, such that this capacity roughly matches the saturation

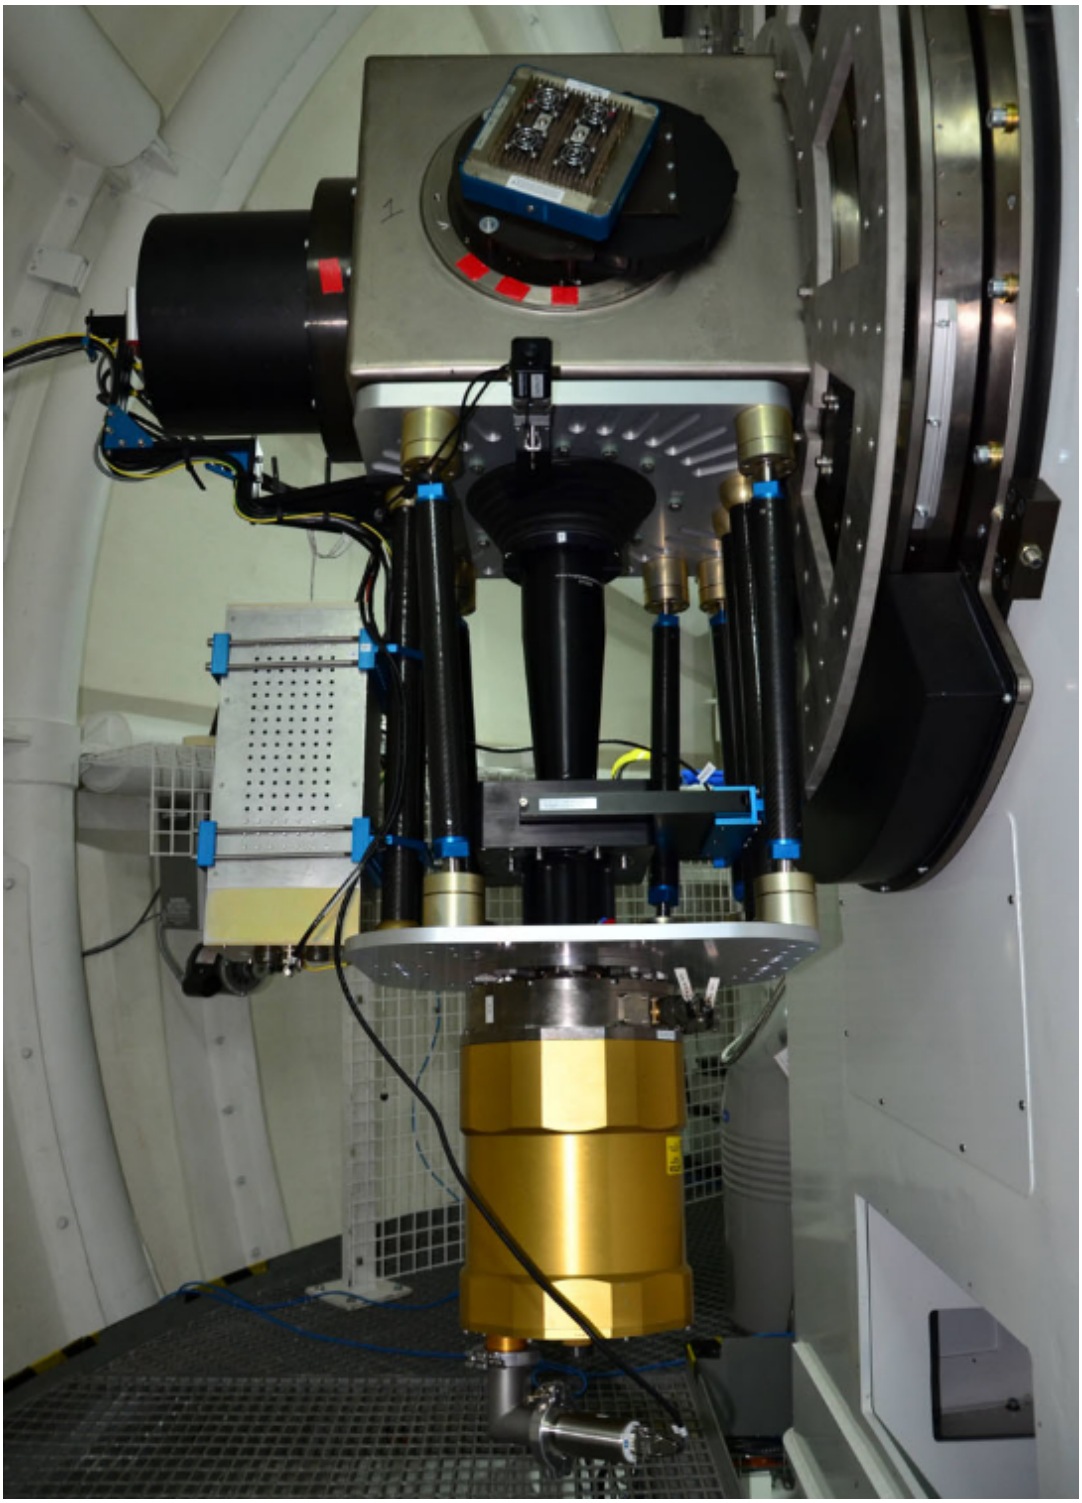

Figure 2.4: Photograph of ULTRASPEC mounted on the Nasmyth cube of the TNT. The cables can be seen following the cable twister and out of the shot to be secured safely to the telescope dome, which is slaved to rotate with the telescope. See Figure 2.3 for a labelled diagram showing the components. Reproduced from Dhillon et al. (2014).

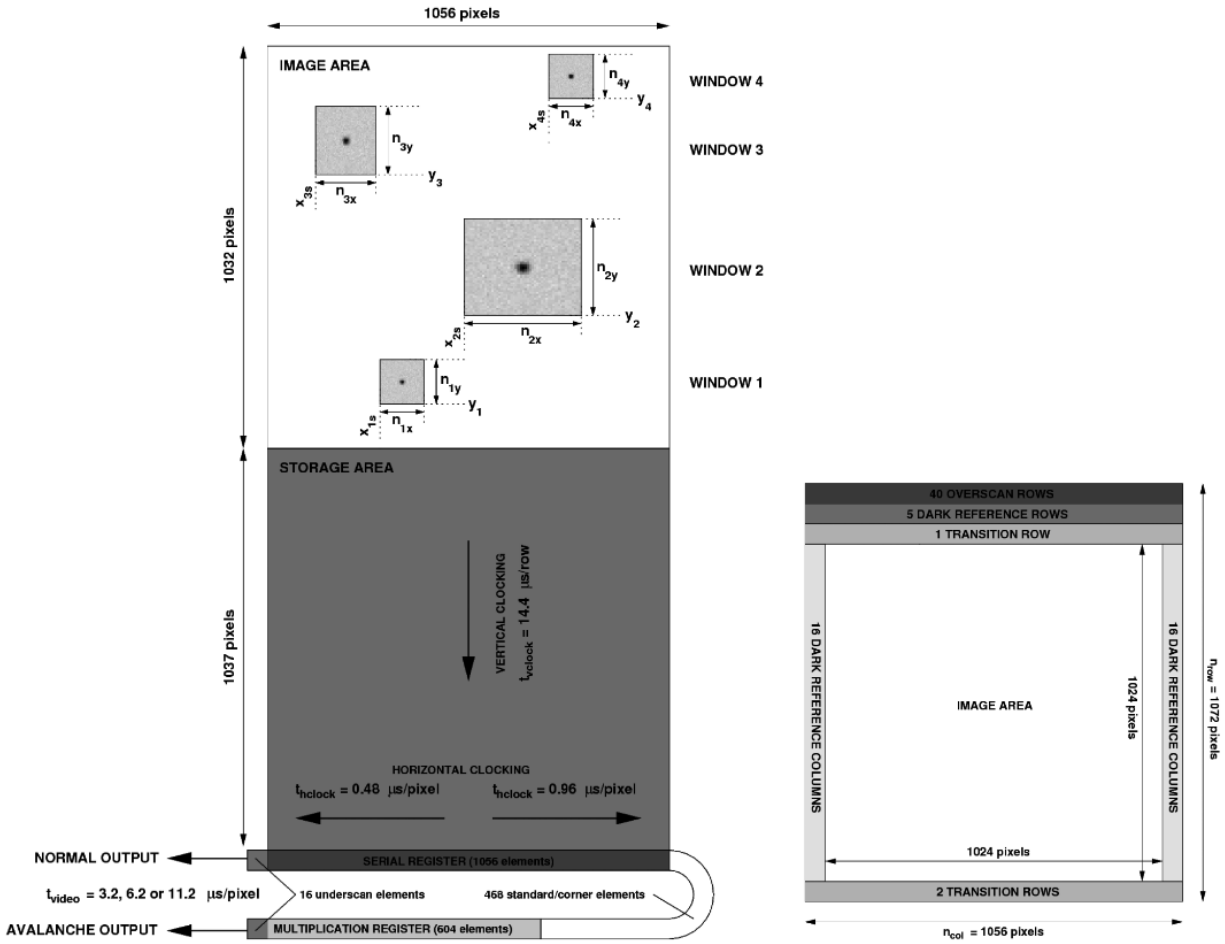

Figure 2.5: Left: schematic of the ULTRASPEC EMCCD, showing some example windows around target stars. The image area and storage area are shown, and pixel clocking times are noted. The serial registers are shown at the bottom, with both the normal and avalanche (see Section 2.2.4) outputs. Right: schematic of an ULTRASPEC image frame, with overscan and dark reference rows in addition to the 1024 × 1024 imaging pixels. Reproduced from Dhillon et al. (2014).

limit of the 16-bit analogue-to-digital converter in the CCD controller (saturation limit = 65,535 analogue-to-digital units, ADU).

The detector is housed inside a 2.5 litre cryostat, which is vacuum pumped to prevent condensation of particles onto the CCD. The vacuum usually lasts up to 12 months, and the 2.5 litre capacity provides enough liquid nitrogen to last around 20 hours. The cryostat is therefore refilled twice a day.

Frame Transfer

One of the key components of the ULTRASPEC EMCCD which enables high-cadence observations is the frame transfer design. This allows the next frame to begin exposing while the previous frame is being read out, significantly reducing the dead time. Essentially, the full CCD is composed of two equally sized square arrays of pixels, one of which is the imaging area, and the other is the storage area, which is covered and does not receive photons from the telescope. A pictorial representation of the ULTRASPEC CCD and typical image frame are shown in Figure 2.5.

After an image has finished exposing, the collected charge is quickly clocked down vertically into the storage area. The next image can begin exposing immediately, meaning the dead time between images is only the time it takes to clock the image into the storage area (around 15 ms). The first image is then read out through one of two outputs (normal, or the EM-gain channel, see Section 2.2.4).

Binning and Windowing

To speed up imaging even further, we can use small windows around the important sources in the field, and can also use on-chip binning of pixels, reducing the read-out time of the image. With windowing, all pixels are clocked vertically and horizontally as normal, but any pixels outside the requested windows are simply dumped away without being digitised. This speeds up the read-out of each image. When binning, the vertical and horizontal clocking is conducted in a way that essentially combines two or more pixels together, clocking the combined charge as if it were from a single pixel. A 2×2 binning setting is often used at the TNT whenever the seeing is worse than ~ 1.5".

Clear Mode

Whilst most observations use the frame transfer and windowing capability to enable higher cadences, occasionally we need to conduct short, full-frame exposures. To allow this there is also an option to run the detector in 'clear' mode, in which the chip is wiped clean of any charge before each short exposure takes place, removing the build up of photo-electrons produced while the previous frame was reading out. The result produces short, full-frame exposures, but with vastly reduced efficiency due to the increased dead time of 13 seconds. Clear mode is used when one needs to control the exposure time in full-frame mode, such as when taking bias frames and flat fields.

No Shutter

One important aspect of ULTRASPEC is the absence of a mechanical shutter. The frame transfer timescale of 15 ms is much faster than any mechanical shutter can operate, although the absence of a shutter does introduce some vertical star trails in the presence of very bright stars, as the pixels are clocked down through the chip. This issue can sometimes be avoided by blocking out the light from bright stars using a focal plane mask. The mask is mounted on a sliding arm at the telescope focal plane, and can be robotically moved in and out of the light path. In addition, the mask is also used in drift mode (see below), and as a fixed shutter, to help block out any light which might enter the light path when taking bias frames and dark frames.

Drift Mode

Even with small windows and high binning factors, there is a $\sim$ 40 Hz limit to the frame rate of normal operation, because of the 14.9ms it takes to move the image into the storage area. To overcome this limit ULTRASPEC can also operate in 'drift mode'. Drift mode uses small windows placed on the border of the image and storage areas. In

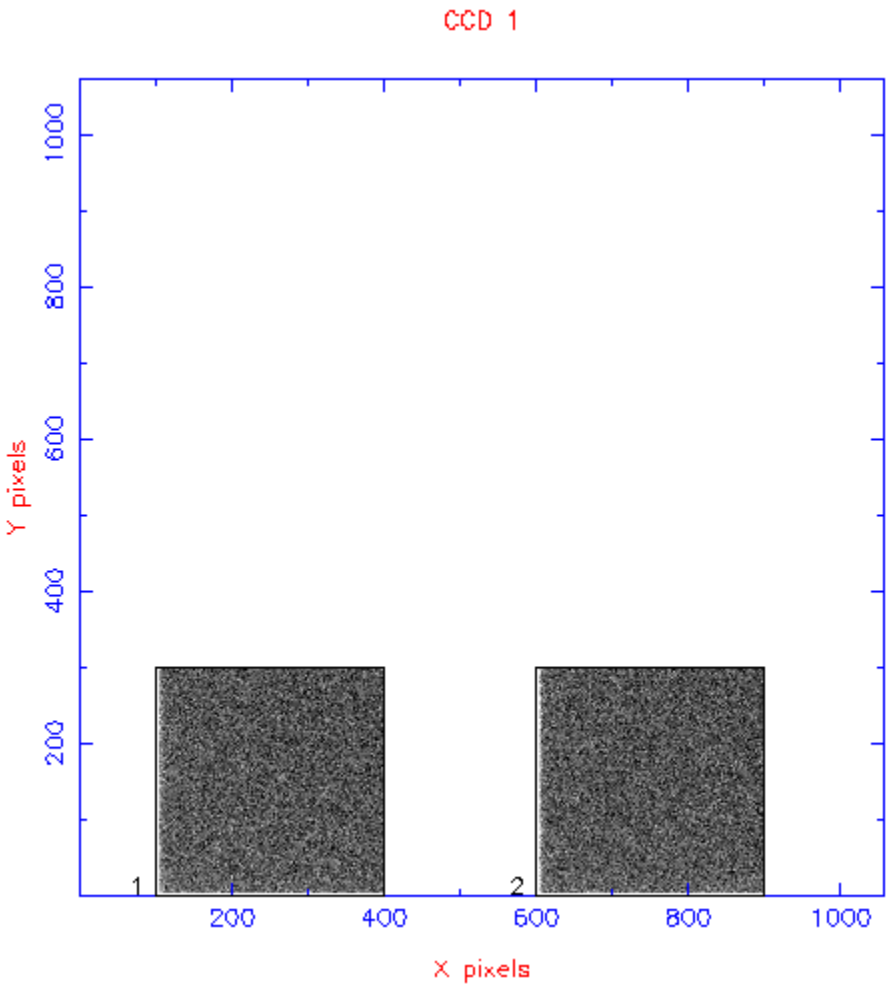

Figure 2.6: An example frame taken in drift mode, showing only the two small windows located at the bottom of the imaging area. The windows are clocked into the storage area before the previous windows have been read out. The stack of images makes its way down the frame transfer region, and the overall dead time is reduced because less clocking is required for each frame.

this mode, only the small windows themselves are clocked into the top of the storage area, rather than the whole frame, reducing the overall clocking time. A number of such windows can be present in the storage area, queueing up behind each other and gradually making their way down to the serial register. In this way very short exposures (<10ms in extreme cases) can be taken, with the dead time being very small (<1ms). An example frame taken in drift mode can be seen in Figure 2.6, and drift mode is explained further in Appendix A of Dhillon et al. (2007).

One key component of the ULTRASPEC instrument that allows drift mode to work is the presence of the opaque, movable mask mounted in the focal plane. This mask can cover some or all of the CCD imaging area by sliding in from the top of the image shown in Figure 2.6. This is used in drift mode to block out all light falling on the chip above the drift mode windows. If this were not in place, the windows would collect light from the sky or other nearby stars whilst being clocked down in to place at the bottom of the image area.

2.2.4 Electron Multiplying Readout

A key challenge when studying faint objects at high speed is that short exposure times often mean very low signal-to-noise ratios. Every frame has added read-out noise, and with low flux from the source, the overall noise in each image could be dominated by the read-out noise. One way to essentially remove the effect of read-out noise is to use an electron multiplying output.

An electron multiplying CCD generates 100's–1000's of times more electrons for each incoming photon than a conventional CCD (Tulloch & Dhillon, 2011). During the read-out process of an EMCCD, photo-electrons are clocked through rows and columns in the same way as a normal CCD. The difference begins when the charge passes through the serial register before reaching the output. In an EMCCD, the charge passes through an extended serial register where the electron multiplying occurs. Here the standard low voltage clocking of $\sim$ 5V is replaced by high voltage clocking using $>$ 40V, which multiplies photo-electrons via impact ionisation (Jerram et al., 2001). This occurs when photo-electrons are clocked into a very deep potential, where they have a small probability of exciting additional electrons within the semiconductor lattice, using the large amount of potential energy they gain when entering the well. This probability is small, but rises sharply with the potential well depth, or clock voltage (which can be controlled

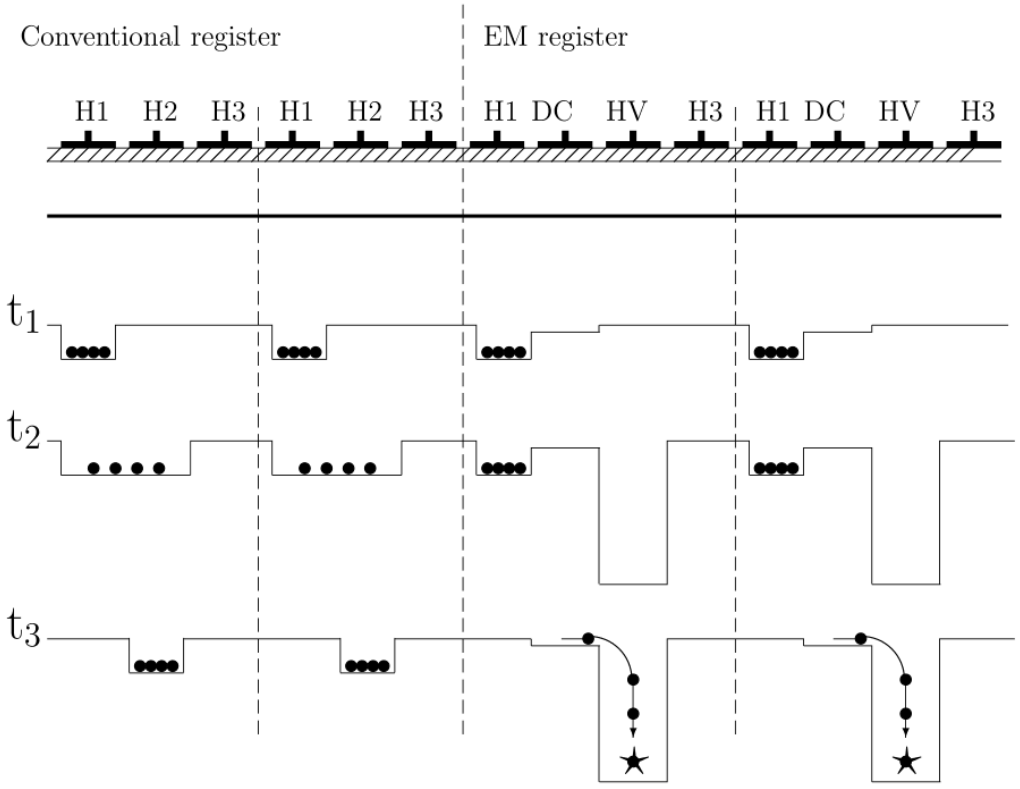

Figure 2.7: Schematic showing part of the serial clocking process for a conventional CCD and an EMCCD. Three steps during the process are shown as t_1, t_2, and t_3. The third stage shows the avalanche multiplication of photo-electrons into the deep well created by the serial phase labelled HV. Reproduced from Tulloch & Dhillon (2011).

externally, thus controlling the amplitude of electron multiplication). A schematic of the process of electron multiplication is shown in Figure 2.8.

Though the probability of a single electron exciting another by impact ionisation is only 1-2%, each pixel is clocked through a large number of stages (604 in the case of ULTRASPEC), producing a large electron-multiplying gain. Any single input electron can thus create a large number of additional electrons for each pixel by the time it reaches the end of the register, as each additional electron can also induce more impact-ionised electrons at each stage of the clocking. The EM-gain is thus also known as 'avalanche gain'. Avalanche gain is described by Equation 2.1.

$$g_A = (1+p)^n, \tag{2.1}$$

where g_A is the avalanche gain after passing through n stages of the high voltage register, each of which has probability p of producing one additional electron from an input electron. The electron multiplication reduces the readout noise by a factor of $1/g_A$, thus with an avalanche gain of > 100 the read-out noise is essentially negligible. This is very useful for read-out noise limited observations. However, the stochastic nature of the avalanche process acts to double the variance in the signal, essentially reducing the quantum efficiency of the detector by a factor of two through a process known as multiplication noise (Tulloch & Dhillon, 2011). Therefore the electron multiplying function actually *reduces* the signal-to-noise ratio in observations which are limited by Poisson noise in the source or sky background.

Because of this, the e2v CCD201-20 used in ULTRASPEC also has a normal output amplifier, as well as the electron-multiplying output, as shown in Figure 2.5. Operation can be switched between the outputs in software, so that a wide range of scientific targets can be observed with maximum signal-to-noise ratios. In practice, we use the normal output most often, and only employ the avalanche amplifier when studying very faint objects at the highest speeds.

As presented in Tulloch & Dhillon (2011), the signal-to-noise ratio (SNR) of a conventional detector can be described by Equation 2.2, whilst for an EMCCD the SNR is given by Equation 2.3.

$$SNR_C = \frac{S}{\sqrt{S + C + D + K + \sigma_R^2}}, \tag{2.2}$$

$$SNR_{EM} = \frac{S}{\sqrt{2 \times (S + C + D + K) + (\sigma_R/g_A)^2}}, \tag{2.3}$$

where S is the mean signal in photo-electrons per pixel, C is the clock-induced charge

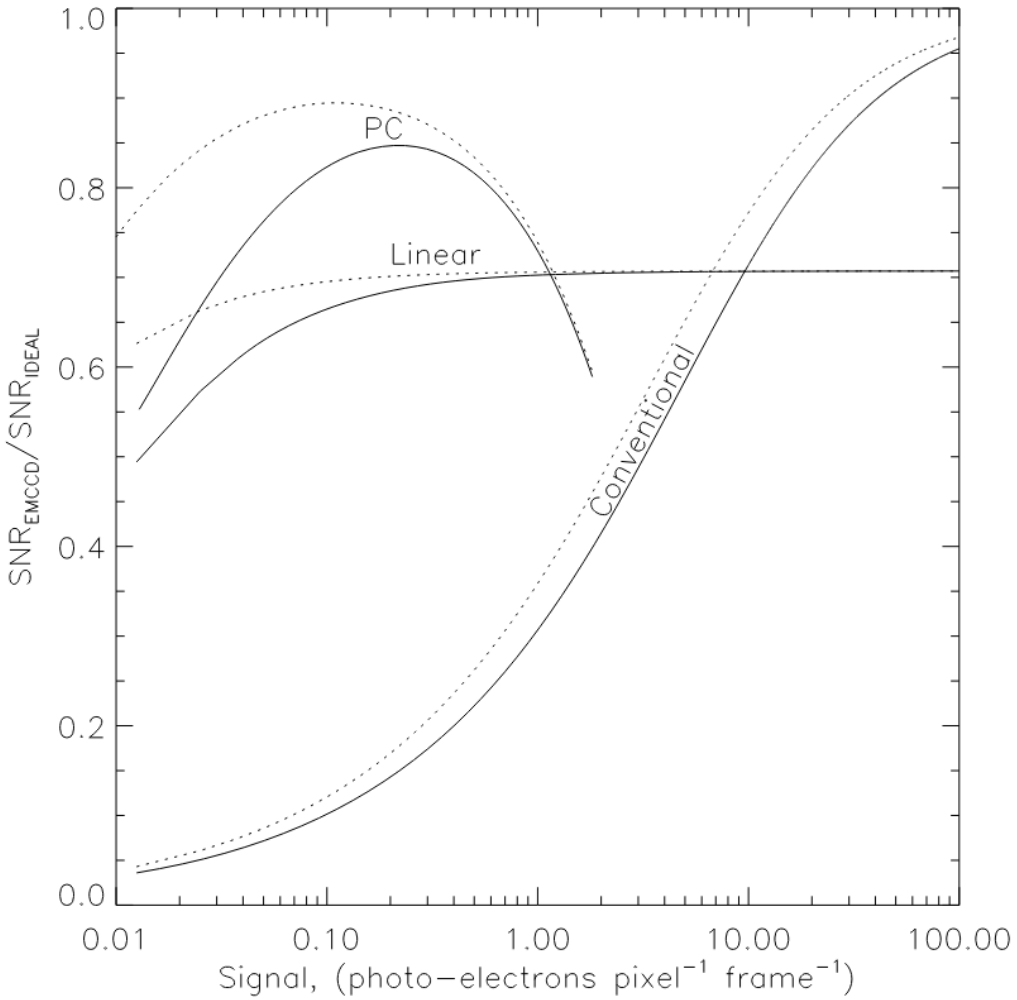

Figure 2.8: Signal-to-noise ratio as a function of input signal, for the different regimes of an EMCCD. PC stands for photon counting, and describes a scenario in which the expected flux from the source is less than one photon per exposure. This essentially avoids multiplication noise, which affects the linear mode. At high incident signal, the conventional (non-electron-multiplying) mode gives the highest signal-to-noise ratio. Reproduced directly from Tulloch & Dhillon (2011).

in electrons (see Section 2.3.3), D is the dark current in electrons (see Section 2.3.4), K is the signal from the sky background in photo-electrons, and σ_R is the read-out noise in electrons. The factor of 2 in Equation 2.3 accounts for multiplication noise.

2.2.5 Data Acquisition System

The data acquisition system of a high-speed camera must be able to cope with the highest frame rates that the detector can be operated at. The ULTRASPEC data acquisition system was especially designed to process frames at rates approaching 200 Hz (Ives et al., 2008).

ULTRASPEC uses a San Diego State University (SDSU) controller built by Astronomical Research Cameras Inc., and is operated by a dedicated Linux PC. The SDSU controller uses digital signal processors to encode observing instructions for the CCD. The signal processors receive XML packets with observing setup requests from the user. In its current setup, ULTRASPEC astronomers use a Python graphical user interface to easily create the necessary XML packets for observations.

2.2.6 GPS Timestamping

In order to conduct useful observations at high speed, extremely accurate time-stamping ($\leq$1ms) of individual frames is necessary. For this reason, ULTRASPEC uses a GPS antenna to collect universal time data from multiple satellites, which are used to regularly synchronise the system clock on the control PC.

Each time an image is started, a dedicated core on the control PC immediately supplies a timestamp and enters it into a first-in-first-out (FIFO) buffer. When the image is being processed by the data handling software, the timestamp is pulled from the FIFO buffer and assigned to the image. This ensures the images and timestamps are always synchronised.

The ULTRASPEC timestamping procedure has been tested with a pulsed LED device, and was found to have an absolute accuracy to within 1 ms. The relative accuracy of the image timestamps with respect to each other will be much better than this.

2.2.7 Data Reduction Pipeline

ULTRASPEC can generate up to 0.6 MB/second, and thus large amounts of data in any single night of observing. The data from each run of exposures is comprised of two files: an XML file describing the instrumental setup and observation details, and a single binary file containing all the images in the run. The ULTRACAM data reduction pipeline (Dhillon et al., 2007) is used to process the data. During observations, the pipeline allows the user to conduct a quick-look analysis of the incoming images, producing a real-time light curve which helps the observers ensure the observations are running smoothly. Alternatively, after the observing is complete, the pipeline can have every single data reduction parameter tweaked in order to conduct the best reduction.

The pipeline operates using an input parameter file, which the user can configure to their needs. The input file includes decisions on the type of extraction requested, for example: fixed vs variable aperture photometry, Gaussian or Moffat profile fitting, the stability of source positions from one image to the next, the size of the aperture for extraction, and the associated calibration frames such as biases and flats. Prior to running the reduction, the user would need to create a file describing the locations on the images of the sources that they wish to investigate - placing apertures around them. The aperture file is read by the reduction process and used to search for the requested sources in each image.

In aperture photometry, 3 differently sized apertures are used. The inner aperture selects the pixels on the source for extraction, and all counts within that aperture are summed to calculate the source flux. The sky counts are measured as the clipped mean of counts in the annulus between the second and third apertures, and this average sky flux is then subtracted from the source flux (after multiplying by the number of pixels inside the first aperture). It is therefore important that these apertures have the correct sizes, such that no source flux is being collected and measured in the outer annulus,

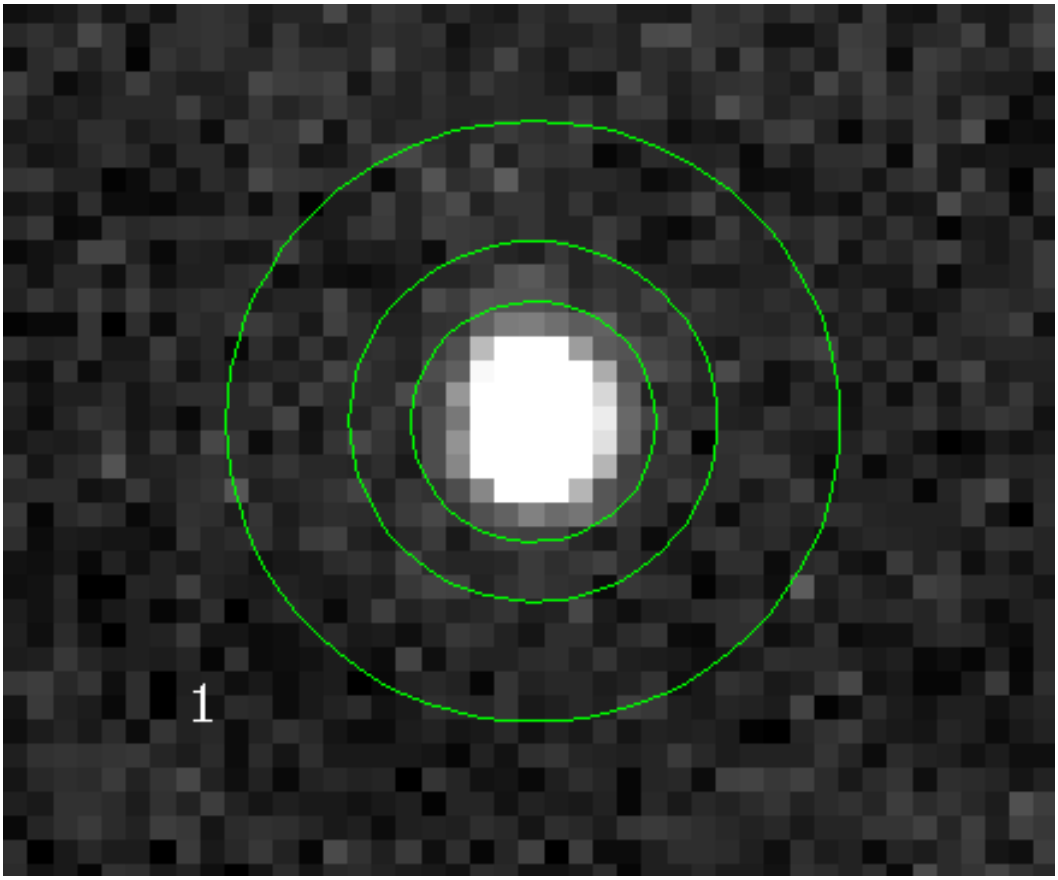

Figure 2.9: Example apertures created around a star by the ULTRACAM data reduction pipeline. The flux inside the inner annulus is extracted as the source flux, and the average sky flux per pixel measured between the middle and outer annuli gives a measure of the sky background.

and not too much sky flux is being collected in the inner aperture, thus adding extra noise. Figure 2.9 shows an example setup, with an inner annulus around the star, and two larger annuli which dictate which pixels are used to estimate the sky background. Usually during the full reduction the apertures can be allowed to vary in size with respect to the changing seeing, although the different apertures around each source must always be the same size as each other for any one image. For a more detailed description of the ULTRACAM data reduction pipeline, see Tom Marsh's webpages at `http://deneb.astro.warwick.ac.uk/phsaap/software/ultracam/html/`.

2.3 Lab-based Detector Tests

The CCD controller software was modified to optimise performance in the lab before shipping to Thailand. In order to provide a standard reference of typical performance which users can expect to obtain at the telescope, a number of quality control tests were executed. This enables observers to choose the mode of operation best suited for their science, and also helps identify any problems.

All the tests in this section were completed by me in Spring 2013 at the University of Sheffield.

2.3.1 Bias levels

On read-out, all CCD frames contain a pedestal level of counts, known as the bias level. The bias is a large positive offset applied to the read-out capacitor in order to ensure the measured values are always positive. Most CCDs employ 16-bit data processing, and without the large positive bias, there would be a high possibility of negative counts being measured due to random read noise. To process a negative count, one of the 16 bits used in data processing would need to be a sign bit (+ or −). With a large positive offset, CCD processing can ignore the need for a sign bit, and use all 16 bits to store data and maximise the dynamic range of the CCD (Howell, 2006).

The measured bias level of a CCD can fluctuate by small amounts over time as electronic components change with temperature and age, but in general the level is stable for a particular read-out mode. Any drastic changes in bias level usually indicate severe problems within the CCD electronics or read-out process. This emphasises the importance of monitoring the bias levels in all the read-out modes, so that bias frames taken at the beginning of each night of observations can be compared with the reference values, and any discrepancies addressed before data collection begins.

We collected multiple bias frames in each read-out mode, measured the mean counts

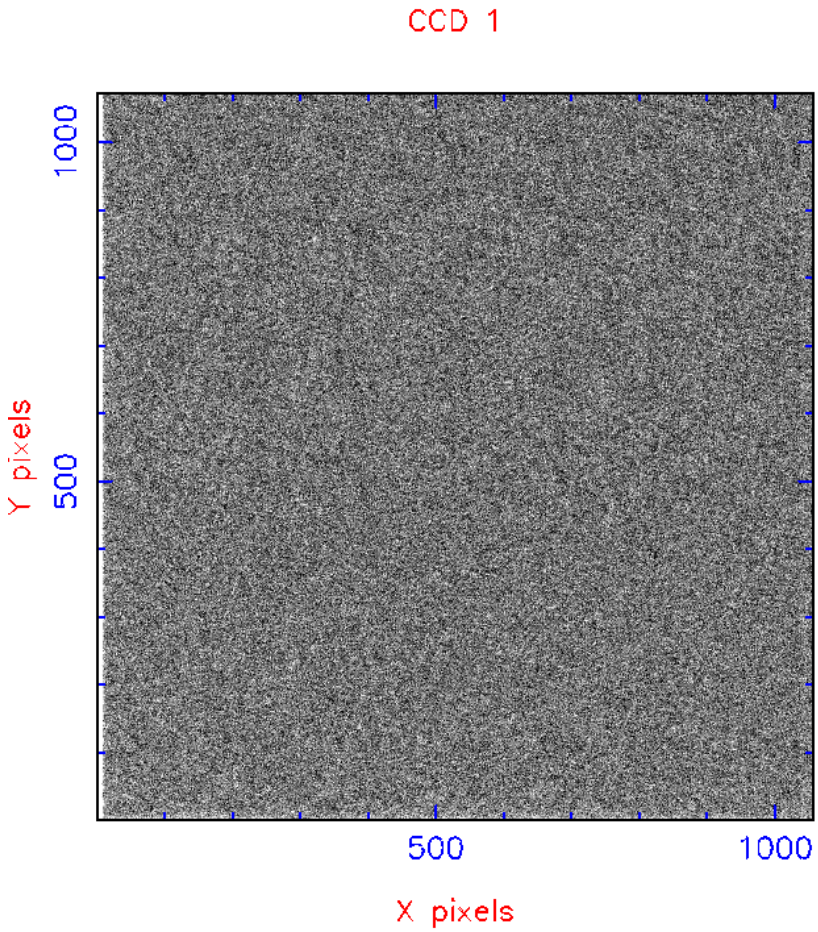

Figure 2.10: Example ULTRASPEC bias frame, made from the combination of 30 individual frames taken in normal read-out mode and slow read-speed.

in a small region of each frame and took the bias level to be the median of all the individual frames (to remove the effect of any cosmic rays). Bias frames were taken with zero added exposure delay in Clear mode (see Section 2.2.3), meaning the true exposure times were zero seconds. We ensured there was no light incident on the chip by placing a metal disk over the CCD window which sits flush with the cryostat, and does not allow any light to pass through. At least 20 frames were taken for each read-out mode, and the mean level was measured within a 100×100 pixel region near the centre of each frame. The measured bias levels for each mode of operation are given in Table 2.2. Figure 2.10 shows an example of a combined bias frame, made up of 30 individual frames. Sometimes

Table 2.2: Mean bias level counts for each mode of ULTRASPEC operation, measured on 2013-05-01.

Output Amplifier	Mode	Binning	Read Speed		
			Slow	Medium	Fast
Normal	Clear	1x1	833	976	1089
Normal	No Clear	1x1	833	974	1092
Normal	Clear	2x2	811	920	845
Normal	Clear	3x3	822	932	863
Normal	Drift	1x1	832	973	1085
Normal	Drift	2x2	810	918	843
Normal	Drift	3x3	821	932	860
Avalanche	Clear	1x1	4396	4833	3695
Avalanche	No Clear	1x1	4394	4841	3695
Avalanche	Clear	2x2	4433	4822	3673
Avalanche	Clear	3x3	4447	4837	3677
Avalanche	Drift	1x1	4393	4834	3695
Avalanche	Drift	2x2	4433	4821	3673
Avalanche	Drift	3x3	4447	4837	3677

bias frames can indicate the presence of pick-up noise due to electrical interference, by showing structure or gradients. There is no sign of structure in Figure 2.10, suggesting that pick-up noise is not present.

It can be seen in the table that the bias levels are largely unaffected by changing the mode of operation (i.e. Clear, No Clear, Drift), which is to be expected, but that the levels do change slightly with different binning factors. This is because bias levels are set by the reset drain voltage[1], which is slightly different for each binning mode so that a maximum dynamic range can be accessed.

2.3.2 Read-out noise and gain

Other parameters of the CCD and read-out configuration which should be checked regularly for discrepancies are the read-out noise and the system gain. Read-out noise is produced in the output amplifier by two key components. One source is the thermal

[1]The reset drain voltage sets the level to which the read-out capacitor should return after each packet of charge is read out.

noise of electrons in the capacitor which transforms the pixel charge into a voltage, and another is thermally-induced variable flicker noise of the electrons in the transistor (Tulloch, 2010). The CCD controller also adds some noise. The thermal noise component scales with temperature, and inversely with bandwidth, so cooling and reading out slowly can reduce this noise contribution. Conversely, the flicker noise component increases with decreasing read time, so there is a trade-off here as to how slowly one should read out the detector. In high-time-resolution astronomy, detectors must be read out relatively quickly, so read-out noise can become a significant factor, especially with short exposure times and thus low count levels. Typical CCDs can have their RMS read-out noise reduced to around 3 electrons, when read out slowly.

The gain of the system describes how many photo-electrons are represented by one analogue-to-digital unit, or ADU, and thus has units of e^-/ADU. The system gain can be controlled manually by adjusting the numerical gain of the video processor which produces the raw output for each image.

The choice of system gain must be made by balancing the effects of two opposing factors, the quantisation noise and the dynamic range. On the one hand, a small gain is beneficial because it reduces the contribution from quantisation noise. Quantisation noise is introduced by the digitising process, and will be significant if the read-out noise in electrons is lower than the gain. For example, if the read-out noise is 2 e^-, and the gain is 4 e^-/ADU, then the digitisation will not adequately sample the read-out noise, and could measure 0 or 1 ADU. This quantisation noise is added in quadrature. The equation describing quantisation noise, σ_q, is given in Equation 2.4.

$$\sigma_q = \frac{N_{fwc}}{2^n \times 12^{0.5}}, \tag{2.4}$$

where N_{fwc} is the full-well capacity of each pixel (usually around 100,000–200,000 for most CCDs, but for ULTRASPEC this is 80,000 e^-), and n is the number of bits in the

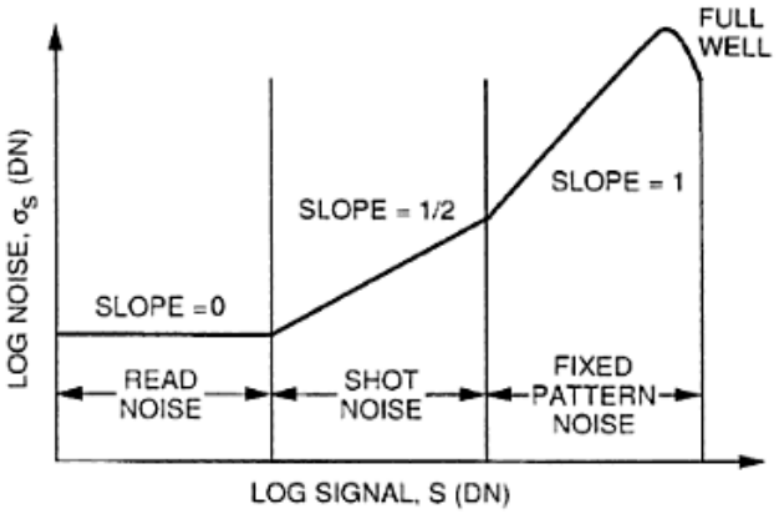

Figure 2.11: Theoretical photon transfer curve reproduced from page 102 of Janesick (2001). The noise in ADU (denoted DN in this example) is plotted against the signal in ADU, on logarithmic axes.

ADC.

On the other hand, a gain that is too small would make inefficient use of the full-well capacity. For example, if the full-well capacity of each pixel is around 150,000 e^-, and the gain is 0.8 e^-/ADU, then the ADC saturating at 65535 (maximum capacity if the ADC has 16 bits) would mean a maximum measurement of $65535 \times 0.8 = 52428$ e^-, only one third of the actual full-well capacity.

The system gain for ULTRASPEC is set to 0.8 e^-/ADU, and the quantisation noise (according to Equation 2.4) is 0.35 e^-. For a deeper discussion of read-out noise and gain see Tulloch (2010).

Read-out noise and system gain can both be measured using the photon transfer curve (PTC) technique. PTCs are created using a number of flat field frames with differing average count levels, from zero counts (at the bias level) all the way up to saturation. For each frame the RMS noise in ADU is plotted against the average counts for a small region within the image, once the bias level has been subtracted. A labelled theoretical PTC is shown in Figure 2.11, whilst an example of the resulting PTC for ULTRASPEC

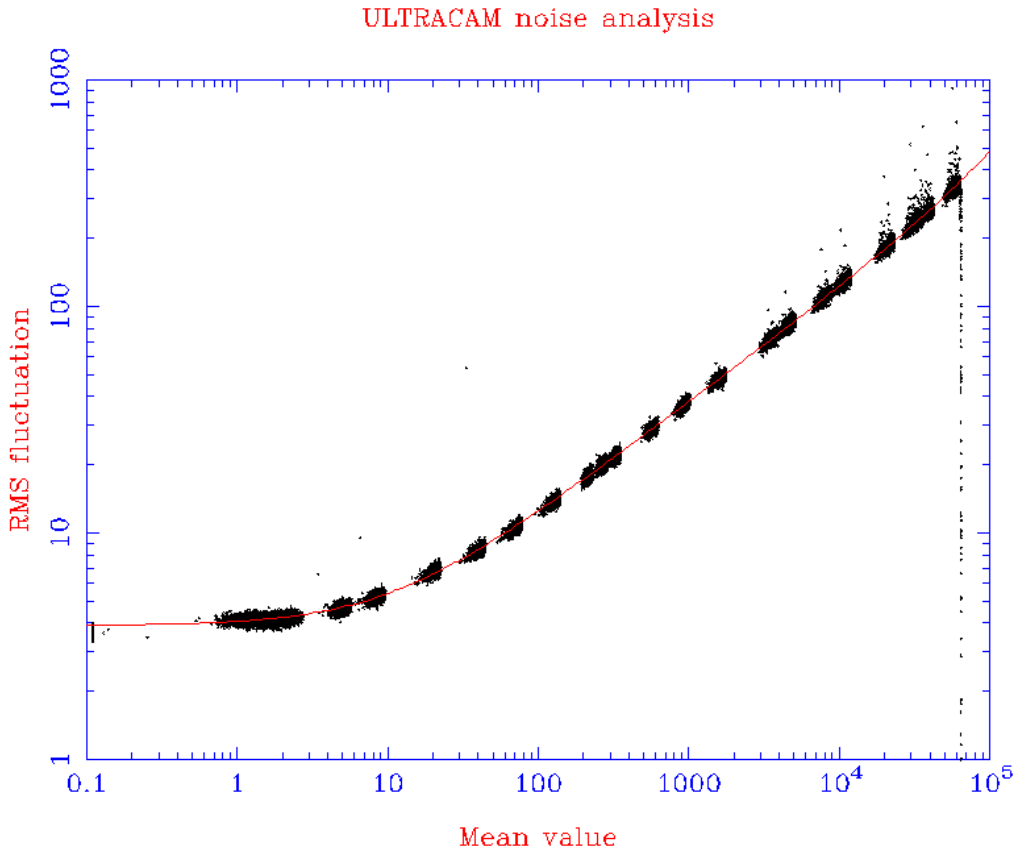

Figure 2.12: Photon transfer curve for the normal output of ULTRASPEC in medium-speed read-out. Three distinct parts are visible in the curve: the flat part on the left indicating the read-out noise, the linear section in the middle which is used to measure the system gain, and the steepening curve at the top right, the shape of which depends on the flat-field noise and shows the saturation level on the extreme right.

is shown in Figure 2.12.

The first section of the PTC describes the read-out noise, and has zero gradient because read-out noise is entirely independent of flux. The linear section in the middle of the plot is dominated by shot noise, and will intersect the x-axis at a value equal to the system gain. From Janesick (2001), the gain g follows from $\log g = \log S - \log\left[\sigma_s^2 - \sigma_r^2\right]$, but during the shot-noise dominated section of the PTC the read-out noise is negligible, which removes the σ_r term. If we set $\sigma_s = 0$, i.e. on the x-axis, the gain is equal to the signal, S.

PTCs also allow for an estimation of the flat-field noise (also known as the fixed

pattern noise), which is due to the slightly different physical sizes of each pixel. This can be estimated from the top section of the curve. Finally, the curve turns over and drops rapidly once the full-well limit is reached, and the CCD becomes saturated (Janesick, 2001).

We produced PTCs for all read-out modes using flat fields taken in the lab, and fitted the curves using the *ncal* routine within the ULTRACAM data reduction pipeline. Tables 2.3, 2.4 and 2.5 show the read-out noise, system gain and flat-field noise estimates found by the PTCs. The values labelled "avalanche" were measured using the avalanche output amplifier, but with the high-voltage gain set to unity, meaning no electron multiplication is expected and one electron entering the serial register should equate to one electron coming out. This enables us to measure the electronic gain of the avalanche output, which can then be used to measure the avalanche gain (see Section 2.3.3).

The system gain is deliberately set higher for avalanche mode (see Table 2.4), since the initial read-out noise in e^- is higher in avalanche mode. This allows a greater dynamic range to be accessed.

As expected, the system gain stays constant for different modes using the same output amplifier, as this is a property solely of the video processor in the SDSU controller and should not be affected by binning or windowing. The read-out noise, on the other hand, does vary slightly, and also has significantly different values in each of the read-out speeds. This is of course expected, as reducing the time taken to read out an image increases the read-out noise. The slight increase in read noise when increasing the binning factor is not surprising either, because when binning the clocking process takes longer, increasing the time between measuring the signal and resetting the voltage level. This gives the voltage more time oscillate and vary, introducing slightly more read-out noise.

Some of the test modes (mainly medium read-out speed with the avalanche output amplifier) suffered from significant pick-up noise in this instance. This can sometimes occur when data cables and power cables sit too close together, and electronic interference

Table 2.3: Read-out noise in e^- for each mode of ULTRASPEC operation as measured by photon transfer curves on 2013-05-01. *Significant pick-up noise was visible in these frames.

Output Amplifier	**Mode**	**Binning**	**Read Speed**		
			Slow	**Medium**	**Fast**
Normal	Clear	1x1	2.3	2.7	4.3
Normal	No Clear	1x1	2.3	2.9	4.5
Normal	Clear	2x2	2.4	3.2	4.6
Normal	Clear	3x3	3.1	2.9	4.6
Normal	Drift	1x1	2.6	2.7	4.0
Normal	Drift	2x2	2.8	3.1	4.3
Normal	Drift	3x3	2.9	3.0	4.2
Avalanche	Clear	1x1	5.5	9.4*	14.4
Avalanche	No Clear	1x1	5.9	11.0*	14.8
Avalanche	Clear	2x2	5.5	8.2	13.2
Avalanche	Clear	3x3	6.1	8.5	15.2
Avalanche	Drift	1x1	5.5	7.4	15.2
Avalanche	Drift	2x2	5.5	8.3	13.6
Avalanche	Drift	3x3	5.7	8.0	15.6

Table 2.4: ULTRASPEC system gain in e^-/ADU measured using photon transfer curves on 2013-05-01.

Output Amplifier	**Mode**	**Binning**	**Read Speed**		
			Slow	**Medium**	**Fast**
Normal	Clear	1x1	0.8	0.7	0.8
Normal	No Clear	1x1	0.8	0.7	0.8
Normal	Clear	2x2	0.8	0.7	0.8
Normal	Clear	3x3	0.8	0.7	0.8
Normal	Drift	1x1	0.8	0.7	0.8
Normal	Drift	2x2	0.8	0.7	0.8
Normal	Drift	3x3	0.8	0.7	0.8
Avalanche	Clear	1x1	1.9	1.6	4.0
Avalanche	No Clear	1x1	1.9	1.6	4.0
Avalanche	Clear	2x2	1.9	1.6	4.0
Avalanche	Clear	3x3	1.9	1.6	4.0
Avalanche	Drift	1x1	1.9	1.6	4.0
Avalanche	Drift	2x2	1.9	1.6	4.0
Avalanche	Drift	3x3	1.9	1.6	4.0

Table 2.5: Percentage flat-field noise of ULTRASPEC measured using photon transfer curves on 2013-05-01. As expected, the flat-field noise does not change significantly with read-out speed or mode, since the noise is generated by the physical differences in pixel size.

Output Amplifier	**Mode**	**Binning**	**Read Speed**		
			Slow	**Medium**	**Fast**
Normal	Clear	1x1	0.3	0.3	0.3
Normal	No Clear	1x1	0.3	0.3	0.3
Normal	Clear	2x2	0.2	0.2	0.2
Normal	Clear	3x3	0.2	0.2	0.2
Normal	Drift	1x1	0.3	0.3	0.3
Normal	Drift	2x2	0.2	0.2	0.2
Normal	Drift	3x3	0.2	0.2	0.2
Avalanche	Clear	1x1	0.3	0.3	0.3
Avalanche	No Clear	1x1	0.3	0.3	0.3
Avalanche	Clear	2x2	0.2	0.2	0.2
Avalanche	Clear	3x3	0.2	0.2	0.1
Avalanche	Drift	1x1	0.2	0.3	0.2
Avalanche	Drift	2x2	0.1	0.2	0.1
Avalanche	Drift	3x3	0.1	0.1	0.1

produces visible structure in the bias frames. An example of this is shown in Figure 2.13. The structure of this noise is variable, and its presence is often temporary. It can sometimes be removed by rearranging the cabling between the detector and the rack PC. During lab testing however, we rarely moved these cables. The appearance of pick-up noise in certain read-out modes is therefore connected to the speed of the voltage swings during read-out, and the resonance between this and the alternating current in the power cables. As discussed later in Section 2.4.9, pick-up noise continued to be troublesome even after commissioning.

2.3.3 Avalanche gain and clock-induced charge

As described in Section 2.2.4, the EM output of the ULTRASPEC CCD can be used in low light-level situations to amplify photo-electrons and virtually remove the effects of read-out noise. The amplitude of avalanche gain (the multiplication factor achieved

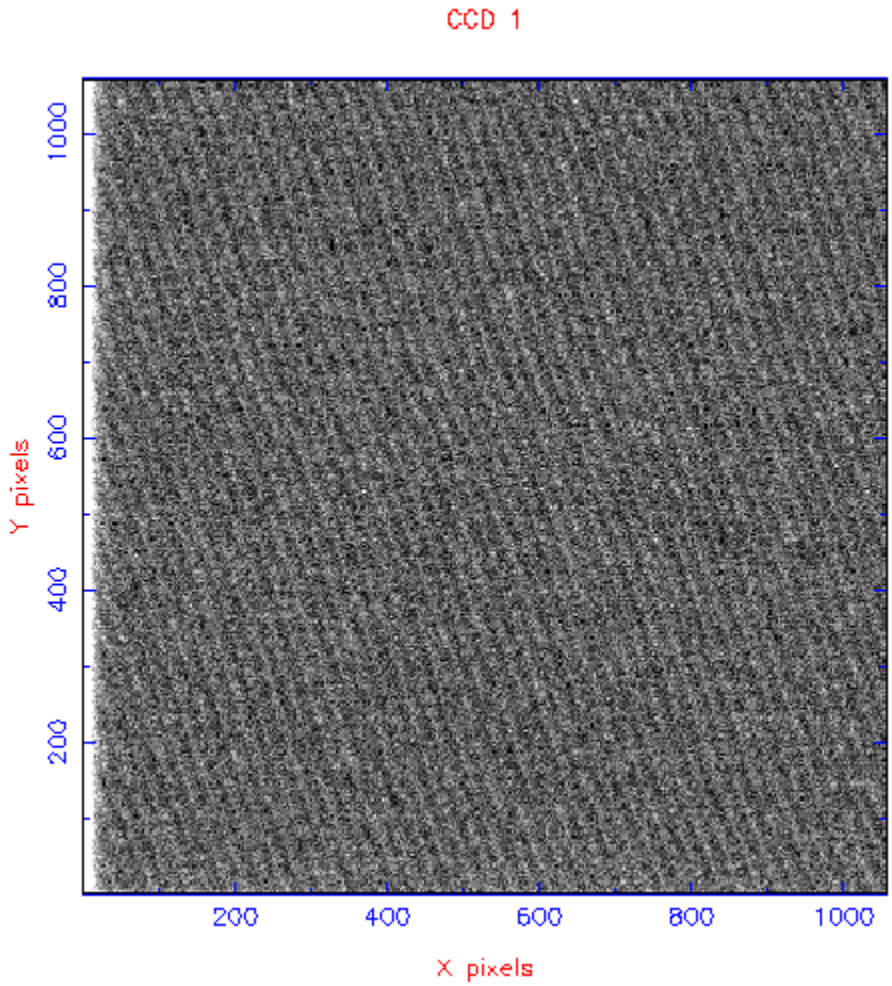

Figure 2.13: Example ULTRASPEC bias frame with medium read-out speed, showing linear structure. This is suspected to be pick-up noise generated because the data cables and power cables are located too close to each other and/or are improperly shielded, causing interference.

when passing through the EM register) is somewhat volatile, and depends heavily on a number of factors including the voltages applied and the temperature of system. As temperature drops, so too does the resistance, allowing the electrons to move faster. However, the lower temperature also increases the excitation bandgap, as the thermal energy of the lattice electrons decreases. Overall the drop in resistance outweighs the increase in bandgap, and the probability of impact ionisation is increased with decreasing temperature (Tulloch, 2010).

It is also thought that the avalanche gain may degrade gradually with time (Tulloch,

2004a). Due to statistical fluctuations, the exact value is impossible to predict directly from hardware specifications, but it can be measured accurately, and overall EMCCDs are quite stable (Mackay et al., 2001). The avalanche gain, g_A, is the ratio of the electronic gain measured through the EM amplifier when using a low voltage (i.e. no avalanche effect), g_{S0}, to the overall system gain, g_S, as shown by Equation 2.5.

$$g_A = \frac{g_{S0}}{g_S} \tag{2.5}$$

Thus the avalanche gain (g_A) is a unit-less multiplication factor which is the average number of electrons output by the EM register for every electron which enters.

Clock-induced charge (CIC, also known as spurious charge) is an additional source of noise, created by the process of impact ionisation. Just like in the EM register, the clocking of photo-electrons from pixel to pixel across the chip can sometimes, though rarely, cause excitation of additional electrons from the lattice. Since it is completely unrelated to photon excitation, CIC is often also considered a form of dark current. However, unlike thermally-excited dark current, the amount of clock-induced charge is independent of exposure time. Once excited, these electrons are clocked out and measured through the output amplifier, making them indistinguishable from photo-electrons (Tulloch, 2010). The probability of generating clock-induced charges is very small. Thus the noise contribution from CIC is negligible in conventional CCDs because it is usually dwarfed by the read noise. However, in EMCCDs where the read noise is essentially zero, CIC becomes important, especially as any clock-induced charges will be amplified in the EM register. Factors which affect the amplitude of clock-induced charge are the temperature, the speed of clocking and the amplitude of the clocking voltages (e2v, 2004; Tulloch & Dhillon, 2011).

To measure the avalanche gain and clock-induced charge, a large number of bias frames were taken in avalanche mode with the high-voltage gain turned on. In these

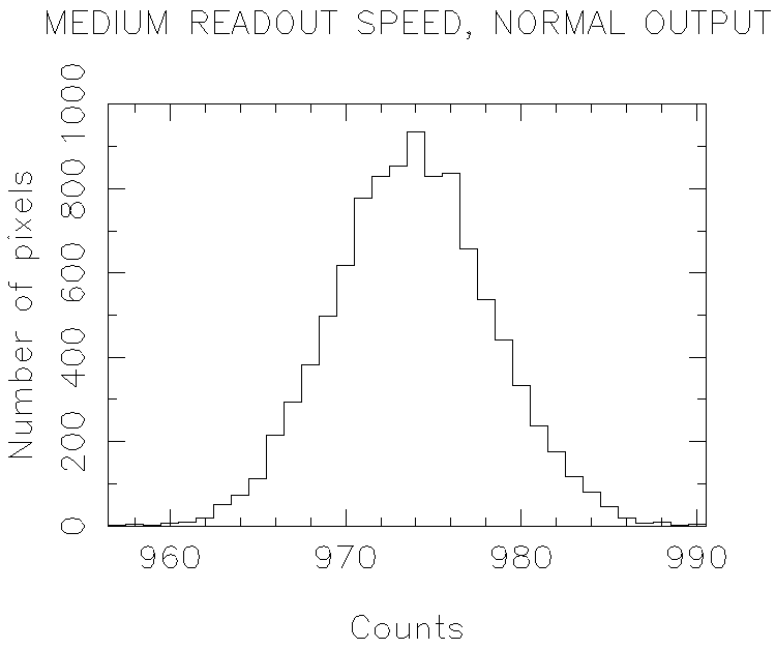

Figure 2.14: Bias frame histogram plotted using normal-output bias frames taken with medium read-out speed. The histogram follows a normal distribution as expected, and the standard deviation of this plot gives an estimate of the read-out noise.

frames, even though no light was incident on the CCD, there is a source of electrons entering the EM output from the clock-induced charge. Plotting a histogram of the resulting distribution of counts from all frames enables the avalanche gain and CIC to be measured. A histogram of an equivalent list of bias frames read out through the normal output would show a bell curve, as seen in Figure 2.14. However, when the EM output is used, two significant artefacts appear in the histogram (see Figures 2.15 & 2.16). The first is a bump on the high-count shoulder of the bell curve, whose size and shape are dependent on the clock-induced charge within the EM register itself (thus not experiencing the full avalanche multiplication of the entire register). Secondly a long, straight tail can be seen extending towards the high count region. This is due to the avalanche multiplication of clock-induced charge produced within the CCD.

Since each pixel is highly unlikely to contain more than one clock-induced electron, the statistical approach of Tubbs (2003) can be used to estimate the expected distribution of electrons excited in the EM-register by impact ionisation (Tulloch, 2010). This is

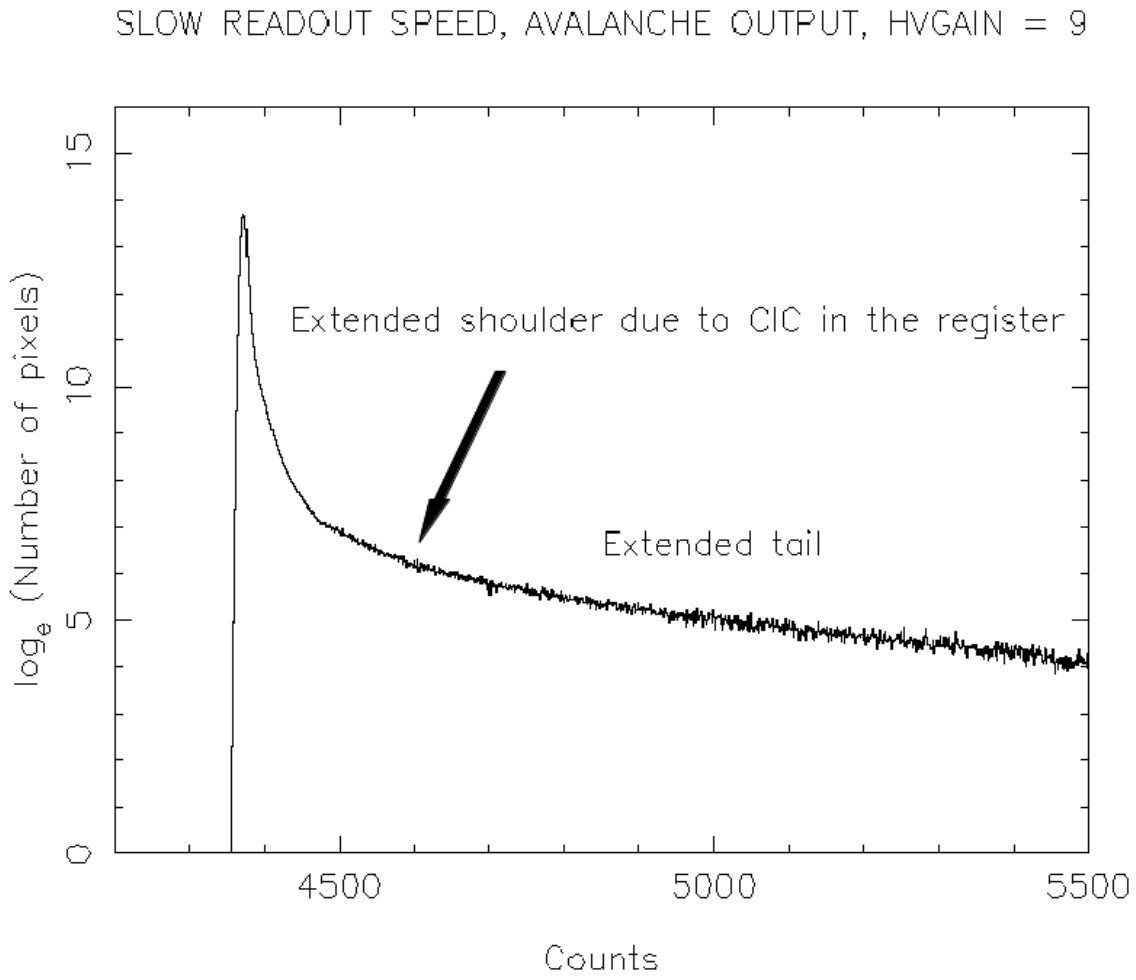

Figure 2.15: Example histogram of multiple slow bias frames taken using the avalanche output. The peak marks the CCD bias level. The extended shoulder due to CIC in the register is labelled, although the feature itself is not easily visible in this instance. Also labelled is the long tail extending to high counts caused by the multiplication of dark current and clock-induced charge in the CCD.

approximated by Equation 2.6 in the case of $g_A \gg 1$.

$$P(n) = \frac{1}{g_A} \exp(-\frac{n}{g_A}), \tag{2.6}$$

where $P(n)$ is the probability of measuring n electrons after a single electron enters the EM register. Taking natural logarithms of both sides of Equation 2.6 gives us Equation

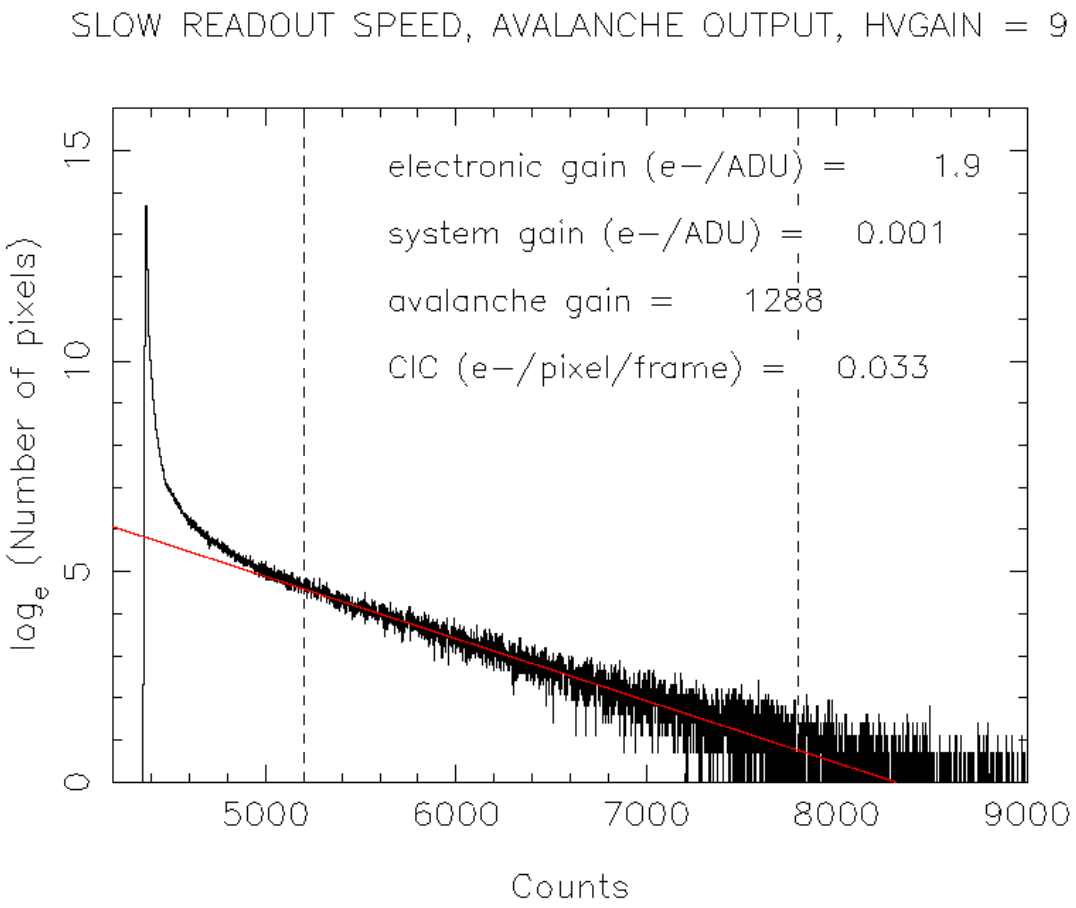

Figure 2.16: The avalanche gain and CIC can be estimated by selecting the linear region of the histogram and fitting a straight line to it. Note the log scale of the y-axis.

2.7, which shows how the gradient of the line in the histogram in Figure 2.15 will be $-1/g_A$, allowing the avalanche gain to be estimated (Tulloch, 2004b, 2010). This is demonstrated in Figure 2.16 for ULTRASPEC data.

$$\ln\left[P(n)\right] = \ln\left[\frac{1}{g_A}\right] - \frac{n}{g_A}. \tag{2.7}$$

The same bias frames were used to measure the CIC, by subtracting the median from the mean as measured in a small region in the middle of the image. The region must have a flat bias profile with no structure, otherwise even a small gradient can disrupt the CIC measurement (Tulloch, 2010).

Table 2.6: System gain values (g_S) measured in avalanche mode on 2013-04-30 in e^-/ADU. *Values in the bottom two rows should be treated with caution, as they incur significant uncertainties due to the small sample size of pixels. Small windows are needed in order for drift mode to function effectively.

Mode	**Binning**	**Read Speed**		
		Slow	**Medium**	**Fast**
Clear	1x1	0.002	0.001	0.004
No Clear	1x1	0.001	0.001	0.003
Clear	2x2	0.002	0.001	0.003
Clear	3x3	0.001	0.001	0.003
Drift	1x1	0.002	0.001	0.004
Drift	2x2	0.002*	0.001*	0.003*
Drift	3x3	0.001*	0.001*	0.002*

Table 2.7: Avalanche gain values (g_A) measured on 2013-04-30. *Values in the bottom two rows should be treated with caution, as they incur significant uncertainties due to the small sample size of pixels. Small windows are needed in order for drift mode to function effectively.

Mode	**Binning**	**Read Speed**		
		Slow	**Medium**	**Fast**
Clear	1x1	1165	1109	1116
No Clear	1x1	1332	1263	1267
Clear	2x2	1266	1241	1238
Clear	3x3	1304	1190	1181
Drift	1x1	1202	1164	1127
Drift	2x2	1148*	1067*	1158*
Drift	3x3	1297*	1590*	2495*

Tables 2.6, 2.7 & 2.8 show the values of system gain, avalanche gain and clock-induced charge for the different modes of operation using the avalanche amplifier output. The values are measured with a chip temperature of 160K.

2.3.4 Dark current

With the read-out noise being so drastically reduced in an EMCCD, other noise sources start to become significant. One such source which can have dire consequences if not

Table 2.8: Clock-induced charge in electrons per pixel per frame, measured on 2013-04-30. *Values in the bottom two rows should be treated with caution, as they incur significant uncertainties due to the small sample size of pixels. Small windows are needed in order for drift mode to function effectively.

Mode	**Binning**	**Read Speed**		
		Slow	**Medium**	**Fast**
Clear	1x1	0.038	0.021	0.018
No Clear	1x1	0.068	0.044	0.027
Clear	2x2	0.051	0.042	0.035
Clear	3x3	0.068	0.062	0.052
Drift	1x1	0.043	0.029	0.021
Drift	2x2	0.059*	0.049*	0.035*
Drift	3x3	0.074*	0.050*	0.028*

minimised accordingly, is dark current. Dark current occurs when electrons are thermally excited from the silicon lattice, instead of being excited by the arrival of a photon. Dark current is read-out and measured in the same way as photo-electrons, and thus it appears in images in the same way and has the same Poisson/shot noise as true photon-induced flux.

As one would expect, dark current can be reduced by cooling the detector, because this reduces the thermal energy of the atoms in the silicon, and thus reduces the probability of them being thermally excited. It is for this reason that almost all astronomical detectors are cooled with liquid Nitrogen which boils at 77K.

Dark current can also be prominent even in appropriately cooled detectors for a short time after the system is powered on, or after saturation of the full well limit of the CCD. When the chip is saturated, many free electrons can be left partially trapped in the silicon lattice, and are released randomly over time, contributing to the higher dark current (Tulloch, 2010). After a system power-on (POn), a similar process occurs, trapping free electrons in the lattice, and dark current measurements in this instance are similar to those made after strong illumination (see Figure 3.12 of Tulloch 2010). Additionally, the dark current is not only in the first image after the POn, but continues

to affect dark current readings for tens of minutes, regardless of how long the system has been powered off for. This is thought to be due to charge becoming trapped in the lattice when the bias voltages swing off and on; a standard procedure whenever the system is initialised (David Atkinson, UKATC – private communication).

We investigated the behaviour of dark current after a POn, along with the dependence on temperature, by taking long exposures (10-25 minutes) without allowing any light to illuminate the chip. The cryostat window was blocked out with the metal disk as was done for taking bias frames. Dark exposures were taken immediately after powering on the system, and every 10 mins subsequently, and also taken with the chip at varying temperatures, after the dark current had settled down. The results can be seen in Tables 2.9 and 2.10. They are also shown graphically in Figures 2.17 and 2.18. It is important to know how the dark current behaves after a system POn, to ensure this noise source is minimised during astronomical operations. The fact that it can take up to 40 minutes after a power-on for the dark current to drop to normal operational levels means that users should avoid powering down and powering up the CCD as much as possible.

In addition, a quick test for dark current linearity was carried out. This involved taking dark exposures of varying lengths, once the dark current had settled down, at a temperature of 160K. Exposure times of 5, 10, 15, 20 and 25 minutes all showed consistent values of around 12 ADU/pix/hour.

2.3.5 Linearity

To assess the responsivity of the chip over a range of flux levels we conducted a linearity test using illumination from a stable lamp source as an artificial star. Exposures were taken in normal, slow read-out mode with steadily increasing exposure times, and we calculated the mean counts in a boxed aperture of 50×50 pixels near the centre of the CCD.

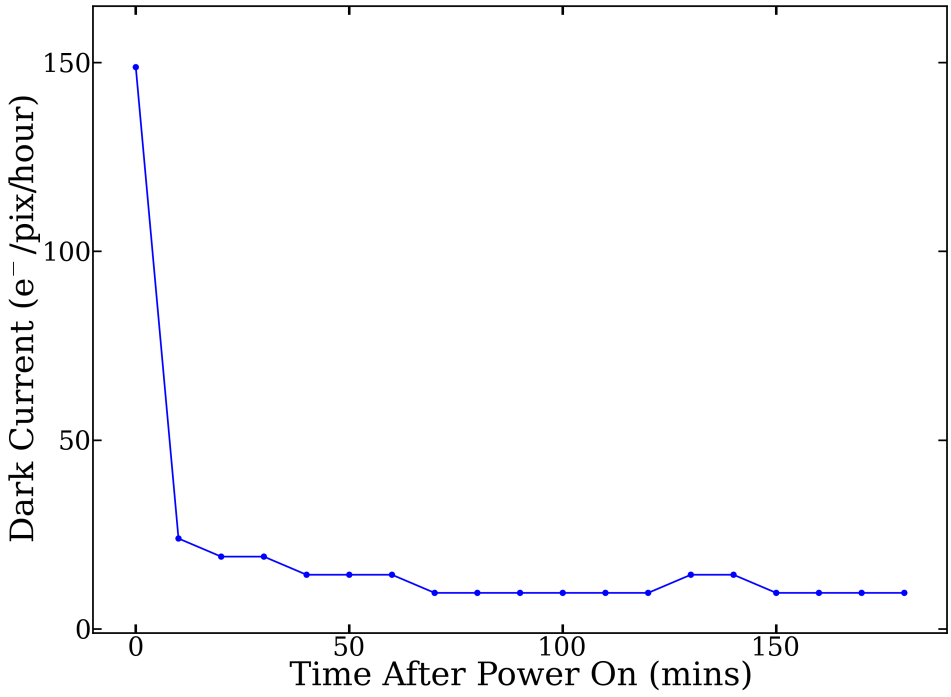

Figure 2.17: Dark current as a function of time after the CCD is powered on, at 160K.

Table 2.9: Dark current as a function of time after powering on the CCD. Measured using 10 minute exposures on 2013-06-07 at 160K.

Time Since POn (Mins, ±2)	**DC (e⁻/pix/hour)**
0	148
10	24
20	19
30	19
40	14
50	14
60	14
70	10
80	10
90	10
100	10
110	10
120	10
180	10

Table 2.10: Dark current as a function of CCD temperature, measured using the mean value of three 15-minute dark frames obtained on 2013-06-08.

Chip Temp (K, ±0.2)	**DC (e^-/pix/hour)**
155	3.2
160	13.8
165	38
170	106
175	288
180	659
185	1389
190	2680

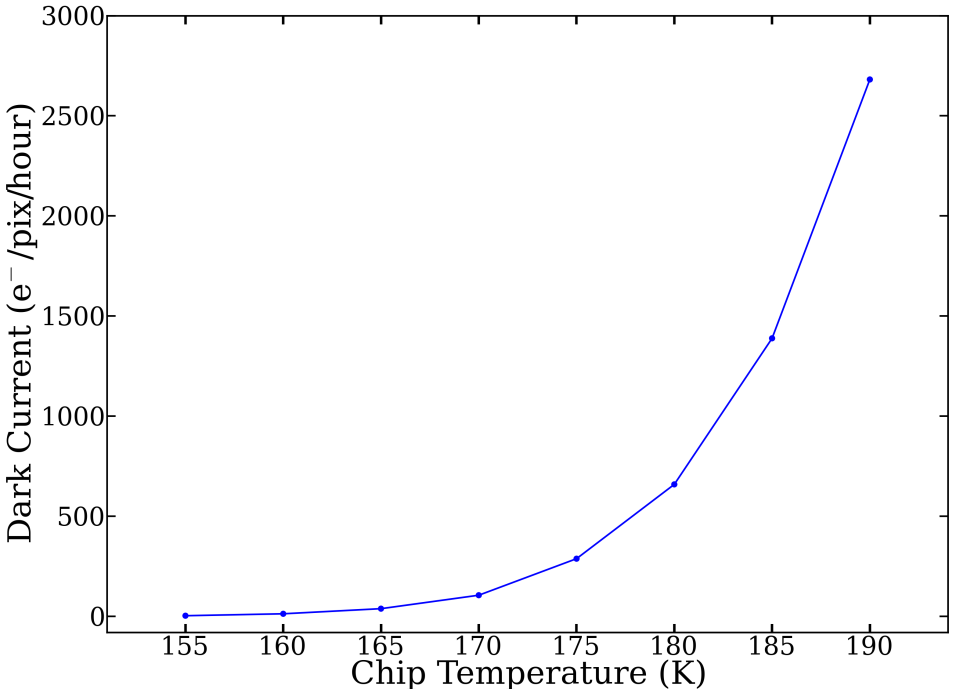

Figure 2.18: Dark current as a function of CCD temperature.

The CCD was found to be linear over the vast majority of its dynamic range (less than 1% variation). At very short exposure times (less than 0.5 seconds), the linearity is affected by the effect of frame transfer, and at high illumination levels (approaching saturation), the response deviates also. Plots showing the flux as a function of time, and the residuals to a straight-line fit, are shown in Figures 2.19 and 2.20.

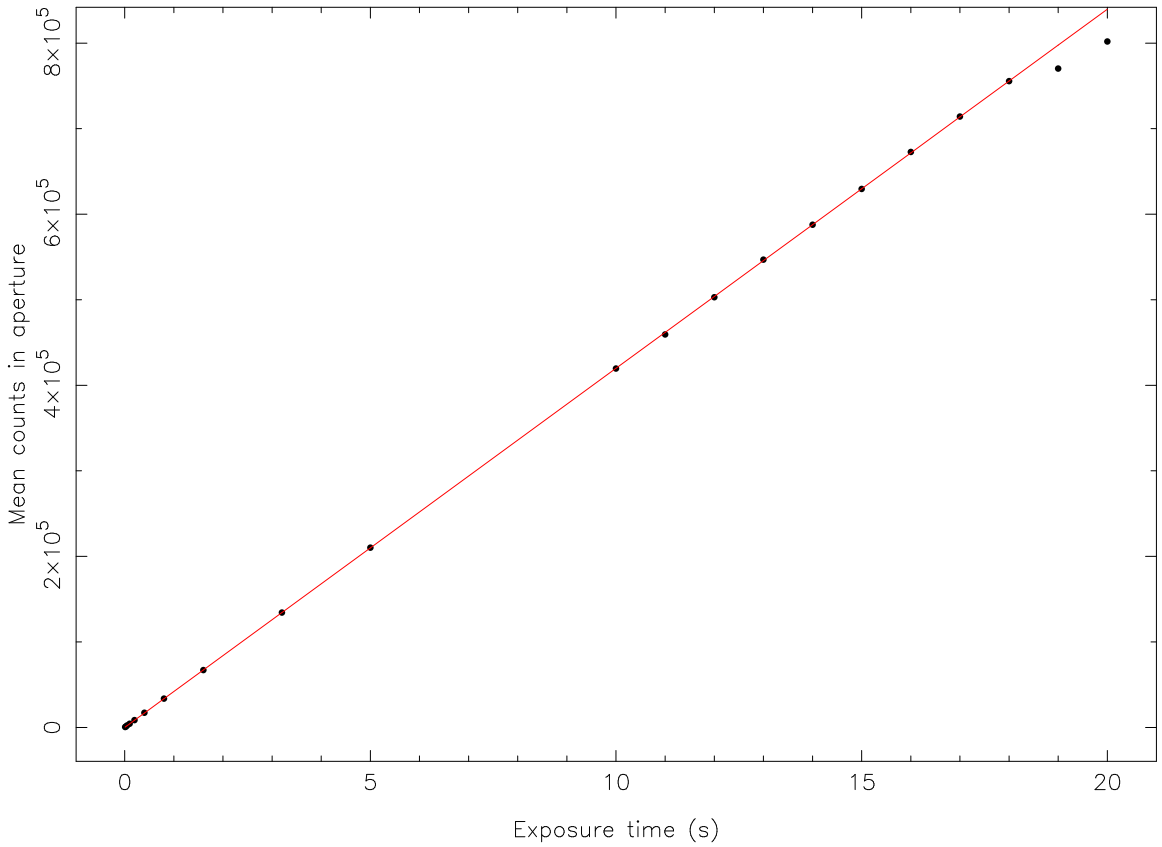

Figure 2.19: Linearity test of the ULTRASPEC CCD.

2.3.6 Pixel defects

Also whilst in the lab in Sheffield, we aimed to create a bad pixel mask to help observers avoid placing their targets on defective areas of the CCD. Such areas could be due to dust or other contaminants on the surface of the chip itself, hot pixels, bad columns or insensitive individual pixels. To identify these areas, a large number of flat fields and dark frames were median-combined create a master flat field frame and a master dark frame. The resulting master flat and master dark images were then visually searched for significantly dark or bright areas respectively.

Only one hot pixel was found in the dark frames, but several hundred dark pixels or regions were found in the flat fields. Many of these were minor variations, under 10% relative to surrounding pixels, and as such were labelled "moderate" defects. Some

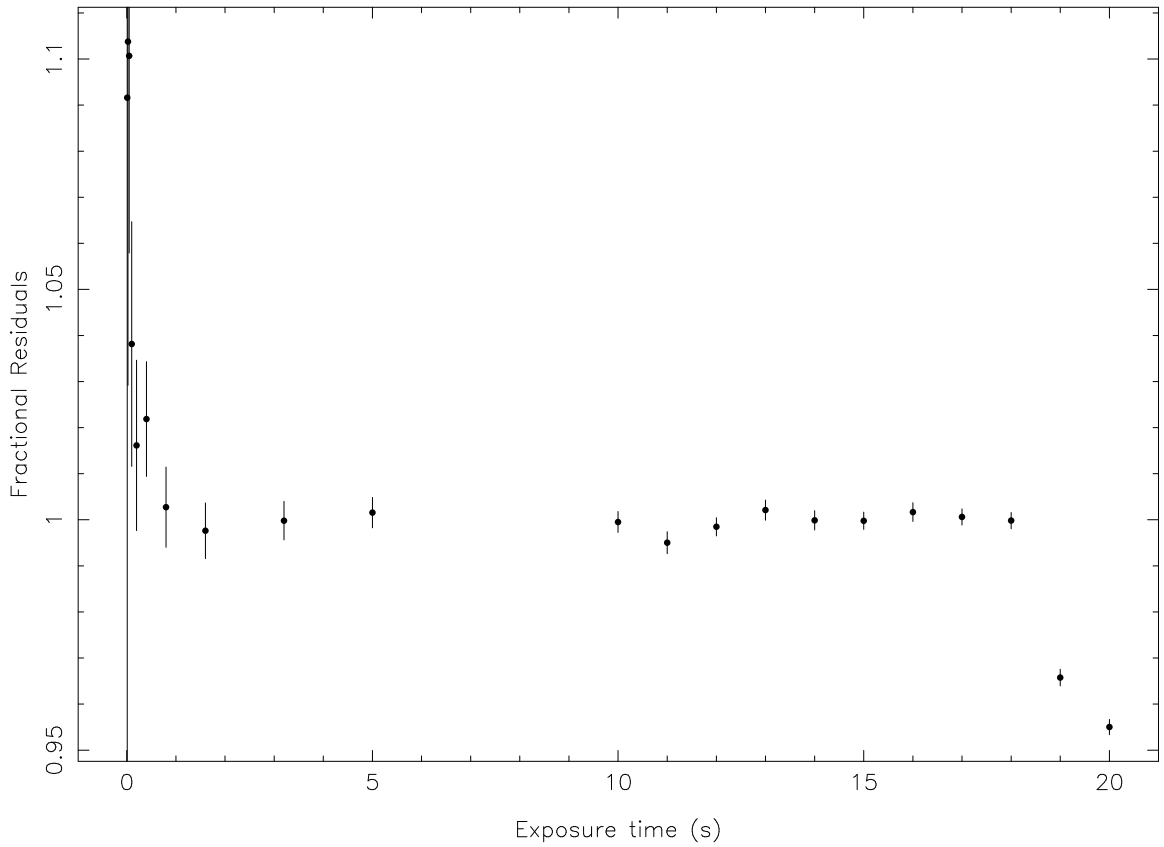

Figure 2.20: The residuals of the straight line fit in Figure 2.19, showing that the chip is linear to within 1% for almost the entire dynamic range.

however, displayed more "disastrous" variations, and were labelled accordingly. Figure 2.21 shows some bad pixels highlighted.

The bad pixel mask was updated during commissioning at the TNT, in case anything had changed during shipping and installation. Some defects were stronger or fainter than before, but very few new defects were seen.

2.3.7 Vacuum Lifetime

Unlike most astronomical cameras, ULTRASPEC uses the first correction lens itself as a seal for the cryostat vacuum, meaning it has what is known as an 'active window'. This improves throughput by removing the need to have a plane-parallel window sealing the cryostat. However, the design is more complicated, because the lens itself must not be

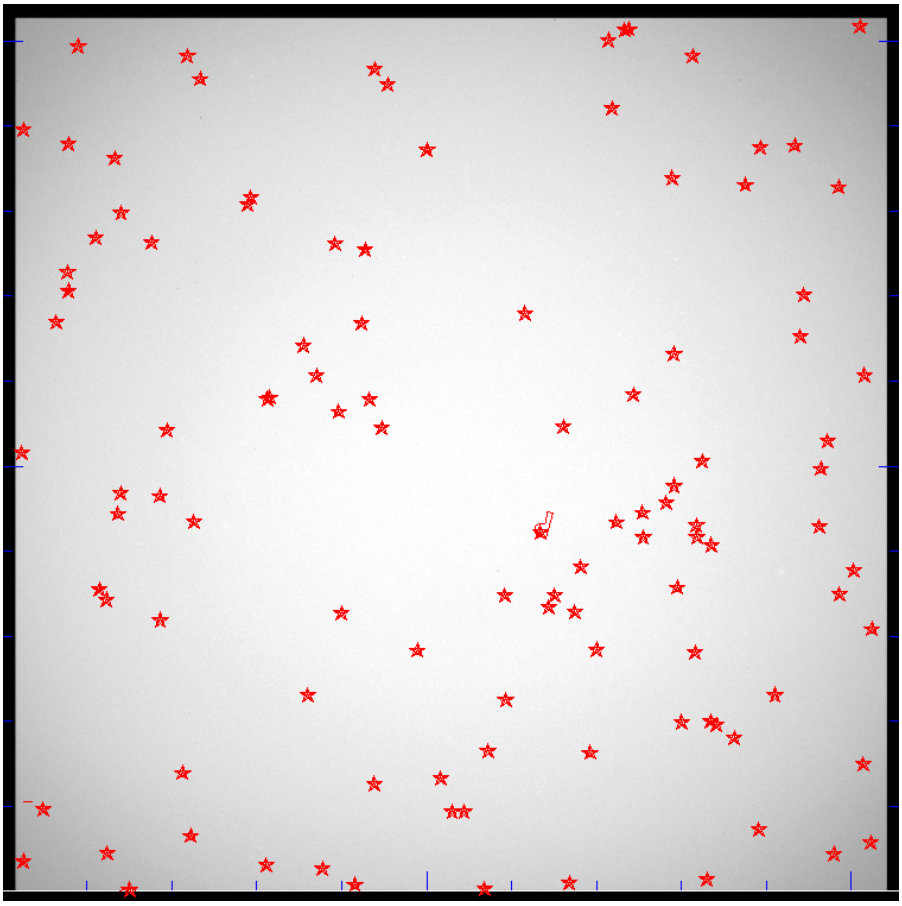

Figure 2.21: Image showing all "disastrous" pixel defects across the chip ($> 10\%$ variation).

able to move with respect to the cryostat if the pressure changes, as this would affect the location of the focal plane and thus degrade the image. Therefore the rim of the lens was polished, and an o-ring is placed between the rim of the lens and the inner wall of barrel 1. A second o-ring is needed between the barrel and the cryostat face.

This is a complex design, and after the delivery of the lenses it was important to test this seal. The clear glass window which was previously being used to seal the cryostat was removed, and the first lens barrel was slotted in its place. With everything else unchanged, any problems pumping down the vacuum would indicate an issue with the seal created by the first lens. The vacuum was pumped down as usual, and thankfully there were no obvious issues. The pressure dropped normally to around 10^{-5} mbar over

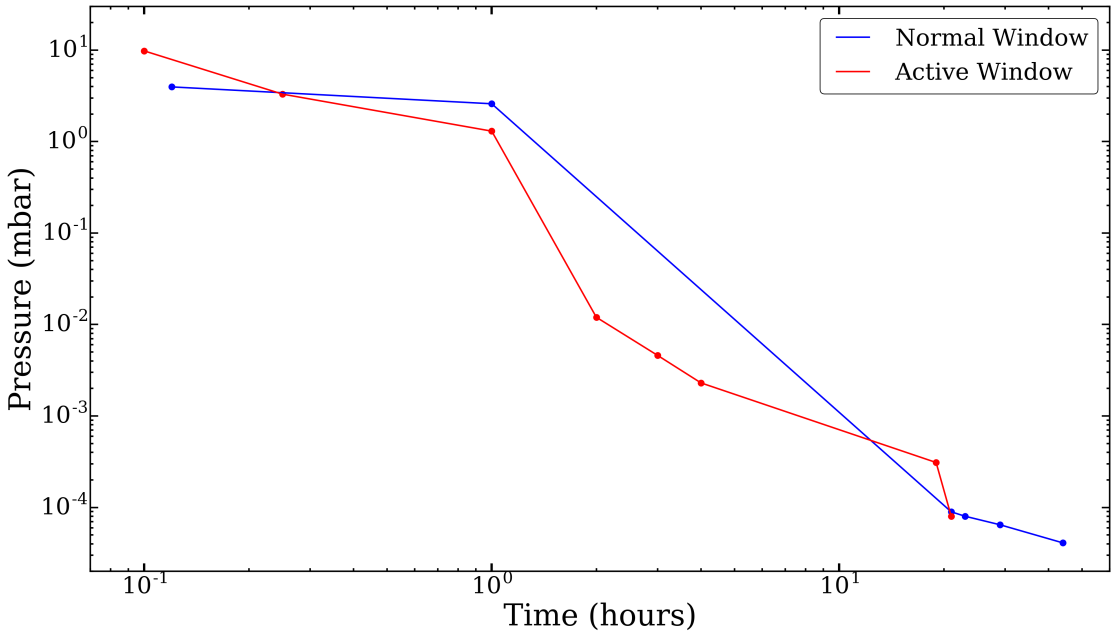

Figure 2.22: Cryostat pressure as a function of time after pumping begins, for the normal and active windows. The similarity of the two curves suggests the seal of the active window is good.

the course of several days. Figure 2.22 compares the vacuum pressures as a function of time after pumping began, for the original window and the active window.

For reference, Figure 2.23 is included to demonstrate how the temperature of the cold finger (thermal contact that connects the chip itself to the liquid nitrogen canister) and the CCD itself drops with time after cooling begins with the first introduction of nitrogen to the cryostat.

2.3.8 Problems in the lab

Before ULTRASPEC was ready to be shipped to Thailand, we encountered three serious issues with the detector which each required opening up the cryostat. The first occurred during the detector testing phase in the lab, when we were experimenting with dark current and different temperatures. We noticed that the temperature readings for the

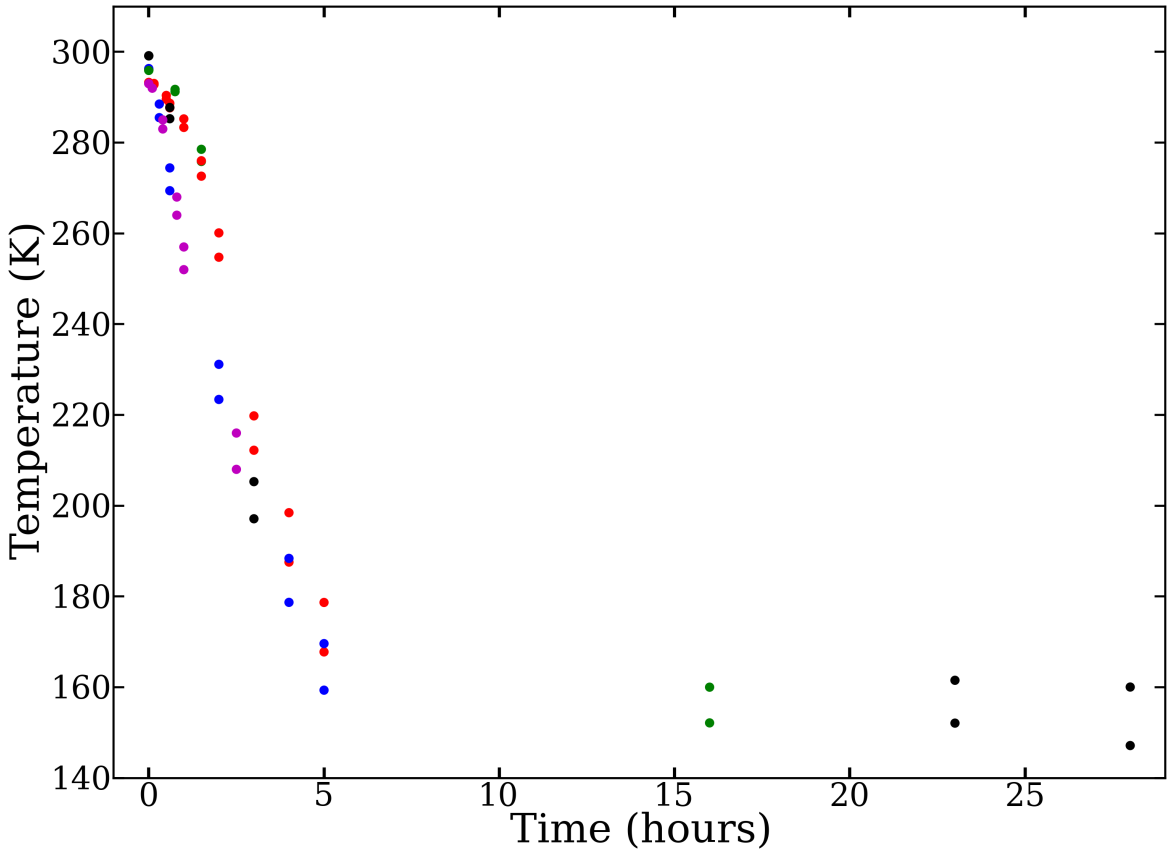

Figure 2.23: Temperatures of the cold finger and CCD as a function of time after cooling begins. Each colour has two points, one for the CCD temperature and one for the cold finger temperature (always cooler than the chip). Different colours represent different cooling sessions.

CCD and the 'cold finger' (the thermal contact between the CCD and cryostat which carries the heat away) were oscillating more and more wildly, and reading incredibly high temperatures (hundreds of degrees Celsius). We ceased operations immediately and transported the detector to the UK Astronomy Technology Centre (UKATC) in Edinburgh to allow the experts to investigate.

It turned out that the temperature sensor cables, and the cables for the heater which maintains the set temperature, had melted away their insulation and were short-circuiting. The original wires which were used should not have been, and new, more robust cabling was installed inside the cryostat, in the clean room at UKATC.

Back in the lab in Sheffield, the detector was now operating at a stable fixed temper-

ature as required, but we then found the dark current to be unacceptably high. Even when cold (at 160K), we were measuring over 400 electrons/pixel/hour, around 40 times higher than expected. We suspected that the thermal contact between the cryostat and the CCD must have been broken during the re-wiring at UKATC. This time we used a clean room in Sheffield to open up the cryostat and check the connections.

No obvious break in the thermal contact could be seen, so instead we checked all the electrical connections. Here we found two short-circuiting pins in the military connector which connects the CCD to the controller. We re-wired this short circuit and replaced all the components within the cryostat. After pumping out the vacuum and cooling the cryostat again, the dark current had returned to the normal levels of around 10 electrons/pixel/hour at 160K.

Finally, just as we believed everything was operating normally, the cryostat was moved again to install the first optics barrel and test the vacuum. During this movement, a small piece of hardware within the cryostat (which was an old, recycled ESO cryostat) came loose and fell onto the cryostat window. The object, which turned out to be a small spring, could clearly be seen through the cryostat window, but miraculously never impacted the CCD chip. This discovery required a third and final visit to the clean room, to open up the cryostat and remove the contaminating object. After this, the instrument was finally ready to ship to Thailand. A photo of one of the clean room operations is shown in Figure 2.24.

Figure 2.24: Photograph of Vik Dhillon investigating the ULTRASPEC CCD inside the open cryostat, in the clean room in Sheffield.

2.4 On-sky Commissioning

This section details the commissioning tests carried out at the 2.4-m Thai National Telescope, at Doi Inthanon, Thailand. These tests were conducted by myself and Vik Dhillon during the initial commissioning run in November 2013. The following text and data analysis is all my own.

2.4.1 Throughput Modelling

Probably the largest and most important part of on-sky commissioning is testing the throughput of the instrument. This is done to make sure the instrument is performing to specifications. The method involves modelling the throughput of the instrument and

Table 2.11: SDSS filter central wavelengths and bandwidths, from Dhillon et al. (2014).

Filter	λ_{cen} **(nm)**	**Bandwidth (nm)**
u'	355.7	59.9
g'	482.5	137.9
r'	626.1	138.2
i'	767.2	154.5
z'	909.7	137.0

telescope, arriving at a theoretical zeropoint. This is then compared to an observed zeropoint, calculated using data from observations of standard stars. We define the zeropoint as being the magnitude of a source from which we would measure one photo-electron per second on the ULTRASPEC CCD.

Theoretical Zeropoints

In modelling the telescope and instrument system, we analysed what fraction of light from a hypothetical zero magnitude star would be detected over the entire optical range of wavelengths. We started with the standard flux per unit frequency from a zero magnitude star in the AB system. This is $f_\nu = 3631 \times 10^{-23}$ ergs s^{-1} cm^{-2} Hz^{-1} in all filters (Fukugita et al., 1996). This value was then converted into a flux per unit wavelength, f_λ, which is measured in units of ergs s^{-1} cm^{-2} nm^{-1}. The flux is different for each filter, and is calculated using Equation 2.8:

$$f_\lambda = \frac{c f_\nu}{\lambda_{cen}^2}, \tag{2.8}$$

where c is the speed of light (in nm/s) and λ_{cen} is the filter central wavelength (in nm). In ULTRASPEC we most commonly use the SDSS u', g', r', i' and z' filters. Their central wavelengths and bandwidths are given in Table 2.11, while Figure 2.25 shows their transmission curves as a function of wavelength.

We next converted to N_λ, the number of photons s^{-1} cm^{-2} nm^{-1} in each filter, using

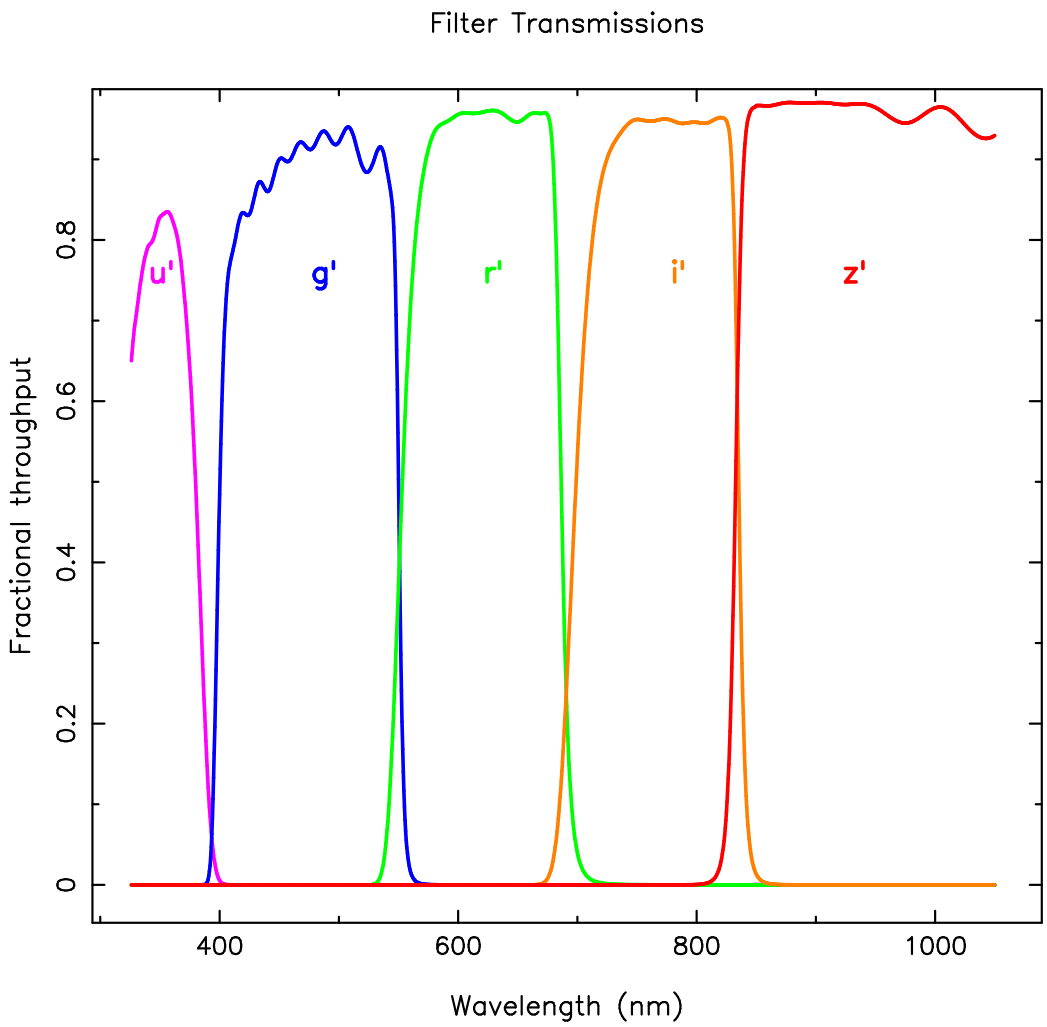

Figure 2.25: Transmission curves for the 5 SDSS filters used in ULTRASPEC as a function of wavelength. Data provided by the Asahi Spectra Company. The data is limited to between 326nm and 1050nm as this is the range within which we have available transmission data for all components of the system.

Equation 2.9.

$$N_\lambda = \frac{\lambda_{cen} f_\lambda}{hc} \tag{2.9}$$

where h is Planck's constant (6.63×10^{-27} ergs/sec).

Next we adjusted for the telescope mirror area, which we assumed to be 90% of the primary mirror area of radius 120cm (i.e. 10% less due to the secondary mirror shadow). This gave us the number of photons s^{-1} nm^{-1}. Finally, we calculated the number of

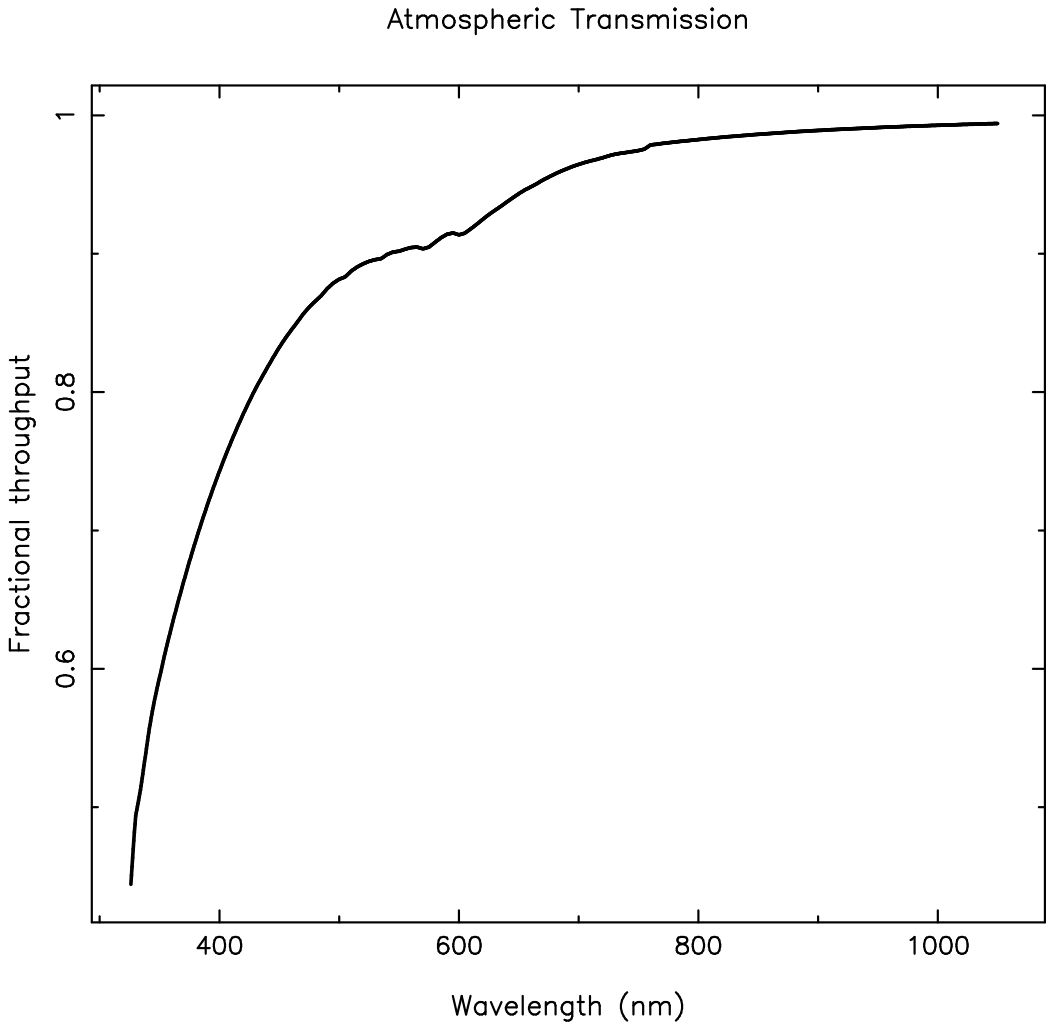

Figure 2.26: Fractional atmospheric transmission as a function of wavelength. These data are for the Roque de los Muchachos Observatory in La Palma (King, 1985), and have been adopted for the Thai National Observatory.

photons per second in each filter by multiplying N_λ by the bandwidth in each filter (given in Table 2.11). Each of the above steps are shown in Table 2.12.

Once the expected photon counts have been calculated, these numbers were then passed through a throughput model of the telescope and instrument, including the atmosphere, instrument optics (lenses and adhesives), filter throughputs and CCD quantum efficiency.

The throughput model is simply the predicted fractional throughput as a function of wavelength, derived by multiplying the individual throughputs of each of the above

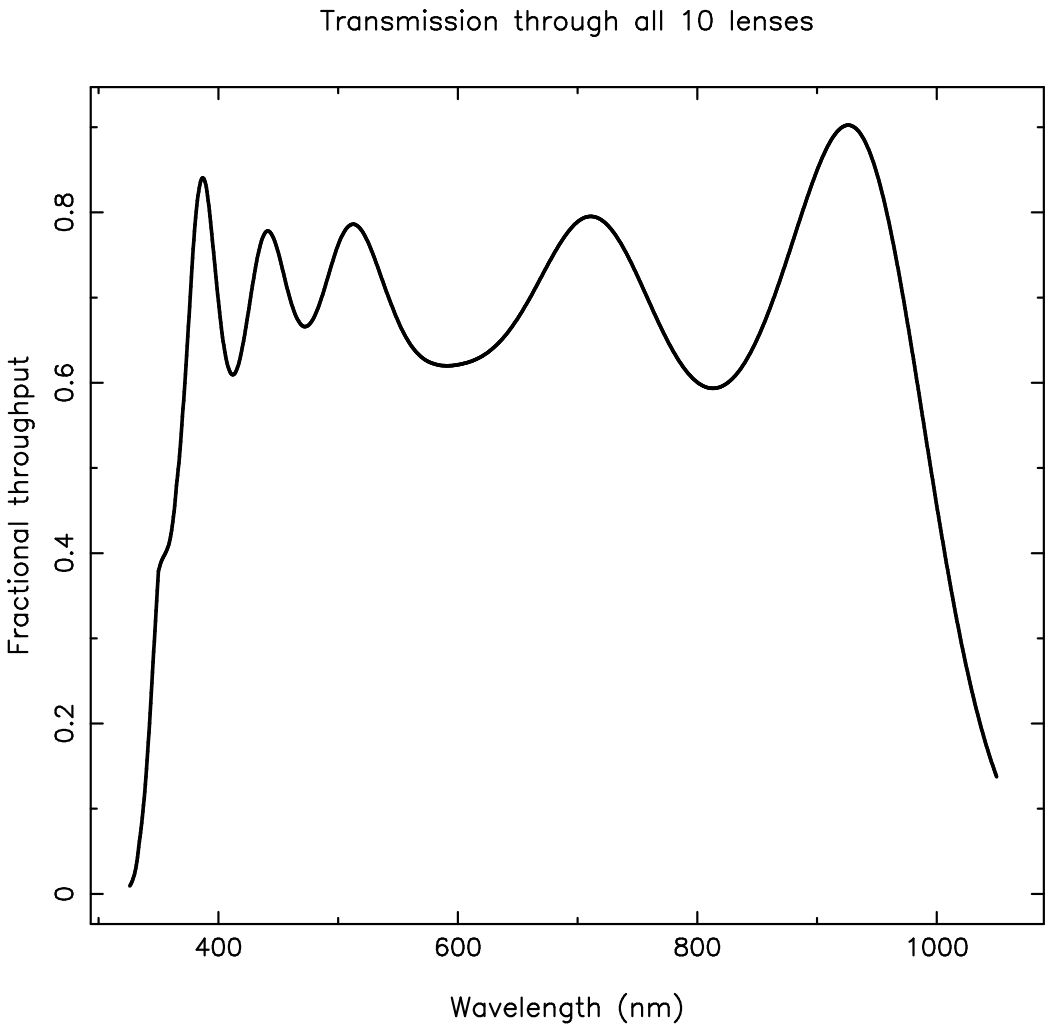

Figure 2.27: Cumulative throughput of all 10 lenses in ULTRASPEC as a function of wavelength. All 10 lenses (including two doublets) are modelled individually, accounting for the anti-reflection coatings on each face, and the internal transmittance of the glass (as provided by Schott UK Ltd.), which depends on the thickness.

Table 2.12: Fluxes and photon count calculations for each filter.

Filter	F_λ (ergs s^{-1}cm^{-2}nm^{-1})	N_λ (phot s^{-1}cm^{-2}nm^{-1})	N_λ (phot s^{-1}nm^{-1})	N_λ (phot s^{-1})
u'	8.61×10^{-8}	1.54×10^{4}	6.27×10^{8}	3.76×10^{10}
g'	4.68×10^{-8}	1.14×10^{4}	4.62×10^{8}	6.37×10^{10}
r'	2.78×10^{-8}	8.75×10^{3}	3.56×10^{8}	4.92×10^{10}
i'	1.85×10^{-8}	7.14×10^{3}	2.91×10^{8}	4.46×10^{10}
z'	1.32×10^{-8}	6.02×10^{3}	2.45×10^{8}	3.36×10^{10}

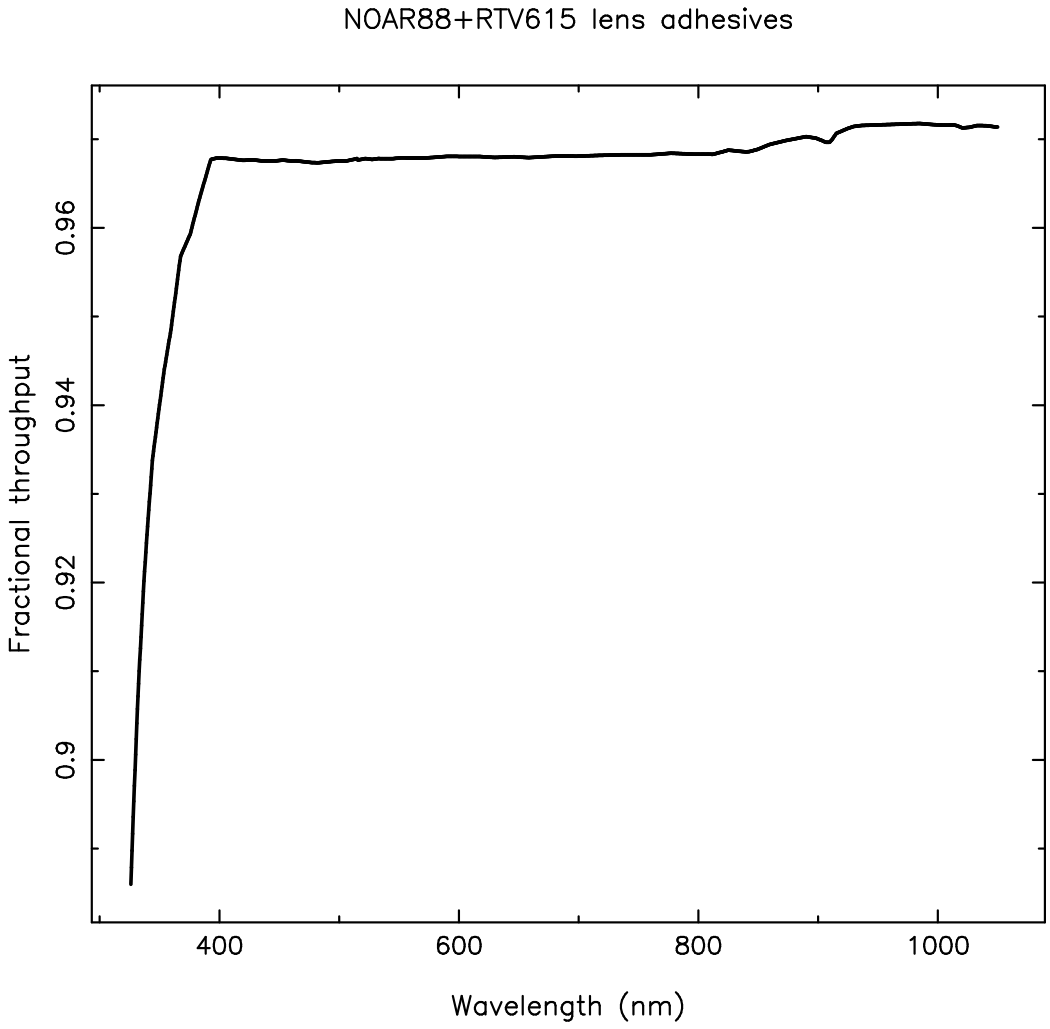

Figure 2.28: Transmission of lens adhesives in ULTRASPEC as a function of wavelength. The two doublet lenses require transparent adhesives to join the individual lenses. The two types of adhesive used were Techsil RTV-615 of thickness 0.12mm, and Norland NOAR88 of unknown thickness, but assumed to be of similar thickness to the Techsil RTV bond.

listed components together. The data for each component was first linearly extrapolated to span the same wavelength steps (0.5nm) and range (326 to 1050nm). The throughput of each component is shown graphically in Figures 2.25, 2.26, 2.27, 2.28 and 2.29. In the absence of a direct measurement, we assumed that each of the four telescope mirrors has a uniform reflectivity of 85% across all wavelengths, and thus our final transmission was multiplied by $0.85^4 = 0.522$. The mean throughputs in each filter are shown in Table 2.13, and the throughput model is shown in Figure 2.30.

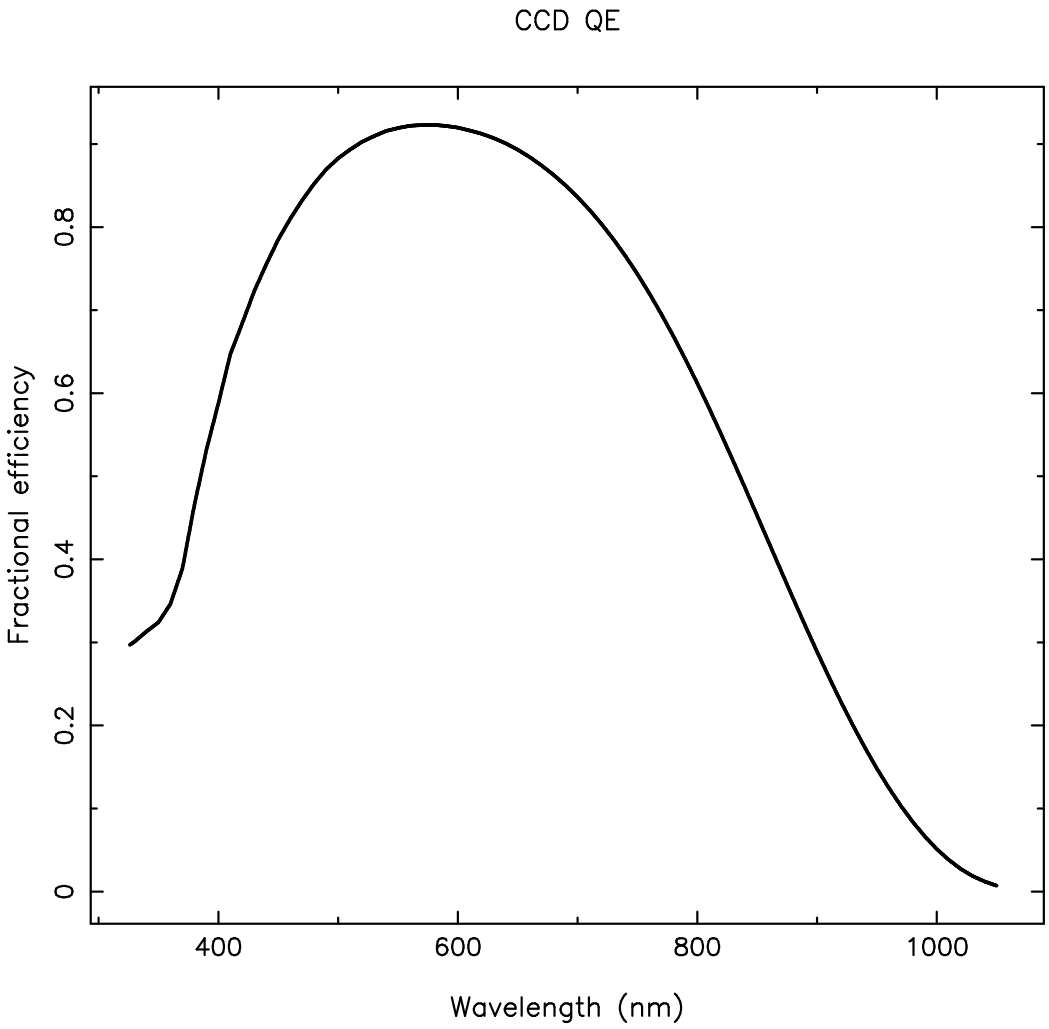

Figure 2.29: Quantum efficiency curve for the ULTRASPEC CCD, as provided by e2v.

Table 2.13: Theoretical throughputs of the entire TNT+ULTRASPEC system and atmosphere, and the resulting expected photon counts from a zero magnitude star in each SDSS filter.

Filter	**Mean throughput**	**Expected photons s^{-1}**
u'	0.0321	1.20×10^9
g'	0.2365	1.51×10^{10}
r'	0.2476	1.22×10^{10}
i'	0.1963	8.76×10^9
z'	0.1007	3.38×10^9

Table 2.13 also shows the expected rate of photo-electrons from a zero magnitude star, after the photons have passed through the throughput model. Using these numbers we were able to calculate a theoretical zeropoint using Equation 2.10.

$$m_0 - m_{ZP} = -2.5 \times \log_{10} \frac{F_0}{1} \tag{2.10}$$

where m_{ZP} is the theoretical zeropoint representing a star from which the detector would receive one photon per second, $m_0 = 0$, and F_0 is the expected flux in photons given in Table 2.13. The final theoretical zeropoints are shown in Table 2.16 alongside the observed zeropoints (see Section 2.4.1).

Atmospheric Extinction

In order to measure the actual zeropoints of the TNT+ULTRASPEC system, we first had to investigate the strength of atmospheric extinction at the site. From our experience observing at Doi Inthanon, there is often significant dust or haze in the lower atmosphere, and the strength of extinction can vary substantially even when no clouds are present. On top of this, thin high altitude cloud is also a common occurrence, sometimes at a level that is invisible to the eye and can only be detected in a transmission light curve.

We measured the extinction on many nights in several different filters, to try to build up an accurate picture of the 'best case scenario' in terms of atmospheric throughput. Plotting the observed flux from a stable comparison star against airmass should show a straight line with a gradient of k, the extinction coefficient, i.e. $F_{comp} = kA + c$ where A is the airmass. Figure 2.31 shows an example of this straight-line plot as observed by ULTRASPEC on 2014-02-02, in the g' filter.

We measured the extinction coefficient for 18 different observing runs, all of which covered an airmass change of at least 0.2. Most measurements are for the KG5 filter, and the results are shown in Table 2.14. Many other observations were analysed, but

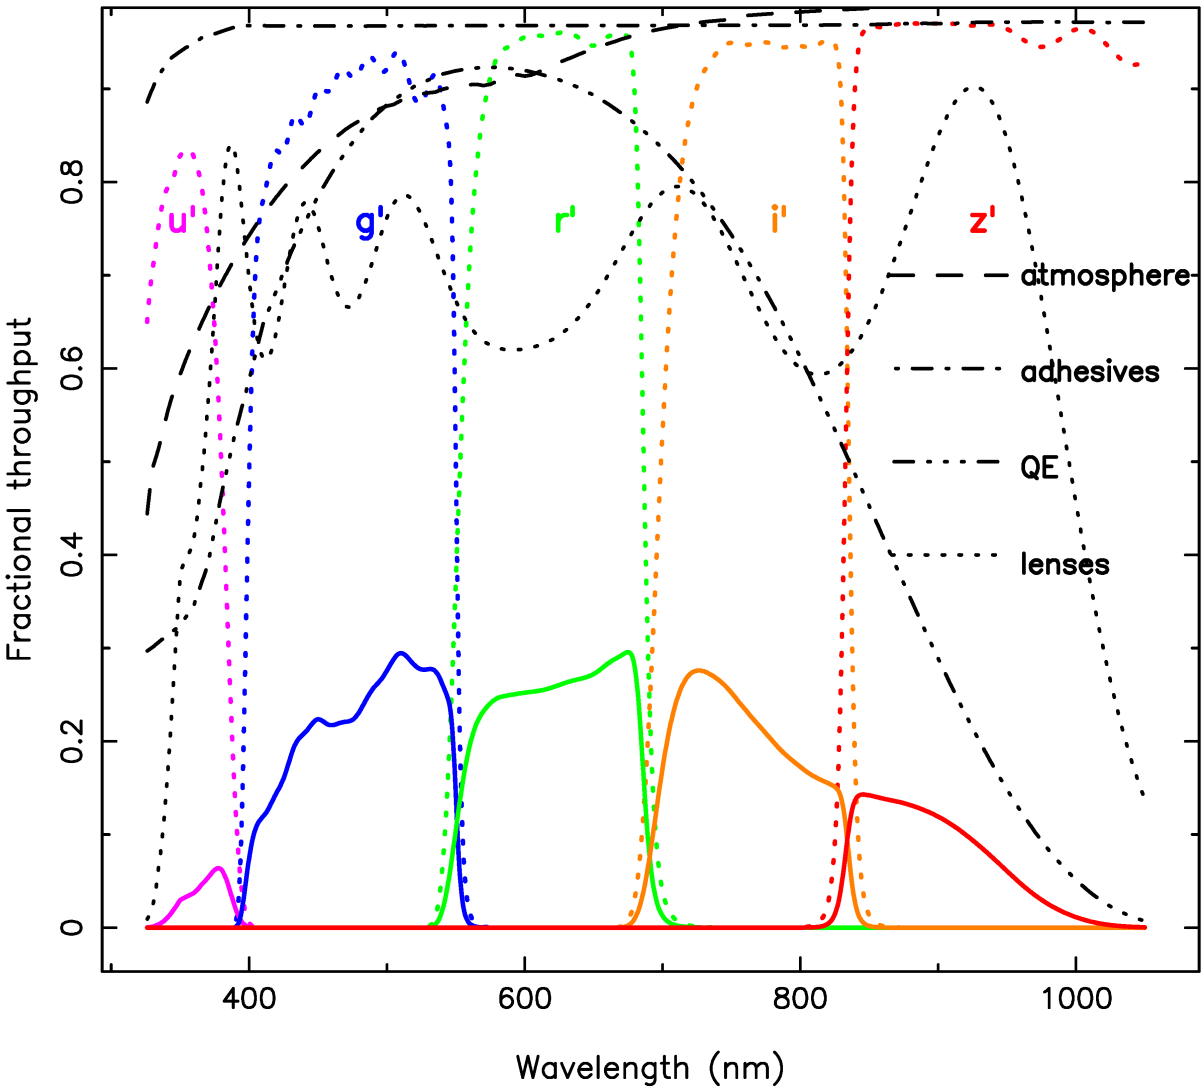

Figure 2.30: Throughput curves of ULTRASPEC optical components. The solid coloured curves show the total throughput of TNT+ULTRASPEC in each filter. This includes the atmosphere, telescope mirrors (not shown), lenses, lens adhesives and CCD quantum efficiency, which are all shown in black.

often the comparison star magnitudes showed transparency variations. These are not included in the table.

As can be seen in Table 2.14, there is substantial variability, demonstrated best by the KG5 filter. There is even large variability within a single night (2016-01-31) where the extinction coefficient changes from 0.11 to 0.18 over the course of a few hours.

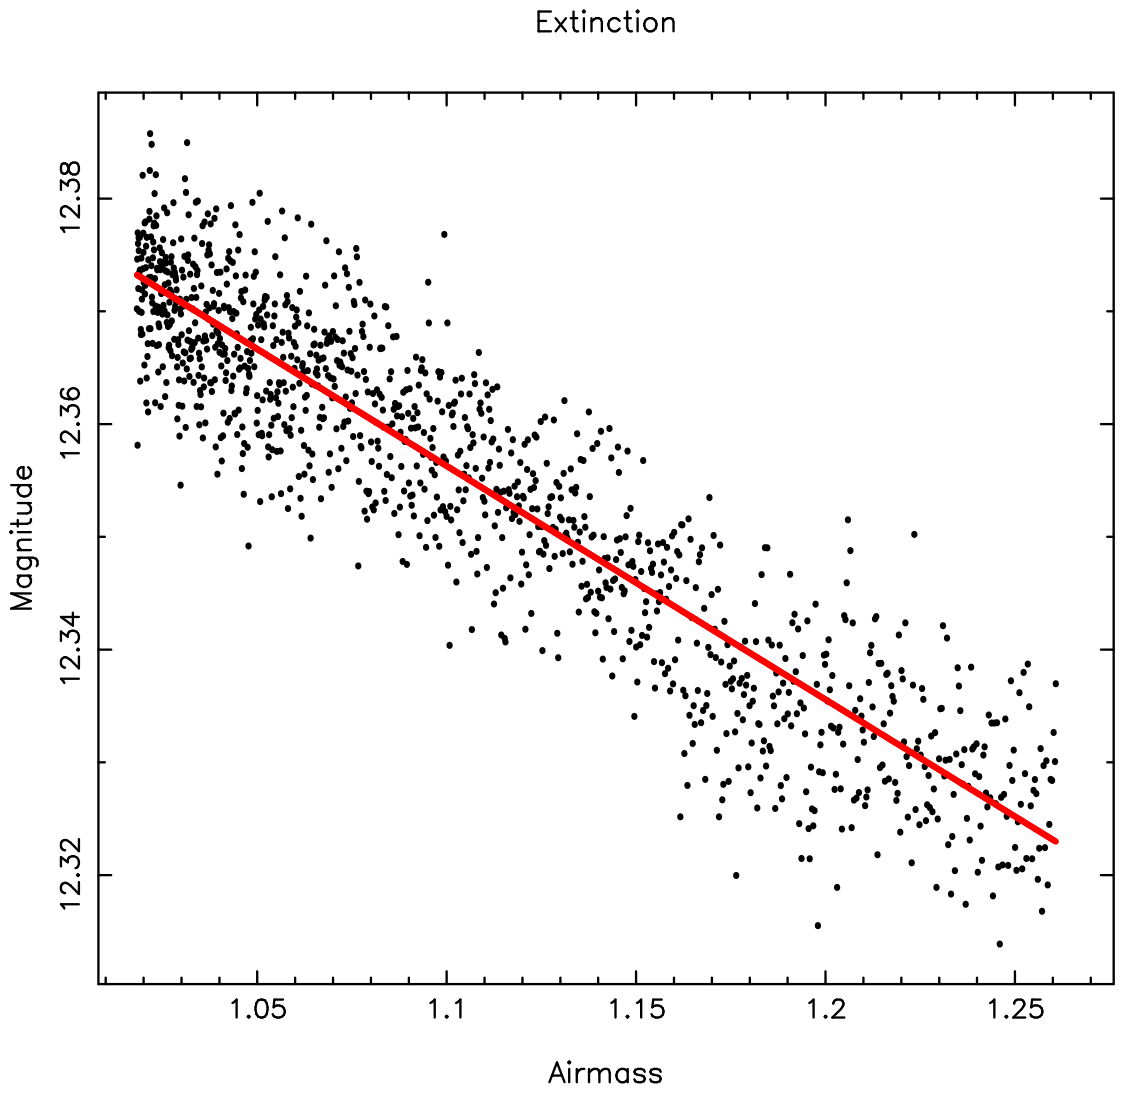

Figure 2.31: Instrumental g'-band magnitude as a function of airmass. A stable comparison star was observed for 2 hours on the night of 2014-02-02, through a range of airmasses. The straight line of best fit is shown in red. In this instance we measured an extinction coefficient of $k_g = 0.19$.

In the context of testing the instrumental throughput, we continue by considering the 'best' values for the extinction coefficients in each filter, in which the atmosphere is most transparent. For KG5 and g' we use 0.11 and 0.16 respectively. For u', r', i' and z' (in which we only have one measurement, or zero in the case of z'), we take the nominal values used at the Roque de los Muchachos Observatory in La Palma[2], which are 0.50,

[2] `http://www.ast.cam.ac.uk/~dwe/SRF/camc_extinction.html`

Table 2.14: Extinction coefficient k values at Doi Inthanon, as measured by ULTRASPEC.

Date	**Target Field**	**Filter**	**k**
2014-01-30	PSR J1023+0038	r'	0.10
2014-02-02	SDSS 1006+2337	g'	0.19
2014-11-27	MASTER 2321+4946	KG5	0.11
2015-01-02	SDSS 2306-0555	KG5	0.13
2015-01-05	ASASSN-13cx	KG5	0.12
2015-02-26	Cze V404Her	KG5	0.12
2015-12-09	V404 Cyg	KG5	0.12
2015-12-09	GALEX 0035+4623	g'	0.17
2015-12-09	WD 1145+017	KG5	0.14
2015-12-10	SDSS 0935+2700	g'	0.16
2015-12-10	WD 1145+017	KG5	0.13
2016-01-28	LSPM 0602+3910	i'	0.08
2016-01-28	WD 0843+516	KG5	0.18
2016-01-29	SSS 130413	u'	0.47
2016-01-30	WD 1145+017	KG5	0.16
2016-01-31	ES Cet	KG5	0.13
2016-01-31	WD 1145+017	KG5	0.11
2016-01-31	WD 1145+017	KG5	0.18

0.09, 0.05 and 0.04 respectively. These appear to match the measured ULTRASPEC values fairly well. In practice, we do not expect the extinction at Doi Inthanon to be this low very often, but our measurements suggest that the extinction can reach such low values occasionally.

Observed Zeropoints

With 'best case scenario' extinction coefficients available, we were then able to measure the observational zeropoints from images of photometric standard stars. In total we have analysed 36 standard star observations, usually in multiple filters, between November 2013 and March 2016. Note: approximately one third of this data reduction was completed by Martin McAllister.

For all standard star observations, the counts were extracted using the ULTRACAM data reduction pipeline. Wide apertures checked by eye in high contrast images were

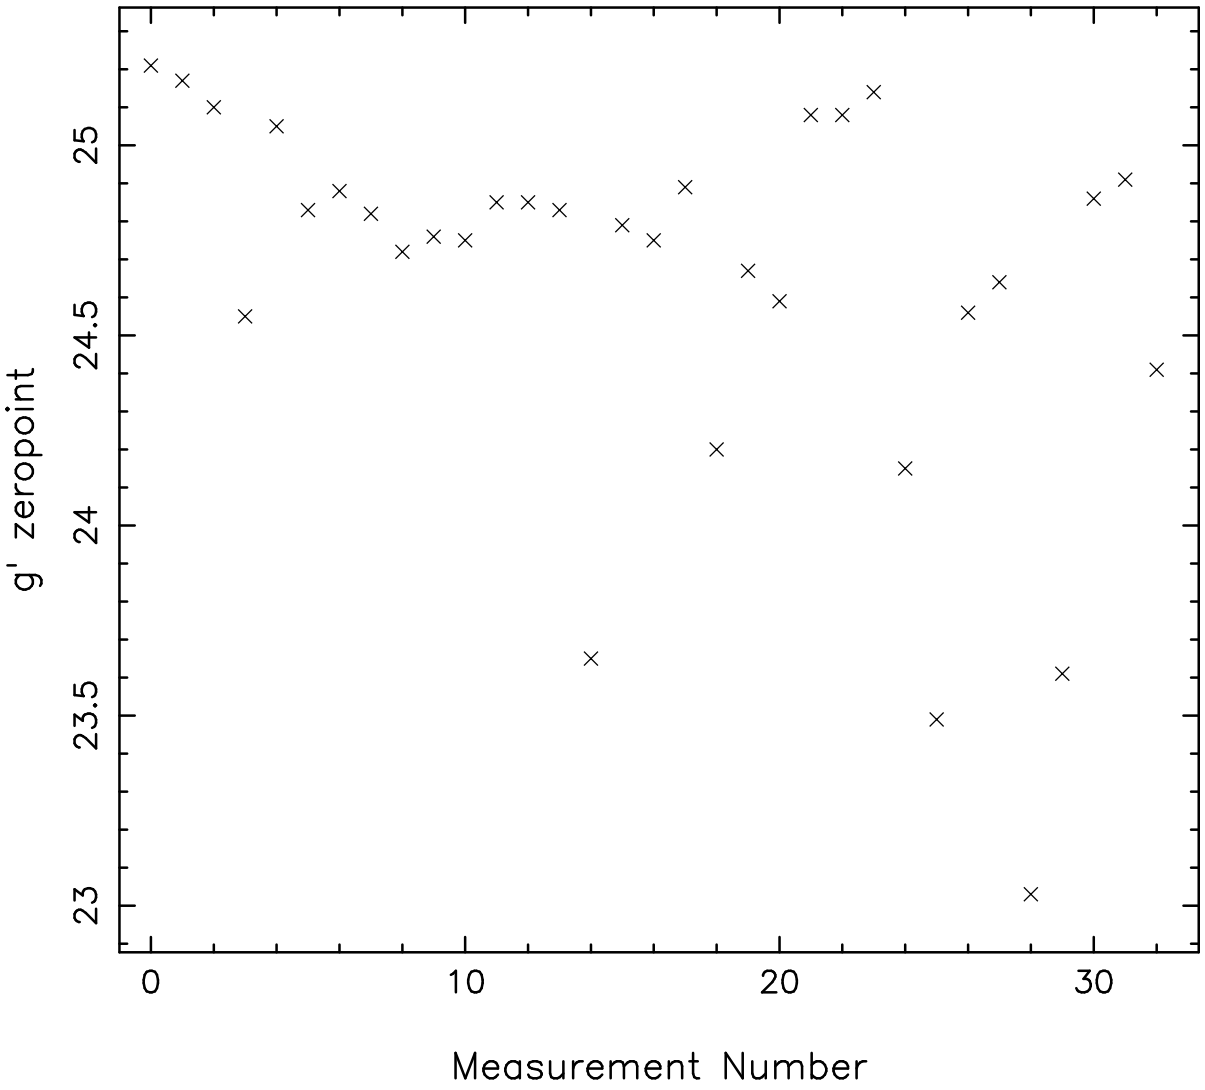

Figure 2.32: 33 measurements of the g'-band zeropoint for TNT+ULTRASPEC, taken from Table 2.15. The x-axis is simply the observation number, and does not scale linearly with time. There is considerable scatter as well as several obviously poor measurements, likely to have been observed through thin cloud.

used to ensure all of the flux from the star was extracted. Bias and flat field frames were used to calibrate the images as usual, and the mean counts per second were measured. The counts were then converted to photo-electrons using the gain (see Table 2.4). The observed zeropoints for each filter were then calculated using Equation 2.11:

$$m_{STD} - m_{ZP} = -2.5 \times \log \frac{F_{STD}}{1} - (k \times A), \tag{2.11}$$

where m_{ZP} is the observed zeropoint, m_{STD} is catalogued magnitude of the standard star, F_{STD} is the flux (in photo-electrons) measured by ULTRASPEC, k is the extinction coefficient for the filter in question, and A is the airmass at which the observations were made. Of course, extinction is a colour-dependent variable, and second-order colour terms could be encorporated here, though we do not attempt such high precision calculations ourselves. Our measured zeropoints for all 35 standard star observations are shown in Table 2.15. Uncertainties in these zeropoints are dominated by the uncertainties in the extinction coefficients (of the order 10%). The resulting uncertainties in the zeropoints are therefore approximately 0.05 mags, 0.01 mags, 0.01 mags, 0.01 mags, 0.01 mags and 0.02 mags for the u', r', i', z' and $KG5$ filters respectively.

Figure 2.32 plots the g'-band zeropoints, showing that there is considerable variability in the observed zeropoints of up to 0.5 magnitudes, as well as clear episodes of poor throughput at certain epochs. We are not aware of any substantial changes to the instrument and telescope setup, and there have been no mirror aluminisation procedures. Thus the throughput of the telescope and instrument are unlikely to be the primary cause of this variability. Instead we attribute it to the presence of cloud or high atmospheric extinction, which we know is often present and can drastically affect the zeropoint measurement. There may be a slight overall trend in zeropoint evolution, with the throughput becoming slightly worse over time. This is not unusual, as the telescope mirror reflectivity is expected to degrade between cleaning or aluminisation procedures.

To test the theoretical throughput model discussed above in Section 2.4.1, we compare the theoretical zeropoints with our best observed zeropoints, since these are likely to have been observed when the sky is clearest, and thus the chosen best extinction coefficients are likely to be most accurate. The 'best' (i.e. faintest) zeropoints are given in Table 2.16, along with the theoretical expectations.

The best zeropoint in g' was measured on the very first night of ULTRASPEC commissioning. Although there weren't many observations on that night covering a range

Table 2.15: Observed ULTRASPEC above-atmosphere zeropoint magnitudes of a star giving one photo-electron per second. The variability is likely due to varying atmospheric extinction and the presence of thin clouds.

		Zeropoint (1 photon/second)						
Date	**Star**	u'	g'	r'	i'	z'	KG5	**Comment**
2013-11-05	Feige 22	22.07	25.21	25.23	24.52	23.43	-	
2013-11-05	SA 95-96	22.29	25.17	25.19	24.58	23.49	-	
2013-11-06	BD+251981	21.87	-	-	-	-	-	Cloudy
2014-01-26	SA 99-447	-	-	-	24.65	-	-	
2014-02-02	Feige 22	-	25.10	25.19	-	-	-	
2014-03-28	SA 98-685	21.83	24.55	24.67	24.21	23.30	25.21	Possible cloud?
2014-03-29	Ru 149D	21.86	-	-	24.22	23.31	-	
2014-03-30	SA 97-249	22.50	25.05	24.99	24.45	23.51	-	Hazy
2014-11-28	SA 101-315	22.25	24.83	24.88	24.30	-	-	Hazy?
2014-12-01	BD+210607	22.34	24.88	24.95	24.38	-	-	Thin cloud?
2014-12-01	SA 97-249	22.26	24.82	24.88	24.29	-	-	Hazy?
2014-12-01	SA 101-207	22.20	24.72	24.79	24.23	-	-	Hazy?
2014-12-01	Feige 34	22.07	24.76	24.88	24.30	-	-	Hazy?
2014-12-02	SA 114-750	22.21	24.75	24.87	24.31	-	-	Hazy?
2015-01-01	SA 92-288	22.32	24.85	24.85	24.27	-	-	Hazy?
2015-01-02	SA 97-249	-	24.85	24.85	24.27	-	-	Hazy?
2015-01-03	SA 93-333	-	24.83	24.83	24.26	-	-	Possible cloud
2015-01-04	SA 93-333	21.49	23.65	23.12	22.73	-	-	Thick cloud
2015-01-05	Hilt 190	22.29	24.79	24.78	24.22	-	-	Hazy?
2015-01-06	Feige 34	22.08	24.75	24.80	24.28	-	-	Hazy?
2015-02-27	SA 97-249	-	24.89	24.82	-	-	25.45	
2015-03-01	BD+210607	-	24.20	24.37	-	-	24.99	Hazy
2015-03-01	GJ 745A	-	24.67	24.63	-	-	25.38	Hazy?
2015-03-03	SA 97-351	-	24.59	24.63	24.13	-	25.36	Hazy?
2015-12-10	SA 100-280	22.63	25.08	25.02	24.46	-	25.67	
2015-12-11	SA 93-333	22.55	25.08	25.05	24.44	-	25.75	
2016-01-31	Hilt 404	22.41	25.14	24.97	24.40	-	-	
2016-03-13	Ross 49	21.31	24.15	24.02	23.53	-	-	Cloudy
2016-03-14	SA 96-83	20.44	23.49	-	-	-	24.25	Cloudy
2016-03-15	SA 97-351	21.52	24.56	24.45	23.85	-	25.21	Cloudy
2016-03-16	SA 97-249	21.53	24.64	24.55	24.02	-	25.19	Possible cloud
2016-03-19	SA 98-978	20.21	23.03	23.22	22.98	-	23.61	Cloudy
2016-03-19	SA 98-653	20.62	23.61	23.63	23.28	-	24.25	Cloudy
2016-03-19	SA 107-351	21.94	24.86	24.71	24.14	-	25.43	Smog?
2016-03-19	SA 110-499	22.19	24.91	24.71	24.13	-	26.00	
2016-03-20	BD+82015	21.16	24.41	24.37	23.90	-	25.37	Cloudy

Table 2.16: Theoretical and observed zeropoints for TNT+ULTRASPEC.

Filter	**Theoretical Zeropoint**	**Observed Zeropoint**	**Difference (magnitudes)**	**Fractional difference**
u'	22.70	22.63	0.07	0.94
g'	25.45	25.21	0.24	0.81
r'	25.21	25.23	-0.02	1.01
i'	24.86	24.65	0.21	0.83
z'	23.82	23.51	0.31	0.75

of airmasses, we attempted to measure the extinction from a 40 minute observation of a field, as it changed in airmass by 0.1. For this data the measured extinction coefficient was k_g=0.16, which is identical to the chosen extinction coefficient in g' based on the full study. This is reassuring and lends further credence to our selected best zeropoints.

As Table 2.16 shows, the instrument appears to be performing as well as expected in the u' and r' filters, but under performs by nearly 20% in the g' and i' filters. The under-performance might be due to systematic errors in the extinction coefficients, or could be attributed to wavelength dependent inaccuracy in the throughput model, most likely manifested in the mirror reflectivity (which was taken to be uniform at 85%). The z' filter under performs by 25%, however the small sample of standard star observations in this filter could be to blame here.

In the future it would be beneficial to match the modelled and measured zeropoints more closely. To do this, we would make further measurements of extinction coefficients, attempting to get reliable values in all filters. We would also aim to acquire actual measurements of the TNT mirror reflectivity, rather than relying on an estimated uniform 85% reflectivity on all four mirrors. In addition, we would continue to observe more standard stars across the visible sky, including at a range of airmasses to help refine the measured zeropoints. In particular we would like to obtain more standard star measurements in the z' filter.

2.4.2 Rotator Centre

To locate the centre of the rotator, assumed to be the centre of the light path, with respect to the CCD, we observed a star cluster at rotator orientations that were 180° apart. Then we measured the positions of 7 bright stars in both images, using a Moffat fit to measure the centroid of the stars. The mean half-way point between the positions of each star from one rotator position to the other gives us the rotator centre. This was measured to be at pixel 524,475, which is only 40 pixels away from the centre of the CCD. This shows there is good alignment of the mechanical axes of ULTRASPEC and the telescope, and will ensure accurate acquisition of targets, since the telescope pointing is centred on the rotator centre.

2.4.3 Flexure

To test for any flexure in the instrument and rotator cube, we observed a star very close to the rotator centre at a range of angles of the rotator, spanning all 360°. Since the star was not located directly on the rotator centre at pixel 524,475, it made a small ellipse around the rotator centre as we moved the rotator. If there was no flexure, or absolutely identical flexure in all directions (unlikely due to the uneven weight distribution of the instrument and cube), the star would trace an exact circle. Instead, as expected we observed an ellipse, as seen in Figure 2.33. This shows that there is a small amount of additional flexure in one direction - it is worst when the rotator angle is set to 220° or 40°. The ellipse has a major axis of 21 pixels and a minor axis of 10 pixels. The difference between these (11 pixels), gives us an estimate of the total flexure in the direction of the major axis. Since each pixel is 13μm and 0.45", this represents a total movement of 0.143mm or 5". This flexure is likely to dominate the small tracking drifts we observe in ULTRASPEC (see Section 3.5.4).

The fitted ellipse in Figure 2.33 is centred on pixel 524,476, which matches well with

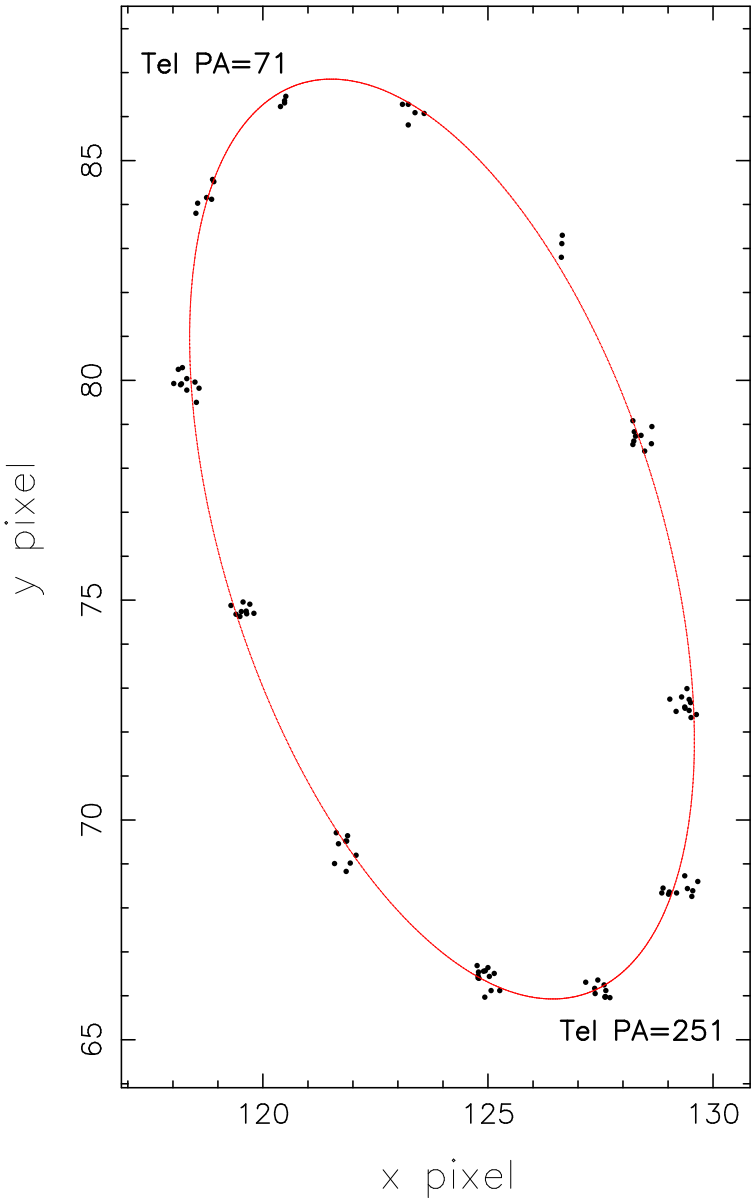

Figure 2.33: x and y pixel positions of a star as it was rotated around the rotator centre. The difference between the major and minor axes of the red fitted ellipse gives the amount of flexure.

the rotator centre measured above in Section 2.4.2. It is not possible to know whether this flexure is due to the flexure of ULTRASPEC or the flexure of the M4 mirror inside the Nasmyth cube. Overall the instrument chassis is very stiff, since the observed flexure is low and does not affect observations.

2.4.4 Tracking

Since the TNT does not have a seperate autoguider, the telescope must be able to track accurately. To investigate the telescope tracking accuracy, and test for the observable effects of flexure, we plot the x and y pixel positions of stars during a 4-hour observing run of a particular field (Figure 2.34). This plot shows that the movement due to tracking errors and/or instrument flexure can reach up to 15 pixels. The maximum rate of positional shift reaches approximately 12 pixels per hour (5.4″/hour).

The majority of this movement appears to come from the flexure in the instrument position, as can be seen more clearly in Figure 2.35. The movement through the observers' meridian, at maximum altitude, produces the largest pixel shift, as the rotator spins the instrument cube around at its fastest rate. The difference in pixel positions before and after the transit match the 11 pixels of flexure measured above in Section 2.4.3.

These plots suggest that, were the positional shifts to become unacceptable for a particular observing campaign, the only ways to improve upon this would be either to address the flexure in the instrument and/or rotator cube, or install an autoguider. Thankfully, the current levels of tracking drift are acceptable for the majority of our observations, which use short exposures to probe time-resolved properties. Assuming an image quality of 0.9" (the median seeing value), exposure times of longer than 10 minutes would begin to show smearing of point sources with the maximum rate of positional shift of 5.4"/hour. The other main imaging instrument at the TNT is the 4K Camera. This can be used for deep imaging as it has a built-in autoguider.

2.4.5 Sensitivity Variations

We noticed immediately during commissioning that there was a strong, striped pattern in the twilight flats of the u' filter. The stripes appeared to be in focus, and did not

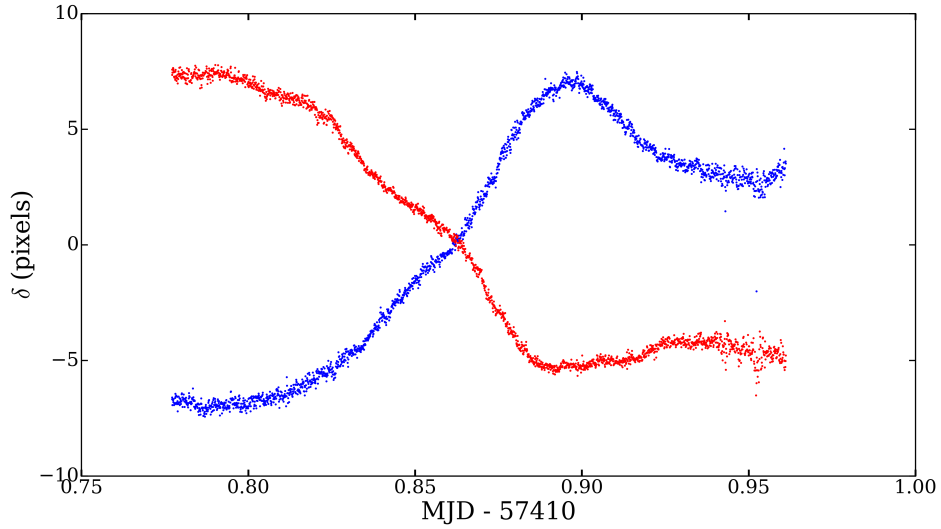

Figure 2.34: Pixel shifts as a function of time, where δ is the difference between the pixel position at a given moment and the mean pixel position throughout the entire observation. Blue represents the x-pixel position while red is the y-pixel position.

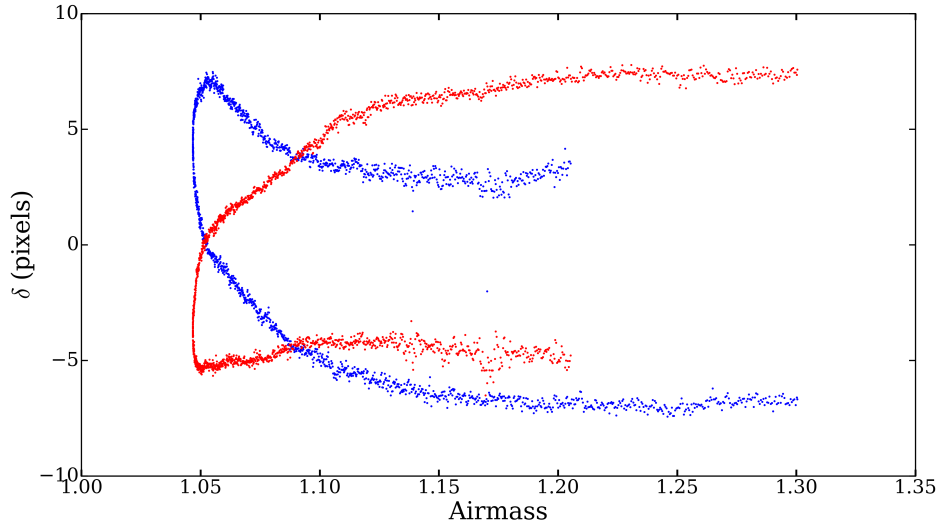

Figure 2.35: Pixel shifts as a function of airmass, where δ is the difference between the pixel position at a given moment and the mean pixel position throughout the entire observation. Blue represents the x-pixel position while red is the y-pixel position.

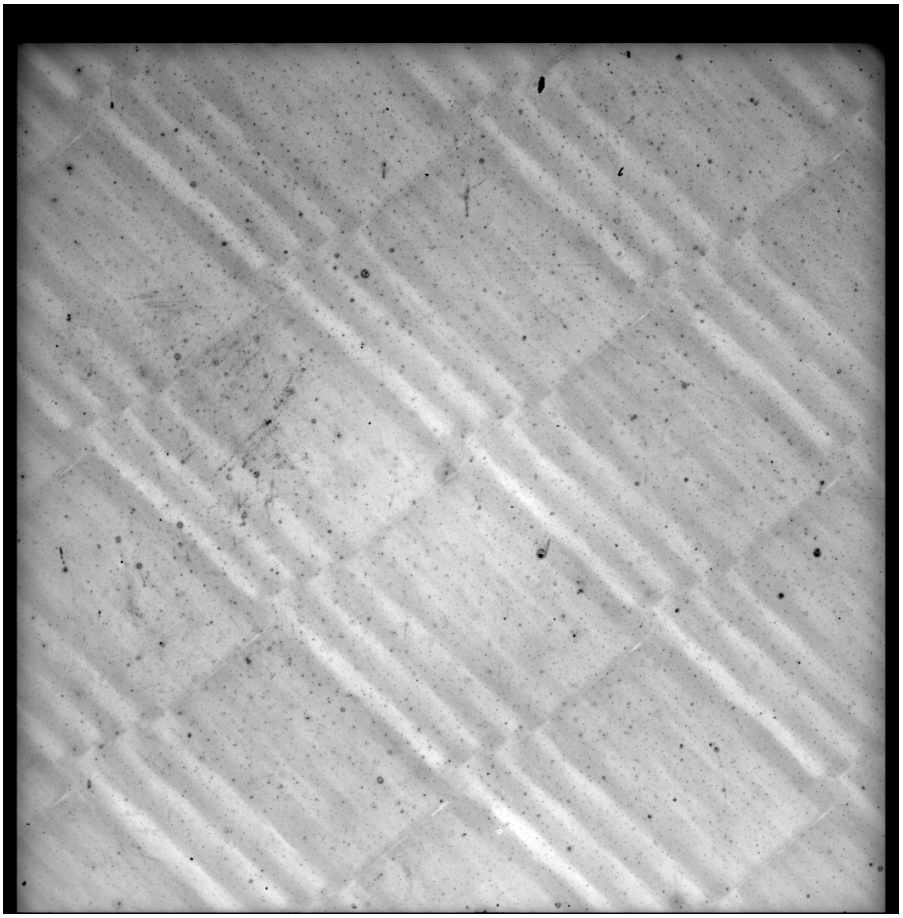

Figure 2.36: Median-combined master twilight flat field taken in the u' filter with ULTRASPEC. The striped pattern appears to be in focus, and must therefore be intrinsic to the CCD, or at least produced close to the focal plane.

change with time. This suggested an origin on the CCD itself, or at least very close to the focal plane, and not related to pick-up noise, which varies in position with time. Figure 2.36 shows the striped pattern in a typical median-combined master u'-band twilight flat.

To test whether this variation in sensitivity is indeed intrinsic to the CCD, we observed a bright star and moved in small steps across the chip, across the striped pattern. An initial test conducted in November 2013 proved to be unusable due to the presence of cloud, however a further test in March 2014 was conducted with clear skies. The results of this test, shown in Figure 2.37, are also somewhat inconclusive. The brightness of the

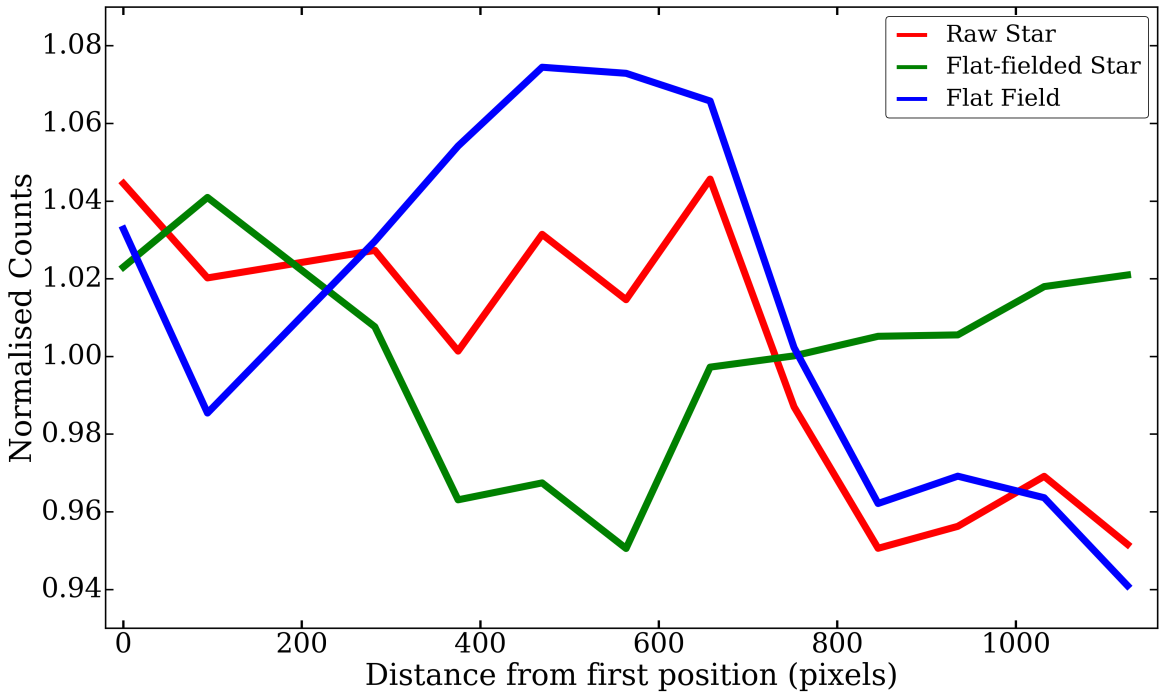

Figure 2.37: Fractional u'-band sensitivity variations with diagonal position across the CCD. Line thickness gives an indication of the Poissonian error. We measured the counts of a star at different positions across the CCD, from bottom left to top right (see Figure 2.38). Assuming a stable atmosphere and intrinsic counts of the star, the red curve should match the blue curve, which shows the counts in the twilight flat field frame, measured in a 5×5 pixel box centred on the same positions as the star. The green curve shows the same test but for frames which have already been divided by the twilight flat. This should show a straight flat line with a normalised flux of 1.

star at different positions on the CCD did not perfectly match the brightness of the flat field at these same positions (as measured by the mean counts inside a 5×5 pixel box). However, the steps made across the chip appeared to be too large, and so the striped pattern was not sampled effectively.

An alternative test to ascertain whether or not the striped pattern could be intrinsic to the twilight flats (e.g. due to scattered light) rather than CCD sensitivity variations, is to compare the twilight flats with deep u'-band images of the night sky. Since we rarely observe in u', deep images of this type are not easily available. However, the original star stepping test, conducted with some cloud present, had a total of 112×12.8 second exposures. Stacking all these frames together gave us a deep night-sky u'-band image

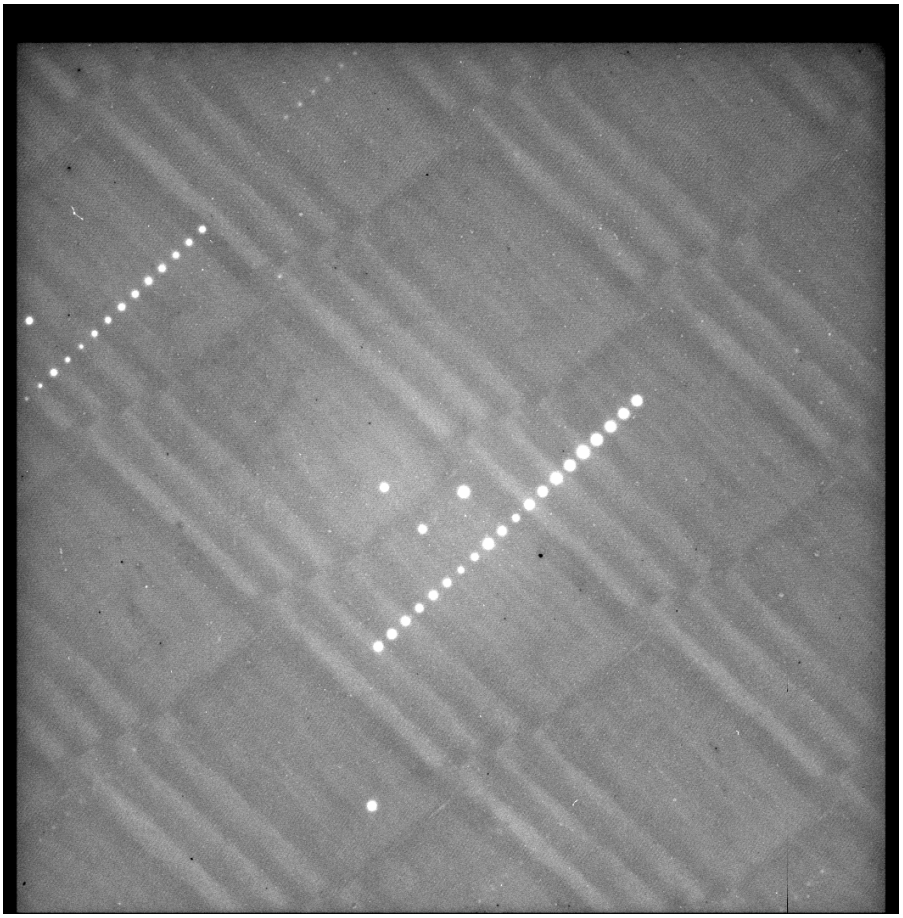

Figure 2.38: Deep u'-band night-sky image, combining all of the ULTRASPEC frames in the original star-stepping test. The striped pattern seen in the twilight flats is clearly visible.

(see Figure 2.38), which clearly shows the same striped pattern as the twilight flats. In both the twilight flat field and the night-sky image, the strength of the sensitivity variations from one small stripe to the next are of the order of 10%. We presume that this variation must be due to some wavelength-dependent CCD sensitivity pattern, since it only appears in the u' and sometimes weakly in the g' and KG5 filters, i.e. at blue wavelengths. It may have something to the do with the CCD thinning process.

It is not clear why the green curve in Figure 2.37 is not flat. There may have been variable seeing conditions, or some minor transparency variations.

To confirm that the twilight flat fields can accurately correct for the striped pattern,

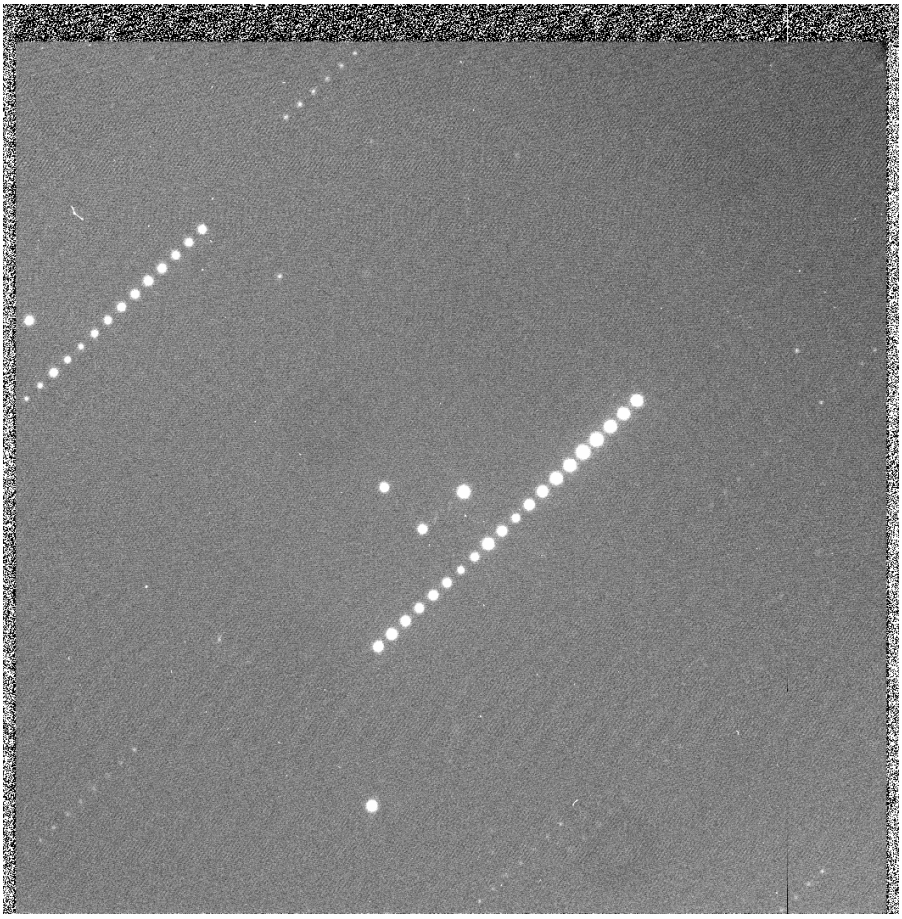

Figure 2.39: Deep u'-band night-sky image divided by the normalised u'-band twilight flat. The striped pattern disappears completely. The remaining structure is due to pick-up noise in the bias frames.

we divided the u'-band night-sky frame by the normalised u'-band twilight flat. The result is a perfectly smooth image (apart from a weak signal due to pick-up noise), as seen in Figure 2.39. This suggests that the twilight u'-band flats are perfectly sufficient for flat fielding u'-band science frames.

2.4.6 Flat Field Structure

In addition to the striped patterns seen in images taken with the u' filter (Section 2.4.5), it was also obvious during the commissioning of ULTRASPEC that there was some strong radial structure in the flat field images. An example twilight flat field image taken in

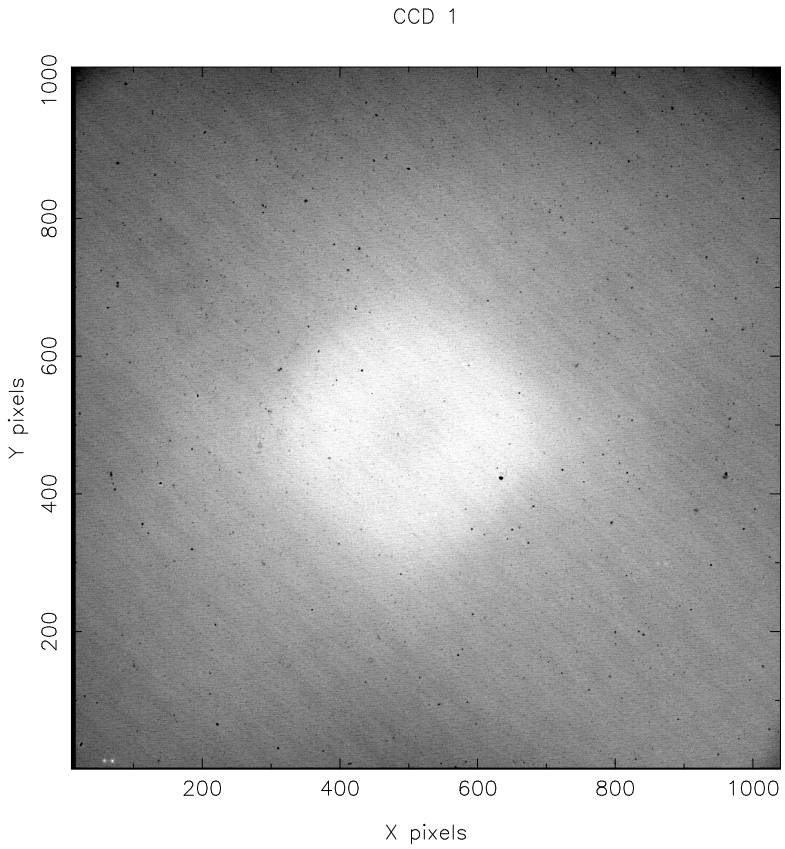

Figure 2.40: Single twilight sky flat in r', showing the bright central region. The variation in brightness from the centre to the edge of the chip is of the order of 10-20%.

the r' filter is shown in Figure 2.40. Figure 2.41 shows a cross-section of the brightness variation across the image. The bright central region in the flat fields is present in all filters, including the less commonly used ones like the 'clear' filter, and the narrow-band filters.

The observed structure is suggestive of scattered light, and hence at first we were worried that the effect might only be present in twilight flat fields. If so, the scattered light would not be present in night time images, and the resulting flat-fielding procedure would not be useful in correcting science images.

To test this, we stacked a deep night-sky image in search of the bright central region, as shown in Figure 2.42. The bright central region is still clearly visible. In Figure 2.43 the same stacked image is shown, but this time divided by the master twilight

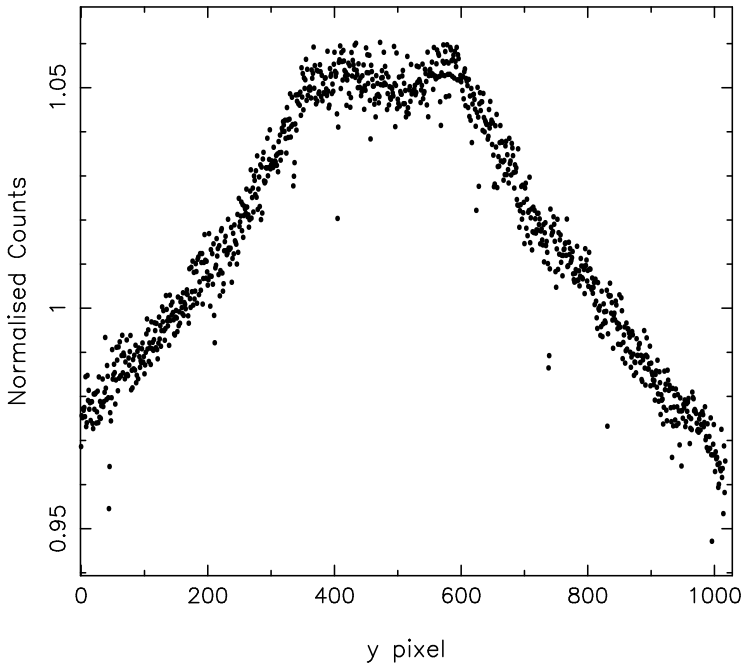

Figure 2.41: Vertical cross-section at column 510 of a normalised, stacked, twilight flat in the r' filter. The behaviour of the bright central region is clearly seen, with a variation of $\sim 10\%$. A diagonal cut-through would show a stronger variation, as the corners of the image suffer from vignetting of the order 10-15%.

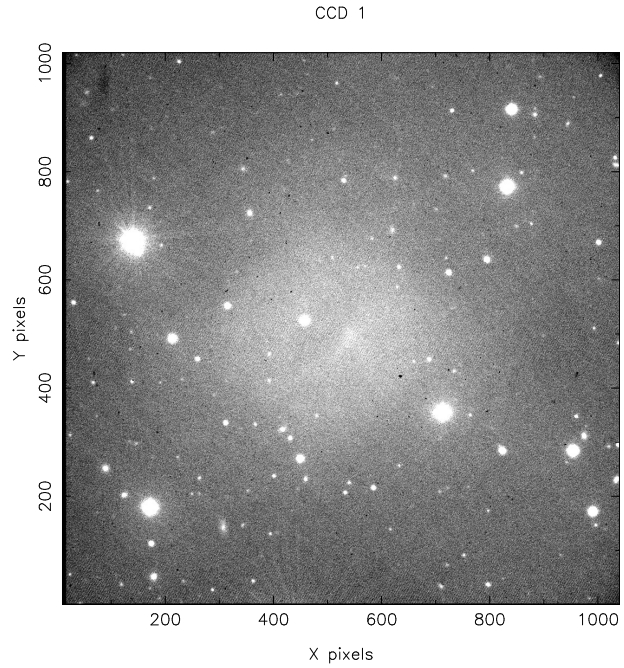

Figure 2.42: Stacked r'-band night-sky image showing that the bright central region persists even at night time.

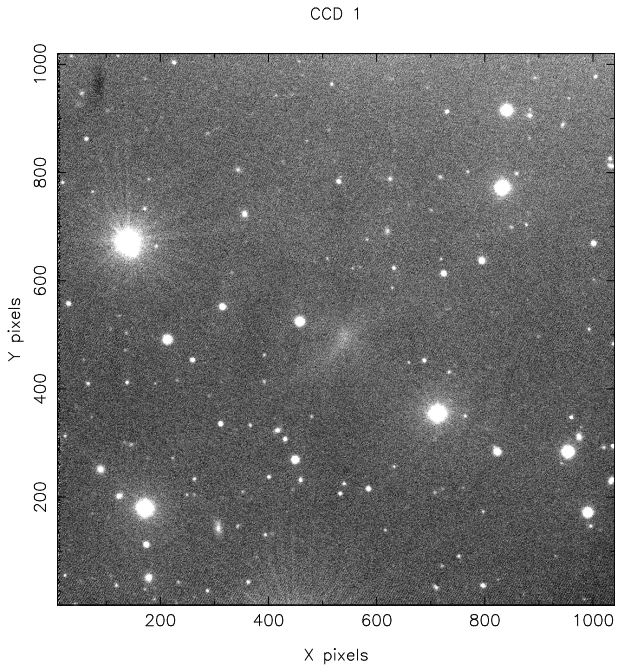

Figure 2.43: Stacked r'-band night-sky image divided by the master flat field. This corrects the bright central region almost perfectly, and the residual bright region at the centre is less than 2% brighter than the rest of the image.

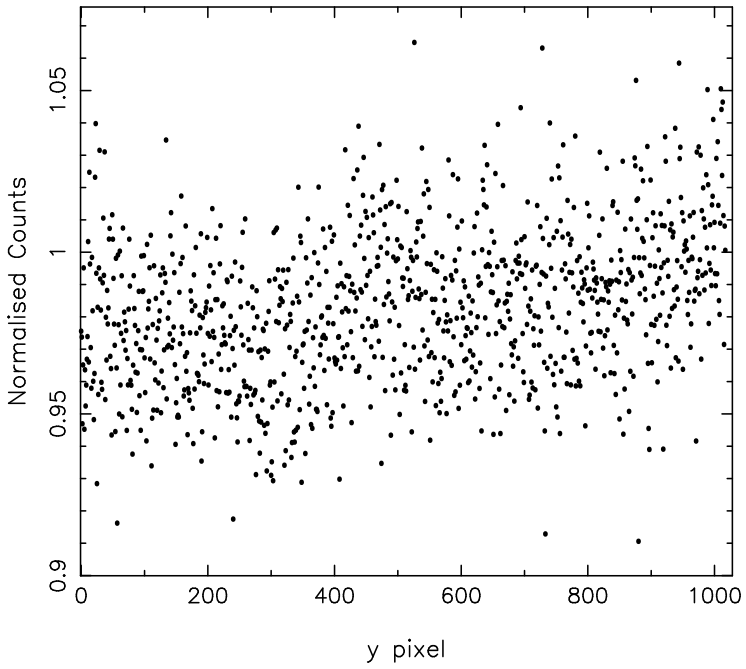

Figure 2.44: Vertical cross-section at column 510 of the normalised, flat-fielded image in the r' filter shown in Figure 2.43. The bright central region has been removed almost entirely, and the image is flat to within 2%, which is less than the intrinsic scatter in the image.

flat. The final image is almost perfectly calibrated, with no strong residual of the bright central region. The resulting image is flat to within 2%, as demonstrated by the cross-section shown in Figure 2.44. This confirms that the bright central pattern is an intrinsic sensitivity function of the telescope and instrument, and this is repeatable and can be calibrated for.

2.4.7 Image Quality

To test the image quality of the instrument, we observed the open cluster NGC 6940, which fills the ULTRASPEC field-of-view with a large number of stars. We imaged the cluster in all five SDSS filters (u', g', r', i', z') to check for any variation in image quality with wavelength. In Figures 2.45 to 2.49 we plot the measured FWHM of stars in the image as a function of radial distance from the centre of the CCD. Each quadrant (top right, bottom right, top left, bottom left) is plotted in a different colour, to investigate the dependence of image quality with field angle.

We find that the image quality is fairly uniform as a function of wavelength, although the u' band images have FWHM around 0.5" larger than the other filters. This is somewhat expected considering the wavelength dependence of atmospheric seeing, although the fractional difference ($\sim$1/3) is larger than would be expected according to the $\lambda^{-1/5}$ dependence (Boyd, 1978). This could be due to extrinsic seeing variations. The other filters all display very similar median FWHM.

As a function of field angle, the image quality tends to deteriorate towards the top left corner in all filters, up to a maximum of 0.4", or $\sim$ 1 pixel, worse than at the centre. This suggests possible warping of the CCD chip in this corner, and in fact this is known to be the case according to the CCD depth measurements taken by the UK Astronomy Technology Centre. The top left corner was claimed to be 0.1mm away from the plane of the other three corners, and a 0.1mm shift is predicted to generate FWHM that are

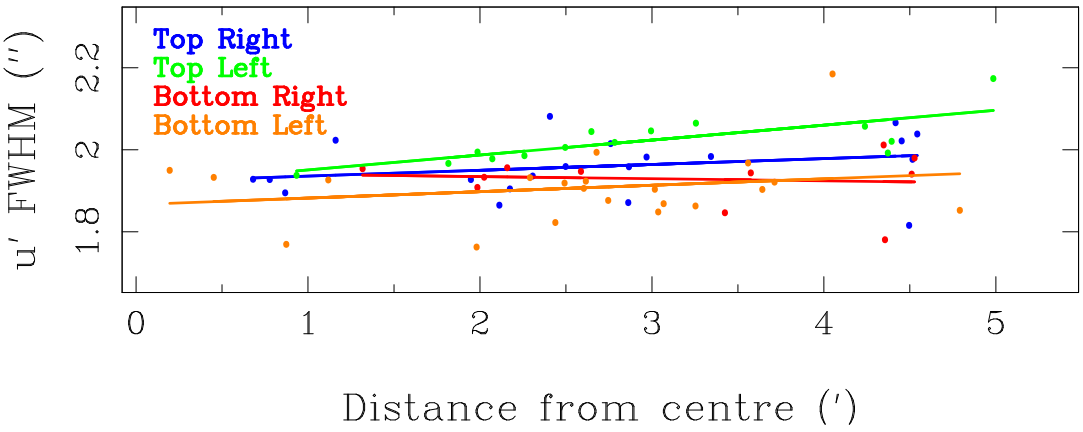

Figure 2.45: Image quality in the u' filter as a function of radial distance from the centre of the ULTRASPEC CCD. Points are individual measurements of stars, and solid lines are linear fits to the points. Different colours describe the FWHM of different quadrants of the field, with the top left quadrant showing the worst image degradation towards the edge of the chip.

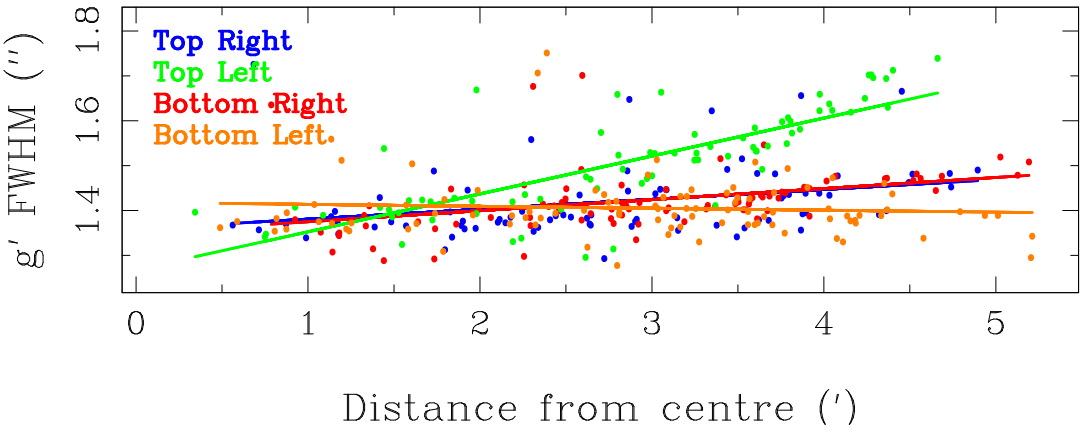

Figure 2.46: Same as Figure 2.45 but for the g' filter.

1 pixel larger. The NARIT optical technician Christophe Buisset confirmed this using a Zemax optical model (Fischer, Tadic-Galeb & Yoder, 2008) of the TNT and ULTRASPEC.

2.4.8 Plate Scale

We measured the plate scale in all five SDSS filters using the images of NGC 6940 described above (Section 2.4.7). Astrometric calibration was performed using `astrometry.net`. The plate scale was measured to be (0.452 ± 0.001)"/pixel in the u', g' and z' bands, and (0.451 ± 0.001)"/pixel in the r' and i' bands. Therefore there is no significant vari-

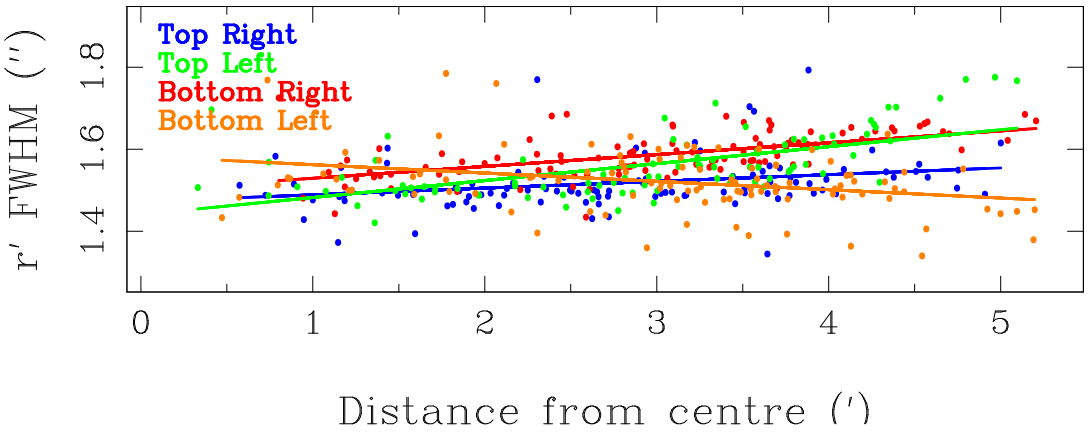

Figure 2.47: Same as Figure 2.45 but for the r' filter.

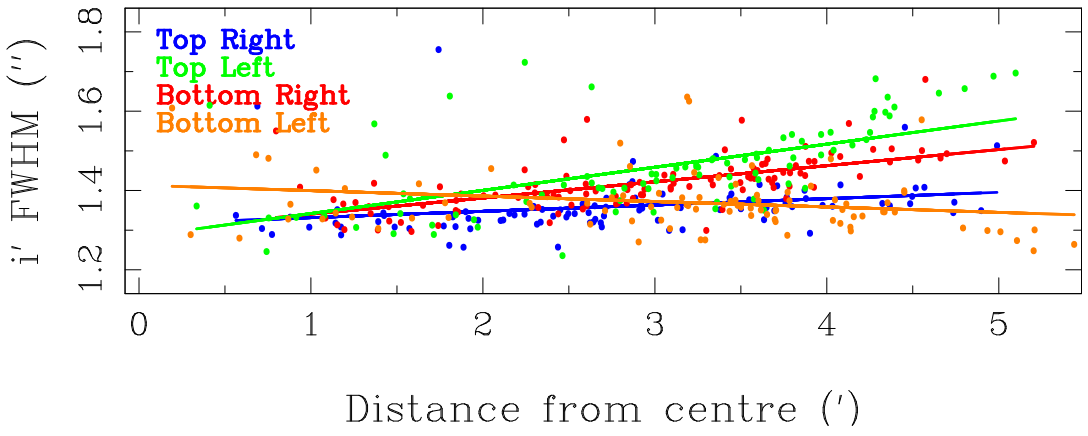

Figure 2.48: Same as Figure 2.45 but for the i' filter.

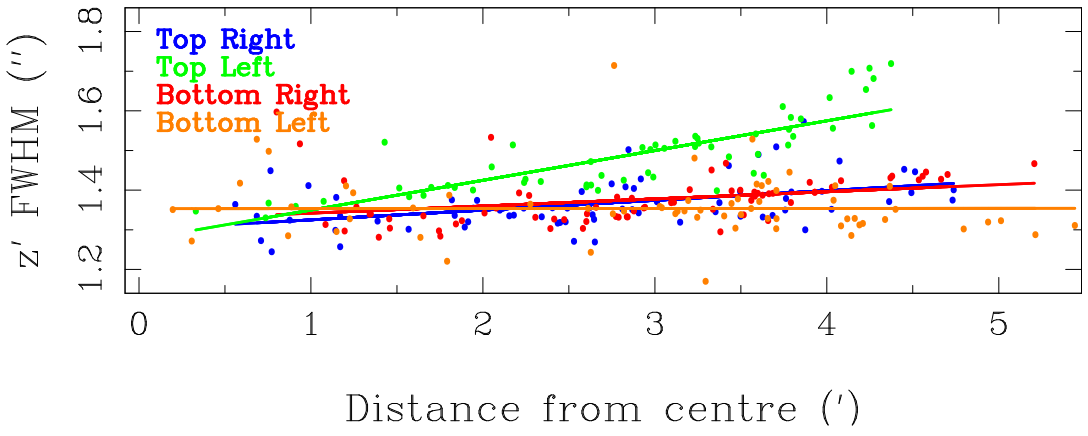

Figure 2.49: Same as Figure 2.45 but for the z' filter.

ability in plate scale with wavelength, as expected. Since the CCD chip has an imaging area of 1024×1024 pixels, the field of view of ULTRASPEC is therefore 463" $\times$ 463", or $7.7' \times 7.7'$.

We also tested for any image shift when cycling through the filters. The centroid position of a bright star near the centre of the CCD was found to be stable to within 2 pixels in all filters, confirming that the filters cause no significant shift in position.

2.4.9 Problems

There have been a small number of problems with ULTRASPEC in its current setup at the TNT. Some were identified immediately during commissioning, whilst others were only spotted later.

Read-out Noise

As discussed in Section 2.3.2, the read-out noise in normal, slow mode in the lab in Sheffield was usually around 2.3 e^-, rising to 2.7 e^- at medium read-out speed, and 4.3 e^- for fast read-out. However, when ULTRASPEC was first installed at the TNT, with the telescope tracking we recorded read-out noise of up to 40 e^-, though usually around 8-12 e^-. This high read-out noise seemed to be much reduced (though still slightly higher than in the lab) when the telescope alt-az motors were turned off, which clearly suggested that there was strong electrical interference being picked up from these motors by ULTRASPEC.

Throughout 2014 NARIT installed deeper and more substantial telescope earthing cables, and attempted more shielding between the telescope and instrument cube. This resulted in the read-out noise dropping to an acceptable level of 3-5 e^- most of the time, although there appeared to be certain angles of the instrument rotator which still produced slightly higher noise levels (5-7 e^-). Since then, David Atkinson has installed special shielded cabling for ULTRASPEC which seems to have brought the noise down

further, to around 2.8 e^- in slow mode. In medium and fast read-out modes, the noise levels often remain high (5-7 e^- in medium, 9-10 e^- in fast) with visible interference patterns. Our current solution is simply to avoid using medium and fast read-outs except when absolutely necessary.

Filter Wheel

Another issue that became apparent in early 2014 was that the filter wheel was, very occasionally, not moving to the correct requested filter. This was only ever noticed when the observed counts in a particular image were significantly higher or lower than expected for the apparent filter. Clearly, this could prove to be a serious issue whenever it goes unnoticed. After extensive testing, the manufacturers of the filter wheel unit (OPTEC) discovered a bug in the firmware version we were using with ULTRASPEC. The bug was thought to be related to the fact that we are using a 6-position wheel when their standard product is a 5-position wheel. With the bug, the wheel controller moves from position 3 to position 1 whenever the move from position 3 to position 5 is requested. After this, all filter positions are incorrectly interpreted as being 2 positions out of place, until the wheel is next initialised. Thankfully, a new version of the firmware was installed on the filter control unit during Autumn 2014, and a spare unit, with the new firmware version pre-installed, was also supplied for emergencies. There have been no further reports about this issue.

Saturation in clear mode

During commissioning, we also saw a serious problem in clear mode when observing in bright/twilight sky. The CCD showed a large saturated area at the bottom of the CCD, while the top area displayed roughly the expected count levels. An example of this behaviour is shown in Figure 2.50. The true nature of this issue was never fully understood, but it is thought to have something to do with reaching the full well capacity

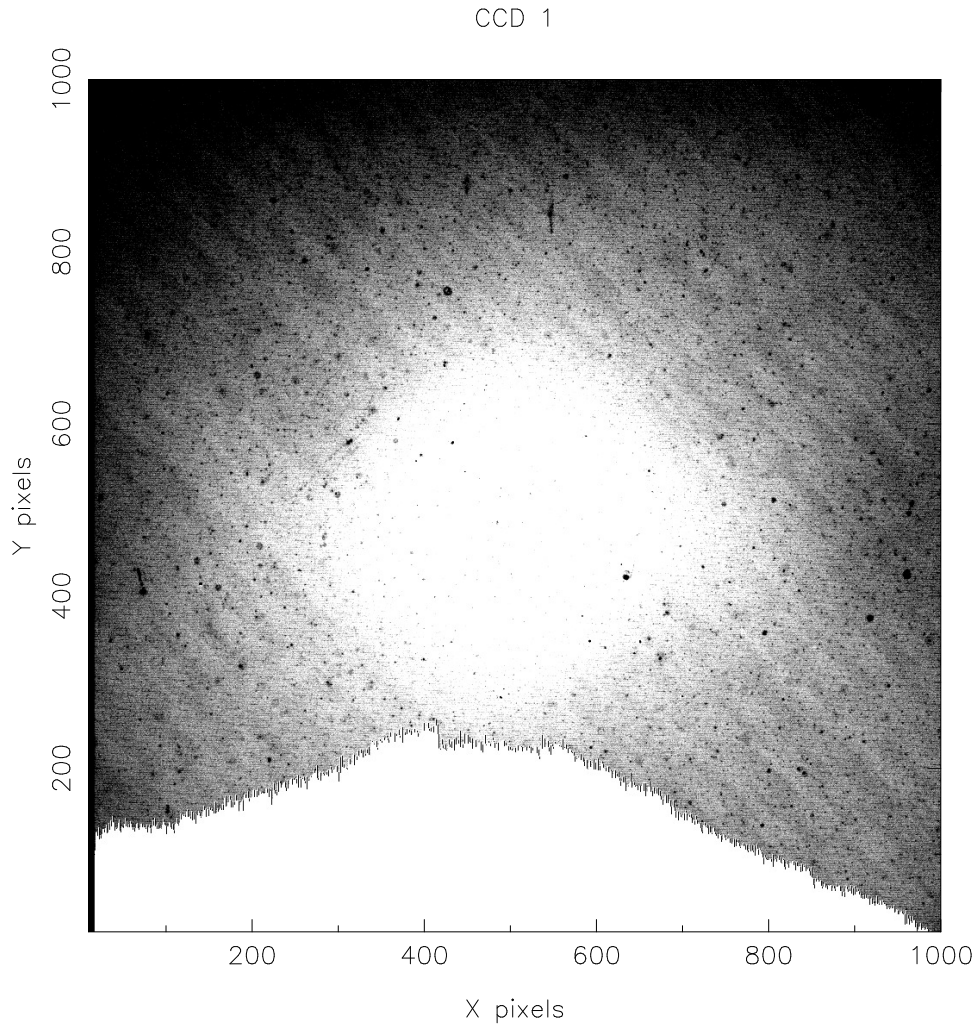

Figure 2.50: Image showing the saturation issue in clear-mode when the chip has high levels of illumination (such as when taking flat fields). The lower region is completely saturated, while the upper region displays the expected count levels. The issue was solved by implementing a double-clear of the CCD.

of either the individual pixels or the charge transfer process (David Atkinson, UKATC, priv.comm.). This problem was solved by introducing a double-clear of the CCD in the low-level controller software. This introduces a 10ms delay when observing in clear-mode, but this should never be an issue because this mode is rarely used for science images and is usually only used for calibration frames.

Time-stamping

A final issue, which could also be disastrous if not spotted, was discovered with the GPS timestamping of frames. It seemed that sometimes a single timestamp would fail to be applied to a given frame (often the second frame in the run, but not always). Instead, this frame would be allocated a null timestamp (sometime in the year 1970), and all subsequent frames would be allocated one timestamp too early, because the time that should have been associated with the original frame is then associated with the following one. It is still unknown why this sometimes occurs, but thankfully it can be easily fixed by manipulating the timestamping information for every frame. A Python script shifts the timestamps and calculates an accurate timestamp for the final frame (which is always missing). Now the end-of-night automated archiving script applies this fix to any data that has suffered from the problem.

2.5 Calibrating the KG5 filter

The following text and data analysis is all my own work, however Martin McAllister conducted some of the data reduction.

In order to obtain better signal-to-noise ratios on faint sources with ULTRASPEC at the TNT, we sometimes use a non-standard broadband filter known as KG5. The KG5 filter essentially encompasses the SDSS u', g' and r' passbands, and allows us to study fainter targets. The throughput of the KG5 and SDSS filters as a function of wavelength are shown in Figure 2.51.

Before proceeding, I must include a note on nomenclature and filter transformations. All filters differ from each other, even those which are designed to cover the same wavelengths and have the same throughputs. For example, the ULTRASPEC filter in the g-band, designated g' as it comes from a slightly different SDSS filter set, will not exactly match the original SDSS g filter used in the Sloan survey. Sometimes the

response of the filter in use is so different from the original catalogue filter response that system transformations are needed to achieve sufficient accuracy when comparing measured magnitudes with catalogue magnitudes. This is especially the case for any high precision photometric studies, where the transformations (usually < 0.2 magnitudes) are larger than any statistical uncertainties (e.g. Bell et al. 2012). In this section I refer to ULTRASPEC filters and magnitudes with primes (e.g. g') and the original SDSS filters and magnitudes from the SDSS catalogue without primes (e.g. g).

2.5.1 Calibrating standard filters

As discussed above in Section 2.4.1, instrumental zeropoints are useful when calibrating the fluxes of observations. As before the zeropoint here is considered to be the above-atmosphere magnitude of a star from which we would detect one photon per second. The zeropoint ought to be stable from night to night, but there may be small changes in mirror reflectivity, lens throughput, and CCD quantum efficiency with time. As mentioned in Section 2.4.1, we know there may be much larger changes in the atmospheric extinction from night to night, which is why we try to observe at least one standard star every night, so that we can, if necessary, calculate a nightly zeropoint.

In Section 2.4.1 I explained that a zeropoint can be measured by observing a standard star and comparing the measured photon counts with the catalogue magnitude, as described by Equation 2.11. In general Equation 2.11 is sufficient, however in the case of the KG5 filter, we do not have well documented or catalogued magnitudes for standard stars. The filter is not part of any comprehensive filter system, and therefore we cannot simply compare the flux from a particular object with the magnitude of a standard or comparison star.

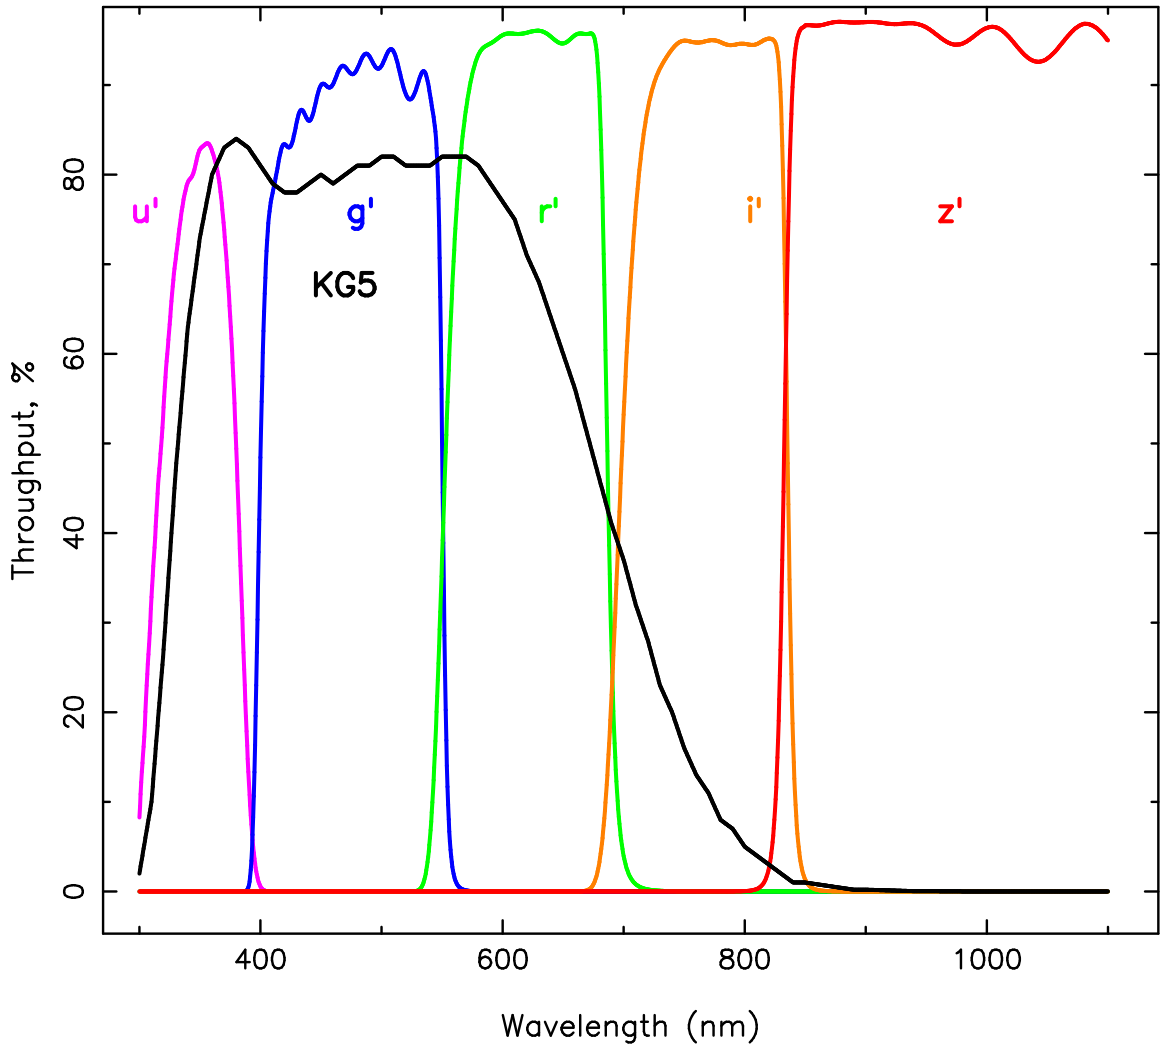

Figure 2.51: Transmission as a function of wavelength for the five SDSS filters and the KG5 filter used in ULTRASPEC.

2.5.2 Theoretical Zeropoints

Our solution to the problem of calibrating KG5 observations is to produce a theoretical zeropoint, based on accurate modelling of the throughput of the entire telescope and instrument system, just as I have done in Section 2.4.1. Recall that one can pass the AB zero magnitude flux of 3631×10^{-23} erg s^{-1}cm^{-2}Hz^{-1} through a model of the instrument optics to predict what flux will be measured by the detector at the end of the light path. The zeropoint magnitude is then produced easily with Equation 2.10.

The problem with this method is that in practice no stars emit perfectly white light.

Instead most show roughly a blackbody spectrum, with numerous absorption lines, with the position of the peak of the blackbody depending strongly on the effective temperature of the star. This of course affects the colour of the star, and thus its magnitude in different filters.

To work around this we use a spectrum for the standard star in question, and fold this through the throughput model instead of the uniform white light (3631×10^{-23} erg $s^{-1}cm^{-2}Hz^{-1}$). Unfortunately, most photometric standard stars have not been studied spectroscopically, and in any case we require a method for calibrating the KG5 filter using any standard star, without having to use the star's spectrum every time we need a KG5 magnitude.

Instead of observed stellar spectra we can use model spectra from spectral atlases based on the effective temperature of the star (e.g. Gunn & Stryker 1983; Jacoby, Hunter & Christian 1984; Pickles 1998; Castelli & Kurucz 2004). However, because we rarely know the star's distance and radius, we cannot simply convert from the flux at the stellar surface (supplied by model atmospheric spectra) to the observed flux at Earth.

2.5.3 Bolometric Corrections

The solution to this problem involves using bolometric corrections to calculate a theoretical KG5 magnitude of any star. Then, Equation 2.11 can be used directly to calculate a zeropoint.

Bolometric corrections (BCs) are the difference between the magnitude of a particular star in a particular filter, and the bolometric (all wavelengths) magnitude of that star, such that $M_{bol} = m_g + BC_g$ for the original SDSS g filter for example. The corrections remove the need to know the distance to and radius of the star. If one knows the BC for a standard star in both the SDSS g filter and the ULTRASPEC KG5 filter, those corrections can be used along with the SDSS catalogue g-band magnitude to predict the

KG5 magnitude, as shown in Equation 2.12.

$$m_{KG5} + BC_{KG5} = m_g + BC_g. \tag{2.12}$$

The formula for calculating a bolometric correction for the g filter, in the AB magnitude system, is shown in Equation 2.13 and follows from the derivations in Bell et al. (2012).

$$\begin{aligned} BC_g = M_{bol\odot} - 2.5\log_{10}\left(\frac{4\pi(10pc)^2 F_{bol}}{L_\odot}\right) \\ +2.5\log_{10}\left(\frac{\int_g \lambda F_\lambda 10^{-0.4A_\lambda} R_\lambda d\lambda}{\int_g \frac{c}{\lambda} f_\nu^0 R_\lambda d\lambda}\right), \end{aligned} \tag{2.13}$$

where $F_{bol} = \sigma T_{eff}^4$ by the Stefan-Boltzmann law, F_λ is the flux at the stellar surface according to the model atmosphere spectrum, A_λ is the appropriate galactic extinction curve, and R_λ is the instrument throughput model as a function of wavelength. All functions of wavelength are in units of Å. We used the following values and units for the physical constants: $M_{bol\odot} = 4.74$, 1 parsec $= 3.0857\times10^{16}$m, $\sigma = 5.670373\times10^{-8}$W m^{-2} K^{-4}, $L_\odot = 3.846\times10^{26}$ W, $c = 2.998\times10^{-18}$ Å s^{-1} and $f_\nu^0 = 3631\times10^{-23}$erg s^{-1} cm^{-2} Hz^{-1}.

The challenge now is to produce an accurate throughput model for the telescope, instrument and detector, as a function of wavelength. Thankfully we have already done this in order to quantify the performance of the instrument during commissioning (see Section 2.4.1). Our model for the system included the throughputs of the atmosphere, telescope mirrors, instrument lenses and adhesives, filter throughputs, and the quantum efficiency of the detector. The largest uncertainties were expected to come primarily from the atmospheric extinction and the mirror reflectivity.

As discussed in Section 2.4.1, we found agreement between the theoretical and the measured zeropoints to within 0.3 mag, but with better agreement in u' and r' than in

g', i' and z'. This leads us to consider that the throughput model might be incorrect as a function of wavelength.

To test this thoroughly, we followed the procedure of Bell et al. (2012) to calculate the difference between theoretical and observed magnitudes (dm) for stars in the open cluster NGC 6940, in all 5 SDSS filters. Stars with different temperatures and colours will have spectra that peak at different wavelengths, and will thus be differently affected by the throughput model. Assuming any differences in the intrinsic filter transformation are small (i.e. from the original SDSS catalogue g-band magnitudes to the observed ULTRASPEC g'-band), any inaccuracy in the throughput model as a function of wavelength should induce non-linearity in a plot of dm against colour $(g - r)$.

If non-linearity was to be discovered in the dm versus colour plots for any or all filters, we could then incorporate a polynomial modifier to the throughput model, allowing us to minimise this non-linearity, and ultimately arrive at a more accurate throughput model. However, as our tests described below suggest, the throughput model is in fact correct, at least to within the limits imposed by our data for NGC 6940, and a polynomial modifier is not needed. The following sections describe this process, and the use of these same bolometric corrections to produce a theoretical KG5 magnitude relation.

2.5.4 Step One: Theoretical Bolometric Corrections

The first step was to build a table of BCs for theoretical stars with a range of colours, so that we can apply the correct BCs to each actual star in the cluster using linear interpolation. We used the Dartmouth isochrones database (Dotter et al., 2008) to build a list of temperature/surface gravity (T_{eff}/log g) pairs for theoretical stars with a range of masses from $0.14M_{\odot}$ to $5.0M_{\odot}$. We set the age and metallicity of the isochrone to match that of the cluster, with the log age being 8.858 and the metallicity being Fe/H = 0.01 (Wu et al., 2009). Figure 2.52 shows the T_{eff}/log g pairs derived from this isochrone.

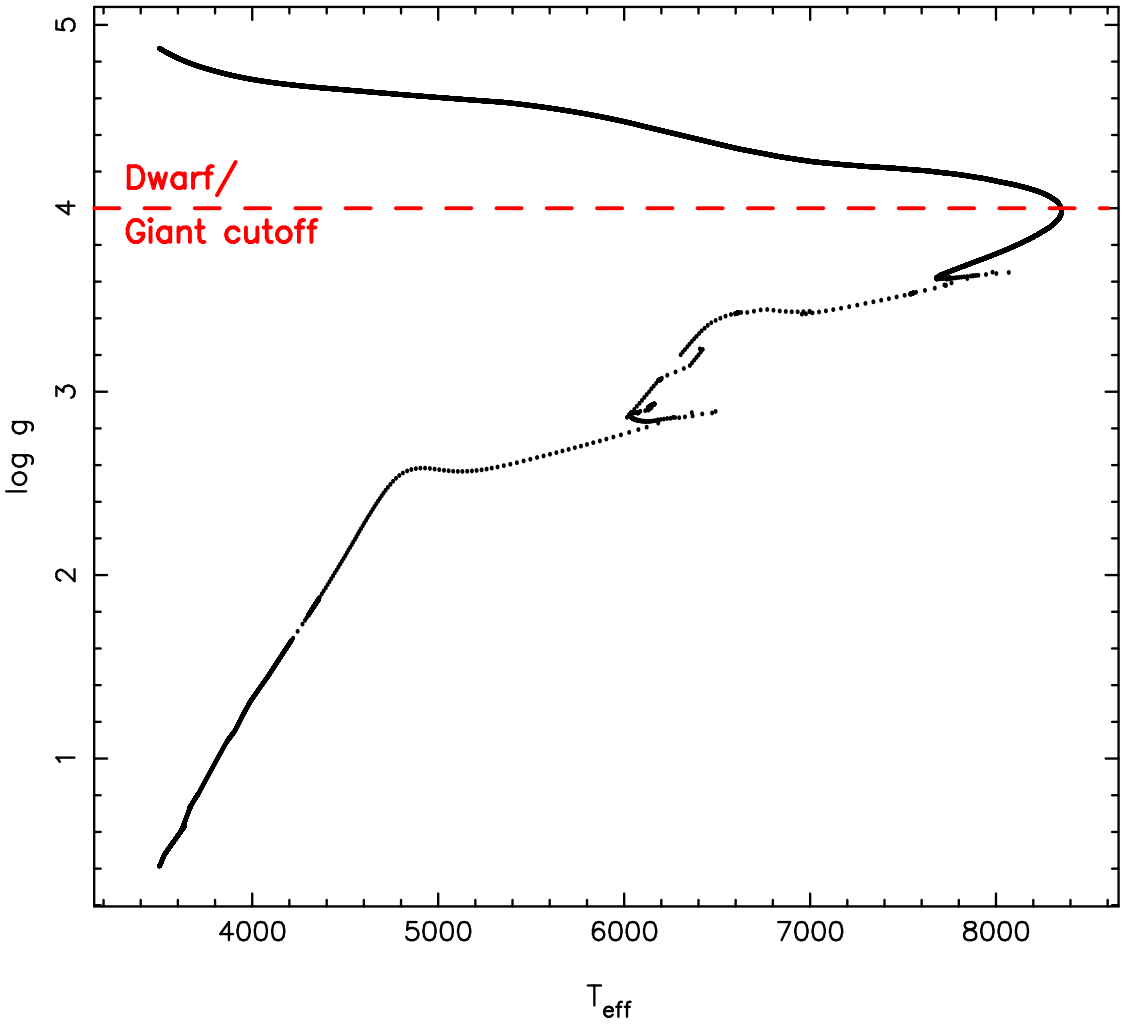

Figure 2.52: Theoretical log g/T_{eff} pairs for the open cluster NGC 6940. We select only the main sequence stars by making a cut at log $g = 4$.

Since the model spectra available are generally constrained to main sequence stars, we rejected those pairs with log $g < 4$, as these values are usually only found in evolved stars.

For each theoretical star in the isochrone we then selected a model atmosphere spectrum from the atlas produced by Castelli & Kurucz (2004). An example spectrum is shown in Figure 2.53. At this stage we incorporated a galactic extinction model, and a reddening measurement of E(B-V) = 0.214 for NGC 6940, into the model spectrum. This allowed us to ignore the A_λ term in Equation 2.13.

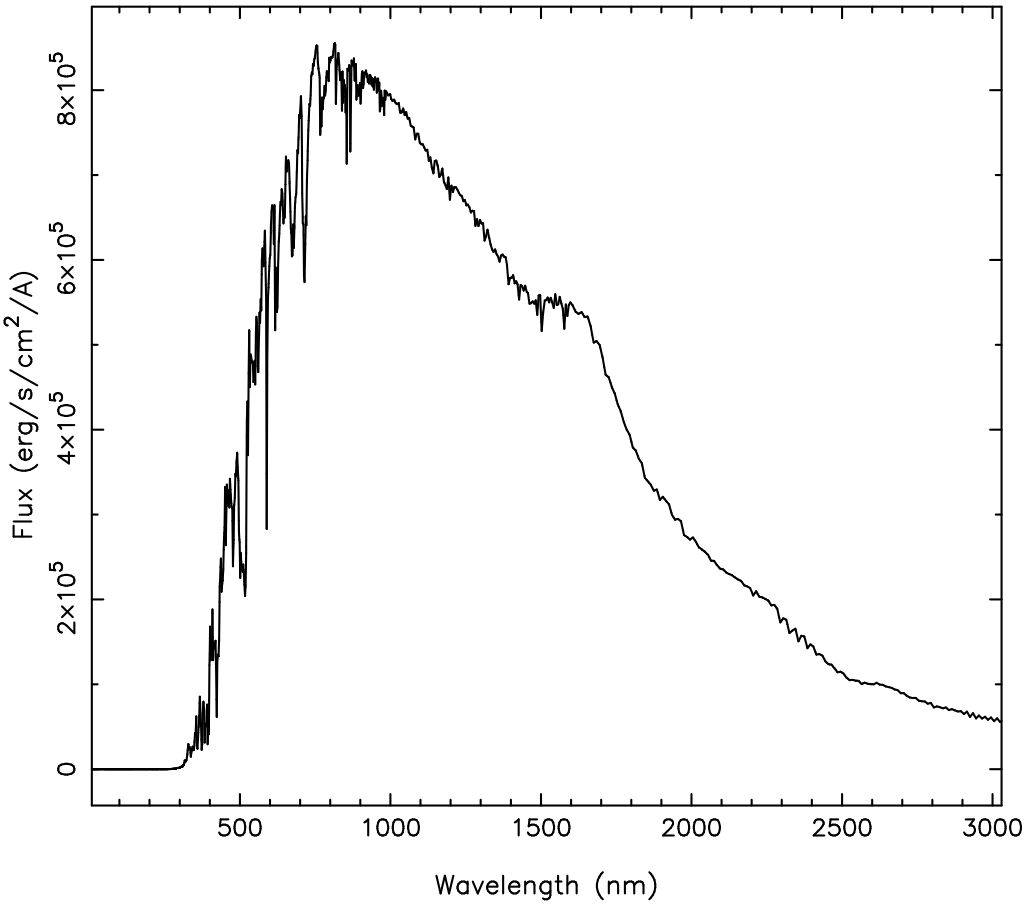

Figure 2.53: An example spectrum of a star with effective temperature 4000K, Fe/H of 0.01 and log g of 4.65, extracted from the atlas of Castelli & Kurucz (2004).

The model atmosphere spectra were then used, along with the throughput model as a function of wavelength, to calculate BCs according to Equation 2.13. BCs were generated for each theoretical star in the isochrone using both the ULTRASPEC throughput model and the SDSS throughput model[3] separately. Using the nomenclature of g' being the ULTRASPEC g'-filter, and g being the original from SDSS, we can use the BCs to state that $g + BC_g = g' + BC_{g'}$. Rearranging gives $BC_g - BC_{g'} = g' - g = \mathrm{d}g$, where $\mathrm{d}g$ is the theoretical transformation between ULTRASPEC g' and SDSS g magnitudes, for the star in question.

The theoretical transformation between ULTRASPEC and SDSS magnitudes is cal-

[3]SDSS throughput curves for each filter were obtained from `http://classic.sdss.org/dr7/instruments/imager/`

Table 2.17: Example transformations between the SDSS and ULTRASPEC photometric systems for a sample of theoretical stars in the open cluster NGC 6940, generated from a Dartmouth isochrone with log age 8.858, metallicity Fe/H = 0.01 and reddening E(B-V) of 0.214. For each filter M, the transformation dM is defined as $BC_{\rm SDSS} - BC_{\rm ULTRASPEC}$.

T_{eff}	log g	du	dg	dr	di	dz
4100	4.135	-0.229	-0.175	-0.012	-0.026	+0.001
5720	4.255	-0.190	-0.087	-0.006	-0.008	+0.002
6799	4.478	-0.160	-0.053	-0.004	-0.003	+0.002
7520	4.601	-0.140	-0.038	-0.002	+0.001	+0.004
8344	4.847	-0.108	-0.022	-0.000	+0.004	+0.012

culated for each of the five filters, u, g, r, i and z, and for each star in the isochrone. We use the abbreviation dM to represent this transformation for any unspecified filter. We now have a table giving the transformations dM in each filter, for a range of stellar effective temperatures. We show some examples of these values in Table 2.17.

2.5.5 Step Two: Check Measured vs Calculated Magnitudes

Next we used real observations with ULTRASPEC to test the accuracy of the transformations between the filter systems. In essence we were testing how well the throughput model can reproduce the observed magnitudes. We used data from the commissioning nights of ULTRASPEC, where we observed the cluster NGC 6940 in u', g', r', i' and z' on the night of 5th November 2013, at an airmass of 1.03. We extracted RA and DEC coordinates, as well as instrumental magnitudes, for as many stars as possible using the Source Extractor program (Bertin & Arnouts, 1996) . For each filter, we then added an extinction correction $k \times A$ where A is the airmass and k is the extinction coefficient taken to be 0.50, 0.11, 0.09, 0.05 and 0.04 for u', g', r', i' and z' respectively.

The star detections were then matched with the SDSS DR9 catalogue (Ahn et al., 2012) such that we could compare the original SDSS magnitudes in the five filters to the instrumental magnitudes measured by ULTRASPEC. We began with between 200 and 600 stars in each filter dataset (least in u', most in g' and r'), but many of these may

have erroneous SDSS magnitudes, or erroneous instrumental magnitudes, as some stars were affected by blending with other nearby sources. To clean our data, we processed each filter dataset as follows:

- Remove any sources flagged by SExtractor as potentially saturated, or with de-blending measures that are overly complex.
- Plot the instrumental magnitude against the SDSS magnitude and remove any sources that are deemed, by eye, to be erroneous (more than $\sim 2\sigma$ from the trend-line). This should remove variable sources and blends.
- Remove any sources that clearly lie away from the trend curve in the plot of instrumental magnitude against uncertainty in instrumental magnitude.
- Conduct a χ^2 analysis for each star. Use the SDSS $g-r$ colours to calculate the predicted magnitudes in other filters using BCs, and compare these with the actual catalogue magnitudes. Remove stars which are measured as very different to their predicted magnitudes ($> 3\sigma$).

With the source list for each filter cleared of outliers, we then calculated the difference between the theoretical magnitude and the measured magnitude (dm) for each star, in each filter catalogue. We took the g-r colours from the SDSS catalogue and used them to calculate effective temperatures for the stars. Equation 2.14 shows the relation between colour and temperature, and is valid only between 3850K and 8500K (Fukugita et al., 2011).

$$T_{eff} \pm 90K = \frac{1.09 \times 10^4}{1.47 + g - r}. \tag{2.14}$$

We rejected any stars with colours so extreme that they do not follow this relation. For those stars with colours which did fall within the valid range, we calculated the

transformation between the SDSS and ULTRASPEC filter systems by linear interpolation with the nearest dM-T_{eff} pairs produced in Section 2.5.4. Then, we added the filter transformation, dM, to the SDSS catalogue magnitude for that star, m, giving us a theoretical ULTRASPEC magnitude for the star in the chosen filter (i.e. $m'_{theo} = m+\mathrm{d}M$).

This theoretical ULTRASPEC magnitude was then compared directly with the measured instrumental (above-atmosphere) ULTRASPEC magnitude to calculate dm - the difference between theoretical and measured magnitudes (i.e. $\mathrm{d}m = m'_{theo} - m'_{meas}$). This process produced a catalogue of stars in the cluster which have, for each filter, a distribution of dm against their SDSS colour, g-r. We plot this in the same form as Bell et al. (2012) for all five filters in Figures 2.54 to 2.58.

If the throughput response model was correct at all wavelengths, Figures 2.54 to 2.58 would all show flat distributions centred on zero. Inspection of these plots suggests that our data are not sufficient to constrain the behaviour of the throughput model. The spread of colours sampled is small, and the u'-band data in particular is very sparse. Given the scatter in the data we have, we cannot investigate the accuracy of the model to better than ± 0.05 magnitudes across a range of 1 magnitude in colour. Nevertheless, to this degree of accuracy, the throughput model appears to be correct, with no strong trends or variance of the transformations as a function of colour.

2.5.6 Step Three: Correct the Throughput Curves

If the throughput model was incorrect, the filter transformations would also be incorrect, and the dm vs colour plots in Figures 2.54 to 2.58 would show some correlation or trend as a function of colour. Figure 7 of Bell et al. (2012) shows this effect.

We might have expected to see similar behaviour to Bell et al. (2012) using our throughput model for ULTRASPEC if we had better data that explored a wider range of stellar colours. If we had seen a trend of dm with colour indicating problems with the

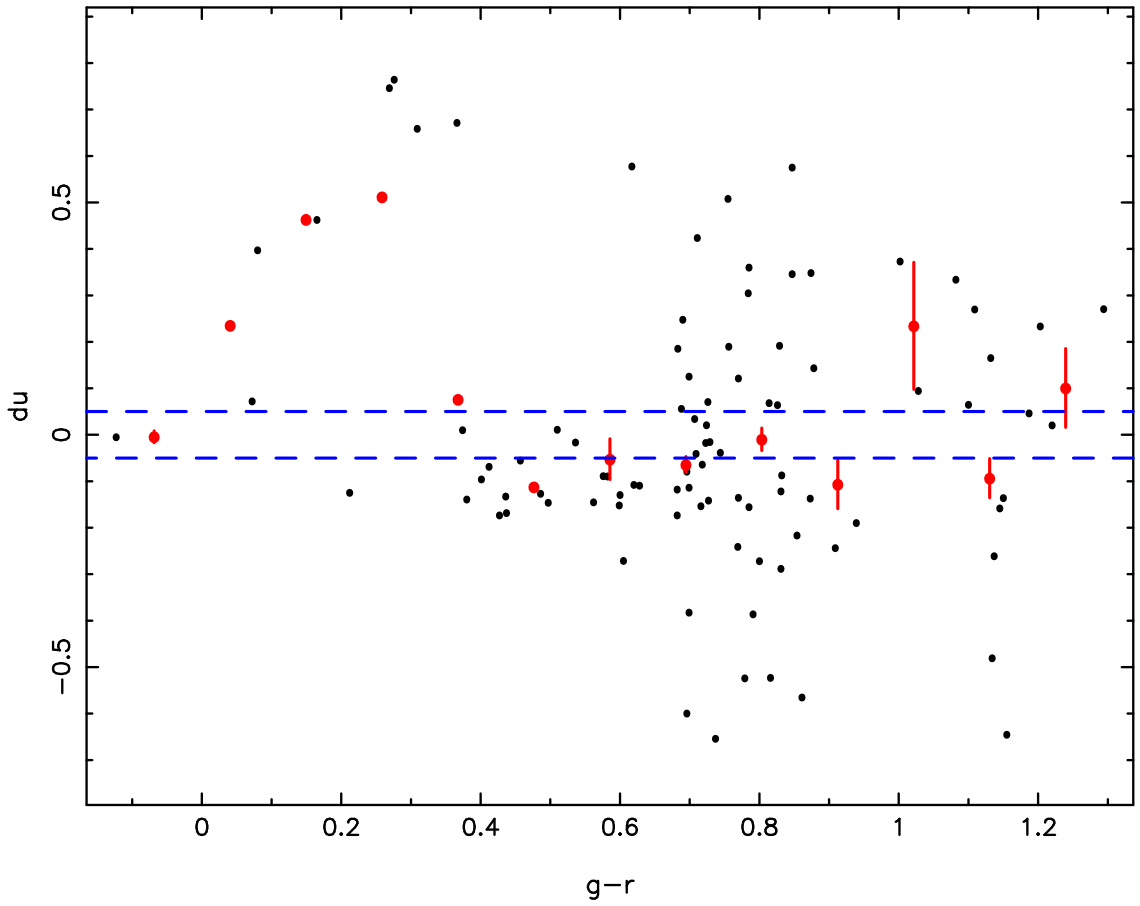

Figure 2.54: Plotting dm against g-r, the difference between the theoretical and measured ULTRASPEC u'-band magnitudes of stars in NGC 6940, as a function of their colour. Individual stars are represented by black points, whilst red points show the mean transform magnitudes of 13 equally spaced bins in colour. The error bars on the binned points are the maximum photometric uncertainty within each bin, divided by the number of points in the bin. The blue dashed lines show the ±0.05 magnitude level with respect to zero.

throughput model, we could attempt to correct for these inaccuracies by incorporating a polynomial correction into the model. Mathematically, the true throughput response function R^*_λ can be thought of as $R^*_\lambda = P_\lambda R_\lambda$, where R_λ is the existing theoretical throughput response function and P_λ is a polynomial function of wavelength.

By computationally minimising the slope, scatter or smoothed variability of the dm vs colour plots, one could arrive at the best values for the polynomial coefficients. These values would describe the appropriate polynomial to modify the original throughput

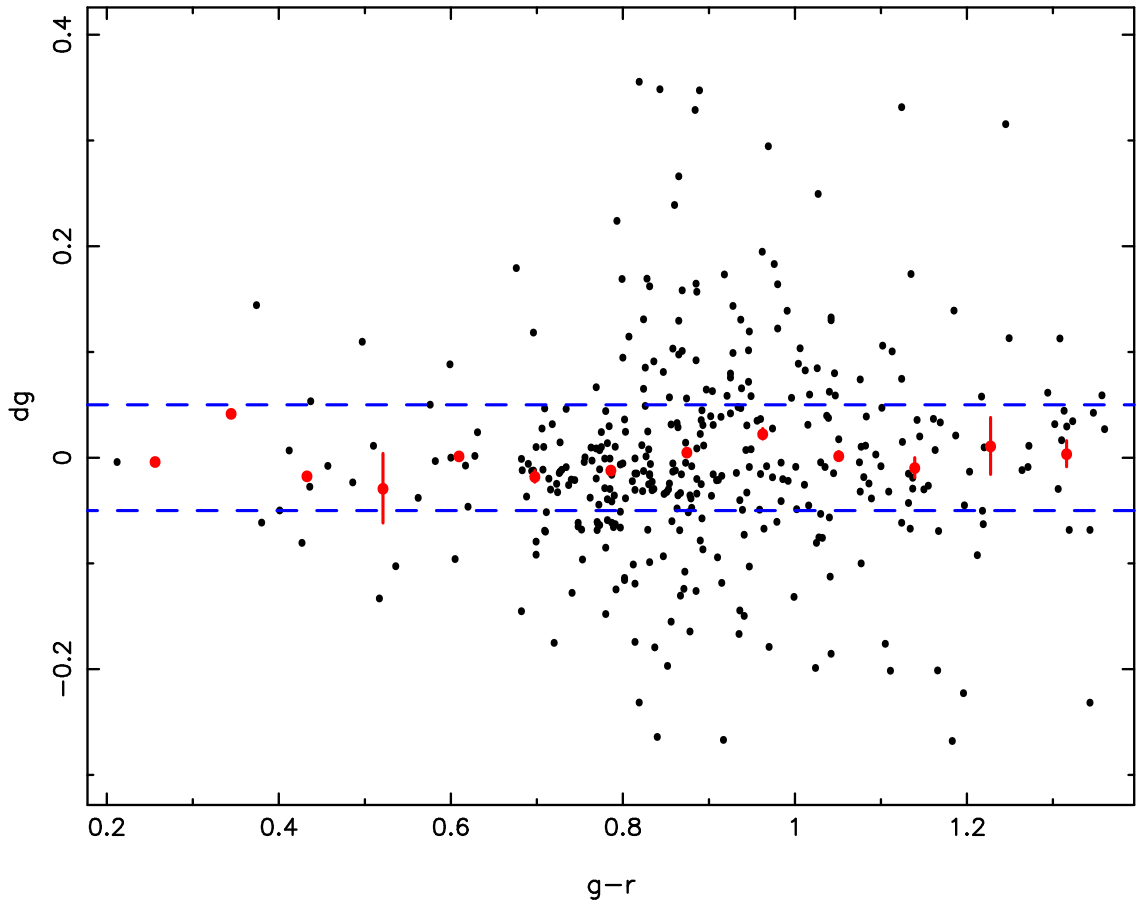

Figure 2.55: Same as Figure 2.54 but for the g' filter.

function such that it best describes the true throughput of the system. However, given that our dm vs colour plots are fairly flat, to within the poor constraints of the data, we chose not to enforce a polynomial modification at this stage, and assumed our throughput model is correct.

Nevertheless, this concept also served as a test for how strongly we could constrain the accuracy of the model with our current data. By modifying the throughput model with a simple linear function (i.e. a polynomial of the form $y = ax + b$), we reshaped the dm vs colour plot. We derived a crude estimation of the limits of the data by folding the throughput model with linear functions of extreme positive and negative gradient. Those functions which cause the dm vs colour plot to become obviously non-flat were then considered the limits of the model. A measurement of the range of dm of

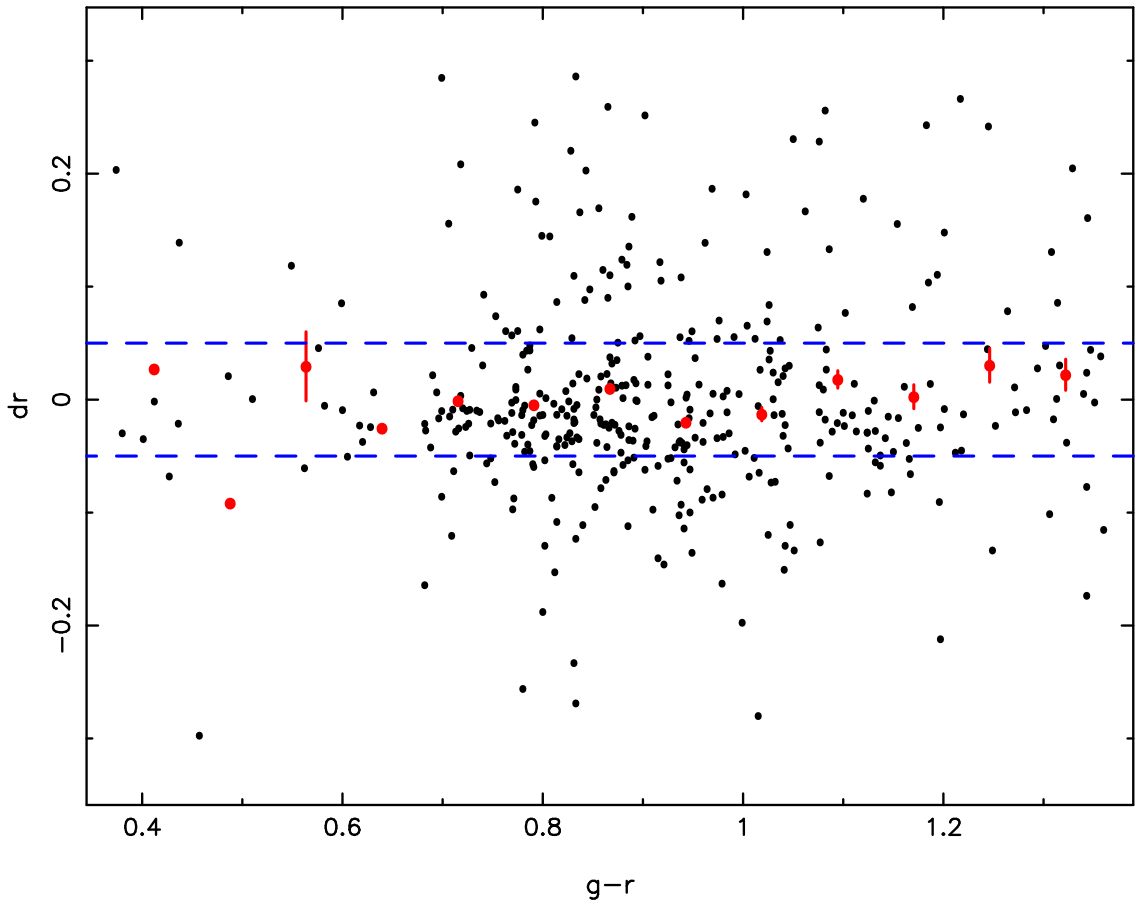

Figure 2.56: Same as Figure 2.54 but for the r' filter.

a typical star at these two extremes gave an estimation of the systematic error in any filter transformations calculated using the existing throughput model.

We investigated the systematic error in this way for each filter, centering the linear modifiers on the central wavelength of the filter, and scaling it to the wavelength range of our original throughput model (3260Å to 10500Å). To quantify the limit of the plot becoming 'non-flat', we overplotted the slope of the distribution, and set the limit of the modifying function gradient to be when this slope crosses the horizontal lines at $dm = \pm 0.05$ mags. An example of this situation is shown in Figure 2.59. Overall we estimate the systematic error to be approximately 0.2 magnitudes in all filters, although we reiterate that this estimation is rather crude.

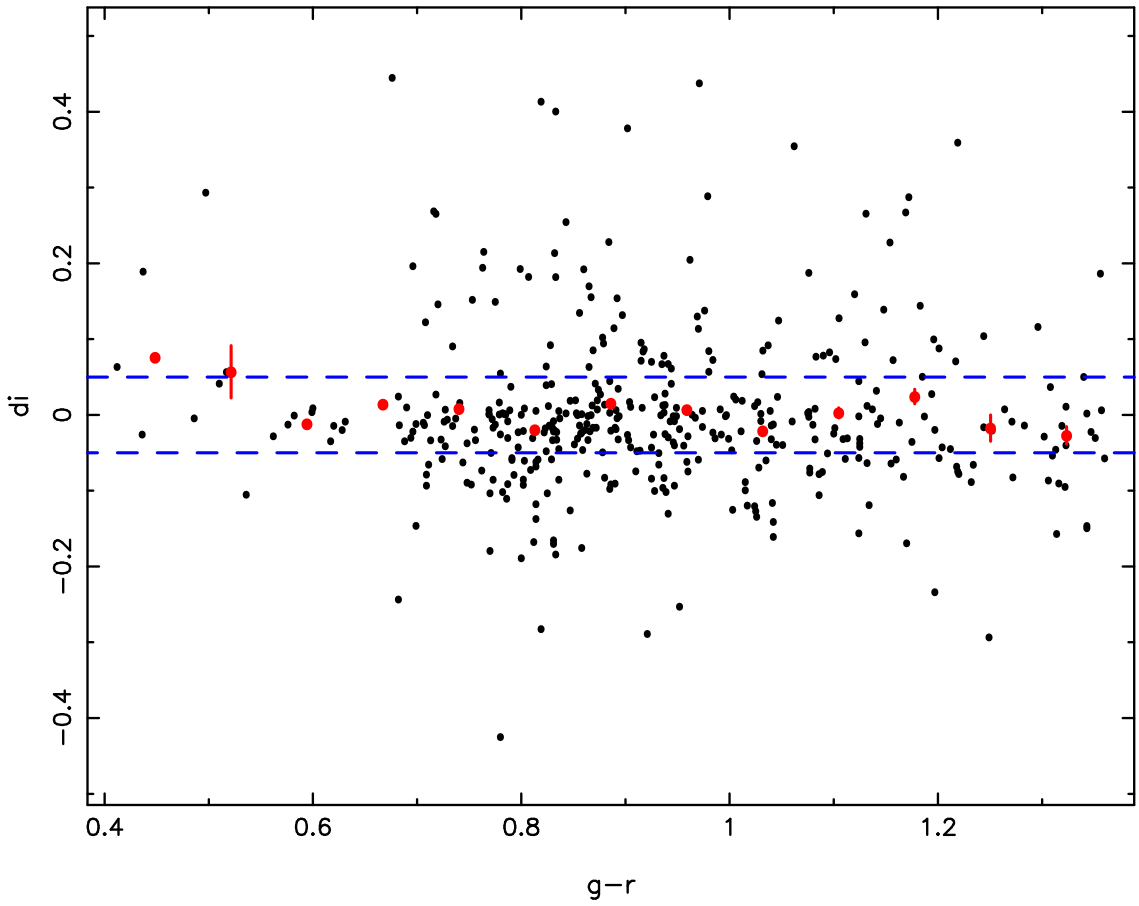

Figure 2.57: Same as Figure 2.54 but for the i' filter.

2.5.7 Step Four: Calculate Accurate Bolometric Corrections for KG5

Accepting that for now our throughput model is correct and cannot be easily improved, we then used the throughput model to calculate theoretical KG5 magnitudes for any given star.

We used the throughput model to calculate BCs for the ULTRASPEC KG5 filter and the SDSS g filter, thus allowing us to predict a KG5 magnitude for any main sequence star. We used the bolometric properties (Equation 2.12) along with the catalogue SDSS magnitude in g, to arrive at a theoretical magnitude in KG5. This magnitude can then be compared with the measured flux of a standard star to arrive at a zeropoint for KG5

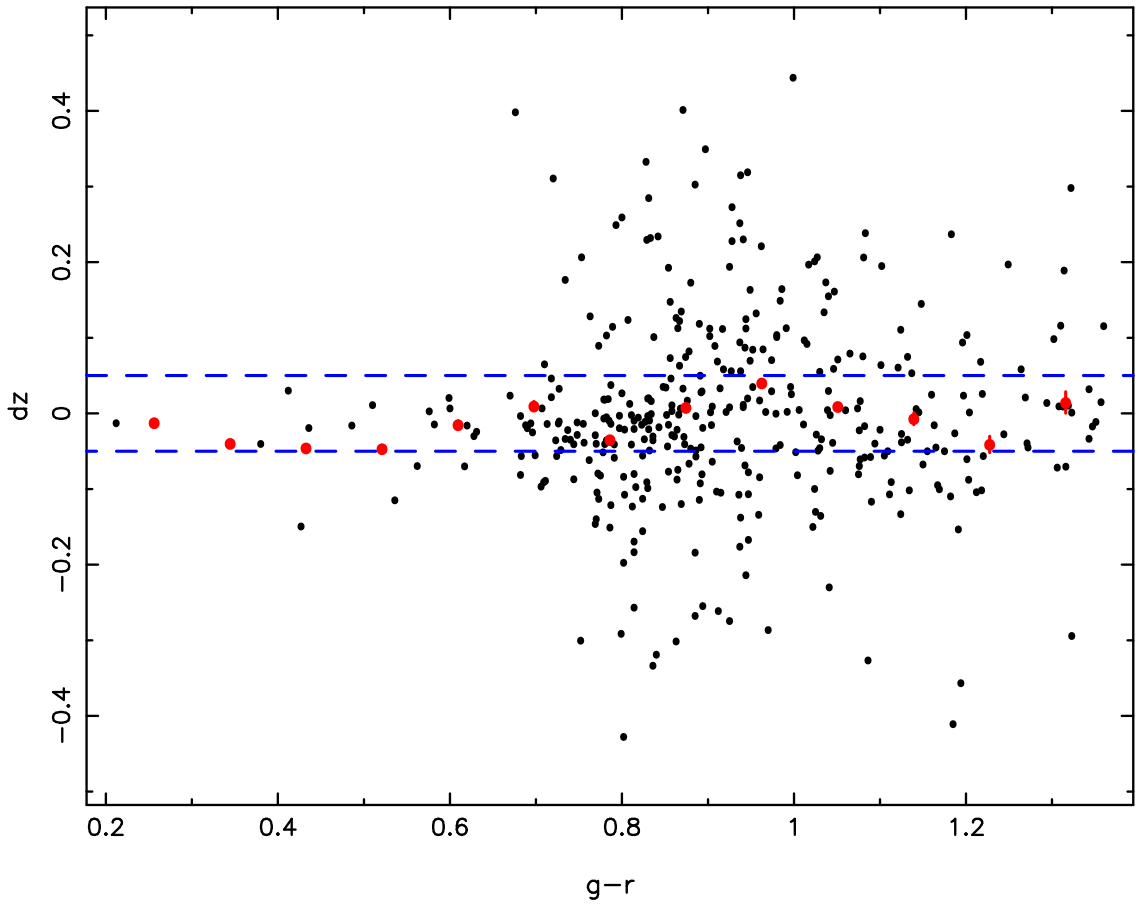

Figure 2.58: Same as Figure 2.54 but for the z' filter.

using Equation 2.11.

To calculate the BCs, we first adopted sensible values for the age, metallicity and reddening for the expected standard stars. Since all standard stars are bright and therefore are likely to be nearby, we assumed there is no significant reddening and let E(B-V)=0. For metallicity we assumed solar values, Fe/H = 0, but noted that altering the metallicity makes very little difference to the calculated BCs (setting Fe/H = 0.1 altered the BCs by 0.01 mags). Most standard stars lie within the galactic disk, and could have ages between 0 and 10 billion years. For simplicity we use a fiducial age of 2 Gyr. However, because some standards are hot (T_{eff} > 8000K) and must therefore be younger, we investigated two scenarios - one with an age of 2 Gyr (log age = 9.301), and another with an age of 500 Myr (log age = 8.699), to allow a broader range of standard stars to

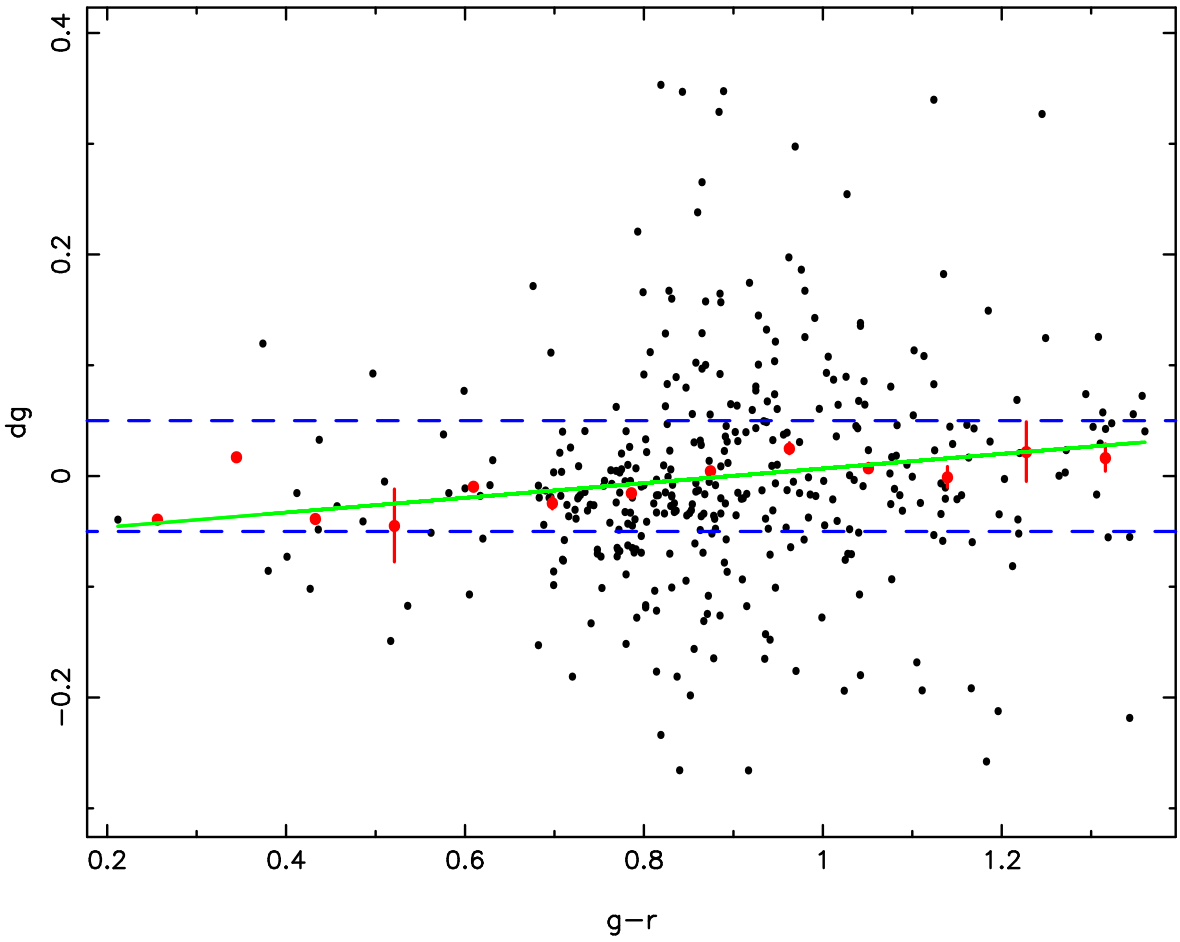

Figure 2.59: Similar to Figure 2.55 but with the throughput model modified by a linear function with a gradient of -3. The green line shows the fitted trendline to the distribution of magnitude differences. The trendline just crosses the blue dashed lines of ±0.05 mags, and we therefore consider this the boundary beyond which the distribution would no longer be flat. This represents one extreme of the range of linear transformation corrections which could still suggest that the throughput model is observed to be correct, given the data we have available.

be used in this calibration. Again as above in Section 2.5.4, we investigated only those stars with log g greater than 4, and in this way only considered main sequence stars. It is therefore important that the tables below (2.18 and 2.19) are only used in conjunction with observations of main sequence standard stars. Using different model spectra, BCs could be generated for other types of stars, such as white dwarfs or giants, but we do not attempt this here.

Tables 2.18 and 2.19 list the predicted KG5 magnitudes for a sample of stars of

Table 2.18: Theoretical ULTRASPEC/KG5 magnitudes for a sample of main sequence stars with SDSS g-band magnitudes of zero. The stars have no reddening, solar metallicities, and an age of 500 Myr.

T_{eff}	**KG5 mag**
3656	-0.804
4277	-0.643
5074	-0.347
5826	-0.193
6366	-0.114
6864	-0.053
7492	+0.010
8115	+0.067
8605	+0.105
8980	+0.124
9177	+0.133

Table 2.19: Theoretical ULTRASPEC/KG5 magnitudes for a sample of main sequence stars with SDSS g-band magnitudes of zero. The stars have no reddening, solar metallicities, and an age of 2 Gyr.

T_{eff}	**KG5 mag**
3534	-0.828
3755	-0.791
4107	-0.700
4553	-0.524
5032	-0.359
5493	-0.255
5863	-0.187
6170	-0.140
6431	-0.104
6617	-0.080
6716	-0.067

differing temperatures, assuming an SDSS g-band magnitude of zero. They can be used along with the known effective temperatures and SDSS magnitudes of standard stars to calculate KG5 zeropoints for ULTRASPEC.

We used Tables 2.18 and 2.19 to predict magnitudes of standard stars in the KG5 filter, in order to measure the KG5 zeropoint. The results are included in Table 2.15.

The zeropoints measured in photometric conditions are all within 0.6 magnitudes of each other, which is a large spread considering the estimated systematic uncertainty of 0.2 magnitudes. We assume this spread is amplified by the variable sky conditions at Doi Inthanon, as discussed above in Section 2.4.1.

2.6 Example Study: Superluminous Supernovae

My main usage of ULTRASPEC has been for the observation of eclipsing CVs at high speed, and searching for optical FRB counterparts. Chapters 4 and 5 are devoted to these studies. However I have also used ULTRASPEC to conduct a small investigation into the limits of short-term variability of superluminous supernovae.

Superluminous supernovae (SLSNe) are a relatively new class of supernovae, with luminosities far greater than their classical counterparts ($M_{peak} < -21$, Gal-Yam 2012). There appears to be three different classes of these events: SLSN-R, SLSN-I, SLSN-II. Although their spectroscopic appearance seems to match well to the classic SN-IIn and SN-Ic classes (hydrogen-rich, narrow emission lines in SN-IIn, and lacking hydrogen, helium and silicon in SN-Ic), their extreme energies ($>7 \times 10^{43}$ ergs/second) suggest the progenitors of SLSNe are likely to be quite different to the classical analogues. It is still unclear what physical mechanisms can power the extreme luminosities in SLSN-I and SLSN-II, but for SLSN-R (which are rarer than the others by a factor of 5), the late light curves match well to large (10's of $M_{\odot}$ of ^{56}Ni) amounts of radioactive decay and are likely associated with the deaths of the most massive stars (Gal-Yam, 2012). In contrast, SLSN-I/SLSN-II appear to be powered by expanding shocks which interact with shells of hydrogen-poor/hydrogen-rich material at large distances, perhaps the result of highly bloated stars or ejected circumstellar material (Gal-Yam, 2012). New, large-scale optical transient surveys, such as LSST, should discover many more examples of SLSNe, and there is still much more to learn (Quimby, 2012).

On 25th February 2015, prompted by an initial investigation by Zampieri et al. (2015b,c) searching for high speed variability in supernovae, we observed the newly discovered SLSN-I CSS141223:113342+004332 (also known as SN 2015bn, Nicholl et al. 2016) with ULTRASPEC prior to its epoch of maximum light. The aim was to search for short-term variability in the light from the supernova, which could reveal information about the density and structure of the ejecta, if variability were observed.

The observations were requested by me but were conducted by Madelon Bours and JJ Hermes. The observations, which were obtained in the r'-band with exposure times of 10 seconds, lasted for 110 minutes. With the seeing at Doi Inthanon varying between 1" and 2", and with some haze present, the data quality is less than optimal.

The light curve for these observations is shown in Figure 2.60, and a Lomb-Scargle periodogram (Lomb, 1976; Scargle, 1982) of these data is presented in Figure 2.61. I found no obvious variability with amplitude greater than the noise, which was of the order 0.04 magnitudes ($\sim$4%). There were no significant periodicities found in the periodogram. By injecting fake sinusoidal signals into the data, I was able to set limits on the amplitude for periodic signal detection in our data. For trial periods of 86 seconds, 14 minutes and 2.4 hours, the maximum amplitudes for periodicities which would go undetected were 0.0025, 0.0035 and 0.0030 respectively, corresponding to 0.5-0.9% of the original light curve signal.

This investigation shows that, in the case of SN 2015bn, at approximately 24 days before peak brightness, the supernova luminosity is constant to within one percent, on timescales shorter than a few hours. Therefore there was no evidence for significant small-scale structure in the envelope surrounding the progenitor.

As an aside, we also observed this particular supernova extensively with *pt5m* (see Chapter 3), with the idea that we might contribute to a long-term light curve campaign. Our multi-colour *pt5m* light curve is shown in Figure 2.62. We also observed short-term light curves of this SN with *pt5m*, with a cadence of two minutes. We saw no signs of

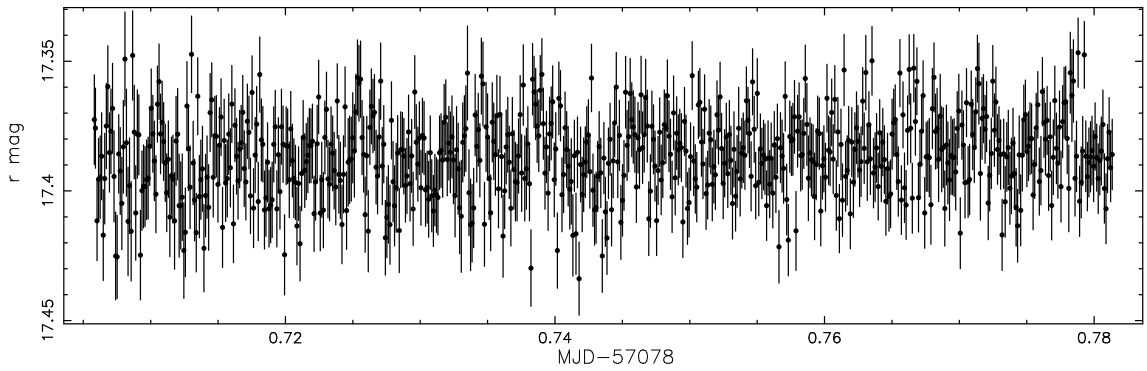

Figure 2.60: Light curve of CSS141223:113342+004332/SN 2015bn at ∼24 days before peak brightness, as observed by ULTRASPEC in the r' filter on 2015-02-25. The light curve shows no signs of variability above the noise.

variability in these observations either.

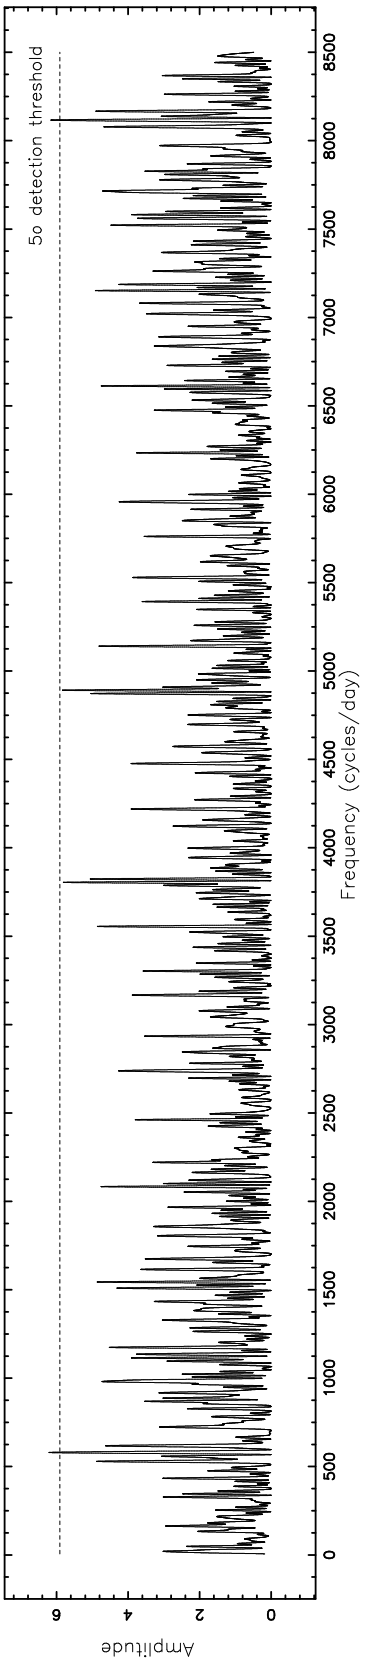

Figure 2.61: Lomb-Scargle periodogram of the CSS141223:113342+004332/SN 2015bn light curve shown in Figure 2.60. The dashed line represents an amplitude 5σ above the mean, and shows that the visible peaks are not convincingly significant.

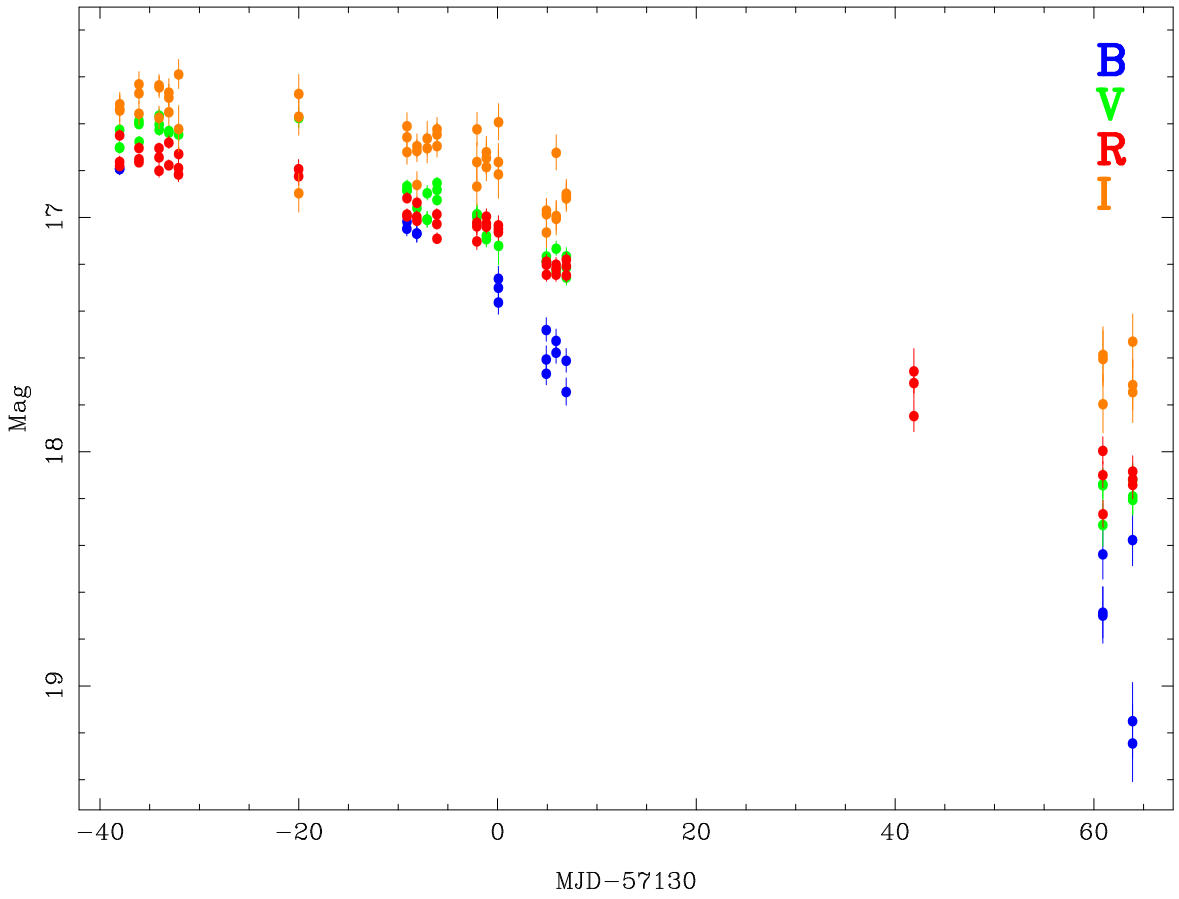

Figure 2.62: 100-day *BVRI* light curve of CSS141223:113342+004332/SN 2015bn as observed by *pt5m* (see Chapter 3). Large gaps in the light curve were due to technical downtime and periods of bad weather.

2.7 Training the Thai Observers

After commissioning the instrument in November 2013, I stayed behind in Thailand for 4 weeks in order to train the local astronomers and telescope operators in the use of ULTRASPEC. Unfortunately, the weather at Doi Inthanon continued to be wet and cloudy for almost all of that time, and we managed to complete very little on-sky training.

Nevertheless, I gave extensive off-sky training to the telescope operators, and continued to support them via email over the rest of the season. I gave a talk about the instrument, and in particular about how to use the ULTRACAM pipeline software. Although there were further questions, queries and minor problems to start with, the local

and visiting astronomers now use ULTRASPEC as the primary instrument on the TNT without any serious issues.

Together with Vik Dhillon, I also helped to write the user manual for ULTRASPEC[4]. This, along with giving the training, was a useful process in helping me understand how to make the most of using the ULTRASPEC.

2.8 Conclusions

I have presented a description of the newly commissioned high-speed imager ULTRASPEC on the TNT. I discuss the detailed analysis of the lab-based and on-sky commissioning data for ULTRASPEC. The instrument is performing well (with less than 1% technical down time), to within the specifications, and the collaboration with NARIT is producing some interesting science (although we show only a small example in this chapter - references to further results from ULTRASPEC on the TNT are given in Chapter 6).

I have modelled the instrumental throughput and compared its predicted performance with observations. I have used the throughput model to calculate bolometric corrections as a function of stellar effective temperature in the KG5 filter. These corrections can be used in conjunction with observations of standard stars to calculate photometric KG5 zeropoints. The procedure for calculating theoretical bolometric corrections was also used to test the reliability of the throughput model with observations of the open cluster NGC 6940. The model appears to be correct to within a systematic error of 0.2 magnitudes.

[4] `www.vikdhillon.staff.shef.ac.uk/ultraspec/userman/uspec_userman.html`

Chapter 3

Using and Improving *pt5m* - a 0.5m robotic telescope on La Palma

In this chapter I outline my work with *pt5m*, a 0.5m robotic telescope on La Palma. The majority of this chapter appears in my published paper on *pt5m* (Hardy et al., 2015b). I have been the primary user of this facility for the past 3 years. As such I have been the first to spot most hardware or software issues with the telescope, and usually I have been able to resolve these issues myself. My contributions to the project, as well as authoring the published manuscript, were focused mainly on designing, writing and testing the transient follow-up software, some of which was completed in partnership with Stuart Littlefair. I also conducted the commissioning and performance tests of the telescope, and a wide range of science observations, totalling over 500 GB of data. The following text is all my own, and/or as published in the paper.

3.1 Introduction

Robotic observing facilities requiring no human input during night time are now commonplace[1], and are becoming ever more popular (e.g. Zerbi et al. 2001; French et al. 2004; Tsapras et al. 2009; Gillon et al. 2011; Brown et al. 2013; Gorbovskoy et al. 2013; Andersen et al. 2014; Schmitt et al. 2014; Strolger et al. 2014). The largest example is the 2m Liverpool Telescope (Steele et al., 2004), providing a fully autonomous, large aperture telescope at a high quality site, with an extensive suite of instrumentation. Such facilities can sharply reduce operation costs as they require much less staff interaction. They can also be programmed to make complex scheduling decisions in order to observe the most appropriate targets, given the observing conditions.

Robotic telescopes can be particularly effective in the field of transient astronomy. For example, using automated alert-listening software, they can respond to unexpected events much faster than human-controlled telescopes. This is especially useful in the follow-up of short-lived transients such as Gamma-Ray Bursts (GRBs). An effective transient follow-up facility requires every aspect of its operation to be conducted as fast as possible. A control system, dome, mount and camera that processes information, responds, moves, and reads-out data quickly are all vital components of a rapid-response robotic telescope. However, overall the leap-frogging of the human interaction stage is most important, since this can delay follow-up by minutes, hours or even days.

The Universities of Sheffield and Durham now boast their own collaborative contribution to the field of robotic astronomy. *pt5m*[2] is a 0.5m robotic telescope located on the roof of the 4.2m William Herschel Telescope (WHT, see Figure 3.1) on La Palma and is hosted by the Isaac Newton Group of Telescopes (ING). It functions as an astrophysics research facility, an atmospheric turbulence profiler, and a teaching resource. *pt5m* op-

[1]For a global list of robotic telescopes see `http://www.astro.physik.uni-goettingen.de/~hessman/MONET/links.html`

[2]The name is an abbreviated form of 'point-five-metre'

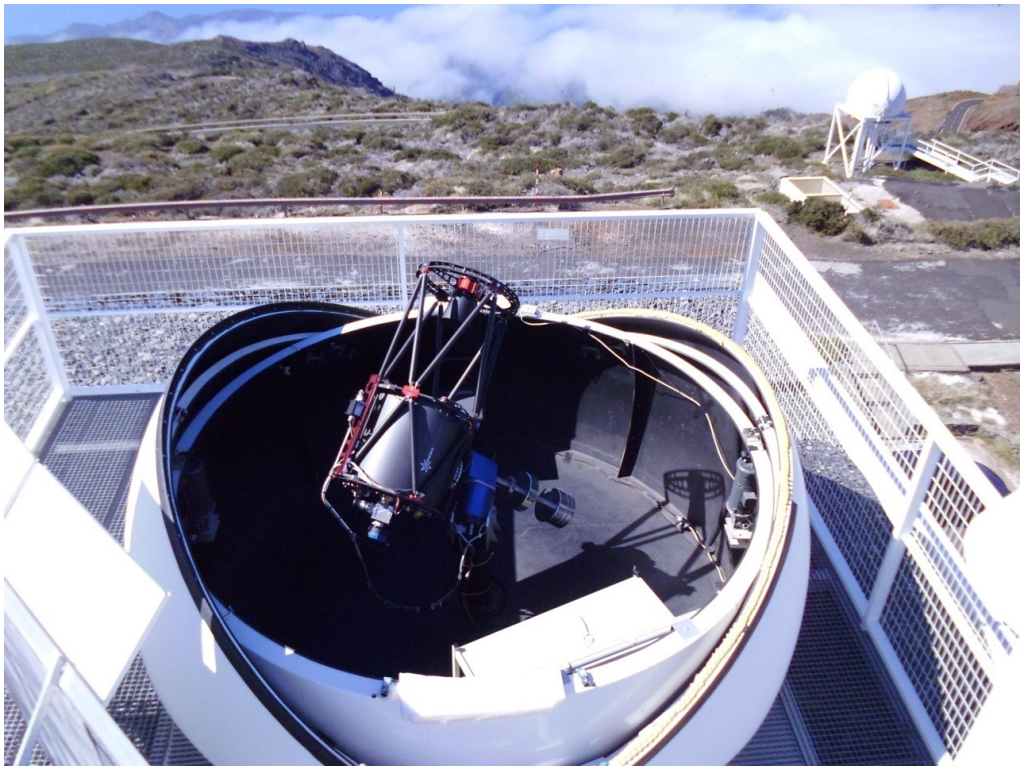

Figure 3.1: Photograph of *pt5m* on the roof of the WHT taken in April 2015. The WHT dome is out of view to the left.

erates autonomously and observes queued targets every night when the weather is good. It has already been used to study pulsating stars (Fritzewski et al., 2016), transiting exoplanets (Garai et al., 2016), solar system objects, interacting binaries (Littlefair et al., 2013; Kupfer et al., 2014; Campbell et al., 2015c), and transients (Gandhi et al., 2014; Hardy et al., 2015a).

One of the greatest advantages in having our own observing facility is the unrestricted access to observing time. Without a Time Allocation Committee or observing proposals, we are free to investigate interesting science as soon as the need for observations becomes apparent. This is particularly useful in the field of transient astronomy, since often we have no idea if or when an interesting source will appear, or a known source will become interesting. This type of setup is also very conducive to long-term monitoring projects, that require just a few minutes of observations per night, over weeks or months. It is

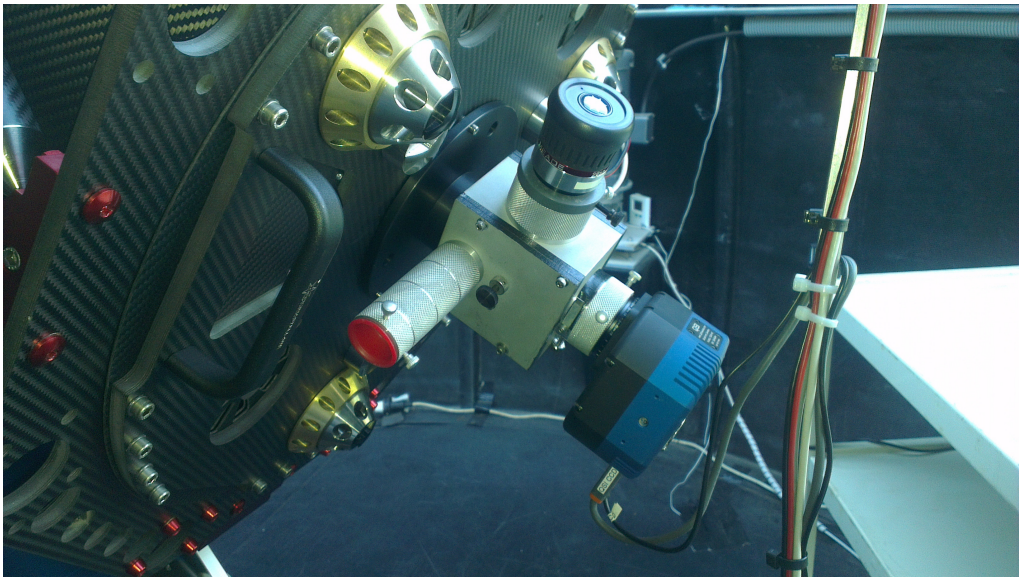

Figure 3.2: Photograph of the Cassegrain focus of *pt5m*, showing the flip mirror mechanism (grey), and the CCD camera (blue) with integral filter wheel (black). In this photo the SLODAR port (pointing upwards) temporarily holds an eyepiece.

difficult to perform observing campaigns of this type at human-operated telescopes. Of course, we must share the observing time between a small collection of astronomers at Sheffield and Durham, as well as with undergraduate students who use the telescope for learning, but in practice *pt5m* is rarely oversubscribed.

This chapter gives a short history of *pt5m* (Section 3.2), and explains in detail its hardware and software architectures (Sections 3.3 and 3.4). I then present an analysis of the telescope performance (Section 3.5), and an example photometric study of the transiting exoplanet WASP-33b (Section 3.6, originally conducted as part of my MPhys project). The final sections describe in more detail the software contributions I have made (Section 3.7), and the scientific results from the use of this software (Section 3.8), before concluding in Section 3.9.

3.2 Brief History of *pt5m*

pt5m was first developed as an atmospheric turbulence profiling facility (SLOpe Detection And Ranging, SLODAR, Wilson 2002) and was originally installed at the South African Astronomical Observatory (Catala et al., 2013). The telescope was then moved to La Palma in 2010 to support the CANARY adaptive optics demonstrator (Morris et al., 2009). For this reason it was sited as close as possible to the WHT, to sample the turbulence in approximately the same beam as the WHT.

In order to fully utilise the telescope when not being used for SLODAR observations, we decided to add a science imager in 2011. We recognised that even a small telescope located at such an excellent site could provide a powerful tool for time-domain astronomy. The science CCD and integrated filter wheel (see Section 3.3.2) were originally mounted on a linear slide to move the science camera in and out of the light path of the SLODAR optics. By 2012 we had implemented the robotic mode (Section 3.4) and started using *pt5m* on a routine basis for time-resolved and transient astronomy. In 2014 the original 0.5m Orion Optics Modified Dall-Kirkham telescope was replaced with a new Ritchey-Chretien telescope (Section 3.3.3). This refurbishment was made to improve throughput and image quality, which were not a high priority for the original SLODAR studies. At the same time the slide mechanism was removed and replaced with a flip mirror that allows light to be diverted to either one of two optical ports. One of these ports hosts the science camera and the other hosts the atmospheric turbulence profiler (see Figure 3.2 and Section 3.3.3).

3.3 Hardware

The following subsections describe the hardware components of *pt5m*. Figure 3.3 shows the components schematically, whilst Table 3.1 summarises the facility specifications.

Table 3.1: Specifications of *pt5m*. Location information is as measured by GPS (WGS84).

Parameter	**Value**
Telescope	
Longitude	28° 45' 38.7" N
Latitude	17° 52' 53.2" W
Altitude	2383 ± 3m
Primary Diameter	0.5m
Primary Focal Length	1500mm
Primary Conic Constant	-1.071
Secondary Diameter	0.148m
Secondary Focal Length	-609.89mm
Secondary Conic Constant	-4.182
Dome Diameter	3.66m
Focal Length	5000mm
Platescale	41.25"/mm
Telescope un-vignetted FoV	9'
Slew Speed	20°/second
Pointing Accuracy (RMS)	0.4'
Tracking Accuracy	12"/hour
Finder FoV	$7.2° \times 5.4°$
Instrument	
Detector Format	2184×1472 pixels
Pixel Size	6.8 μm
Pixel Scale	0.28"/pixel
CCD FoV	$10.2' \times 6.9'$
Full Well Depth	83000 e^-
Readout Noise (binned 1x1)	10.4 e^-
Readout Noise (binned 2x2)	14.4 e^-
Readout Noise (binned 3x3)	18.5 e^-
Gain	1.3 e^-/ADU
Flat field noise (grain)	0.15%±0.05%
Dark Current	8 ± 4 e^-/pixel/hour
Dead Time (binned 1x1)	8 seconds
Dead Time (binned 2x2)	5 seconds
B-band Zeropoint	21.91 mag
V-band Zeropoint	22.06 mag
R-band Zeropoint	22.16 mag
I-band Zeropoint	21.05 mag

3.3.1 Dome and Mount

pt5m is enclosed in a 12-foot Astrohaven clam-shell dome. The use of a clam-shell design is vital for rapid follow-up work, as there is no delay whilst the dome rotates to catch up with the telescope. The disadvantage of such an open design is the reduced protection from wind buffeting compared with a fully enclosed dome.

The telescope mount, a direct-drive Astelco Systems NTM-500, is of German-Equatorial design. This mount design requires a 'pier-flip' when tracking a target across the observer's meridian from East to West. The pier flip introduces a discontinuity in the x-axis of time-series observations of up to several minutes, as the mount must wait for the target to pass through the observer's meridian before reacquiring. The pier flip also rotates the field of view through 180°, placing the target and comparison stars on different parts of the chip, with different levels of vignetting. This results in a small step in the differential photometry which, in principle, can be corrected for with flat fields. The only major negative effect of the pier flip is when it occurs in the middle of an interesting event, e.g. during an eclipse.

In its implementation in *pt5m*, the mount has an RMS pointing accuracy of $\sigma = 0.4$' around the mean pointing position, and an average tracking drift of 12"/hour. These values are larger than expected from the manufacturer's specification, most probably due to flexure in the telescope and roof-top foundations, and small errors in the polar alignment.

Although it is stable during normal use, the telescope and mount both suffer from vibrations in high winds (8-10 m/s), as the clam-shell dome leaves them exposed. They also vibrate when the main WHT dome motors are engaged, as the facility's roof-top foundations are not fully isolated from the building. We are in the process of erecting substantial wind baffling around the dome, particularly on the South-West side where the wind is forced between the WHT dome and the rest of the building.

An advantage of this particular mount is its fast slewing speed of 20°/second[3]. This allows new targets to be acquired very quickly, which is essential for rapid follow-up work, as well as reducing overheads between observations.

3.3.2 Science Camera

The 'Quantum Science Imaging' QSI-532ws camera uses a Kodak KAF-3200ME CCD detector. The science camera is equipped with an integral 5-position filter wheel, which at present holds the following Astrodon filters: Johnson B, Johnson V, Cousins R, Cousins I, and Hα 5nm. The camera employs a thermoelectric cooling system, and is maintained at -20°C when in use. The CCD specifications are listed in Table 3.1. Of particular note is the low dead-time of the CCD (5 seconds in 2x2 binned mode), which reduces overheads and improves the cadence of time-series observations. The CCD camera uses a shutter, and takes timing information directly from the control PC (see Section 3.3.5). The timestamps are moderated by the Network Timing Protocol (NTP) which has a typical accuracy of < 100ms. As well as a science camera, *pt5m* also has an ASTROVID StellaCam finder camera mounted on the telescope.

3.3.3 Telescope

The 0.5m optical tube assembly (OTA) is a custom built $f/10$ Officina Stellaire Pro RC 500, made almost entirely of carbon fibre, resulting in a light-weight, rigid and low thermal-expansion structure. The OTA includes primary mirror covers, primary mirror heaters to prevent the formation of dew (these are not used), and fans for cooling the primary mirror. The telescope focus is adjusted by moving the motorised secondary mirror.

A tertiary mirror is housed in a *Van Slyke Instruments Flipper* (Figure 3.2). It uses

[3]In fact, with a smaller payload, the mount could slew at up to 100°/second.

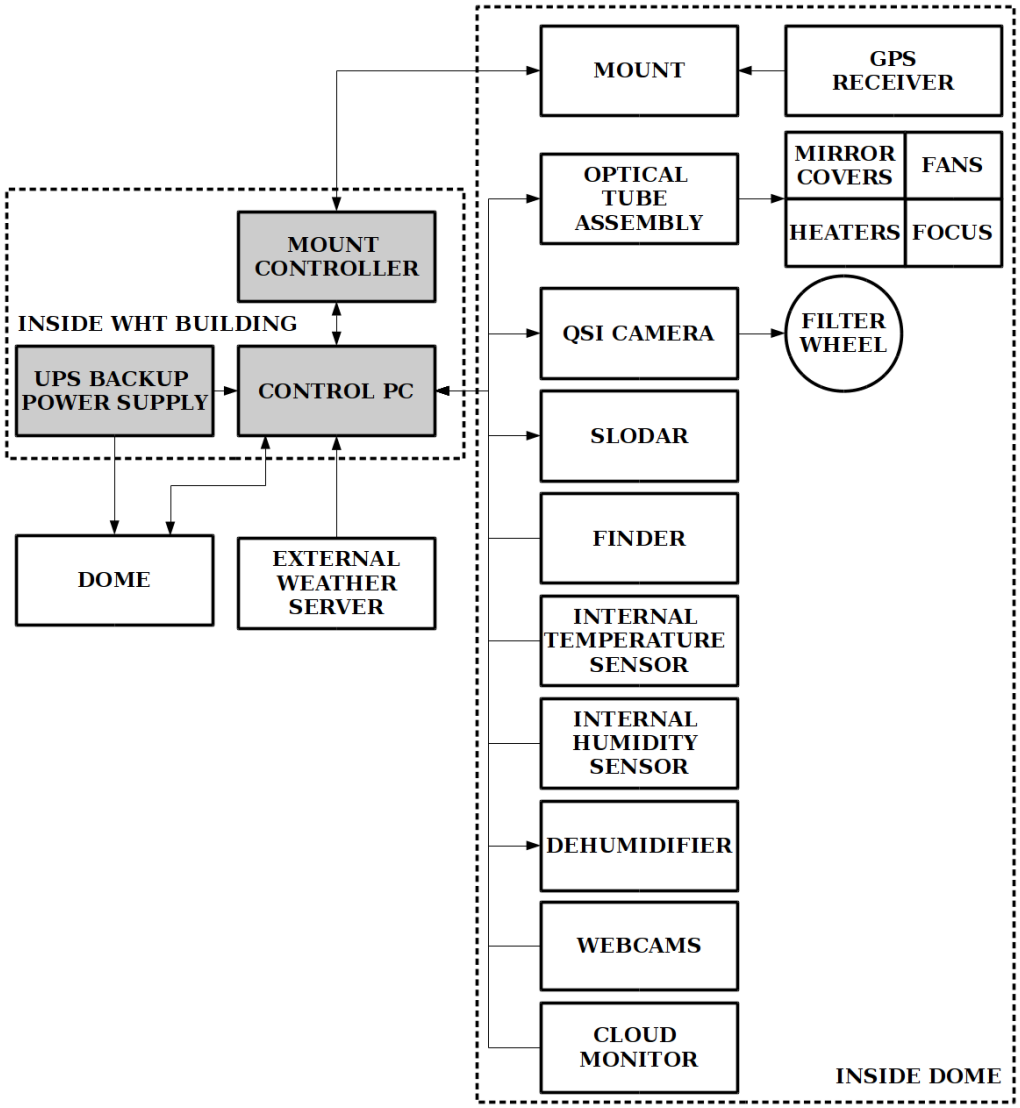

Figure 3.3: Schematic showing the *pt5m* hardware components and their connections.

a mirror on a hinged support that flips up and down. This allows the light beam to either pass straight through to the filter wheel and CCD or to be folded by 90° into the port dedicated to atmospheric turbulence profiling.

3.3.4 Weather System

In order to keep *pt5m* safe from weather damage, a reliable weather station, with backup systems available in case of failure, is essential. Since *pt5m* is hosted by the ING,

we are able to benefit from access to their weather stations[4]. These include the main weather system for the whole observatory located on a mast close to the Jacobus Kapteyn Telescope (JKT), as well as a local mast located on the WHT. In addition, *pt5m* employs its own internal sensors, including a humidity and temperature sensor located in the dome. The Conditions Monitor (see Section 3.4.4) monitors the weather information and decides when the dome should be open or closed. The dome is forcibly closed if this weather data becomes unavailable, or returns exactly the same values for a period of more than 20 minutes (we assume this only happens if the ING weather server has crashed and is not supplying reliable data).

A dehumidifier has been installed in the dome, which turns on automatically when the humidity in the dome is above the safe observing limit. This ensures the unit is only engaged when the dome is closed, and maintains the internal humidity at a safe level when the external humidity is high.

The dome is also equipped with an *AAG CloudWatcher* cloud detector. This is mounted on the internal rim of the dome and sees the entire sky when the dome is open. The unit uses an infrared sensor to measure the temperature of the sky. We have calibrated this sensor against images taken by the Gran Telescopio Canarias all-sky camera over the course of a year. We found a linear relationship between the infrared sky temperature and cloud coverage in oktas, where 0 represents a completely clear sky and 8 is completely overcast. We tried correcting the sky temperature for seasonal variations using a polynomial based on the measured ground temperature. The corrected data describes the same trend as the uncorrected data, and thus we found no evidence for seasonal variations. We use this cloud coverage data as a guide to determine whether or not a night is photometric, although we do not currently use this information when scheduling observations robotically. Since the cloud monitor is mounted on the inside rim of the dome and therefore requires the dome to be open in order to measure the sky

[4] `http://catserver.ing.iac.es/weather/index.php`

temperature, it cannot be used to make a decision on whether or not to open or close the dome.

3.3.5 Electronics

All of the *pt5m* hardware is controlled by a dedicated Linux PC on site. The machine serves as an interface between the different information inputs (e.g. weather data) and outputs (e.g. dome control). Both the control PC and the dome motors are connected to an uninterruptible power supply (UPS), so that the dome can always close even in the event of a power failure. All of the *pt5m* hardware components are connected to remote-control power switches, enabling them to be turned on and off over the network. The control PC, UPS, mount controller and remote power switches are all located inside a purpose-built enclosure in the WHT elevator control room, instead of inside the dome. By locating the electronics inside the WHT building, we reduce the possibility of damage by sudden adverse weather conditions.

pt5m is also equipped with three physical emergency stop buttons. The first stops the telescope and mount, engaging the brakes immediately to stop any potential damage if the telescope is about to crash into the pier, for example. This button is mounted on the telescope mount. The second button is mounted on the interior dome wall, near to the access hatch. This button shuts off power to the dome motors to prevent accidents when climbing through the hatch. The third button is located inside the WHT elevator control room, and is used to shut down the telescope and close the dome in the event of software malfunction or severed connectivity to the external network.

3.4 Software

[5]All of the software used to control *pt5m* is written in Python and runs on the Linux control PC (Section 3.3.5). *pt5m* can be used in either "manual" or "robotic" mode. In manual mode, a human user controls the facility in real time by sending a series of instructions via a command-line terminal that has been opened on the control PC. This mode is used for testing and teaching purposes, and can be used to control *pt5m* locally or remotely over the internet. In robotic mode, *pt5m* operates autonomously, opening and closing the dome, monitoring the weather, taking and processing calibration data, and selecting and observing science targets.

A schematic showing the interactions between the software components of *pt5m* is shown in Figure 3.4.

3.4.1 Hardware Control: Daemons

The main hardware components of the system each have a dedicated software process or "daemon" that runs in the background. These daemons allow the hardware components to be initialised independently from one another. The "Python Remote Objects" (Pyro) library is used for communication between the daemons and other software processes – each daemon is a Pyro server. This server/client model also allows multiple software components to query and send commands to a piece of hardware without clashes and without having to wait for a response.

The *pt5m* hardware daemons are as follows:

- Power: Communicates with the remote-controlled power switches (Section 3.3.5), used to power the individual hardware components on and off. This daemon runs at all times.

[5]Note - My contributions to *pt5m* software development are discussed in Section 3.7. The other software components discussed here were developed by my *pt5m* collaborators, although I have instigated several bug fixes for some of these programs.

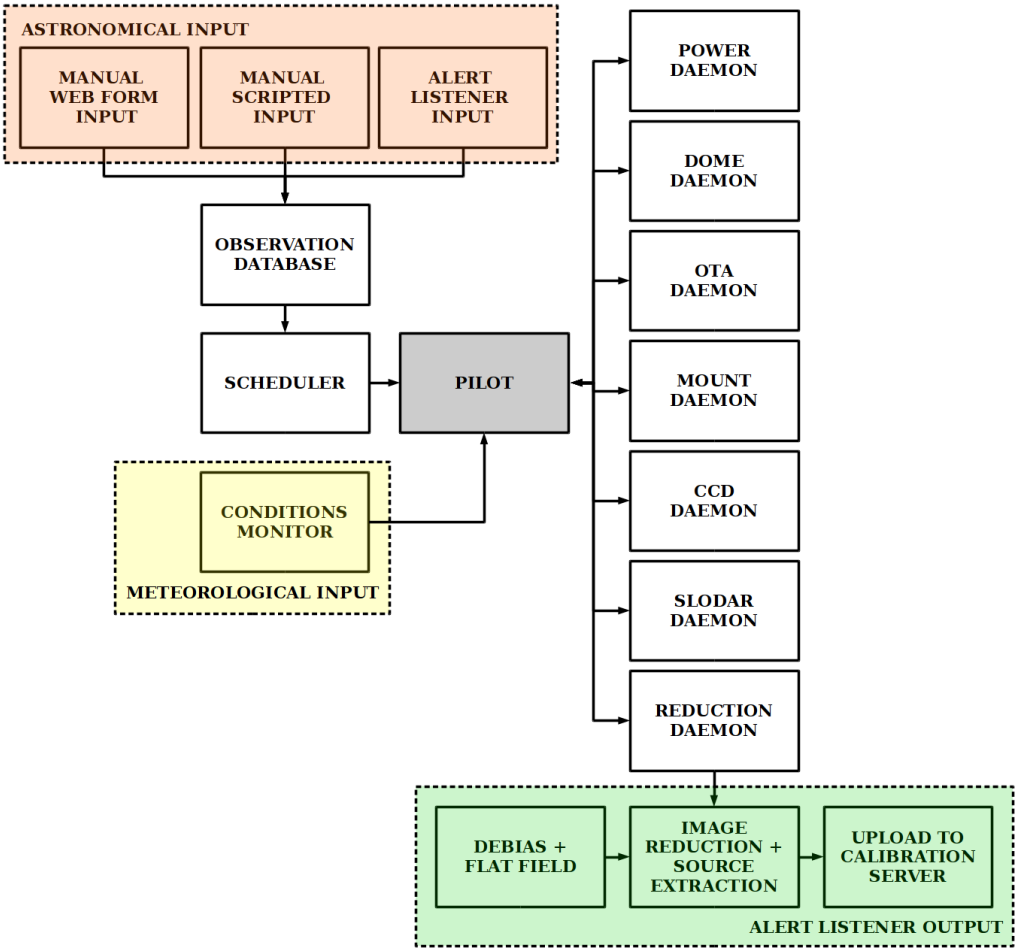

Figure 3.4: Schematic showing the software components of *pt5m* and their interaction.

- Dome: Controls opening and closing of the dome via an Arduino interface. It constantly checks the current weather conditions via the Conditions Monitor (Section 3.4.4). It will automatically close the dome (or refuse to open it) if any of the Conditions Monitor flags are set to 1, if the external power fails or if the network connection to the outside world fails. This daemon runs at all times.

- OTA: Controls the OTA's built-in electronics via an RS232 interface. Functionality includes opening/closing the primary mirror cover, enabling/disabling the mirror heaters, and focusing.

- Mount: Communicates with the telescope mount controller via an internal network.

This daemon controls the pointing and tracking of the telescope. It also constantly monitors a wide range of status information that is available from the mount and makes this available to other software processes.

- CCD: Communicates with the CCD camera via a USB interface. While the daemon is running the thermoelectric cooler is enabled and the CCD is ready to take images. The daemon maintains a queue of requested exposures (each with its own filter, exposure time and binning settings) which are processed sequentially. When this daemon is shut down at the end of each night the CCD temperature rises in a controlled manner before the CCD is powered off.

3.4.2 System Control: Pilot

The 'pilot' program is the software process that is responsible for autonomous control of *pt5m* in robotic mode. It gathers information from all of the daemons and other processes, and issues commands to them. The pilot is started automatically each day in the late afternoon. It remains dormant during daytime, and begins preparing for the night's observations during twilight.

The evening-twilight start-up procedure begins when the Sun reaches 10° above the horizon, and is as follows:

1. Sun at 10°: power on CCD and mount.

2. Sun at 7°: take bias and dark frames (dome still closed).

3. Bias frames complete; open dome and then open mirror covers.

4. Sun at 5°: point telescope to a blank field in the sky and take flats.

5. Focus the telescope using a bright star.

6. Process flats and biases into master frames if enough individual frames were taken in good conditions, i.e. no cloud presence detected during flat field observations.

7. Optionally, observe flux standards.

8. Park the telescope and wait for the end of twilight.

Focusing is conducted automatically on a bright Gliese catalogue star. The half-flux diameter (the radius at which half of the total encircled energy is contained) is measured at the current focus position, and then measured again at a position several millimetres away from the original position of the secondary mirror. The behaviour of the half-flux diameter away from the best focus is linear with respect to the position of the secondary mirror (focus). These two measurements, along with a predetermined function for the half-flux diameter behaviour, which does not vary with time, are all that is required to calculate the best focus position. If the auto-focusing program fails (because a bright star could not be found, or the half-flux diameter measurements do not behave as expected), the telescope returns to the focus used on the previous night.

The pilot is considered to be in the night-time state while the Sun is below $-10°$ altitude. During this time, its behaviour is governed by the "scheduler" (see Section 3.4.3). The pilot periodically queries the scheduler and, if the scheduler suggests a new target should be observed, the pilot slews the telescope and adds the specified exposures to the CCD queue. Once the pilot verifies that the telescope is tracking on the target, the CCD exposures begin. The system then remains in this state, taking exposures, until one of the following occurs:

- all exposures in the CCD queue have been taken and the observation is complete,

- the scheduler returns a new, higher priority target to observe,

- conditions force the dome to close, or

- the end of the night is reached.

If the scheduler returns no valid targets, the telescope is moved to the park position. However, the dome remains open so that the system is ready to respond immediately when a new target is requested.

In the event of bad weather, the dome is closed immediately by the dome daemon (see Section 3.4.1) without the intervention of the pilot. This is an essential precaution to protect the instrumentation from wind and water damage as the weather conditions can change quickly. The pilot then responds as follows: if the telescope is currently observing a valid target it will pause the queue of CCD exposures and leave the telescope tracking. If the conditions improve after a short time then the pilot simply reopens the dome and resumes the CCD exposures. If/when the current observation becomes invalid (e.g. the target sets below the horizon), the pilot parks the telescope and clears the CCD queue.

The morning-twilight shutdown procedure is:

1. Sun at $-9°$: point telescope to blank field and take flats.

2. Sun at $0°$: park telescope, close mirror covers, close dome.

3. Power off mount and CCD.

3.4.3 The Scheduler

The scheduler is a Python script that is interrogated every 15 seconds by the pilot during the night. The script returns a decision on what should be done, either to continue observing what is currently being observed for the next 15 seconds, or to move immediately to a new target.

Each new target (also referred to as a 'job' or 'pointing') is described by an ASCII text file that defines the position of the target, the CCD exposures required, the priority (an integer from 0 to 5, where 0 is the highest priority), and various other constraints

on the required observations, as shown in Figure 3.5. Only one target per job is allowed. Each job is given a unique ID number, which is also the name of the ASCII file. The jobs define the 'queue', and all of the awaiting job files are stored in a single directory on the control PC. When a job is complete, the file is moved from the queue directory to a 'completed' directory. Alternatively, if the job is aborted, deleted, expired or interrupted, it is moved to a correspondingly named directory.

```
# Pointing     Object   RA(J2000)  Dec(J2000)  Priority  ToO  Flip  Sun Alt (deg)  Min.Alt.(deg)  Min.Time(sec)  Max.Moon  User        StartUTC             StopUTC            guide
#-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------
  11490      ASASSN14gu  104.032   33.6016        4       0    1       -15.0          35             9000          B        lkh   2015-02-19T18:00:00   2015-02-20T08:00:00     1
#
#
#
#TYPE    FILTER   BINNING   EXPTIME   NUMEXP
#-----------------------------------------
SCIENCE     V        2        150       200
DARK                 1        150         5
```

Figure 3.5: Example ASCII text file defining a pointing. The first variable is the unique pointing ID, the priority ranges from 0 (highest priority) to 5 (lowest priority), ToO = Target of Opportunity and is a flag used for targets that need override status, Flip is a flag used to indicate whether the observations can tolerate a pier flip, Sun Alt is the maximum acceptable altitude of the Sun, Min.Alt. is the minimum acceptable altitude of the target, Min.Time is the minimum amount of time the telescope should spend on the observations in order for them to be useful, Max.Moon takes values B, G or D for bright, grey or dark time, StartUTC and StopUTC define the time period in which the observations may begin, but once initiated they can extend beyond StopUTC, and guide is a flag used to indicate whether autoguiding should be used. The pointing file also defines the exposures requested, showing the filter, binning mode, exposure time and the number of exposures. The different types of exposures available are SCIENCE, FOCUS, DARK, SKY and BIAS.

Table 3.2: Flags used by the *pt5m* scheduler. Each flag can only take the value 0 (legal) or 1 (illegal). The artificial horizon refers to a map of minimum observable altitude for any given azimuth, which ensures that *pt5m* never tries to observe targets which are vignetted by the WHT dome or building. See Figure 3.5 for an explanation of the terms StartUTC, StopUTC, Sun Alt, Min.Alt. and Min.Time.

Flag name	**Flag=0 if:**
time	The time now is between StartUTC and StopUTC
minalt	The target is above Min.Alt. now
mintime	The target will still be above Min.Alt. after Min.Time has elapsed
altart	The target is above the artificial horizon now
altartmint	The target will still be above the artificial horizon after Min.Time has elapsed
moonphase	The Moon is not too bright for the required observations
moondist	The target is not too close to the Moon
sunalt	The Sun is below Sun Alt now
sunmintime	The Sun will still be below Sun Alt after Min.Time has elapsed
flip	The telescope will not need to conduct a pier flip, or the job allows a pier flip to be conducted

Each time it is polled, the scheduler runs through every job in the queue and calculates a set of flags for each job, with a result of 0 indicating the flag is legal and 1 indicating it is illegal. An explanation of each of the flags is listed in Table 3.2. The job is legal only if all the flags are 0. The scheduler then updates the priority of each job in the queue, as follows:

- If the job is defined as a Target of Opportunity (ToO=1 in Figure 3.5), its priority is given by the user-defined integer priority (Priority in Figure 3.5) plus airmass/100, where the airmass of the target is calculated at a point in time mid-way between the current time and the minimum end-time of the job (Min.Time in Figure 3.5). Hence a priority 2 job with a current airmass of 1.234 becomes 2.01234.

- If the job is not a ToO (ToO=0 in Figure 3.5), its priority is given by the user-defined integer priority plus airmass/10. Hence a priority 2 job with a current

airmass of 1.234 becomes priority 2.1234.

- If any of the scheduler flags are non-zero, add 10 to the priority. Hence an illegal, non ToO, priority 2 job with a current airmass of 1.234 becomes 12.1234.

- If the priorities of two jobs are equal after the above updates, the job that entered the queue first has the higher priority.

The scheduler takes the job with the highest priority (equal to the lowest numerical value) in the queue and compares it with the updated priority of the job currently being observed. On the basis of this comparison, the scheduler decides what to do next, as follows:

1. If the current observation is still the highest priority in the queue, and legal, continue observing.

2. If the current observation is no longer the highest priority in the queue, and the highest priority job in the queue is *not* a ToO, continue observing. This rule prevents jobs from constantly being overridden, and usually allows all of the requested observations to be obtained. Note that the exception to this rule is for priority 5 observations. These are "queue fillers" and can be overridden immediately by any job in a higher-priority band, regardless of its ToO status. Queue filler jobs will also be resumed after interruptions, when conditions allow, which helps keep the telescope active even on quiet nights with few high priority jobs in the queue.

3. If the current observation is no longer the highest priority in the queue, and the highest priority job *is* a ToO, interrupt the current observation by aborting the current exposure and immediately slewing to the new target. This rule guarantees that time-critical observations can be obtained. Note that if the current observation is also a ToO, it will not be interrupted unless the interrupting ToO lies in a higher priority

band: for example, if the current observation is a ToO of priority 3.01234, it would not be overridden by a ToO job of priority 3.01000, but would be overridden by a ToO job of priority 2.01234.

In this way the scheduler is simple and flexible. By setting appropriate values for the priority, ToO and start/stop times, it is possible to accommodate a wide variety of scientific projects, from high-priority time-critical jobs to low-priority long-term monitoring jobs. The drawback of the scheduler, as currently implemented, is that it does not look forward more than one job into the future. It therefore might observe a short high-priority job at the expense of a longer job with only a slightly lower priority. One solution would be to look forward to the end of the night and construct a schedule based on minimising some statistic, e.g. $\Sigma_n(t_{obs} \times p)$, where t_{obs} is the length of time spent observing a job of priority p, summed over all of the jobs observed, n. The drawback of this approach is that we would then lose some control over which jobs are observed, and therefore we have decided not to implement it.

3.4.4 Conditions Monitor

The Conditions Monitor is a Python script that reads the weather information provided by up to three on-line sources from the ING, and is updated every 60 seconds. These sources record the meteorological conditions at three different locations at the Observatorio del Roque de los Muchachos: the weather masts at the JKT, Isaac Newton Telescope (INT) and WHT. The various parameters that are read from each mast are listed in Table 3.3, along with the limits defining when the dome should be closed. In order to prevent repeated opening and closing of the dome when conditions are close to any of the limits, we define different limits for when the dome should be re-opened following a closure due to bad weather. These re-opening limits are also listed in Table 3.3. The parameters are more lenient when the dome is already open compared to when the dome

Table 3.3: Weather parameters considered by the *pt5m* Conditions Monitor, with different limits applicable depending on whether the dome is open or closed. The final two parameters (Max Sun and Max Moon) do not affect the opening and closing of the dome, but dictate when observing can commence.

Parameter	**Limit (open)**	**Limit (closed)**
Windspeed (m/s)	10.0	9.0
Windspeed (km/h)	36.0	32.4
Relative humidity (%)	70	63
Internal humidity (%)	70	63
Dry/Wet sensor	DRY	DRY
Temperature (Celsius)	1.0	2.0
Age of weather data (seconds)	700	700
Max time for parameters to be unchanged (seconds)	1200	1200
Max time without contact to Durham server (seconds)	300	30
Max Sun altitude (°)	-10	-10
Max Moon altitude for dark time (°)	-5	-5

is closed, as the latter is likely to be because of recent bad weather, and some additional contingency is needed.

The Conditions Monitor usually takes its input from the JKT mast. This is the most exposed location on the edge of the Caldera on La Palma, and hence usually reports the most extreme meteorological parameters. By using these parameters we are therefore taking the most conservative approach to opening and closing the dome, which is essential for a robotic telescope.

Occasionally there is a problem with the JKT mast, and one or more of the meteorological parameters reported by the system is either unavailable, corrupted or not updated. In this case, the Conditions Monitor automatically switches to taking its input from one of the other masts. If this mast is also in error, the Conditions Monitor will shut the dome. The Conditions Monitor also continuously checks that the internet link between *pt5m* and the UK is working. If the link is broken, the Conditions Monitor will shut the dome. Once again, this is adopting a conservative approach: with no internet link, there is no possibility of us logging in remotely to shut the dome in case of a prob-

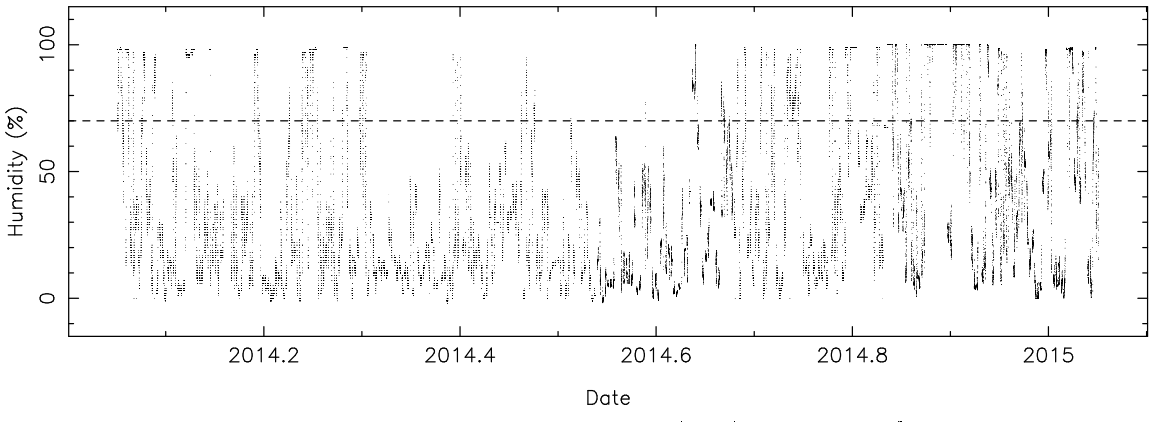

Figure 3.6: Humidity values as measured by the Conditions Monitor during 2014. The open dome humidity limit of 70% is shown with a dashed line.

lem. The system also notes the timestamps of all weather data packets, and treats the data as unreliable if it is more than 700 seconds old, or if the values don't change for more than 20 minutes.

The Conditions Monitor works by setting flags. All of the parameters listed in Table 3.3 have an associated flag, which is set to 0 when the parameter is legal (i.e. it is safe to open the dome) and 1 when the parameter is illegal (i.e. it is not safe to open the dome). Only when all of the flags are 0 can the dome be opened, and if any of the flags are subsequently set to 1, the dome will close. When observing locally, and when sure that it is safe to do so, it is possible to override this system and allow the dome to remain open even if one or more of the flags are set to 1.

We have analysed the fraction of time that the dome remains open during night time hours, given the current weather parameters. Between 2013-01-20 and 2015-09-17 (971 nights), 21% of night time experienced high humidity, 17% experienced high windspeeds, and 10% experienced low temperatures. These fractions will overlap with each other, as often high humidity is coupled with high winds and/or low temperatures. Overall, 32% of night time was lost because of adverse weather, while 5% was lost due to incomplete or corrupted meteorological information. Figures 3.6, 3.7 and 3.8 show the weather station values of humidity, wind and temperature for a year-long period from early 2014.

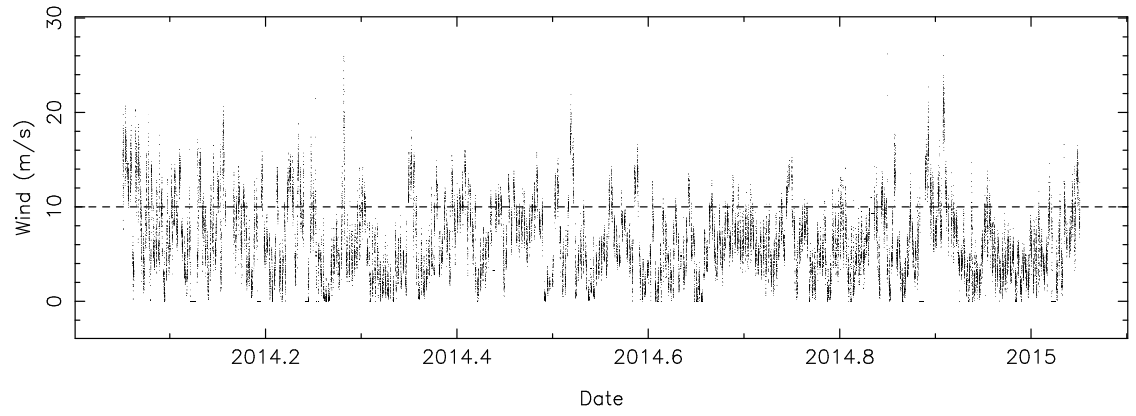

Figure 3.7: Wind speeds as measured by the Conditions Monitor in 2014. The open dome windspeed limit of 10 m/s is shown with a dashed line.

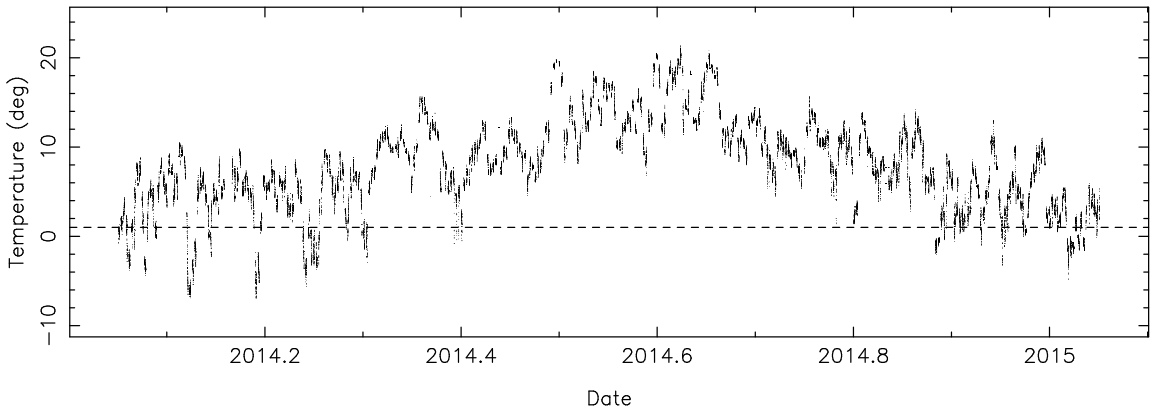

Figure 3.8: Temperatures as measured by the Conditions Monitor in 2014. The open dome temperature limit of 1.0°C is shown with a dashed line.

Were we to be more liberal with our weather limits, e.g. setting the humidity limit to 80%, the windspeed limit to 11 m/s, and the temperature limit to 0°C for both open and closed dome scenarios, we could expect the fractions of lost night time to drop to 14%, 11% and 6% for humidity, windspeed and temperature respectively. Despite this, we choose to maintain our current weather limits, as the risk of equipment damage outweighs the marginal gains made in observing time.

3.4.5 Observation Database

A Structured Query Language (SQL) database of all jobs submitted to *pt5m* is hosted at the University of Sheffield, thus allowing observation requests to be submitted whilst the telescope is offline. Tables in the database include a list of *pointings*, the *exposures* associated with each pointing, and an *observing log*, which keeps track of the science images associated with each observation.

Submission of jobs to the SQL database is performed either via a web form[6] which interfaces to the database through PHP scripts, or by custom-written Python scripts which can be run anywhere, and interface to the same PHP scripts. The PHP scripts send the job details to the SQL database, which writes the ASCII files (Figure 3.5) to be picked up by the scheduler. Access to the database requires a username and password, which we administer directly to collaborators.

A flag in the database keeps track of the status of the pointing: this flag can be set to indicate that observations are pending, completed, that the pointing has been deleted, the observations were interrupted, or the pointing's validity expired before it was observed. The pointing ID is used as a key to identify pointings in the *exposures* and *observing log* tables. The observing log has an entry for each science frame obtained with the telescope, including relevant information from the FITS headers.

3.4.6 Alert Listener

In order to follow up the fastest-decaying transients, the human processing stage of an alert system needs to be bypassed, as this can lose valuable time. The Alert Listener was developed specifically to perform transient follow-up, and is composed of three elements: software to listen for alerts and schedule observations, software to automatically reduce the data, and software to publish the data.

[6] `http://slittlefair.staff.shef.ac.uk/pt5m/pt5mQueue.html`

The listening software monitors transient alerts via VOEvents (Seaman, Allan & Williams, 2008), and previously also via `skyalert.org`. It then analyses the alert information, and decides whether or not the event is worthy of follow-up observations with *pt5m*. Initially, new cataclysmic variables, new AGN-variability, and new unknown transients were followed up, and now also any new Gamma Ray Bursts which have also triggered immediate radio follow-up with the AMI Large Array radio telescope (Staley et al., 2013). Transients classified as anything else are not followed up at this stage. If follow-up is requested, the software dictates what observations are required (e.g. a light curve for new cataclysmic variables, or $BVRI$ colours for unclassified objects) and, using the magnitude given in the alert packet (if available), sets an appropriate exposure time.

The data reduction component of the Alert Listener is a software daemon (Figure 3.4) that automatically processes every science frame as soon as it has been read out. The following tasks are performed:

- De-bias, dark-subtract and flat field using the master calibration frames (Section 3.4.2). If no master frames are available from that night, the daemon searches backwards through previous nights until a master frame is found, up to a maximum look-back time of 30 nights.

- Using the SExtractor package, extract a catalogue of all stellar sources in the image.

- Match the extracted catalogue to published *Vizier* catalogues and conduct an astrometric calibration, allowing for rotation, shift and stretch.

The daemon also records this activity in a nightly log, including any failures to complete the tasks.

The third component of the Alert Listener automatically publishes the transient data. Each extracted catalogue associated with a transient alert is automatically uploaded to

the Gaia Science Alerts Calibration Server[7] which calculates and publishes a calibrated magnitude for the transient. We use this server because we have been providing photometric follow-up for the Gaia Science Alerts project (see e.g. Campbell et al. 2015c), and it is a useful repository for transient data. Although the server is designed for use with Gaia events, non-Gaia transients can also be uploaded and processed.

A recent upgrade to the Alert Listener has involved implementing an option to build stacked frames for each filter before uploading. The resulting frame is then uploaded, providing a more precise magnitude measurement, or a deeper upper limit, in the case where the source is not detected.

A detailed description of the transient follow-up and automatic data reduction software is given below in Section 3.7. Results from the implementation of the follow-up software follow in Section 3.8.

3.4.7 Autoguiding

Since *pt5m* does not have a separate autoguider, we have implemented autoguiding using the science frames. For each new image, the catalogue of sources extracted by the data reduction daemon (Section 3.4.6) is compared with that of the first science frame in the sequence. A minimum of 11 stars are needed for an accurate comparison; any fewer than this and the autoguiding process is skipped for that image. The median offset between the catalogues is then used to adjust the telescope position. The offset is applied during the readout of the following image. Guiding can be enabled or disabled at any time and is automatically re-established following a pier flip. The downside to our autoguiding implementation is that when particularly long exposures are needed, the guiding is unable to correct for tracking errors on shorter timescales than the exposure time. Equally, when very short cadence observations are required ($<$ 10 seconds including exposure time and read out), the data reduction process can sometimes cause minor delays. However,

[7] `http://gsaweb.ast.cam.ac.uk/followup`

observations of such short cadence are very rare. Section 3.5.4 describes the performance of the telescope tracking with and without autoguiding enabled.

3.5 Performance

All of the performance data given in this section refer to the final configuration of the telescope described in Section 3.2. The performance tests used in the analyses below were all conducted and analysed by myself.

3.5.1 Detector

The CCD readout noise, gain and flat field noise were measured from bias frames and dome flats using the photon transfer curve technique (Janesick, 2001) and are shown in Table 3.1. The gain is unchanged when binning, but the readout noise depends on the binning factor. The flat field noise is the same in all filters, to within the measurement uncertainties, suggesting the dominant source of this noise is variations in the physical size of the pixels, rather than some wavelength dependent response.

The dark current was measured at night with the dome, mirror and camera shutters closed. Exposure times of 20 minutes were used, and the dark current was measured as the median count level after the median bias level (measured only minutes earlier) had been subtracted. The detector shows negligible dark current (8 ± 4 e^-/pixel/hour), but this tends to be concentrated in hot pixels and hence it is important to use median-combined dark frames when reducing data.

By observing a star field with a range of exposure times in quick succession, we were able to analyse the linearity of the detector. We found that the detector performs linearly to within 0.5% at light levels of up to 45000 counts/pixel, but departs from linearity by up to 2% at higher light levels.

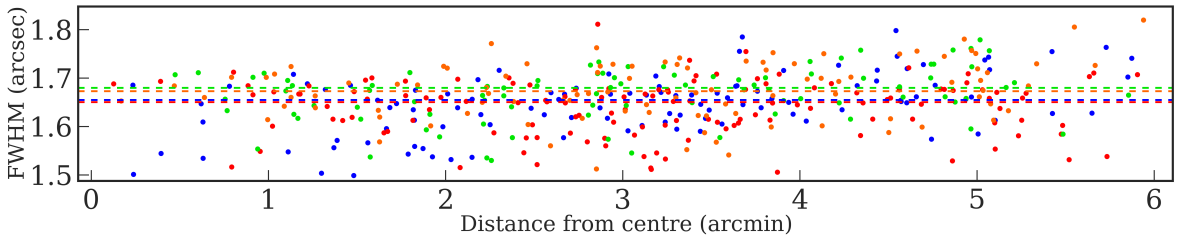

Figure 3.9: FWHM in the R-band as a function of radial distance from the centre of the CCD, measured from an image of the open cluster IC 4996. The standard deviation of the FWHM is 0.06" in this image. The four colours represent four different quadrants of the CCD; blue is top right, red is bottom right, green is top left and orange is bottom left. The horizontal lines represent the median FWHM in each quadrant.

3.5.2 Optics

To measure the performance of the telescope optics, we observed a star cluster filling the entire field of view in all five available filters ($H\alpha$, B, V, R, I). We measured the full width at half maximum (FWHM) of stars across the image. Our best measurements had a FWHM of 1.3". This is worse than expected given the median seeing at the site of 0.69" (Wilson et al., 1999) and the intrinsic optical performance of the OTA (which is close to diffraction limited at optical wavelengths, according to the manufacturer's tests). The most likely contributions to this discrepancy are the local seeing due to the roof-top location of *pt5m*, and the effects of wind shake.

The image quality appears to be largely independent of wavelength; the difference in FWHM from filter to filter is less than 10% (approximately 2σ). The image quality degrades by approximately 0.05" from the centre to the corners of the CCD, as shown in Figure 3.9, in accordance with the expectation from ray tracing of the telescope optics.

Flat-field images show signs of vignetting, with the drop in flux from the bright central region to the darker corners of a typical sky-flat being 10-15%. The vignetting pattern is centred slightly above chip centre, suggesting that the science camera is mounted just off centre of the optical axis.

3.5.3 Zeropoints

The photometric zeropoints in the four broad-band filters were measured using the standard stars SA101-315 and SA102-620. Each star was observed twice on two separate photometric nights, at different airmasses. We corrected for extinction using the coefficients given in the ING Observers Manual, 1995[8]. The above-atmosphere zeropoints representing the magnitude of a star giving one photo-electron per second in each filter were then calculated. The final zeropoints, representing the mean of all four standard star measurements in each filter, are shown in Table 3.1.

Figure 3.10 shows the 5σ limiting magnitudes that *pt5m* can achieve as a function of exposure time and moon brightness, calculated using the measured zeropoints.

3.5.4 Pointing, Tracking and Guiding

The pointing accuracy of *pt5m* is 0.4' (RMS), with marginally worse pointing at low elevations and high declinations. The tracking accuracy of *pt5m* was measured from observations following the same targets for 4 hours. The worst measured tracking drifts were up to 30"/hour for certain areas of the sky (at low elevations and at high declinations), but on average tracking drifts of 12"/hour were seen. This drift is larger than expected according to the mount specifications, most probably due to small errors in the mount's polar alignment, and flexure in the telescope and roof-top platform. The drift limits exposure times to a maximum of 5 minutes. With the introduction of autoguiding (Section 3.4.7), positional fluctuations of only 1.3" (RMS) from the original position are seen.

[8] `http://www.ing.iac.es/astronomy/observing/manuals/ps/general/obs_guide.pdf`

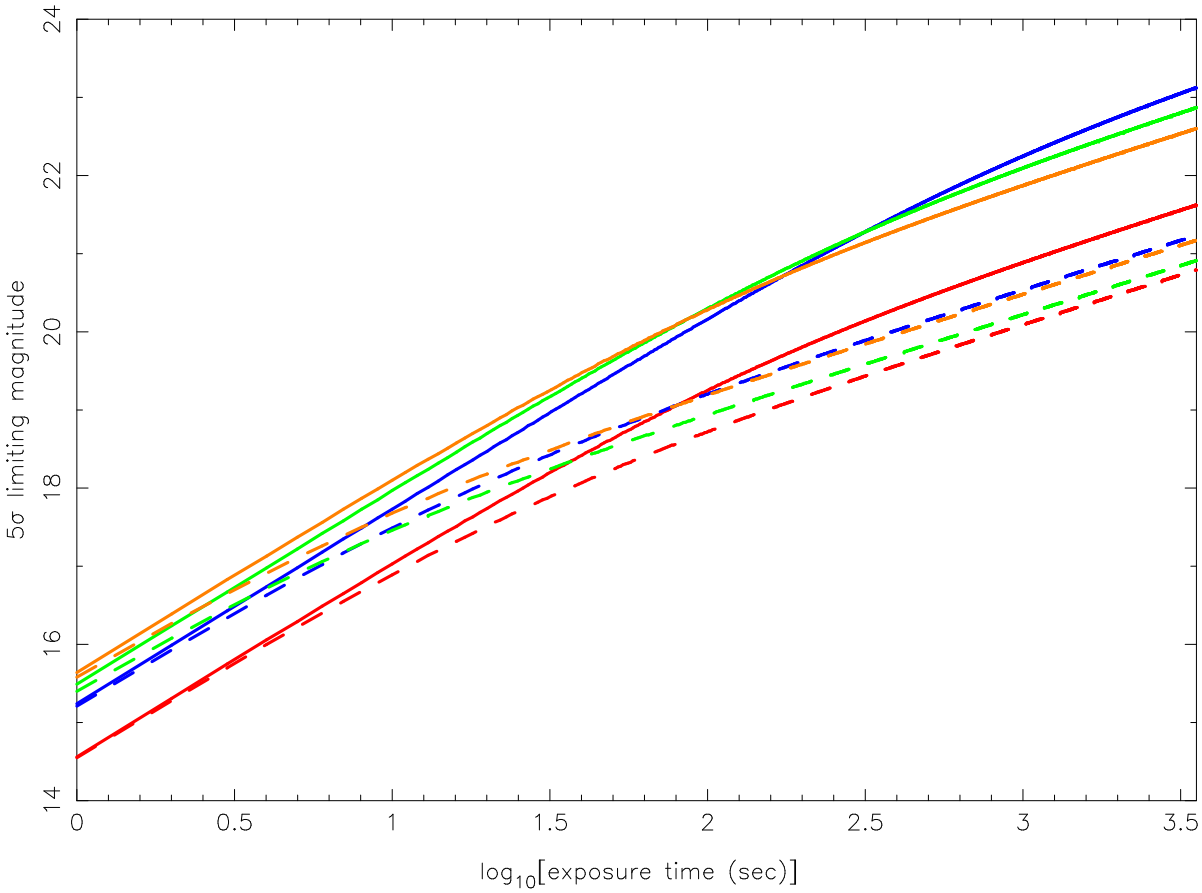

Figure 3.10: 5σ limiting magnitudes of *pt5m* as a function of exposure time. The blue, green, orange and red curves show the results for the B, V, R and I filters respectively. Solid lines show the results for dark time and dashed lines for bright time. The calculations assume seeing of 1.5", airmass of 1.0 and 2x2 binning. This figure was produced by Vik Dhillon.

3.5.5 Response Times

It takes less than two seconds from the announcement of a Gamma-Ray Burst as a VOEvent for the Alert Listener (Section 3.4.6) to submit a job for follow-up observation to the *pt5m* queue. The pilot checks the scheduler for the highest priority job every 15 seconds. We estimate a maximum slew time of 30 seconds, assuming that the telescope has to conduct a pier flip and move to the opposite side of the sky. Adding these together we can expect to be tracking on target less than a minute after an alert is announced.

Since the publication of the *pt5m* manuscript (Hardy et al., 2015b), we have received

two observable, real-time, live GRB triggers during periods of good weather. In both instances, the first *pt5m* images were taken of the GRB field just 50 seconds after the trigger was released by Swift. The release of the trigger is often a minute or two behind the first detection of the bursts, so our observations began within a few minutes of the burst. A more detailed analysis of the successful GRB afterglow detections is given in Section 3.8.3.

3.5.6 Residual Images

In October 2013, we serendipitously observed an apparent transient event whilst monitoring the active blazar 4C+38.41 (Gasparrini, 2013). The bright ($B\sim16$) transient was observed in the same field, but was located several arcminutes away and was not thought to be associated with the blazar. The transient appeared to fade quickly (0.2 magnitudes/minute) over the course of around 15 minutes, as seen in Figure 3.11, and moved slowly (0.5"/min) with respect to background stars. We could not identify the transient as a known comet, asteroid or minor planet, and the coordinates did not coincide with the orbit of any known man-made satellite. The object appeared to have a stellar point spread function, as shown in Figure 3.12. Bizarrely, the transient reappeared two weeks later, 18" away from the original position, and this time appearing as a double source.

After exhausting all possible astrophysical origins, we later discovered that this fast-fading transient signal was in fact the residual image of a bright star used for focusing just minutes earlier. Photo-electrons appear to have been trapped in impurity sites in the substrate of the chip. After the CCD was read out and cleared, these trapped electrons gradually leaked out during the next few exposures, creating a spurious signal in the resulting images.

This residual image phenomenon is not uncommon in infrared detectors, but is rare in optical detectors (Rest et al., 2002). It transpires that our Kodak KAF-3200ME CCD

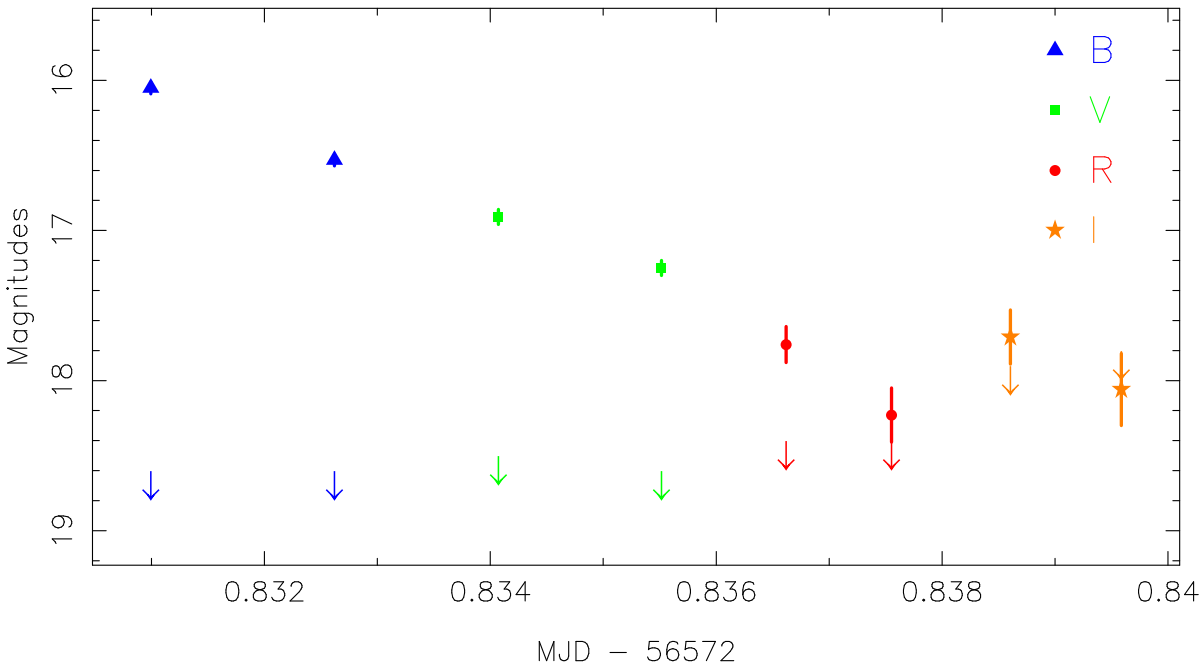

Figure 3.11: Light curve of the first transient source observed on 2013-10-07. The downward arrows represent the 5-σ limiting magnitudes of each image.

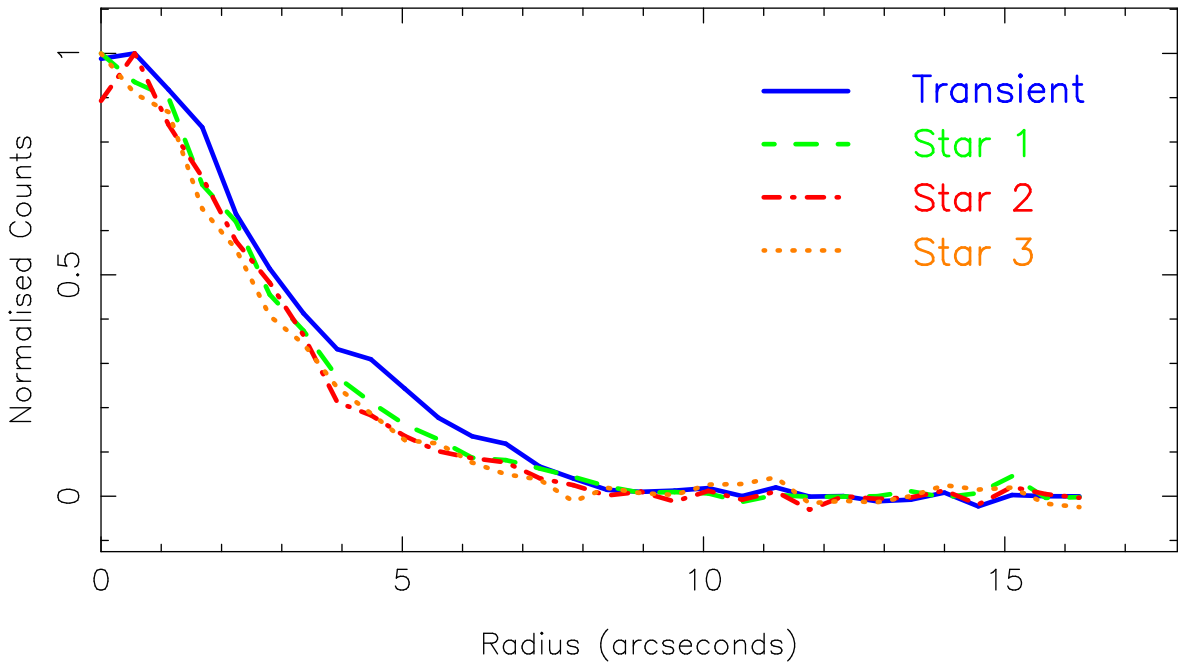

Figure 3.12: Radial profiles of the mysterious transient and three comparison stars in the field. The transient profile is similar to those of the nearby stars, suggesting an origin above the atmosphere.

is known to exhibit this behaviour[9]. Clearly, care must be taken when analysing images taken immediately after exposures containing bright stars.

3.6 Example Science: Transiting exoplanet WASP-33b

Note - this study was conducted by myself as part of my MPhys project report and included in the MNRAS publication (Hardy et al., 2015b). More recent observational campaigns, which focus directly on testing and demonstrating the automated transient follow-up software are discussed in Section 3.8.

We used *pt5m* to conduct an extensive study of the transiting exoplanet WASP-33b (Collier Cameron et al., 2010). This hot Jupiter is unusual in that it orbits an A-type star which itself exhibits δ-Scuti pulsations. The planet is also observed to be hotter and more bloated than most other exoplanets. In addition, there have been tentative claims of a resonance between the period of orbit and period of pulsations, suggesting a tidal link (Herrero et al., 2011).

We observed a total of 15 transits of WASP-33b in 2011 and 2012. We also collected several hours of out-of-transit data in order to study the stellar pulsations. The observations were phase folded using our own ephemeris, binned and combined using inverse-variance weighting. The combined transit light curve is shown in Figure 3.13. We model the planetary transit using the geometrical analysis of Sackett (1999), finding a planetary radius of $1.603 \pm 0.014\ R_J$ and an orbital inclination of $87.04 \pm 0.26°$. Uncertainties were estimated using the 'prayer-bead' technique of residual permutation (Désert et al., 2009), in which false light curves are generated by shifting the residuals between each datapoint and the best fit light curve along point by point. If the number of data points in the light curve is N, then $N - 1$ false light curves will be generated.

[9] `http://canburytech.net/QSI532/RBI.html`

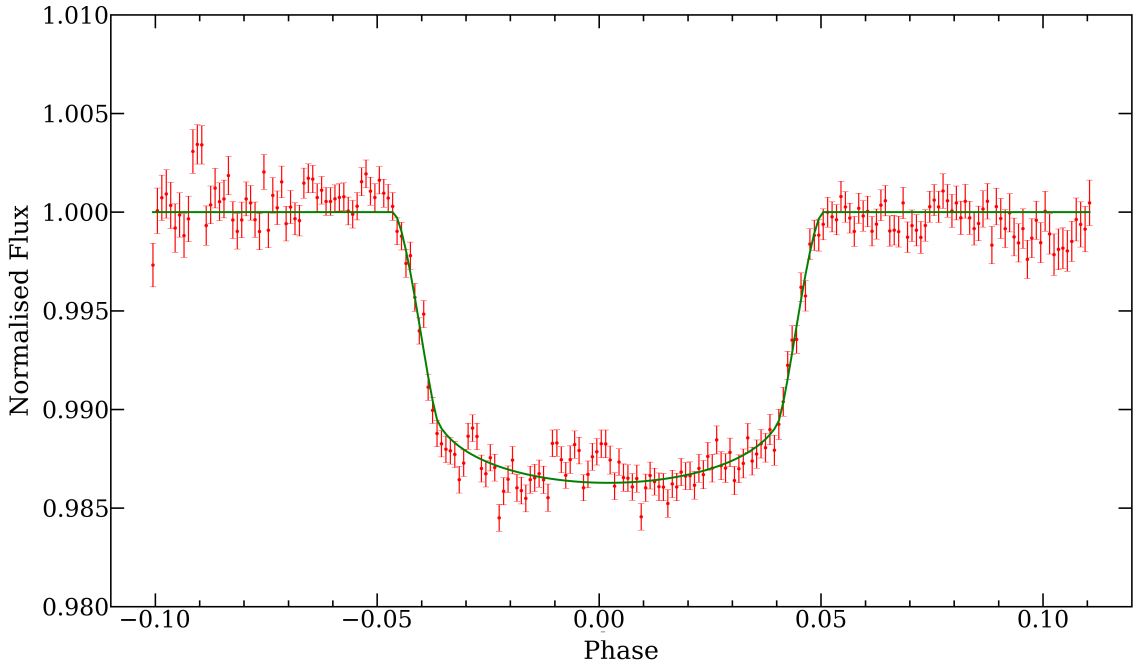

Figure 3.13: Transit light curve of WASP-33b, comprised of 15 individual *pt5m* transits. The observations have been phase-folded, binned and combined using inverse-variance weighting. The solid curve shows our best fit model. See text for details.

These light curves are then fit individually by the model, and the standard deviation of measured parameters in all of these fits gives an estimate of the final uncertainty in the fit.

Our measured inclination matches well with previous studies (Collier Cameron et al., 2010; Herrero et al., 2011; Smith et al., 2011; Kovács et al., 2013; von Essen et al., 2014), but the planetary radius is found to be up to 3σ larger. This highlights the fact that WASP-33b is substantially bloated in comparison to other exoplanets. Our measurement of the radius is plotted on a mass-radius diagram and compared to theoretical models of exoplanetary structure in Figure 3.14. The measured radius is almost 20% larger than expected from planetary models (Fortney, Marley & Barnes, 2007), although it should be noted that these models are not tailored to such a close-in exoplanet as WASP-33b.

The origin of an enlarged planetary radius could be attributed to a number of factors, including tidal heating (Bodenheimer, Lin & Mardling, 2001), Ohmic heating (Batygin

& Stevenson, 2010) and kinetic heating (Guillot & Showman, 2002). The circular orbit of WASP-33b (Smith et al., 2011; Deming et al., 2012) means that significant tidal heating is unlikely, but Ohmic and kinetic heating may play a role. However, we propose an additional and simpler reason for its enlarged radius; the young age of the system. The age of WASP-33 could be as young as 10 Myr (Moya et al., 2011), making it one of the youngest known exoplanet hosts. The size of an exoplanet is expected to decrease with time after formation, as the planet cools and contracts (Fortney, Marley & Barnes, 2007). In addition, standard planetary models usually use arbitrary starting conditions, with their main scientific interest being the state of the planet some Gyrs after formation, by which time their starting conditions have little effect (Fortney & Nettelmann, 2010). In this sense, it is not surprising that WASP-33b is found to be bloated with respect to these models, as it is much younger than the planets the models are designed to simulate.

Our phase-folded light curve (Figure 3.13) shows a mid-transit bump, which is present in several of the individual transit light curves. The star spot explanation for this phenomenon proposed by Kovács et al. (2013) is unlikely, as the host is an A-type star. Instead, we propose that the bump may be due to a pulsation in the host star resonating with the orbital period, and therefore is still visible after phase folding. However, von Essen et al. (2014) find no conclusive evidence for such star-planet interactions, and we found no pulsation periods in resonance with the orbital period in power spectra of the out-of-transit light curves. The nature of the bump remains unclear.

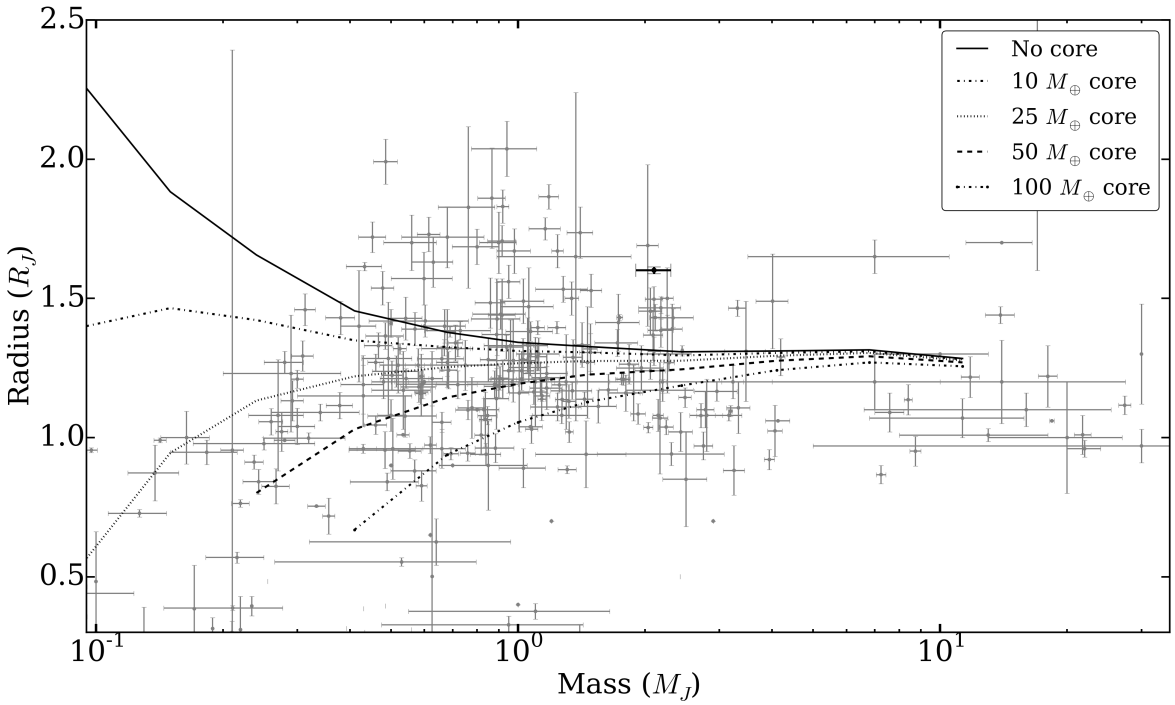

Figure 3.14: Mass-radius diagram showing the known exoplanets with measured or estimated masses and radii, as of July 2015 (Schneider et al., 2011), plotted as grey points. WASP-33b is shown in black, with the planetary mass measured by Lehmann et al. (2015). Also shown are planetary models of a H/He planet of age 300 Myr and a scaled orbital separation of 0.02 au, with five different core masses (Fortney, Marley & Barnes, 2007).

3.7 Transient Follow-up and Data Reduction Software

In Section 3.4.6 I briefly discussed the transient follow-up software. This section will expand on this and explain in more detail the software I have developed for *pt5m*, all of which was written in Python.

3.7.1 Job Submission

`submitdaily.py`

At first, the only way to submit a job/observation/target to the *pt5m* queue was to use the web form. This works for a single job needing no repetition and when the observations are not urgent. However, often we need to submit targets for long-term monitoring, adding a new job to the queue every night. We also need to be able to submit urgent jobs automatically, without human input.

The solution to this problem was to build scripts which could submit jobs to the queue, based on an input file describing one or more jobs and their requested exposures. The scripts would then be run daily or weekly at a certain time using the Linux 'cron' capability.

The daily job submission script reads in a file, with each line representing an individual job and including all the parameters needed to define the job (see Figure 3.5 for examples of the required parameters). Some parameters which do not depend on the job request, such as the start time, end time and user, are hard-wired into the script rather than given in the input file. Many of the key components of the job submission script were used later in the automated transient listening scripts (Sections 3.7.3 and 3.7.6).

Recently, I have made an addition to this script, to enable certain objects to be flagged as Gaia transients. The script can then include these objects in the SQL transient database on the *pt5m* control PC. This allows *pt5m* observations of Gaia transients to be automatically noticed, reduced and uploaded to the Gaia Science Alerts (GSA) server by the transient follow-up scripts (see Section 3.7.5 below).

3.7.2 Building Master Calibration Frames

`make_cals.py`

Science frames taken with *pt5m* need to be corrected for bias and dark current and flat-fielded, to improve the accuracy of the images. We needed to automate the process of producing master calibration frames (biases, flat fields and darks) in order to develop an

automated data reduction pipeline (see Section 3.7.4). The calibration-frame building script is run by the pilot at the end of evening twilight. It logs all output to a daily log file.

Bias Frames

The first stage of calibration is to subtract the bias level from the science images. The bias level is a constant voltage applied to the read-out capacitor on a CCD to ensure that the digital output can never be negative due to read-out noise. See Section 2.3.1 for more detail on CCD bias.

Usually, several individual bias frames are taken and median-stacked together to produce a master bias. This removes any spurious noise sources such as cosmic rays, and reduces the read-out noise per pixel. Since the bias will be subtracted from every science frame (and thus its noise will add in quadrature), we want to keep the read-out noise as low as possible. Individual bias frames are usually zero-second exposures taken with the CCD shutter closed and the lowest throughput filter in the beam (this is the $H\alpha$ filter in the case of *pt5m*). Every night, the *pt5m* pilot takes 5 bias frames in each of the three binning modes (1x1, 2x2, 3x3), and combines them to form master biases for each binning mode, as long as at least three frames exist for each mode.

Dark Frames

For many astronomical cameras, dark frames are no longer necessary because the CCD is cooled to below $-100°$C. However, since the CCD used by *pt5m* is only cooled to $-20°$C, dark current is a major issue that must be corrected for. Dark current consists of thermally excited electrons which are indistinguishable from photo-electrons. The dark current brings its own source of shot noise, and often dark current is concentrated in certain pixels, known as hot pixels. See Section 2.3.4 for more detail on dark current.

Dark frames are taken with the shutter closed, but unlike bias frames, they must

have long exposures, to build up an accurate indication of the dark current. Ideally, darks should have the same exposure time as the science frames in use. In practise, final dark frames can be divided or multiplied by the necessary constant in order to match the exposure time of the science frames. The *pt5m* pilot takes 5 × 60-second dark frames in each binning mode, every evening after the bias frames. The calibration script subtracts the master bias from these frames, then makes a median stack of the 5 dark frames in each mode, again only if there are at least 3 frames available.

Some astronomers use raw dark frames which have not had the bias level subtracted. This essentially skips one step in the process, and allows you to remove the bias level and dark current with one frame. However, we opt to keep these steps separate, and always produce master bias and master dark frames. This allows more flexibility, in case a user does not want to remove the dark current, and allows us to more easily track the bias level and dark current with separate frames.

Flat Fields

Ideally flat fields are images of a bright, uniform source. Flat fields are needed to remove the variation of pixel sensitivity across the chip. Some pixels perform better than others, and some regions of the chip might receive more incident light than others, due to vignetting of the optics, for example. This small and large scale structure in pixel sensitivity can be accounted for by dividing a science frame (after bias and dark subtraction) by a master flat field.

Flat field images are usually taken during morning or evening twilight, when the sky is bright. Often several flats are taken in order to reach very high signal-to-noise ratios (SNRs). If the pixel to pixel variations are of the order of 1%, we need a flat field SNR of at least 100, otherwise the shot noise may be larger than the pixel variations. This would mean a signal of at least 100,000 is needed. The twilight sky is also uniform (as long as there are no clouds), and thus serves as a good flat field. Every night *pt5m*

automatically attempts to observe 5 sky flats in each filter in the evening and morning twilight, while pointing at blank regions of the sky (i.e. containing no bright stars). Between each image the telescope makes a small offset, so that if there are any stars in the field they will be found in different locations in each image and will then disappear when the frames are median stacked. Five evening flats may not be sufficient to reach the required SNR, but since the twilight sky fades quickly, it can be difficult to obtain more than five images of sufficient signal. Some astronomers use dome flat (images of the smooth inside of the telescope dome) as back-up flat fields when necessary.

The calibration script first checks that there are no clouds in the sky, by investigating the most recent cloud monitor value. If the cloud monitor oktal is above zero, the script skips the production of flat fields for that night, noting this in the log file. However if the sky is clear and the cloud monitor returns zero, the script will build master flat-field frames for each filter that was observed, if at least 3 frames are available. The combined flat-field images are normalised by dividing every pixel by the mean pixel value of the whole image.

3.7.3 Listening to SkyAlert

`skyAlertListener.py`

`eventProcessing.py`

`jobSubmission.py`

The first transient follow-up software we developed was designed to listen to transient alerts from CRTS which were announced via the website `skyalert.org`. This service, which is no longer operating, allowed users to design their own RSS feeds, selecting their own parameters. We set up a feed for all CRTS transients which had been classified by human interpretation and identified with a likely source class (e.g. CV, SN, AGN, high proper motion star, etc.) based on its historical light curve and similarity to other

objects. My aim was to pick up these events, select those of interest (CVs, AGN and those labelled 'other', i.e. unknown sources), and automatically submit follow-up observations to the *pt5m* queue.

To ensure it avoids repeating observations on past events, the script compares all events in the stream to those which have already been analysed and saved in an SQL database hosted on the control PC. For every new event, the script builds a portfolio of information, extracting the object name, RA, DEC, event type and an event magnitude from the full event information string.

Next the script decides what to do with each event. It checks for four key factors: whether the coordinates of the event are observable from La Palma (i.e. if the declination is above -20°), whether the event is recent enough to still be worth observing (if the time between now and when the event was first discovered as a transient is less than 10 days), whether the event type matches any of the classes that we are interested in, and finally whether the event magnitude is bright enough for us to observe it with *pt5m*. Only if all of these conditions are met will the event be flagged as observable.

The script decides what observations will be required, based on the event magnitude and classification. It selects a maximum Moon brightness based on the target magnitude, and also uses this value to dictate the exposure times required, based on the measured zeropoints of the telescope. If the event is a CV, a light curve is requested in order to search for eclipses. The script requests the V filter, and calculates an appropriate exposure time in order to reach a signal to noise ratio of 20 (within an allowed range of 30 to 300 seconds). The script chooses a priority setting of 5 (lowest priority, but can be resumed after an interruption), and submits the required observations directly into the *pt5m* queue. If the event is an AGN, the script requests two images of the target in both the $H\alpha$ and R filters, with exposure times set according to the magnitude. The queue priority is set to 2, since the observations will not take very long to complete and follow-up within a short time frame is useful. If the event is classified as 'other', the

script requests two images in each of the B, V and R filters, also with a priority setting of 2.

In all cases, the requested observations are submitted to the queue immediately, so that they can be observed at the next available opportunity, which might be right away. Whenever the script triggers follow-up observations of an event, it sends an automated email to me with information about the alert.

3.7.4 Automatic Reduction of Science Frames

`pipeline.py`

One of the initial goals of the transient follow-up software was to be able to release immediate results from *pt5m* observations of new transients. For example, if a transient was detected and classified by a wide-field survey, and *pt5m* was able to conduct immediate follow-up observations confirming the presence or nature of the event, the results might be useful to other follow-up teams preparing observations. The *pt5m* pipeline allows for automated data reduction and the immediate publication of results.

Prior to being processed by the pipeline, images are first astrometrically calibrated using the Starlink package STILTS. This matches the stars in the image to those in the published *Vizier* catalogues (such as UCAC4, 2MASS, USNO, SDSS), finding an astrometric solution and adding this to the FITS header. This solution describes the RA and DEC for every pixel in the image.

The pipeline operates on every science frame taken by *pt5m* in real time. The pipeline first searches for the appropriate calibration frames that match the science frame's binning and filter settings. If a matching bias, dark or flat field image cannot be found in that night's data, it will look back one night at a time, searching for a suitable frame. If the script is unable to find a suitable calibration within the last 30 nights of data, it gives up and works instead on the raw image. If suitable calibration frames have been

found, the script applies these to a temporary copy of the science image (subtracting the bias, scaling and subtracting the dark frame, and dividing by the normalised flat field).

The key stage of the pipeline script is to use SourceExtractor (SE, Bertin & Arnouts 1996) to extract a catalogue of all detected sources in the image. For a more detailed explanation of how SE operates, see Chapter 5. The resulting catalogue, listing all extracted sources within the image along with their RA and DEC coordinates and instrumental magnitudes, is saved as a FITS catalogue. These catalogues are then ready to be analysed by the transient follow-up script described below.

3.7.5 Automatic Upload to GSA Server

`transientFollowup.py`

`gaiaSubmitModular.py`

Once the catalogue of sources has been extracted (Section 3.7.4), we next want to measure a magnitude for the target. To do this, we take advantage of the photometric calibration facility provided by the Gaia Science Alerts (GSA) collaboration[10]. The calibration server will match an astrometrically calibrated catalogue of sources with a published photometric catalogue, in order to identify the filter used and measure the intrinsic zeropoint of the image. With a given location (and search radius) for the target in question, the calibration server will then measure a calibrated magnitude for the object, based on the photometry of all stars in the image.

The script begins by analysing the headers of all new images, comparing the unique pointing/job identification number with the list of ID numbers in the SQL database of transients. If there is a match (i.e. the image being considered is associated with a transient event), the script then searches for the corresponding source catalogue for the transient image which should have been produced by the pipeline. Sometimes the

[10]http://gsaweb.ast.cam.ac.uk/followup

pipeline is still processing the image, and the script therefore waits for a maximum of 20 minutes for the corresponding catalogue to be generated, before giving up and moving on to the next image. At very high cadence, images can be produced almost as quickly as they can be processed, and if the network connection is slow or intermittent, it can take some time to astrometrically calibrate the image by comparison with online astrometric catalogues. When the matching catalogue has been found, the script submits it directly to the GSA server for calibration.

If the GSA server does not already have an entry for the transient in question, the script will automatically generate an entry based on the coordinates supplied in the original event alert. In addition, the script will delete any data points from the server which have been associated with an incorrect filter. For example, occasionally the server will match the source catalogue to the SDSS r filter instead of Johnson R, and we remove these points automatically to avoid any future confusion.

The GSA server always returns information about the uploaded data point, including the measured zeropoint, the number of sources used in the calibration, the measured magnitude of the target and an uncertainty on this magnitude (which is set to be -1 whenever the source is undetected, in which case the given magnitude is a detection limit). The datapoints and resulting light curve are available publicly on the GSA webpage[11] as soon as the data has been uploaded. Figure 3.15 shows an example light curve of a CV which was automatically followed up and reduced by *pt5m*. This CV is discussed in more detail in Chapter 4. As well as the key information being saved in an SQL database, every step of the procedure in this script is saved to a log file.

Recently, the process of uploading data to the GSA server has been upgrading, allowing frames to be stacked before submission to the server. This allows more precise photometry to be measured, and deeper magnitude limits in the case where the source is not detected. Additional upgrades have been discussed, including the option to automat-

[11] `http://gsaweb.ast.cam.ac.uk/followup/list_of_alerts`

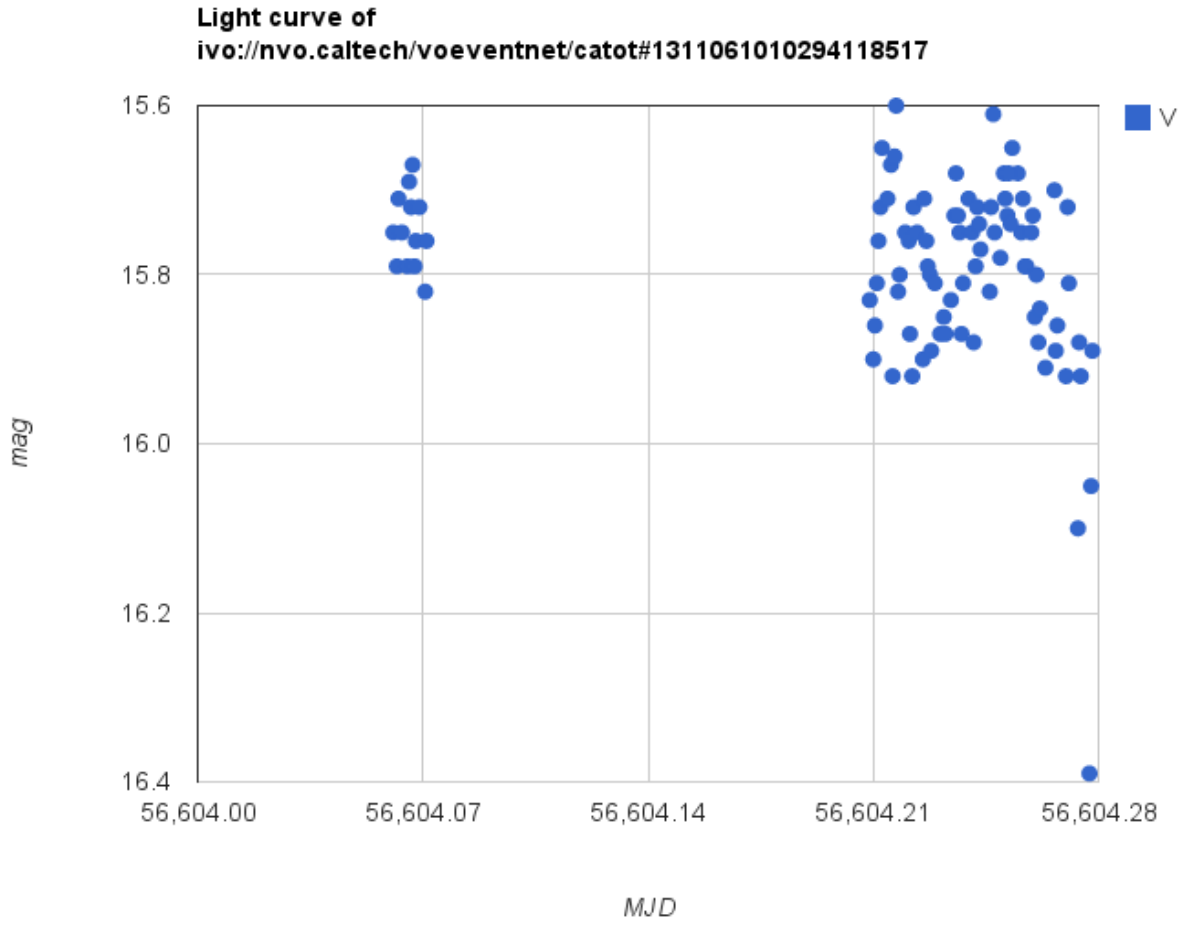

Figure 3.15: Automatically generated V-band light curve of the CRTS transient with identifier 'catot#1311061010294118517'. This object was automatically observed by *pt5m* after being classified as a CV by CRTS, and the images were automatically reduced and uploaded to the GSA calibration server by my scripts. The system turned out to be eclipsing, and the beginning of an eclipse can be seen in the final data points. See Chapter 4 for more details on this particular object.

ically delete points with bad data, perhaps based on the calibration zeropoint returned by the server. If for example the returned zeropoint is several magnitudes off from its usual value, this could indicate a serious inaccuracy in the photometry, caused for example by clouds or by severe blurring of the image, perhaps due to wind shake, dome vibrations, or tracking errors.

3.7.6 Listening for Swift/AMI VOEvents

`listenForVOEvents.py`

`eventProcessing.py`

`jobSubmission.py`

In 2013 we began a collaboration with '4 Pi Sky' (Staley & Fender, 2016) with the aim of optically following up radio transients. It had become clear that the long-awaited real-time detection of radio transients from LOFAR (the LOw-Frequency Array for Radio astronomy, Fender et al., 2006) was still far from materialising. Instead we decided to attempt simultaneous follow-up of gamma-ray bursts (GRBs) at both optical and radio wavelengths, with *pt5m* and the Large Array of the Arcminute Microkelvin Imager (AMI), located in Cambridgeshire. Simultaneous optical and radio fluxes, or detection limits, can help constrain energetics and emission mechanisms in GRBs (van der Horst et al., 2015).

The 4 Pi Sky team aim to follow-up new GRB alerts whenever possible with the AMI Large Array, hoping to detect and characterise radio afterglows. Tim Staley built a broadcaster to announce VOEvents (Virtual Observatory Events – a universal protocol for transient alerts), which could be picked up by the AMI telescope, and also by *pt5m*. I built a script to listen for these events, which are broadcast from a 4 Pi Sky webserver.

The script operates with 'factories' to handle individual connections to the host port supplying the VOEvents. Every time an event is received, or a time-out limit is reached, the factory closes the connection to the host port and attempts to make a new connection. This prevents connections hanging and the receiver missing any important packets or messages.

If a message includes the string 'ami-request', we treat the message as a Swift GRB trigger which is requesting follow-up with AMI and *pt5m*. In this case, we handle the event in a similar way to our earlier scripts, by building a portfolio with the important

event data. The script then checks for observability, that the alert is not more than 10 days old, and that the positional uncertainty is small enough that we expect to find the burst inside the *pt5m* field-of-view.

If the GRB follow-up trigger passes all of these tests, we immediately submit a job to the *pt5m* queue with ToO override status, so that we can begin observations as soon as possible. The script requests an R-band light curve of 60 $\times$ 60-second exposures. The job enters the queue with a priority setting of 1, and with no limit on the Moon brightness. We also add an additional job to the queue for the following night, which is only valid when the target is close to transiting the meridian. This is because radio observations are often attempted at $\sim$24 hours after the burst, and an optical detection or limit at the same time is also useful. The radio observers at AMI can then attempt a +24 hour follow-up observation when the target is close to transit, expecting that *pt5m* will also be trying to observe the target at the same time.

Finally, the script adds the event to the transients SQL database, and sends an automated email with information on the event to the *pt5m* and 4 Pi Ski teams.

3.7.7 Calibrating GRB Images and Sending to 4 Pi Sky

`transientFollowup.py`

`sendToLofar.py`

Initially we had hoped that our *pt5m* observations of GRB afterglows could be reduced automatically, with the results being announced to the community immediately. However, this is very difficult, primarily because the first and fastest alert trigger from the Swift Burst Alert Telescope (BAT) does not supply a sufficiently accurate position. The directional uncertainty in the initial trigger is usually of the order of several arcminutes. This is small enough to usually be found within the *pt5m* field-of-view, but too large for our automated reduction procedure to detect and analyse the GRB source. On top of

this, the GRB optical afterglow, if there is one, is usually quite faint, and can sometimes only be detected by stacking many images together.

One possible solution to this might be to await a more refined position from the Swift X-Ray Telescope or UV/Optical Telescope, which is often (but not always) supplied some minutes after the initial burst trigger. Once this position is available, *pt5m* could reduce the images automatically. It could then send a message to the *pt5m* managers, and perhaps even to the wider community if certain conditions are met, e.g. that an afterglow has been detected. We have not attempted to programme a complex follow-up procedure such as this, but it is something to keep in mind.

Instead, we currently calibrate any observations of GRBs and send the files directly to the 4 Pi Sky servers. As before, every new science frame is checked to see if the target was triggered by Swift/AMI as a GRB follow-up. If so, the image is calibrated (bias subtracted and divided by a flat field), and every morning the *pt5m* control PC sends any GRB follow-up images from the previous night to the 4 Pi Sky server.

3.8 Transient follow-up observations

My primary observing project with *pt5m* has been the hunt for eclipsing CVs, and this is discussed in detail in Chapter 4. However, since the development of the transient follow-up software, I have also ensured that *pt5m* has been used to obtain useful observations of a number of transient events, mainly by following up detections from surveys and monitoring the Astronomer's Telegram feed[12]. I have also supervised the automated follow-up of GRB events, manually reduced the data and published several notices regarding our *pt5m* observations.

[12] `http://astronomerstelegram.org/`

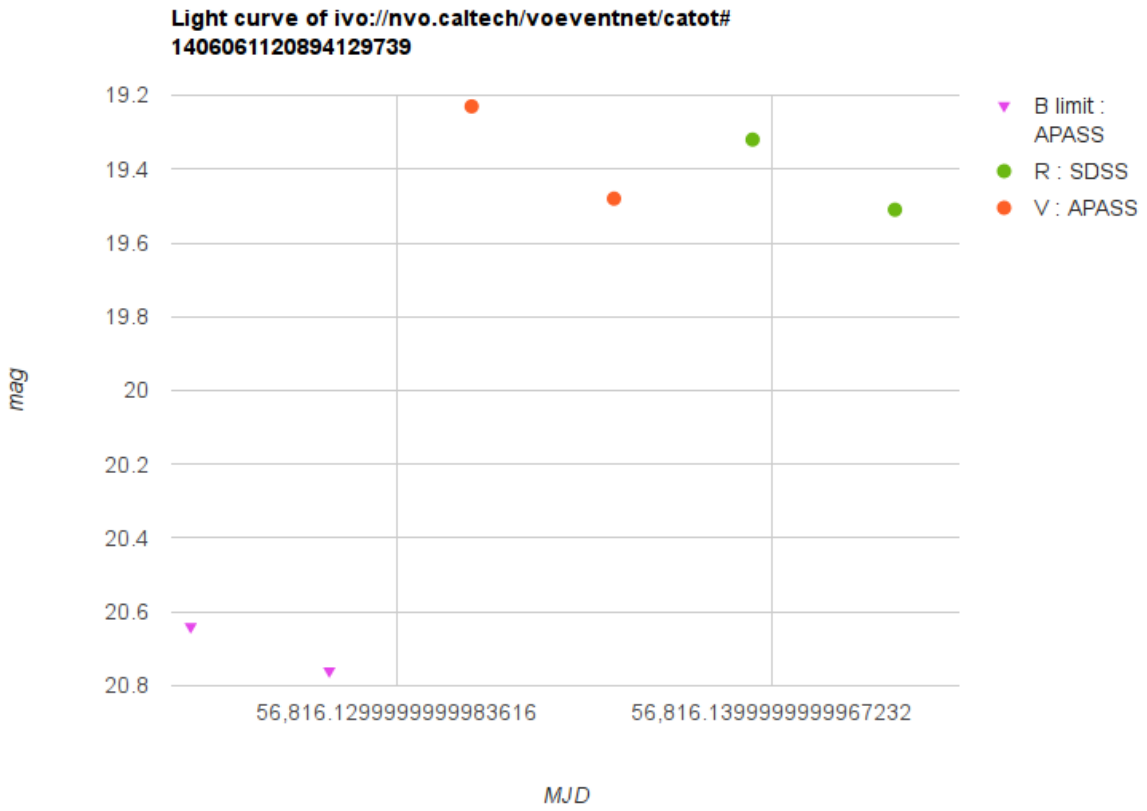

Figure 3.16: Light curve of CSS140606:164803+122120, which was automatically observed by *pt5m* and reduced by the transient follow-up software. The B-filter values are magnitude limits, suggesting that the object is quite red or highly reddened. Perhaps the original transient was a flare from an M-dwarf star or a quasar, and these data show the source in quiescence.

3.8.1 CRTS

My first transient follow-up campaign began in 2013 with observations of CRTS targets which had been classified as AGN, CVs or 'unknown'. The aim was to obtain BVR colours of unknown objects, such as CSS140606:164803+122120 - see Figure 3.16. I hoped that a reliable colour measurement might help classify the sources. In the end, very few objects were observed before the skyalert event stream was closed down.

For new CVs, I was searching for eclipses or superhump periods. A good example is CSS130828:205253-023953, also known as ASASSN-13cg (see Figure 3.17). The *pt5m* observations showed possible superhump behaviour in the automatic photometry taken

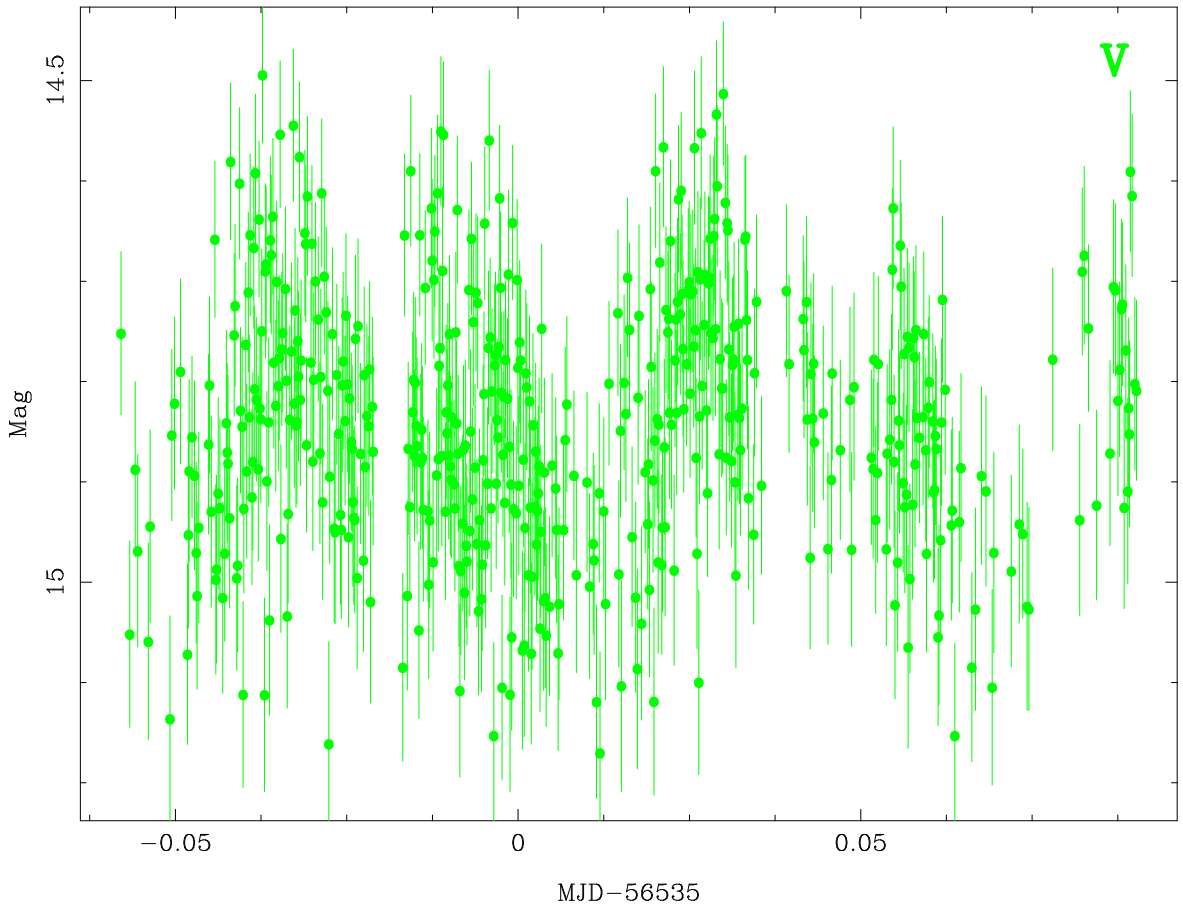

Figure 3.17: Light curve of CSS130828:205253-023953, which was automatically observed by *pt5m* and reduced by the transient follow-up software. The light curve shows possible superhump behaviour indicating this object was correctly classified as a CV.

during outburst. This agrees with the observations reported by the amateur observers network *vsnet* (alert 16302).

3.8.2 Gaia

We have been a member of the Gaia Science Alerts (GSA, Wyrzykowski 2016) working group for several years, and have always aimed to perform follow-up observations of interesting Gaia transients with *pt5m*. Even before the satellite was launched I observed a number of test targets, detected by other surveys but followed-up by the GSA working group in order to test the Gaia photometric calibration server. We contributed multiple

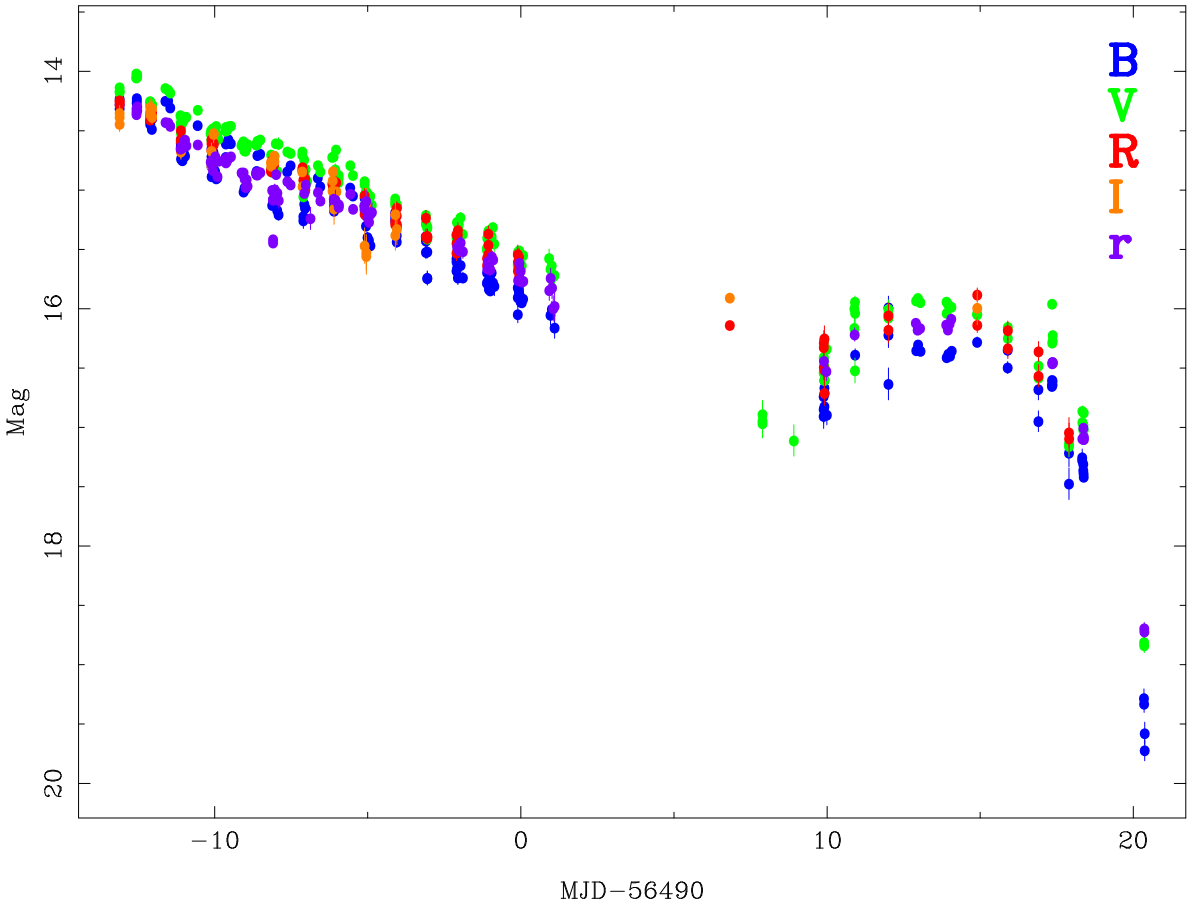

Figure 3.18: Multi-colour light curve of ASASSN-13ax observed with *pt5m* for the GSA collaboration, to help test the calibration server. Some points in this light curve, e.g. in the SDSS *r*-band, come from other observers in the GSA collaboration. The object is a CV, and the light curve shows the decline from superoutburst (Copperwheat et al., 2013; Kato et al., 2014a).

BVRI datapoints to the follow-up of the super-outbursting CV ASASSN-13ax, the full GSA light curve for which is shown in Figure 3.18. We also conducted multi-colour observations of the variable blazar ON325 (Figure 3.19), and the Type Ia SN ASASSN-14bb (Figure 3.20).

Our tests found some minor issues in the calibration server, such as some originally mismatched target coordinates, the need for the option to delete erroneous datapoints automatically, and the need for the option to enter new events into the database automatically.

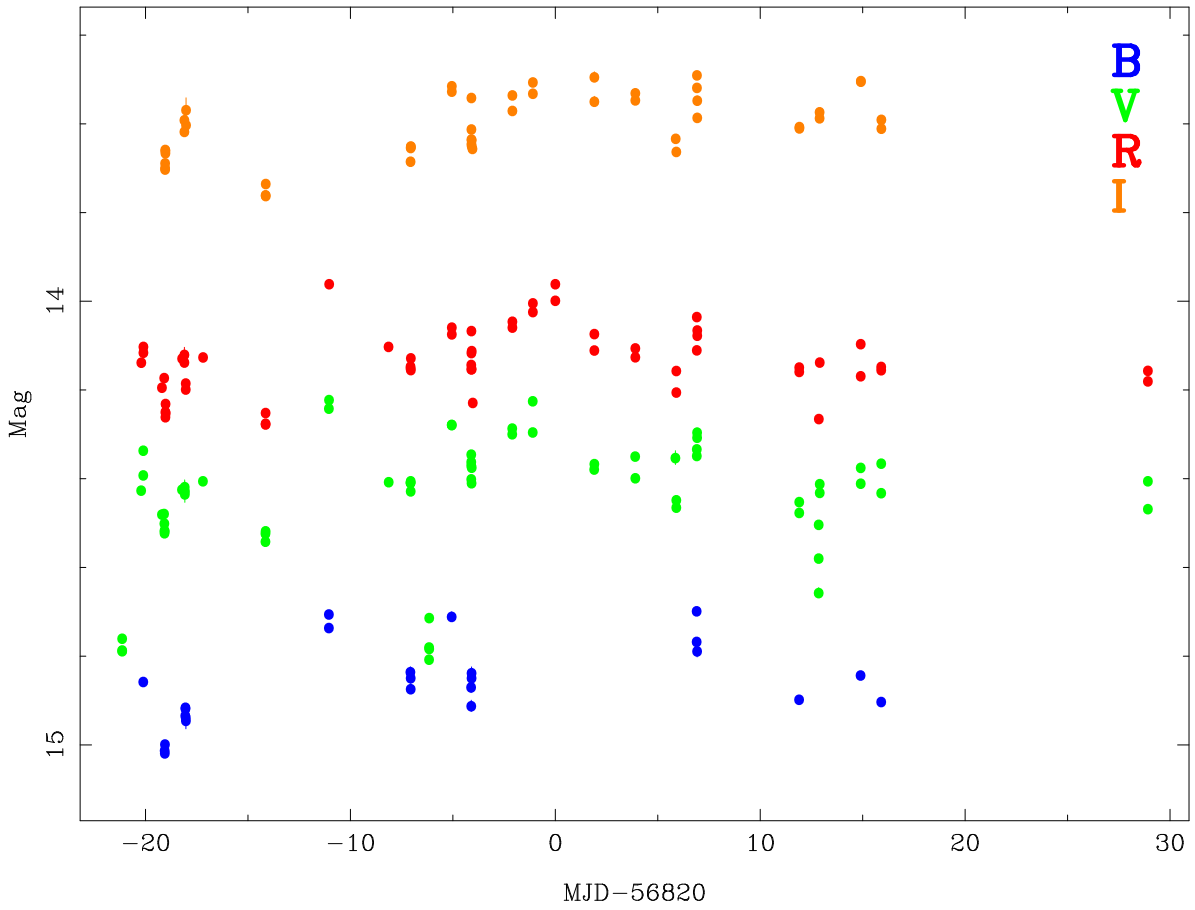

Figure 3.19: I observed *BVRI* colours of the BL Lac object ON325 for the GSA collaboration, to help test the calibration server. Some points in this light curve come from other observers in the GSA collaboration, but most are supplied by *pt5m*.

Since the full release of Gaia transients began in October 2014, I have been following up any objects of interest, particularly CVs and SNe when requested by the collaboration, whenever these objects are bright enough and visible to *pt5m* (see Rixon et al. 2014 for the first Gaia transients follow-up notice). For example, I helped refine an orbital period for the new AM CVn system Gaia14aae (Campbell et al., 2015c), and helped confirm possible supernovae candidates such as Gaia15aap (Campbell et al., 2015a). I also observed Gaia14acr and Gaia14acq, which had quickly returned to their quiescent brightnesses, suggesting they may have been stellar flares.

In total *pt5m* has contributed over 15,000 datapoints to the GSA calibration server, most of these submitted using my scripts. This is more than any other observatory in

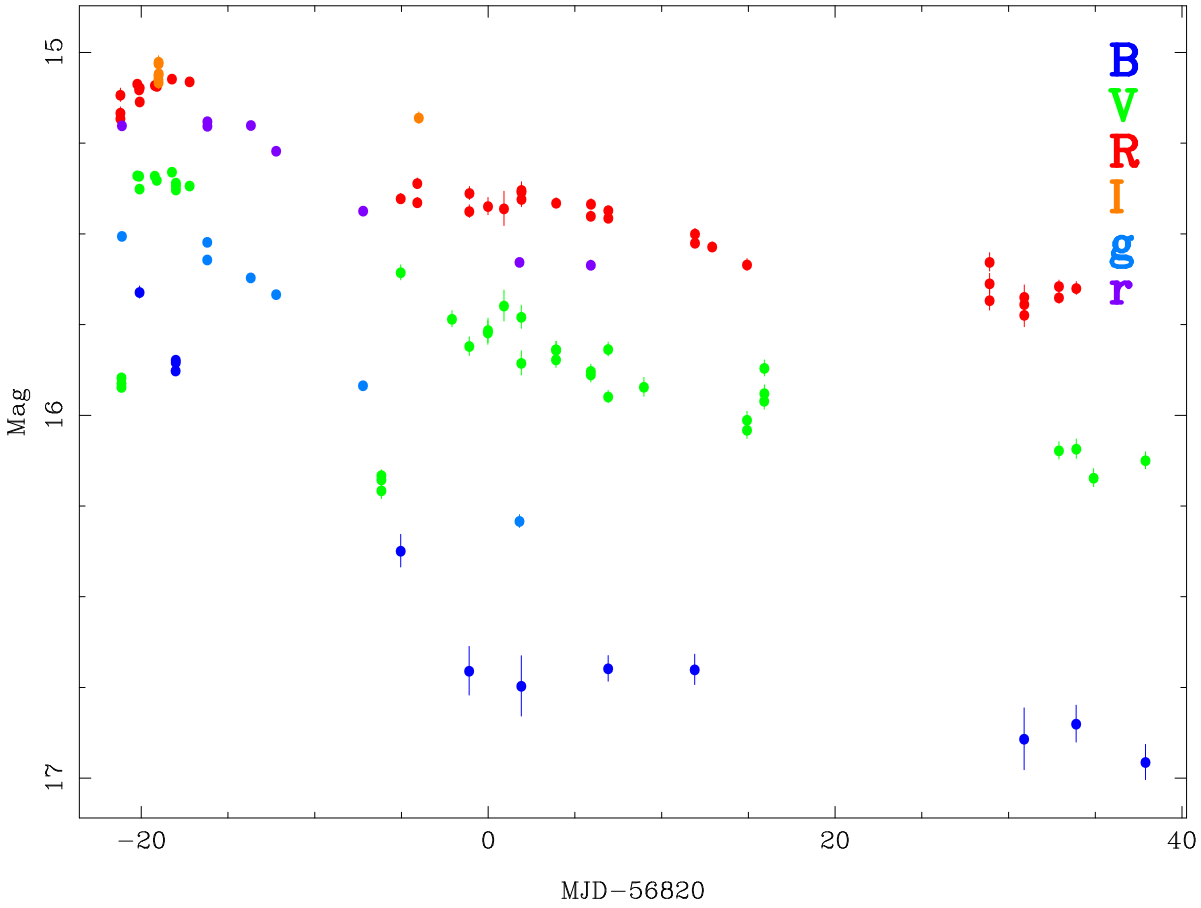

Figure 3.20: Multi-colour light curve of the Type Ia SN ASASSN-14bb, observed for the GSA collaboration to help test the calibration server. Some points in this light curve come from other observers in the GSA collaboration, but most are supplied by *pt5m*.

the collaboration. The *pt5m* team is continuing to monitor the Gaia transient stream and observe any objects we can.

3.8.3 GRBs from *Swift*

Although Swift GRB triggers arrive frequently (∼1 per week), it took many months before a suitable trigger arrived during night time, while the sky was clear, and at an observable location. *pt5m* has made several observations of GRB fields, sometimes several hours after the initial burst, and usually finding no optical counterpart at that time, to detection limits of around $R \sim 21$–22. However, as described below, I have

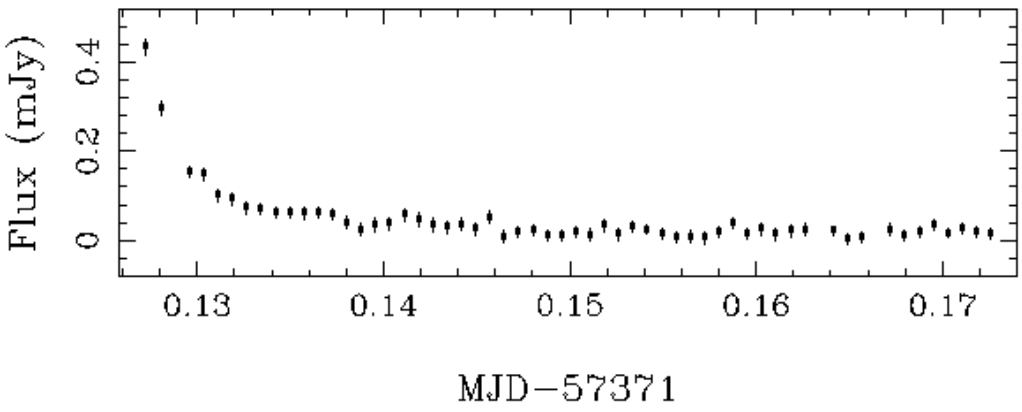

Figure 3.21: Decaying R-band light curve of the optical afterglow of GRB151215A, as observed by *pt5m*. Observations began just 53 seconds after the announcement of the burst, 74 seconds after the original burst detection.

now detected several optical counterparts and have successfully shown that *pt5m* can immediately follow-up a burst, beginning exposures on target 50 seconds after the release of a Swift trigger.

I have published two GRB Coordinates Network circulars, one announcing an optical counterpart detection and showing the decaying light curve of GRB151215A (Figure 3.21, Hardy et al. 2015d), and one providing early (few minutes post burst) non-detection and limits of GRB151118A (Hardy et al., 2015c). Unfortunately neither of these GRBs were simultaneously followed up with AMI, but another afterglow detection, GRB140907A, was observed by *pt5m* and AMI at an overlapping epoch of 9-10 hours post burst. I measured the afterglow (Figure 3.22) with a magnitude of $R = 20.52 \pm 0.29$, in 100 stacked 30-second exposures. The source was not detected in the AMI observations, but the lower limit was consistent with the expected flux from a long GRB (Rob Fender, priv. comm.).

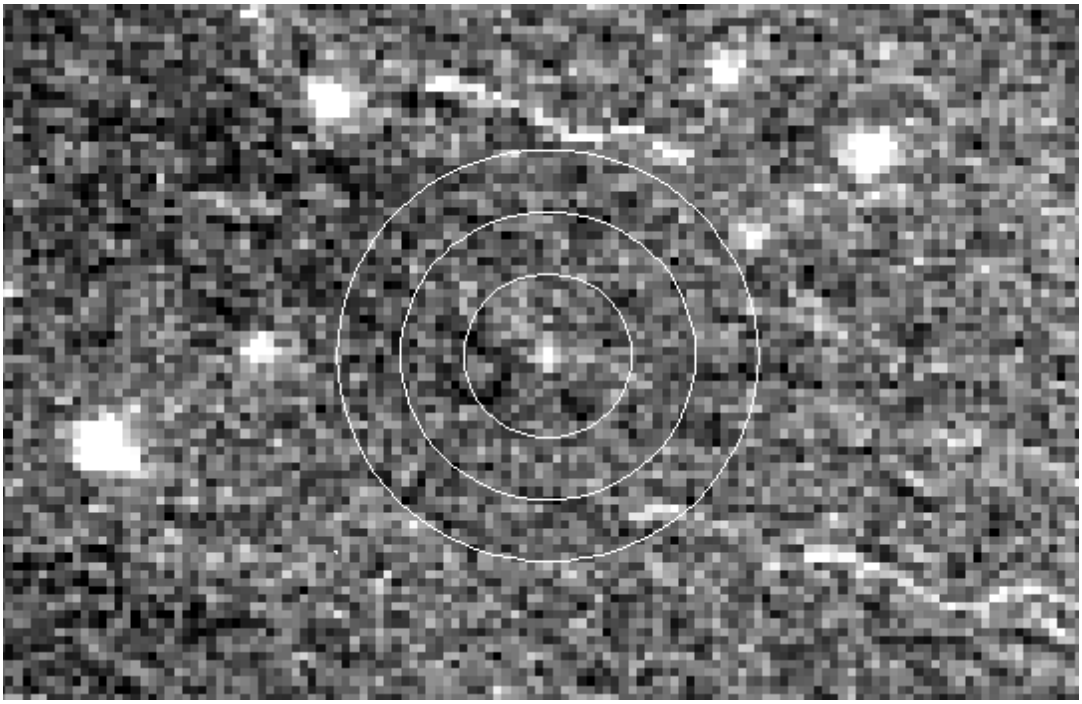

Figure 3.22: Optical afterglow of GRB140907A (Krimm et al., 2014) as observed by *pt5m*, visible inside the white circle. The afterglow was measured to be $R = 20.52$ approximately 9 hours after the initial Swift detection, which is consistent with other reports (Kuroda et al., 2014).

3.8.4 V404 Cyg

One transient object of particular interest that I followed extensively with *pt5m* was V404 Cygni. This black hole X-ray binary (Khargharia, Froning & Robinson 2010 and references therein) had been in quiescence since its discovery during an outburst in 1989. This lasted until 15th June 2015 when it began another bright outburst, first spotted by the Swift BAT (Barthelmy et al., 2015; Negoro et al., 2015; Sivakoff et al., 2015; Martí, Luque-Escamilla & García-Hernández, 2016). The June outburst lasted around 50 days, with observations at all wavelengths showing its excited state (Sivakoff et al. 2015 and references therein), and also showed slow and fast variability (Hynes, Robinson & Morales, 2015).

X-ray binaries can sometimes enter violent outbursts, via similar methods to that seen in cataclysmic variables (see Chapter 4). In a black hole binary, the black hole

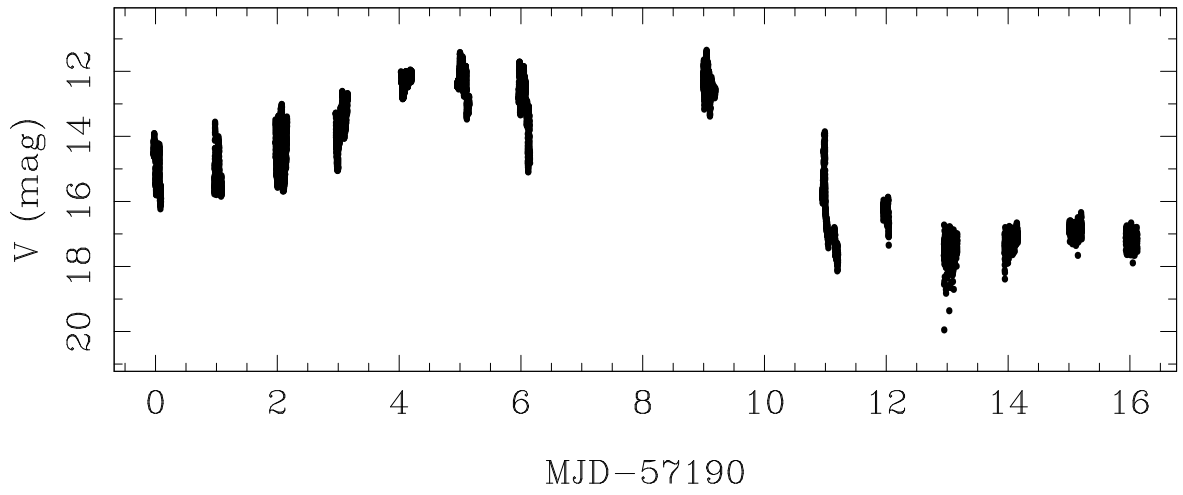

Figure 3.23: Multi-night V-band light curve of V404 Cyg during the June 2015 outburst, as observed by *pt5m*.

accretes matter from its companion via an accretion disk. If the mass accretion rate is low enough, the accretion disk will be unstable. The long-lived quiescent intervals can therefore be interrupted by violent outbursts when the gas is ignited (King, Kolb & Szuszkiewicz, 1997). These so-called X-ray novae can occur with rise times of a few days, and reach extreme luminosities before decaying over the course of a month or so (Życki, Done & Smith, 1999).

pt5m received an automatic trigger requesting observations of V404 Cyg via the 4 Pi Sky VOEvents broadcaster (Section 3.7.6). My scripts triggered observations immediately, and they began 5 hours after the initial Swift detection, as soon as the object was visible from La Palma. I observed the system for as long as possible on subsequent nights over the next two weeks. I also continued monitoring it once per night or so, over the following 50 nights until it settled back into quiescence. I published a light curve in an Astronomer's Telegram (Hardy et al., 2015a), showing the source brightening over several nights, and with increasing amplitude of variability. A longer duration light curve than that published in the Telegram is shown in Figure 3.23. I have supplied all of our *pt5m* observations to Greg Sivakoff, who is putting together a large collaborative paper with data from hundreds of observers at all wavelengths.

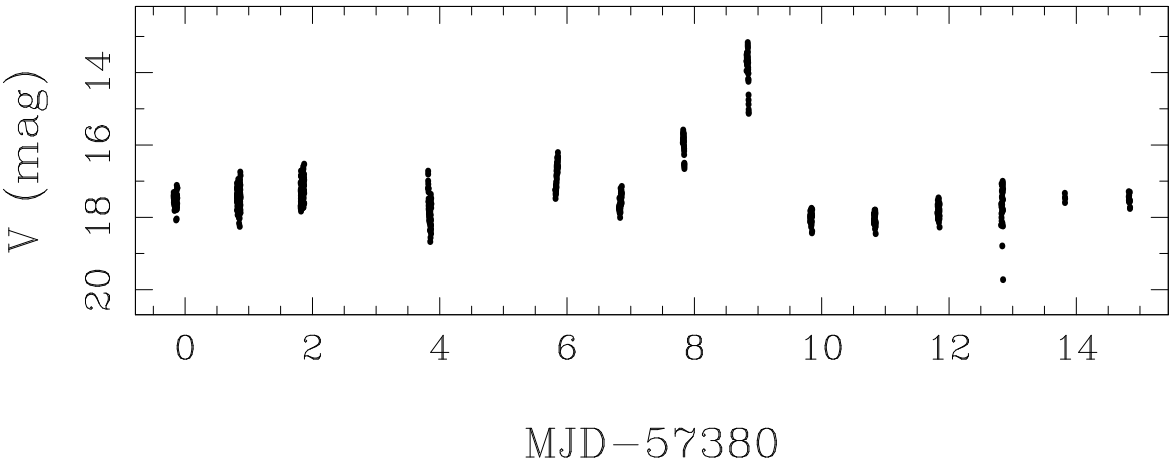

Figure 3.24: Multi-night V-band light curve of V404 Cyg during the December 2015 outburst, as observed by *pt5m*.

In addition to the long term observing campaign with *pt5m*, I also observed V404 Cyg with ULTRACAM during the June 2015 outburst. These observations were published in Gandhi et al. (2015, 2016) and Dallilar et al. (2016).

V404 Cyg's subsequent quiescent state was relatively short-lived, lasting only a few months. On 23rd December 2015 Swift BAT was again triggered by an outburst of V404 Cyg (Barthelmy, Page & Palmer, 2015). This heightened state was less luminous than in June, and lasted only a couple of weeks. Again, we observed the source with *pt5m* over the Christmas and New Year period, collecting long-term and short-term light curves which were published in another Astronomers Telegram (Hardy et al., 2016a). These light curves are shown in Figures 3.24 and 3.25. V404 Cyg appeared to return to relative quiescence in early January 2016 (Li & Kong, 2016), however due to its proximity to the Sun it could no longer be observed.

My *pt5m* observations will have been helpful to the community, by encouraging or discouraging further follow-up on larger telescopes. When the December 2015 outburst of V404 Cyg began to fade in the optical much faster than the June outburst, my *pt5m* observations showing this trend may have helped others avoid wasting valuable observing time on larger telescopes, if they were particularly interested in observing the system in

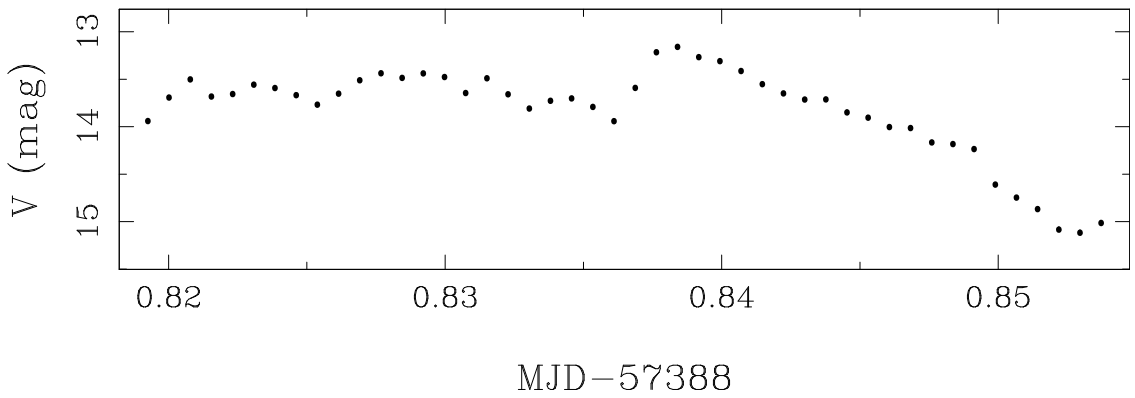

Figure 3.25: *pt5m* V-band light curve of V404 Cyg on 2016-01-01. These data were temporally coincident with sub-mm and X-ray observations, which also found V404 Cyg to be in a brightened state (Tetarenko et al., 2016; Motta et al., 2016).

outburst.

3.8.5 Astronomer's Telegram

Since the development of *pt5m* I have monitored the Astronomer's Telegram (ATel) feed and kept track of any interesting developments of transients in the northern hemisphere. Listening to the transient events via Astronomer's Telegram cannot be automated, because the telegram format is fluid rather than fixed. Nevertheless, it has been worthwhile to manually follow-up occasional events announced via ATels with *pt5m*.

For example, the residual images discussed in Section 3.5.6 were observed during the follow-up observations of the active blazar 4C+38.41 (Gasparrini, 2013), which I was alerted to via the ATel calling for multi-wavelength collaborative observations. I observed the blazar in the *BVRI* filters on 54 occasions over the course of 85 nights. The resulting light curve is shown in Figure 3.26. 4c+38.41 appears to be a very active blazar, having shown several bright outbursts in recent years (Arkharov et al., 2015).

However, probably the most interesting object that I studied following an announcement via an ATel was CSS160603:162117+441254 (Drake et al., 2016). This object had previously been classified as a short period W UMa star with a period of 0.207852

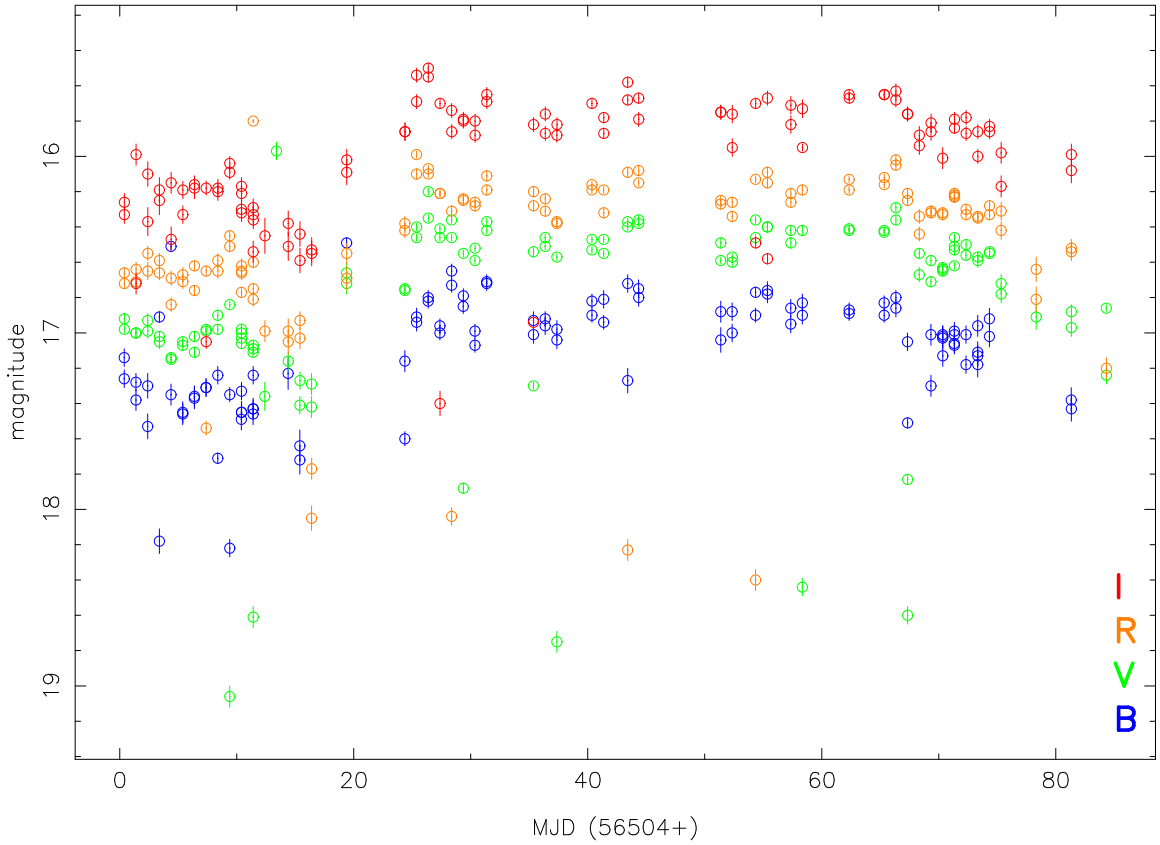

Figure 3.26: Multi-colour light curve of 4C+38.41 during its July 2013 outburst, as observed by *pt5m*.

(Palaversa et al., 2013), showing strong ellipsoidal variability (Lohr et al., 2013). The phase-folded CRTS light curve containing a total of 11 years of observations is shown in Figure 3.27.

Ellipsoidal modulation in systems like CSS160603:162117+441254 is caused by the elongation of the component stars filling their Roche lobes. When we see the system at phase zero (with one star eclipsing the other), we see a minimum, whilst at phases 0.25 and 0.75 we see the maximum visible surface area of the Roche-lobe filling stars, and thus a maximum in the light curve. If one star is larger or brighter than the other we will see different depths of minima at phases 0 and 0.5.

In June 2016 Drake et al. (2016) announced that CRTS had observed a brightening of around 1.5 magnitudes in the source, triggering their transient detection pipeline. If

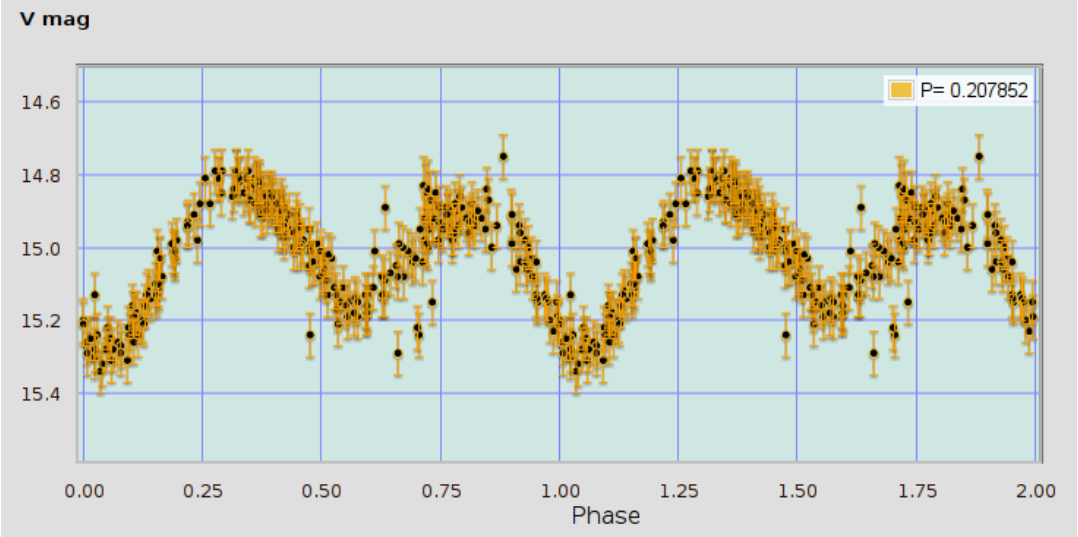

Figure 3.27: 11-year CRTS light curve of CSS160603:162117+441254 phase-folded on the period of 0.207852 days. The ellipsoidal variations with an amplitude of 0.5 magnitudes are visible, with a deeper primary minimum and shallower secondary minimum.

the object was indeed a contact binary system of two main sequence stars, the outburst could have represented a rare stellar merger event (Tylenda et al., 2011). However, Drake et al. (2016) also noted that the short orbital period and GALEX magnitudes (Martin et al., 2005) mean that this system was more likely to be a cataclysmic variable.

I observed the source with *pt5m*, first on the night beginning 2016-06-06, and on several subsequent nights. The object showed a very clear eclipse signal (confirmed independently by Zejda & Pejcha, 2016), indicating that the system was almost certainly an outbursting CV. This was confirmed soon after with time-resolved spectroscopy which measured the binary component masses to be approximately 0.9 and 0.4 $M_{\odot}$ (Thorstensen, 2016). Phase-folded light curves from five nights of observing CSS160603:162117+441254 with *pt5m* are shown in Figure 3.28. The system has faded from outburst over the course of about 10 days, and the light curve has gradually returned to its original ellipsoidal modulation seen in quiescence.

The stark difference between the outburst and eclipse light curves is surprising. Such strong ellipsoidal modulation is uncommon in CVs, because the light tends to

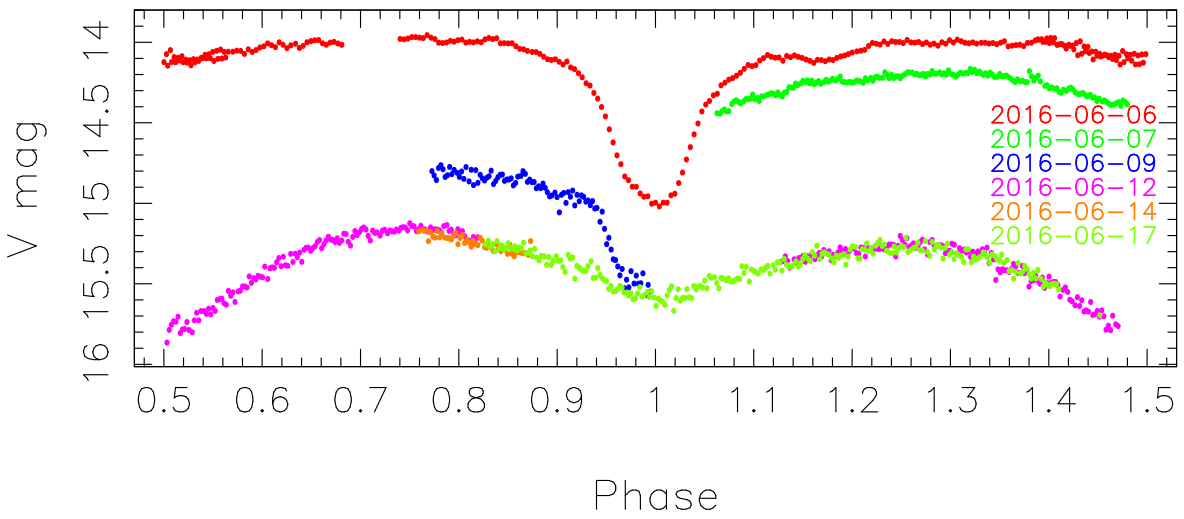

Figure 3.28: V-band light curves of CSS160603:162117+441254 observed with *pt5m* over several days following its initial detection by CRTS. The system appears to be fading and returning from a classic disk-dominated dwarf nova eclipse light curve during outburst, to its quiescent ellipsoidal modulation light curve.

be dominated by the white dwarf rather than the secondary star. In the case of CSS160603:162117+441254 we see a light curve that is dominated by the secondary star ellipsoidal modulation, with very little contribution from the white dwarf or accretion disk. This is evidenced by the small difference in primary and secondary minima, and any eclipse of the white dwarf or disk is fairly weak. There may even be no eclipse of the white dwarf or disk at all, and the difference in minima could be due to irradiation of the inner face of the secondary star. However, this is unlikely because the light curve is dominated by the secondary star, and the white dwarf flux is so low.

Contrastingly, in outburst the secondary minimum (which occurs when the donor star is behind the white dwarf) is barely visible, and is much smaller than the primary eclipse. This behaviour suggests that the disk is bright only in outburst, since in quiescence the light is dominated by the donor star.

Because the disk is barely visible in quiescence, the eclipse may actually be grazing (i.e. not completely edge on), and the disk could have grown in size during outburst, in order to produce the stronger eclipse signal.

Further spectroscopy and photometry is needed to understand this source fully. I will continue observing the object with *pt5m*, and we may seek UV or X-ray observations in order to investigate whether the system may be magnetic.

3.9 Conclusions

In this chapter I have described the architecture and operation of *pt5m*, and presented the results of the performance testing I have carried out. I have also described the additional software I developed to enable immediate and automated transient follow-up observations and data reduction.

I have demonstrated the capabilities of *pt5m* with a number of interesting observations, including a full study of the transits of the exoplanet WASP-33b, and transient follow-up observations of supernovae, GRBs, AGN, X-ray binaries and CVs. I have shown that *pt5m* is a powerful facility for time-resolved astronomy, and is particularly well suited to rapid follow-up of transient events.

I have also used *pt5m* extensively to conduct a search for eclipsing CVs, which is described in detail in Chapter 4.

Chapter 4

High-Speed Observations of Cataclysmic Variables

This chapter describes my search for new eclipsing cataclysmic variables with *pt5m*, as well as the high-speed follow-up of these new systems which I obtained with ULTRASPEC to investigate the structure of their eclipses. I have submitted this chapter for publication in MNRAS with 14 co-authors (Hardy et al., 2016b). The following data collection, reduction, analysis and text are all my own, except where otherwise stated.

4.1 Introduction

Cataclysmic variables (CVs) are binary systems in which an evolved white dwarf (WD) accretes mass from a (usually) late-type, main sequence companion. The accretion tends to be via an accretion disk, or directly onto the poles of the WD in the case of a strong magnetic field. In all cases, the donor star loses mass via Roche lobe overflow, and steady mass transfer is driven by the loss of angular momentum. For the remainder of this chapter we focus our attention on non-magnetic CVs only. A schematic diagram showing an imaginary CV system is shown in Figure 4.1. The four main types of CV

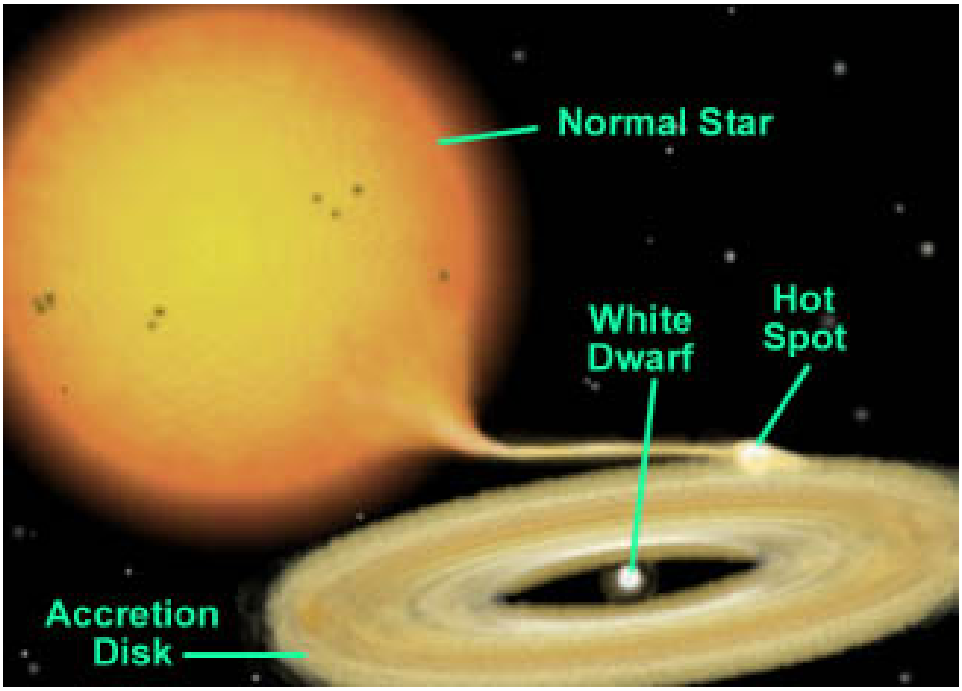

Figure 4.1: Artists impression of a CV system, showing the white dwarf, donor star, accretion disk and hot spot (also known as the bright spot), where the accretion stream impacts the disk. Image credit: K. Smale, obtained from *https://heasarc.gsfc.nasa.gov/docs/objects/cvs/cvs.html*

are dwarf novae (DNe), nova-like variables, classical novae and recurrent novae. For extensive reviews of CVs see Warner (1995) and Hellier (2001).

Types of CV

In dwarf novae, the accretion disk occasionally goes into dwarf nova outburst. This happens when disk instabilities cause the viscosity in the disk to change, allowing more material to be accreted onto the WD, and causing the system to become brighter by $1-10$ magnitudes within one or two days, before slowly fading again after several days (Smak, 1984). Dwarf nova outbursts can occur quasi-periodically on timescales of weeks to years, and each outburst can last days to weeks. DN outbursts are now the main process by which CVs are discovered, due to their detection as transient sources in wide-field surveys.

In nova-likes, a higher rate of mass transfer maintains a bright accretion disk which dominates the system brightness, essentially being in a state of constant dwarf-nova outburst. The term 'nova-like' comes from their similarities to remnant stars in old novae, although nova-likes are systems in which nova eruptions have not been seen. The non-magnetic nova-likes can be divided into three non-exclusive classifications: UX UMa stars, VY Scl stars and SW Sex stars. The former are considered the default or standard for nova-likes, whilst VY Scl stars exhibit deep 'low' states in which their luminosities are much reduced, for durations of days to years (La Dous, 1993; Dhillon, 1996). SW Sex stars have periods of 3-4 hours and show distorted radial velocity curves and seemingly high orbital inclinations (usually showing eclipses), as well as strong, possibly flared accretion disks with extended bright spots (Thorstensen et al., 1991; Dhillon, Smith & Marsh, 2013; Tovmassian et al., 2014).

Classical novae are systems which have been discovered by a nova eruption, in which an external shell of hydrogen on the surface of the WD ignites in a thermonuclear explosion. Nova eruptions are much more luminous than dwarf novae, and many novae have been discovered in other galaxies, with mass ejections of up to $10^{-4} M_{\odot}$ (Darnley et al., 2012). It is assumed that if we could watch them for long enough (perhaps tens of thousands of years), nova-likes would eventually show novae eruptions, although this is not certain.

Some novae have been seen to repeat, or show evidence for previous eruptions (e.g. by the presence of nova shells). These are known as recurrent novae. The rate of recurrence of a given system is dependent on how much hydrogen is needed to be accreted before ignition, and this itself depends on the mass and radius of the WD, and the accretion rate (Starrfield, Sparks & Truran, 1985). Those recurrent novae which have been discovered already presumably have short recurrence timescales driven by high accretion rates and masses close to the Chandrasekhar limit (Darnley et al., 2012).

In addition to the types above, another key class which falls under the umbrella of

cataclysmic variables are the double degenerate systems known as AM CVn binaries. These systems have very short orbital periods (typically less than 60 minutes), and tend to show helium in their spectra (Nelemans, 2005). AM CVn systems have evolved helium white dwarfs as their donor stars, and therefore must have very tight orbits in order for their Roche-lobes to be filled. They may be the progenitors of some channels of Type Ia supernovae (i.e. the double-degenerate channel, Maoz & Mannucci, 2012).

Eclipsing CVs

If the orientation of the orbital plane of a CV is coincident with our line of sight, eclipses may be observed. CVs often have complex eclipse structure, revealing different components of the system, including the accretion disk, the WD and the bright spot, where the in-falling material from the donor star encounters the accretion disk. In certain systems these features are evident and clearly separated from one another, but in others they are not present at all, are very weak, or blended together. For example, in some systems the WD ingress and bright spot ingress might occur simultaneously, blending together into a single feature. Other systems may have a weak bright spot that cannot be identified above the noise level. In most CV systems there is the added stochastic variability component known as flickering, which is due to random fluctuations in the accretion rate. Some systems show varying rates of accretion and disk radius, which can be observed as changes in the timing and depth of the bright spot eclipse. All of these factors will affect the structure of an eclipse light curve (Savoury et al., 2011).

When a CV of the dwarf nova type displays clear WD and bright spot features in its high cadence light curve, this light curve can be modelled allowing precise measurements to be made of the component masses and radii, in addition to other system parameters. In such cases, this information can be extracted without any form of spectroscopic study, e.g. Littlefair et al. (2008); Savoury et al. (2011); McAllister et al. (2015). The validity of this method has been successfully verified using time-resolved spectroscopy (Tulloch,

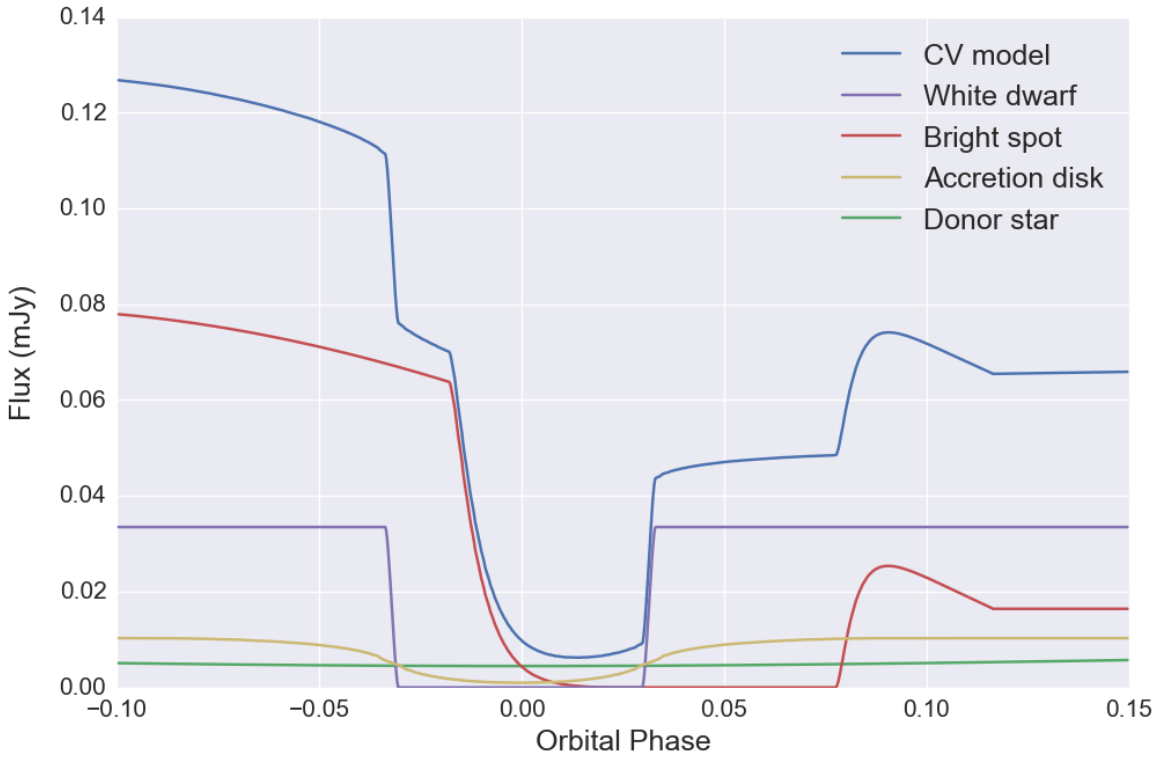

Figure 4.2: Artificial CV eclipse light curve, showing the different components of the model, which all contribute to the total light curve. An eclipse light curve showing clearly separated structure like this is ideal for modelling. Figure supplied by Martin McAllister.

Table 4.1: Criteria for modelling CV eclipse light curves, grouped by their importance.

Importance	**Criterion**
Vital	White dwarf ingress and egress visible
	Bright spot ingress and egress visible
	WD and BS ingresses and egresses separated temporally
Important	High signal-to-noise ratio
	No flickering or only weak flickering
Useful	Only mild contribution from accretion disk

Rodríguez-Gil & Dhillon, 2009; Savoury et al., 2012; Copperwheat et al., 2012). Figure 4.2 shows an example of an ideal CV eclipse suitable for modelling, with the different components shown separately. Table 4.1 lists the criteria of the features required in an eclipsing CV light curve in order to conduct accurate and precise system parameter modelling.

Using eclipses to study CV evolution

Eclipsing CVs are useful in improving our understanding of CV evolution. Theoretical and empirical studies of the WD mass distribution in CVs, and the properties of CV donor stars, both rely on the limited database of precisely measured CV system parameters. The average WD mass is considerably higher in CVs than in pre-CV systems and single WDs (see Figure 4.3). This was considered in detail by Zorotovic, Schreiber & Gänsicke (2011), who found no sign of WD mass evolution as a function of orbital period, despite Savoury et al. (2011) earlier arguing that WDs may in fact grow in mass via accretion (contrary to nova outburst theory, e.g. Townsley & Bildsten 2004), because of the increased WD mass for CVs at shorter orbital periods (see Figure 4.4). However, the study by Savoury et al. (2011) provides only weak support for this claim, since they were unable to show conclusive evidence that CVs at shorter orbital periods have systematically higher WD masses than those at longer periods. An alternative explanation for the WD mass discrepancy is rapid mass accretion during CV formation (Schenker & King, 2002). This is of relevance to the ongoing discussion regarding Type Ia supernovae progenitors, since WDs which slowly grow in mass could conceivably explode as supernovae once they reach the Chandrasekhar limit (Starrfield, Iliadis & Hix, 2016).

Wijnen, Zorotovic & Schreiber (2015) used binary population synthesis models to simulate the observed WD mass distribution among CVs. They found that neither WD mass growth by accretion, nor a phase of rapid mass transfer at the formation stage, could explain the observed mass distribution. Schreiber, Zorotovic & Wijnen (2016) and Nelemans et al. (2016) have independently considered the impact of selective angular momentum loss due to early nova eruptions in low white dwarf mass CVs. Both studies found that this process can plausibly explain the mass distribution discrepancy between CV and pre-CV populations, eliminating the need for the WDs to grow in mass with time (thus casting further doubt on the theory that CVs could be the progenitors of

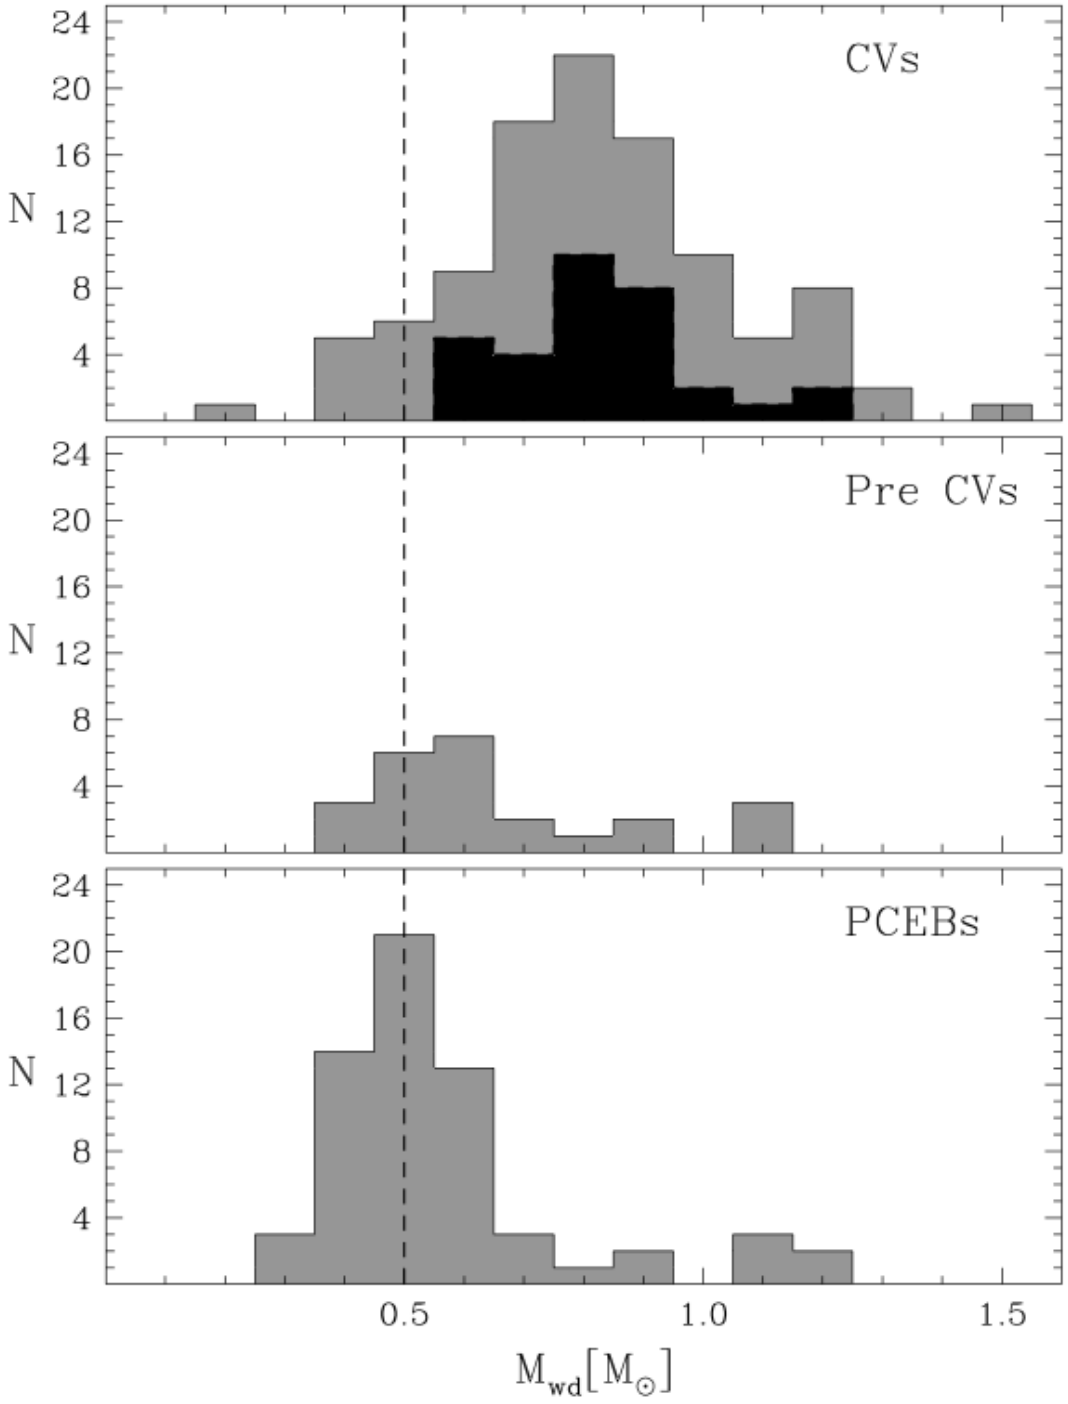

Figure 4.3: White dwarf mass distributions for CVs, pre-CVs and post common envelope binaries. WDs in CVs are found to be significantly more massive than WDs in detached systems. The darker panels represent more accurately measured masses then the paler sections, with measurements made through eclipse modelling or spectroscopy, rather than superhumps. Reproduced from Zorotovic, Schreiber & Gänsicke (2011).

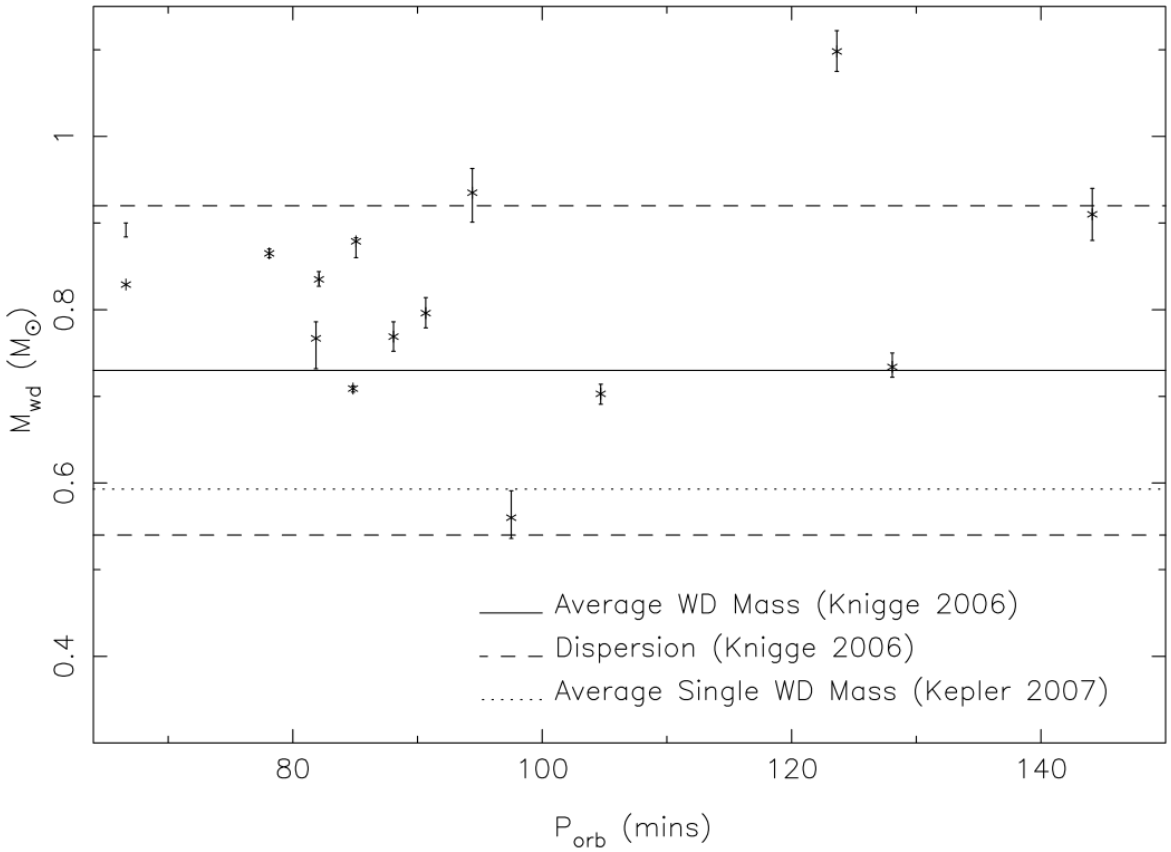

Figure 4.4: White dwarf masses in short period CVs, reproduced from Savoury et al. (2011). CVs at short periods were found to have higher WD masses than the average for all periods, and higher WD masses than in single WDs. This led to the suggestion that WDs could grow in mass with time, since CVs evolve to shorter periods with time. Horizontal lines for mean WD masses come from Knigge (2006) and Kepler et al. (2007).

SN Ia). However, these studies are empirical and rely on existing measurements of WD masses which are limited to only 35 objects, some of which have poorly constrained masses (Zorotovic, Schreiber & Gänsicke, 2011; Schreiber, Zorotovic & Wijnen, 2016). An expansion of the sample of CVs with precisely measured masses will help refine the models of Nelemans et al. (2016) and Schreiber, Zorotovic & Wijnen (2016), and could finally reject CVs as possible SNIa progenitors (Wijnen, Zorotovic & Schreiber, 2015).

Aside from the properties of the WD, the donor star mass and radius values in CV populations can also be very useful in CV evolution studies. To fully understand the evolution of a CV with time, we need to understand the processes of angular momentum

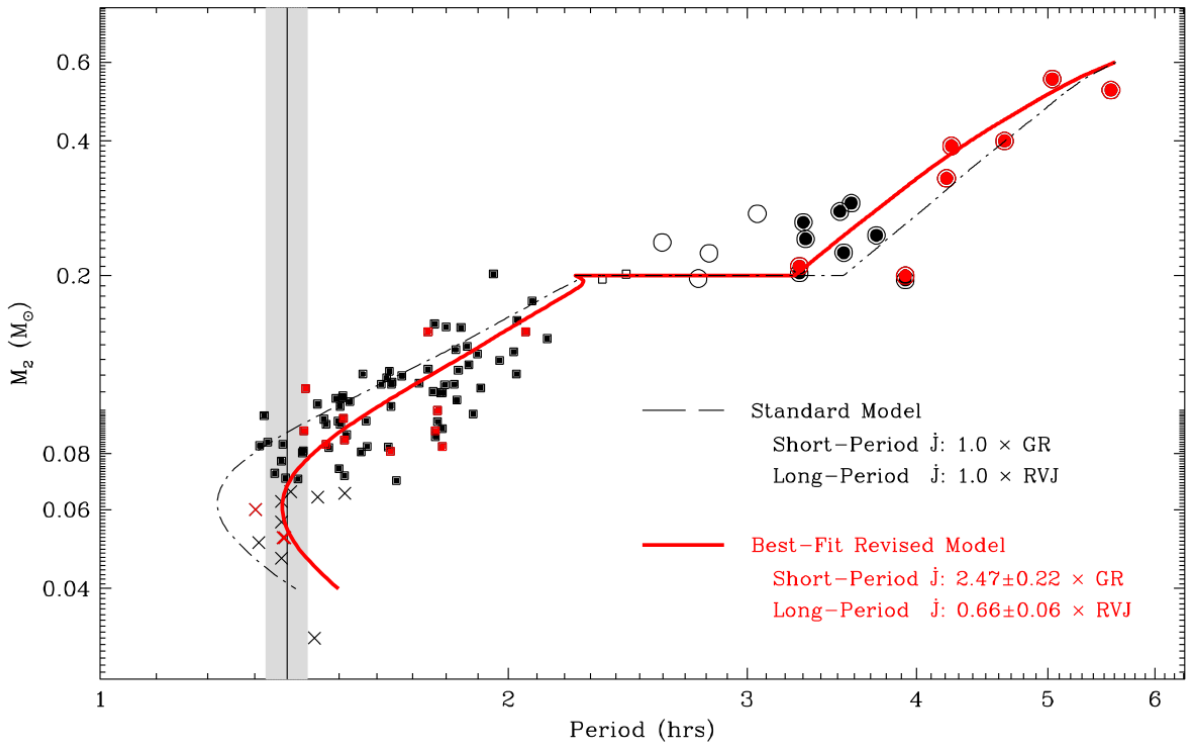

Figure 4.5: Semi-empirical CV donor sequence by Knigge, Baraffe & Patterson (2011). The red symbols are values from eclipsing CVs, which provide the most precise measurements. The standard model of angular momentum loss (black, dashed), via gravitational radiation and magnetic braking, shows some discrepancy with the observational results. Empirical mechanisms for angular momentum loss are incorporated in the best fit model in red.

loss which drive the steady flow of gas from the donor star to the WD. Much work has gone in to producing a semi-empirical donor sequence which describes the mass, radius, effective temperature and spectral type of CV donor stars as a function of orbital period (Smith & Dhillon, 1998; Patterson et al., 2005; Knigge, 2006; Knigge, Baraffe & Patterson, 2011). The donor sequence (see Figure 4.5) is one of the best ways of tracing long-term mass transfer rates, and therefore also rates of angular momentum loss. The alternative measure, the WD temperature, can only give a snapshot measurement of accretion rate, and does not allow any insight into the historical mass transfer rate.

These empirical studies of CV evolution have so far required an additional unknown mechanism for angular momentum loss, that cannot be fully explained by the exist-

ing mechanisms of gravitational radiation and magnetic braking (Knigge, Baraffe & Patterson, 2011). In addition, with only around 30 systems with precisely measured parameters, the empirical donor sequences still require considerable refinement with more measurements over a range of orbital periods. In order to do this, and identify the missing mechanism for angular momentum loss, further constraints are required through additional donor mass and radius measurements. Finding and characterising new eclipsing CVs will help this process.

Search for eclipsing CVs

In this chapter I present a search for new eclipsing CVs, predominantly in the northern hemisphere. Similar studies of CVs in the southern hemisphere have been conducted by Coppejans et al. (2014), Woudt et al. (2012), and in the preceding papers by the same group. In addition, I also present a large collection of high speed observations of existing eclipsing systems, many of which already have published photometry, but of a lower time resolution. For all eclipsing systems I provide an evaluation of their suitability to be used in light curve modelling. These observations are part of an ongoing campaign to derive CV system parameters (McAllister et al. 2015 and references therein), and will form the observational basis of a future paper detailing the modelling of these systems.

4.2 Candidate Selection

Dwarf novae routinely enter phases of brighter emission known as outbursts, due to thermal instabilities causing changes in viscosity of the accretion disk, dumping large amounts of material onto the WD in a short period of time (Osaki, 1996; Buat-Ménard, Hameury & Lasota, 2001; Lasota, 2001). New CVs are being detected in outburst at a rate of at least one per day by current wide-field transient surveys. As of January 2016, the Catalina Real-Time Transient Survey (CRTS, Drake et al., 2009) had detected almost

1300 new CVs, the All-Sky Automated Survey for SuperNovae (ASAS-SN, Shappee et al., 2014) had found over 600, and the Mobile Astronomical System of TElescope-Robots (MASTER, Lipunov et al., 2010) had discovered almost 600. The Gaia astrometric surveying satellite, launched in 2013, is also now discovering transient events, including CVs, at a rate of several per day (Wyrzykowski et al., 2012; Blagorodnova et al., 2016). On top of these, many hundreds have been identified in the Sloan Digital Sky Survey (SDSS, e.g. Gänsicke et al., 2009; Szkody et al., 2011) and by a vast network of amateur astronomers; see for example *vsnet* (Kato et al., 2004) and the Center for Backyard Astrophysics (Patterson et al., 2013).

Our aim was to find new eclipsing systems amongst the CVs found in these surveys. Starting with a list of nearly 2000, we first rejected all systems which were not observable from our main facility in La Palma (see Section 4.3). We also excluded all systems with a quiescent magnitude greater than $V \sim 19$. We used the 0.5m telescope *pt5m* to search for eclipses (see Section 4.3 and Hardy et al. 2015b) which is unable to reach sufficient signal-to-noise ratios to detect eclipses in systems fainter than this.

Finally, we rejected any systems which have already been studied at sufficiently high time resolution to rule out eclipses. These include systems near the celestial equator which may have already been studied with time-resolved photometry (e.g. Coppejans et al., 2014), and systems which have been closely followed by amateur observers. The international network of amateur astronomers studying CV superhumps, *vsnet*, have studied many systems whilst in outburst, sometimes at high enough cadence and for long enough to determine an orbital period and rule out eclipses. Superhumps are hump-like variations in the light curve of outbursting dwarf novae, with a period slightly longer than the orbital period. They occur when the tidal torques of the secondary star entice the heated accretion disk into an elliptical geometry, with a projection effect that produces the periodic modulation (Whitehurst, 1988; Hellier, 2001). Characterising superhump periods can allow CV mass ratios to be estimated, but not to the high level

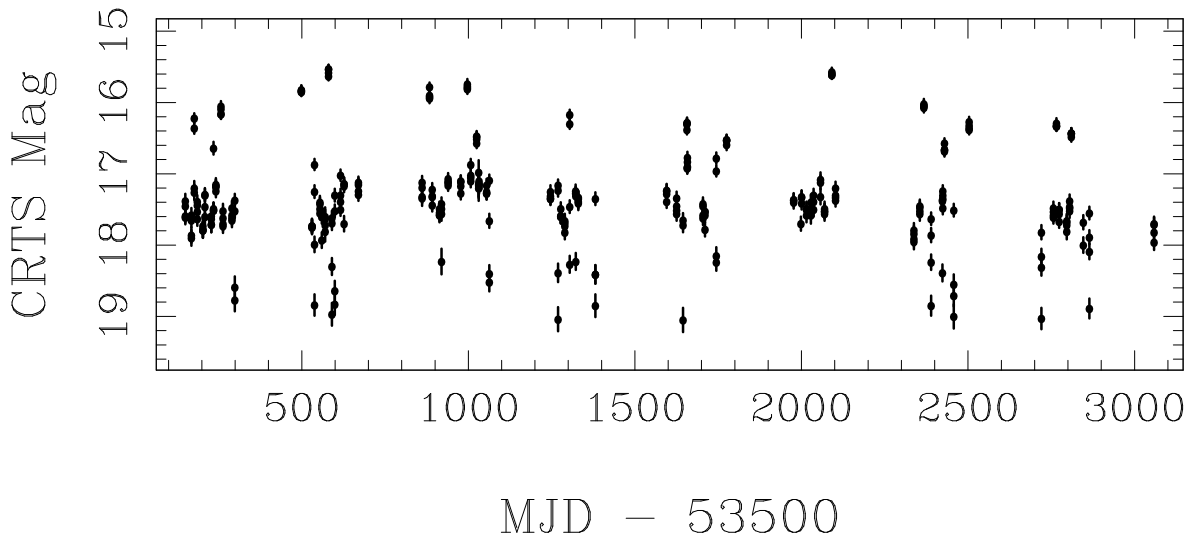

Figure 4.6: CRTS light curve of CSS131106:052412+004148 created from publicly available data (Drake et al., 2009). Numerous dips in brightness indicate the possibility that the system is eclipsing. The brighter measurements are dwarf nova outbursts.

of precision achieved through light curve modelling (e.g. Kato et al., 2015 and references therein).

At the end of this process we were left with around 80 systems eligible for follow up. This number has grown gradually since our observing campaign began in 2014. We began by observing the brighter targets, and those which were already suspected to be eclipsing according to other studies. This could be either from spectroscopy suggesting that the accretion disk is close to edge-on (e.g. Breedt et al. 2014), or from long-term light curves provided by CRTS, for example, that show regular dips or faint measurements. An example of such a system, CSS131106:052412+004148, is shown in Figure 4.6.

4.3 Observations

For new CVs, initial observations were made with *pt5m*, a 0.5m telescope located at the Roque de los Muchachos Observatory on La Palma (see Chapter 3). The telescope is robotic and sits on the roof of the control-room building of the 4.2m William Herschel Telescope (WHT).

If eclipses were seen and a period could be measured, we then observed the systems at high time-resolution using either ULTRACAM on the WHT (Dhillon et al., 2007) or ULTRASPEC on the 2.4m Thai National Telescope (TNT, Dhillon et al. 2014 and Chapter 2). Both are imagers utilising frame-transfer CCDs for high cadence, low dead-time (of order 10ms) observations. ULTRASPEC has a single channel, whilst ULTRACAM has three beams, and can therefore observe in three filters (e.g. u', g', r') simultaneously.

In Section 4.6 we also present new observations of CVs which were already known to be eclipsing. Along with the WHT and the TNT, these observations have also been conducted at other telescopes, such as the 3.6m New Technology Telescope (NTT) at La Silla or the 8.2m Very Large Telescope (VLT) at Paranal, both in Chile. A select few systems were also observed with the 11.1×9.8m (effective circular aperture of 9.2m) South Africa Large Telescope (SALT) and its SALTICAM instrument (O'Donoghue et al, 2006). An extract from the full log of all of our observations is presented in Table 4.2. The full version can be found in Appendix B. We use the most accurate timing system available; Barycentric Modified Julian Date (BMJD), under the Barycentric Dynamical Timescale (TDB), which takes account of relativistic time dilation.

In all observations, bright stars near to the target were used as comparisons, and we performed differential photometry in order to account for any transparency fluctuations. The fluxes of the target and comparison stars were extracted using variable aperture photometry with the ULTRACAM pipeline reduction software (Dhillon et al., 2007). For particularly faint targets, the optimal photometric extraction technique of Naylor (1998) was used. We calculated the magnitudes of the comparison stars by extracting their full flux using large, fixed apertures, then flux calibrated this using average extinction coefficients and instrumental zeropoints. The flux of the target was then calibrated using the calculated magnitude of the comparison star. Light curves are therefore plotted as calibrated flux in mJy as a function of time.

We supply average out-of-eclipse magnitudes in Table 4.2, primarily as a guide to

show when systems are in outburst, or to demonstrate which systems show strong variability. However, since the zeropoints measurements (based on observations of photometric standard stars) are made less frequently than the science observations, these magnitudes should always be considered estimates. Systematic uncertainties of 0.2 magnitudes or more could be introduced by variable extinction or high altitude clouds.

We measured extinction coefficients from observations of stable, non-variable stars taken over a range of airmasses on clear, photometric nights. We therefore note that our method only enables accurate flux calibration for observations conducted in similar conditions. The atmosphere tends to be stable and clear in La Palma and in Chile, but in Thailand there is frequently low altitude haze, as well as thin, high altitude clouds, which strongly affect the extinction coefficients and zeropoints. Therefore all measured magnitudes from the TNT are to be treated with caution.

Most observations made with ULTRASPEC at the TNT were conducted in the standard SDSS filters, but sometimes we used a non-standard *KG5* filter. This filter encompasses the SDSS u', g' and r' pass bands, and can therefore obtain better signal-to-noise ratios on fainter targets than individual SDSS filters. The obvious disadvantage here is that it is not easy to convert the measured counts into a calibrated flux. Our method for flux calibration of the *KG5* filter is described in Chapter 2.

Table 4.2: Journal of Observations. Some object names have been shortened. Start(MJD) is the start time of each observing run, given in MJD(UTC). Mid-eclipse times are given in BMJD(TDB) and the number in parentheses is the uncertainty in the last digit. T_{exp} is the exposure time in seconds and is supplied as $T_{blue}/T_{green}/T_{red}$ for the three beams of ULTRACAM. ΔT is the duration of the observing run in minutes. Mag. is the estimated out-of-eclipse magnitude. This is an extract showing only one object. The full table containing all 74 objects can be found in Appendix B.

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T$_{exp}$ (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
ASASSN-13cx	56943.91536	56943.9710(3)	150	137	*pt5m*	V	16.1
	56968.81315	56968.8218(3)	150	206	*pt5m*	V	18.0
		56968.9015(3)					
	56973.81141	56973.8401(7)	150	406	*pt5m*	V	18.3
		56973.9186(7)					
	56989.65725	56989.68988(3)	3.9	60	TNT/USPEC	g'	18.9
	56993.59267	-	3.9	53	TNT/USPEC	*KG5*	17.9
	57014.83182	57014.8596(2)	120	69	*pt5m*	V	16.2
	57015.86791	57015.89471(7)	120	69	*pt5m*	V	16.1
	57023.51317	57023.54121(3)	3.9	57	TNT/USPEC	*KG5*	18.5
	57024.56187	57024.57665(3)	4.9	46	TNT/USPEC	g'	18.3
	57027.56687	57027.60339(3)	5.9	69	TNT/USPEC	*KG5*	18.9
	57198.12288	57198.13413(15)	12/4/4	31	WHT/UCAM	$u'/g'/r'$	18.4/18.4/18.3

4.4 Results: New Eclipsing Systems

All data in this section were reduced and analysed by me.

We present 13 CV systems which were either unknown to be eclipsing, were suspected eclipsers, or whose eclipsing nature has only recently been published. Most discoveries were made with *pt5m* observations over several nights, in the Johnson V filter. Light curves for each and every object were generated and checked by eye for signs of eclipses, requiring at least two data points to be considered as a candidate eclipse. When eclipses were found, mid-eclipse times were measured by fitting a Gaussian to the eclipse profile. An example Gaussian fit to an eclipse of MASTER OT J003059.39+301634.3 (the first eclipsing CV we discovered) is shown in Figure 4.7.

Ephemerides were then obtained by a linear fit to the eclipse times. When fitting for an ephemeris, we always selected the orbital cycle zero-point (T_0) which minimised the co-variance between P and T_0. Occasionally we omitted some mid-eclipse times from

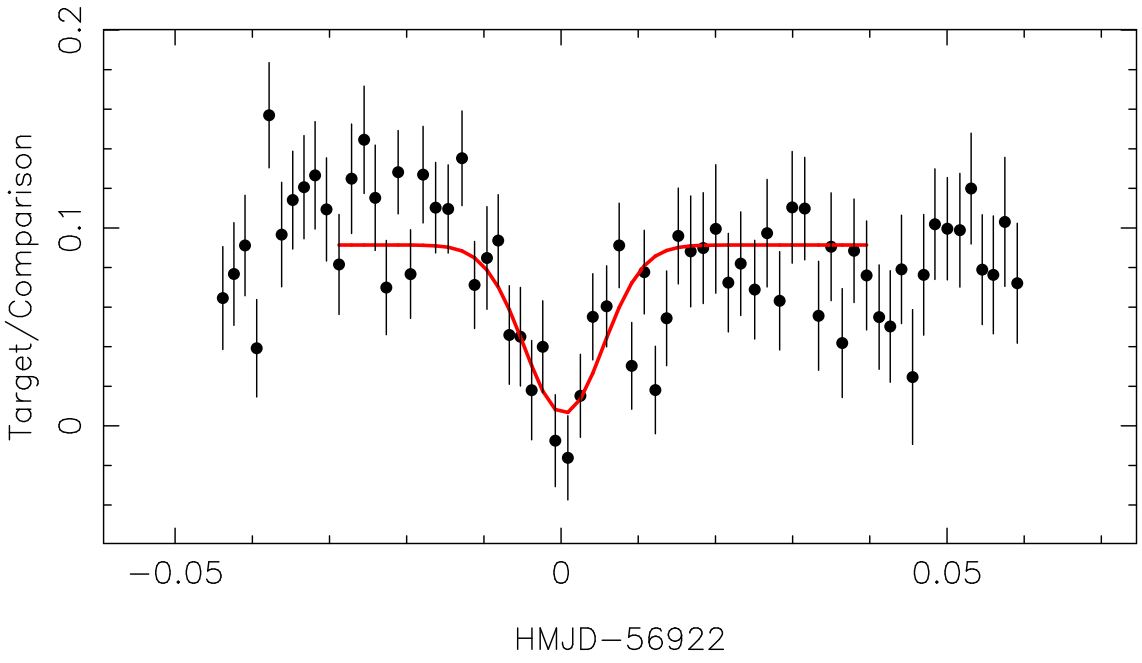

Figure 4.7: Gaussian fit to a MASTER OT J003059.39+301634.3 eclipse observed with *pt5m* on 2014-09-21. The mid-time of the eclipse is measured to be HMJD=56922.0003± 0.0010.

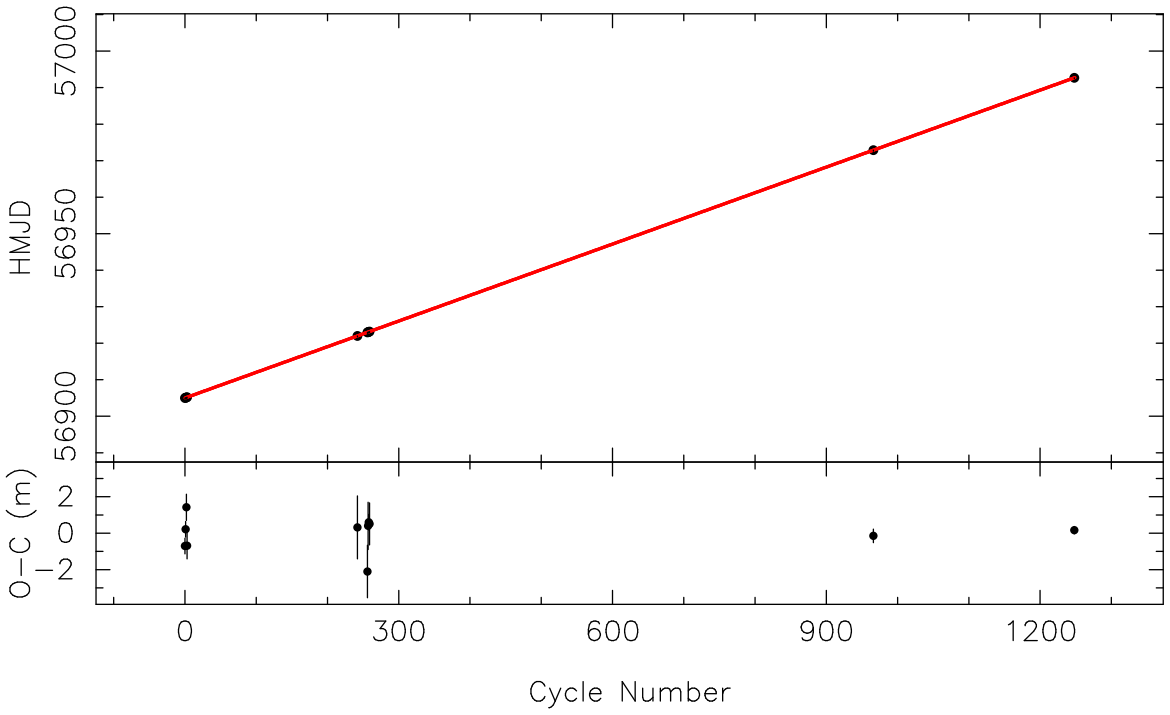

Figure 4.8: Measured eclipse mid-times of MASTER OT J003059.39+301634.3 plotted against the orbital cycle number. The red linear fit gives us a period measurement of $0.07026247 \pm 0.00000003$ days. The O−C residuals from the linear fit are plotted in the lower panel in units of minutes.

the ephemeris fit, if for example the system was in outburst, there were orbital cycle ambiguities, or the times had very high errors (often the case with *pt5m* observations). An example linear fit is shown in Figure 4.8.

Occasionally this method is subject to failure, if there are period aliases, or when there is an ambiguity in the orbital cycle number. We therefore conducted additional checks on period measurements using phase-dispersion-minimisation (PDM, Stellingwerf 1978). To do this we took all the light curve data, including eclipses and out-of-eclipse observations, normalised it and produced a PDM periodogram which shows troughs at potential orbital periods. We phase folded the data on those periods to test which were valid. Visual inspection of the phase-folded light curves was sufficient to confirm or reject reliable orbital periods for all systems. An example phase-folded light curve of MASTER OT J003059.39+301634.3 is shown in Figure 4.9. The derived ephemerides of

all systems are presented in Table 4.2.

Table 4.3: Ephemerides and suitability for modelling of eclipsing systems. Ephemerides are measured and quoted in BMJD(TDB). Some ephemerides have been calculated using additional mid-eclipse times from the literature. In these cases the relevant references are supplied. Evaluation of the potential for light curve modelling is given as 'Y' for yes, 'N' for no, and 'M' for maybe/more data are needed.

Object	T_0	P	**Additional Times**	**Modelling Potential**
1RXS J180834.7+101041	54926.3108(1)	0.06949063(5)	Southworth & Copperwheat (2011)	N
ASASSN-13cx	57104.30636(1)	0.07965006(1)	-	M
ASASSN-15au	57148.04781(4)	0.06894973(2)	-	M
ASASSN-15bu	57080.98512(1)	0.0768262(9)	-	M
CSS080227:112634-100210	56294.94350(3)	0.077421572(3)	Southworth et al. (2015)	M
CSS080306:082655-000733	56444.203767(8)	0.059764424(1)	-	Y
CSS080623:140454-102702	55488.728318(8)	0.059578972(2)	-	Y
CSS081220:011614+092216	56713.570050(8)	0.065843019(3)	Coppejans (priv.comm.)	M
CSS090102:132536+210037	56207.945385(5)	0.0623849134(7)	Southworth et al. (2015)	Y
CSS090419:162620-125557	56551.813968(9)	0.075442719(5)	-	Y
CSS090622:215636+193242	56877.82014(2)	0.0709293(2)	-	Y
CSS100218:043829+004016	56658.417503(1)	0.0654948(1)	-	M
CSS110113:043112-031452	56396.181193(3)	0.0660508721(6)	-	Y
CSS110513:210846-035031	56767.65993(3)	0.15692657(3)	-	Y
CSS111003:054558+022106	56852.23197(9)	0.12097138(4)	Miszalski et al. (2016)	Y
CSS111019:233313-155744	56866.13556(2)	0.04285020(1)	-	M
CSS111101:233003+303301	57015.4708(2)	0.1559784(6)	-	N
CSS131106:052412+004148	57126.37256(2)	0.17466647(2)	-	M
CzeV404 Her	56871.91730(4)	0.09802125(2)	-	M
GALEX J003535.7+462353	56489.69294(4)	0.17227491(1)	Wils et al. (2011)	M
GY Cnc	55938.26369(1)	0.175442402(1)	Kato, Ishioka & Uemura (2002)	Y
HS 2325+8205	55661.080122(7)	0.194334533(1)	Pyrzas et al. (2012)	Y
HT Cas	55550.16974(3)	0.073647175(1)	Feline et al. (2005)	N
IY UMa	54679.998816(9)	0.0739089285(3)	Steeghs et al. (2003)	Y
MAST003059.39+301634.3	57159.48849(5)	0.07026247(3)	-	M
MAST192328.22+612413.5	56852.10211(9)	0.1676465(3)	Kennedy et al. (2016)	N

Table 4.2: ...continued.

Object	T_0	P	**Additional Times**	**Modelling Potential**
MAST232100.42+494614.0	57010.6753(1)	0.2123715(9)	-	M
MLS101226:072033+172437	57023.1777(2)	0.1504066(3)	-	N
MLS120517:152507-032655	56821.89967(5)	0.06438163(7)	-	N
SDSS J040714.78-064425.1	56332.1340(1)	0.17020393(2)	-	N
SDSS J075059.97+141150.1	56306.51099(1)	0.093165496(2)	Southworth et al. (2010, 2015)	Y
SDSS J090103.93+480911.1	55229.666015(7)	0.0778805326(6)	-	Y
SDSS J090403.49+035501.2	-	-	-	N
SDSS J092009.54+004245.0	56145.35967(3)	0.147875676(7)	-	N
SDSS J092444.48+080150.9	-	-	-	N
SDSS J093537.46+161950.8	55670.30839(6)	0.06405514(3)	Southworth et al. (2015)	N
SDSS J100658.40+233724.4	56396.98203(6)	0.18591303(2)	Southworth et al. (2009, 2015)	Y
SDSS J115207.00+404947.8	56082.544698(2)	0.0677497017(2)	Southworth et al. (2010); Savoury et al. (2011)	Y
SDSS J125023.84+665525.4	54614.43088(1)	0.0587356818(6)	Dillon et al. (2008)	M
SDSS J152419.33+220920.1	56287.823054(3)	0.0653187310(8)	Southworth et al. (2010)	Y
SDSS J155531.99-001055.0	55267.794254(4)	0.0788455518(3)	Southworth et al. (2007)	M
SDSS J155656.92+352336.6	55633.75890(3)	0.088091490(4)	-	M
SSS100615:200331-284941	56874.0223924(8)	0.05870445(6)	-	Y
SSS120402:134015350512	-	-	-	N
SSS130413:094551-194402	56809.688700(9)	0.065769292(3)	-	Y
V2051 Oph	55314.156237(4)	0.06242785751(8)	Vrielmann, Stiening & Offutt (2002); Baptista et al. (2003); Qian et al. (2015)	Y
V713 Cep	56532.192754(5)	0.085418509(1)	-	Y

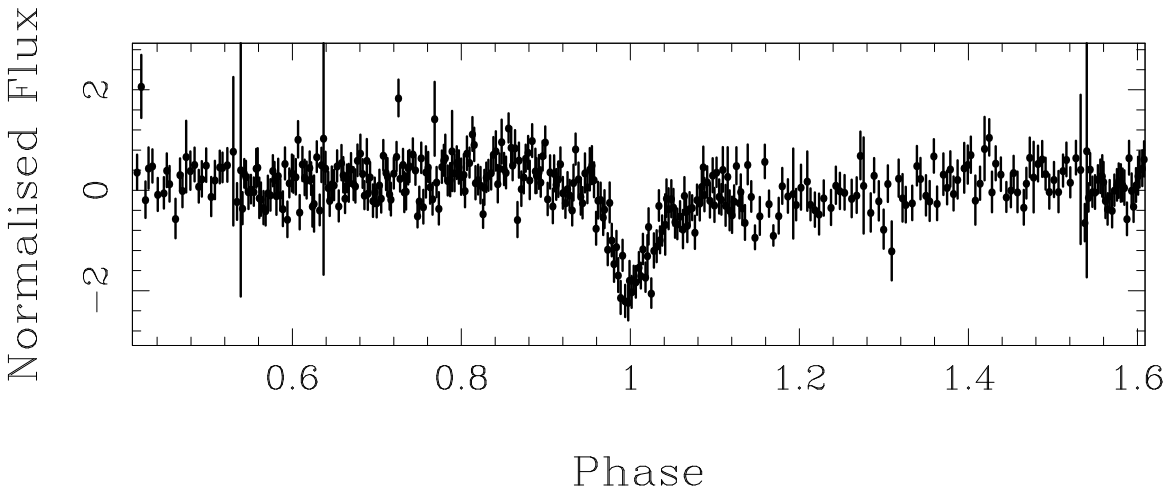

Figure 4.9: Phase-folded light curve of MASTER OT J003059.39+301634.3, comprising 13 individual light curve observations and binned by a factor of 5. The clear eclipse at phase one and lack of eclipse-like signals at phases other than one gives us confidence in the measured period.

Once a reliable period had been derived, we then attempted to investigate the eclipse structure by observing the system at high time resolution. Details of these observations are included in Table 4.2. In high cadence observations, often the Gaussian fit to an eclipse profile is unsuitable, because the different components of the eclipse are clearly visible and the shape can no longer be approximated by a Gaussian profile. In this case, we use the mid-point of the WD eclipse as the mid-eclipse time, determined from measurement of WD mid-ingress and mid-egress. We supply the mid-eclipse times in Table 4.2 in case this information is useful to future orbital period variability studies (e.g. Parsons et al. 2010b; Bours et al. 2016).

We briefly discuss each system below, and comment on their suitability for light curve modelling. A summary of these judgements is also included in Table 4.2.

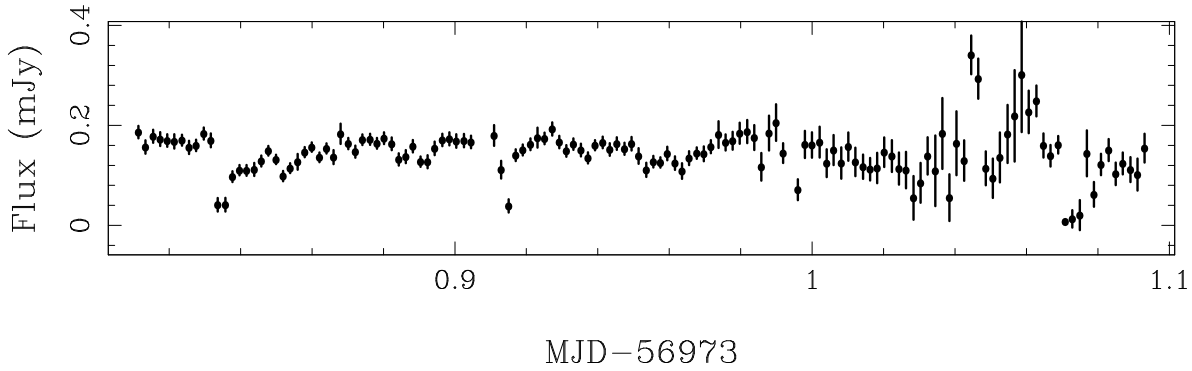

Figure 4.10: ASASSN-13cx light curve observed in quiescence on 2014-11-12 with *pt5m*. Up to 4 eclipses are visible, with the last part of the observations degraded by passing clouds.

ASASSN-13cx

This object was discovered by ASAS-SN[1] with an outburst to $V = 15.5$ magnitudes, from $V = 18$ as recorded in the NOMAD catalogue (Zacharias et al., 2004). ASASSN-13cx also showed previous outbursts in the Catalina Sky Survey (CSS, part of CRTS), as well as fainter measurements indicating eclipses. During an outburst in 2014, *vsnet* were able to confirm the eclipsing nature and estimate a period (Kato et al., 2015). We observed multiple eclipses of ASASSN-13cx with *pt5m* and later at high speed with ULTRASPEC and ULTRACAM.

We refined the initial orbital period estimates supplied by Kato et al. (2015) using 12 eclipses (see Table 4.2). Figure 4.10 shows a 6 hour light curve observed with *pt5m*, displaying several eclipses, while Figure 4.11 shows a light curve observed with ULTRACAM. At high cadence the eclipse structure can be seen with clearly separated egress for the WD and bright spot. The ingress of the bright spot is less clear, and in this state it may be difficult to model these light curves to extract system parameters. This system will continue to be monitored.

[1]References for new ASAS-SN transients are not always available, as they are often not published in the literature. Instead transients are listed on the ASAS-SN webpage at `http://www.astronomy.ohio-state.edu/~assassin/transients.html`

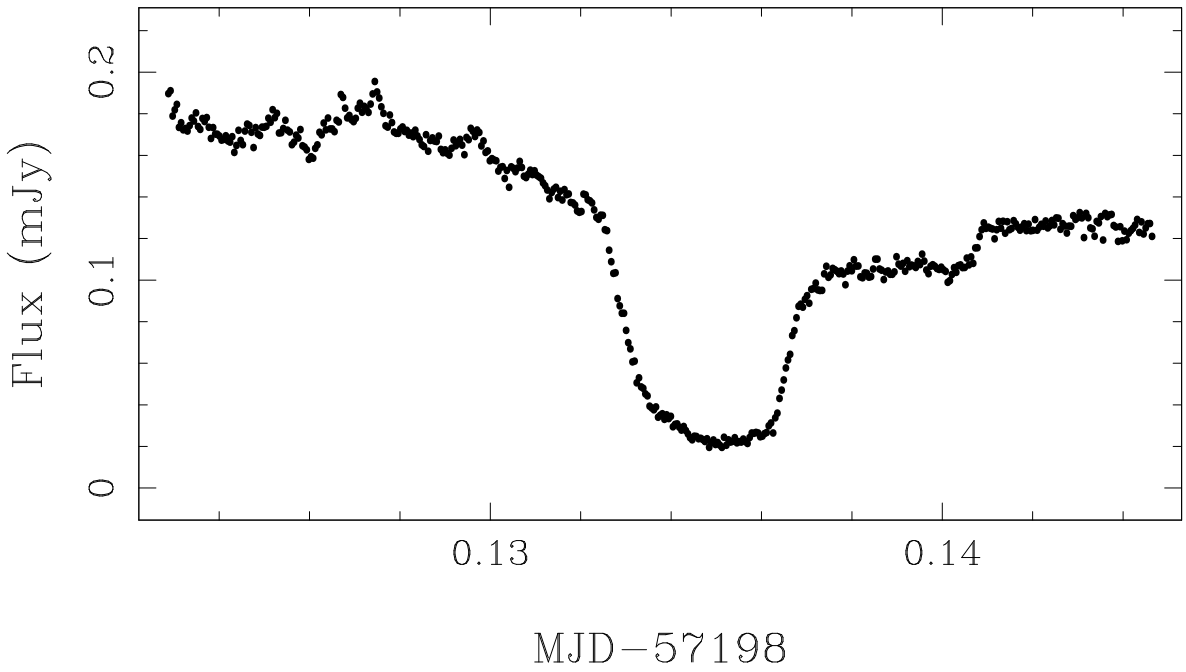

Figure 4.11: ASASSN-13cx r'-band eclipse observed on 2015-06-24 with ULTRACAM. The bright spot egress and WD egress are clearly separated, but the ingresses are less clear. In this light curve and all subsequent densely-sampled light curves, error bars have been omitted for clarity.

ASASSN-15au

This object was discovered by ASAS-SN during an outburst to $V = 15.3$ magnitudes, from a quiescent SDSS magnitude of $g = 17.8$ (Ahn et al., 2012). It also shows some variability in archival CSS and Mount Lemmon Survey (MLS, part of CRTS) light curves, including outburst and eclipse-like faint measurements. We found multiple eclipses in *pt5m* light curves (see Figure 4.12), and attempted high speed observations with ULTRASPEC (see Figure 4.13).

The observations at high cadence show signs of eclipse structure, including weak features that could be the bright spot ingress and egress (Figure 4.13 shows the clearest example of light curve structure). However, these features are not clear enough for this system to be a candidate for light curve modelling in its current state or without further study.

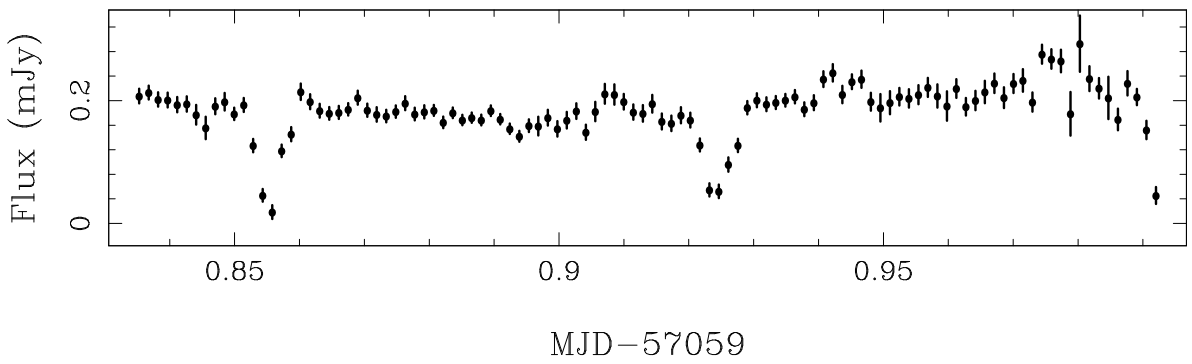

Figure 4.12: ASASSN-15au light curve from *pt5m* observed on 2015-02-06, showing three eclipses.

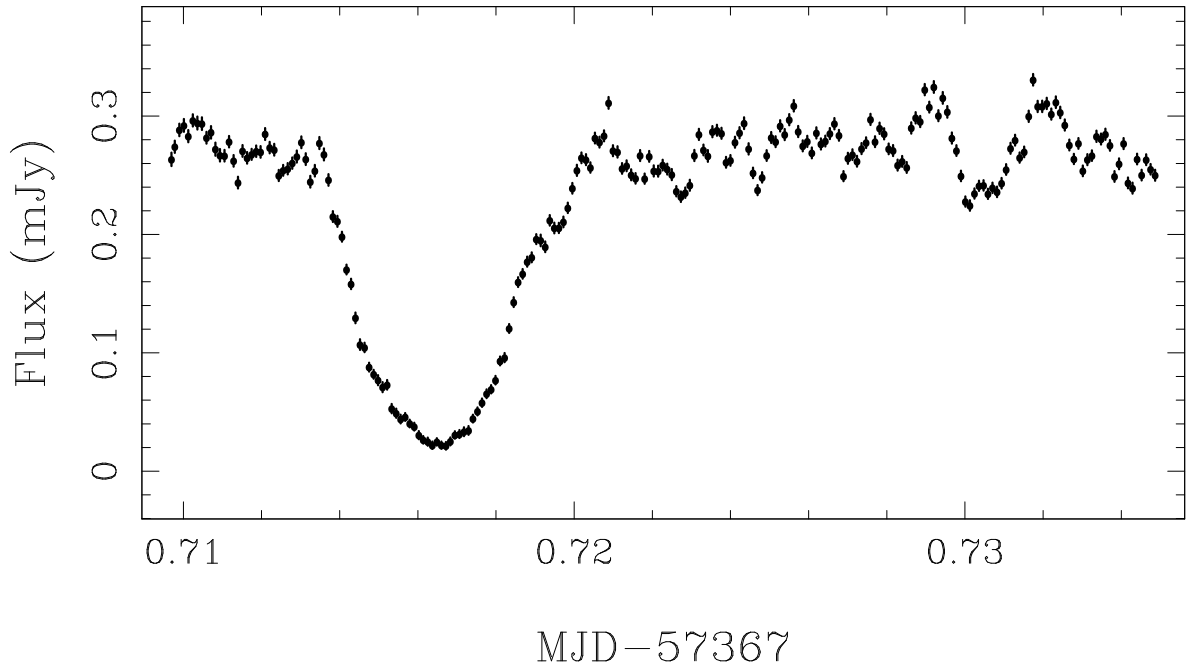

Figure 4.13: ASASSN-15au eclipse observed with ULTRASPEC on 2015-12-11. Close inspection reveals possible bright spot features, but these may also be attributed to flickering.

ASASSN-15bu

ASASSN-15bu was discovered by ASAS-SN via a dwarf nova outburst. It also showed previous outbursts in the CSS light curve, as well as multiple dips signalling eclipses. *vsnet* (alert 18228) used the CSS data to find a coherent orbital period from the eclipses of 110 minutes.

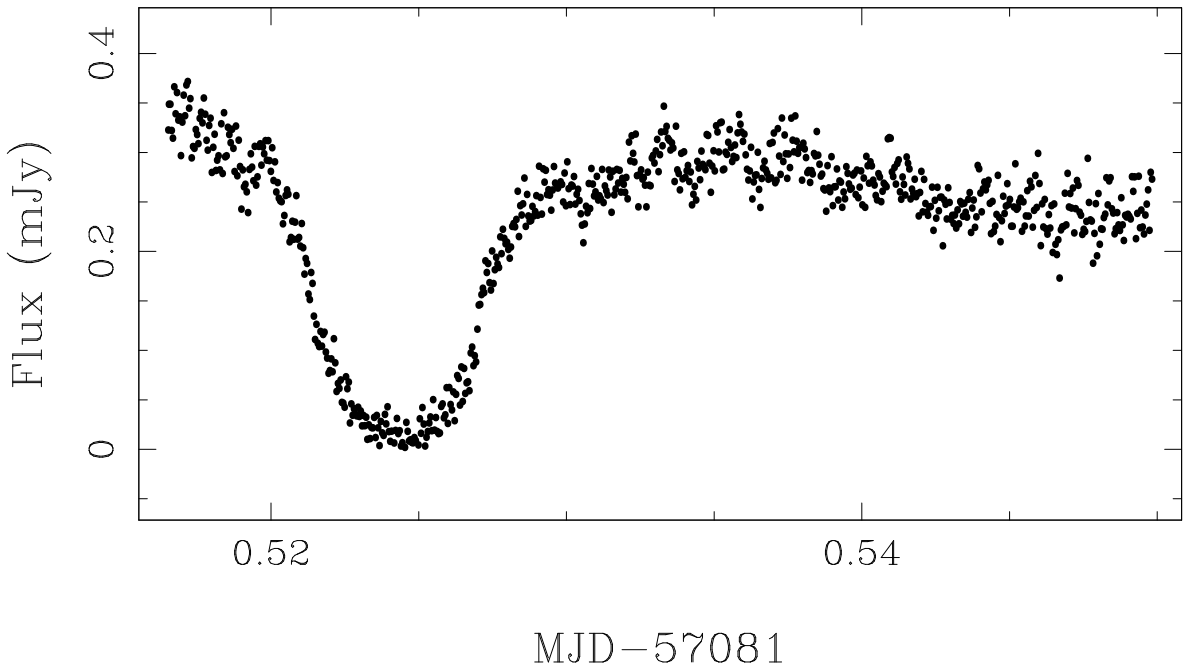

Figure 4.14: ASASSN-15bu eclipse observed with ULTRASPEC on 2015-02-28 during in bright time, which adds considerable noise. Nevertheless it appears that there is no strong bright spot feature.

We confirmed the eclipsing nature with 4 ULTRASPEC observations. Figure 4.14 shows the clearest eclipse observations available, in which no clear bright spot features are visible. We intend to study this system with a larger telescope in future, but for now it cannot be modelled.

CSS110113:043112-031452

This system was discovered via a super-outburst in CRTS[2], from a quiescent magnitude of $V = 19$-20 to $V = 15$. The *vsnet* collaboration discovered eclipse signals and estimated an orbital period of 1.58 hours (alerts 12614, 12661), but no detailed studies have yet been published.

We confirmed the eclipsing nature with high-speed observations of 12 eclipses in varying conditions. Sometimes the bright spot signal is hidden in noisy data or blended

[2]References are often not available for CRTS discoveries, though all transients are available online at `http://crts.caltech.edu/` and the discovery of CVs in particular is discussed in Drake et al. (2014).

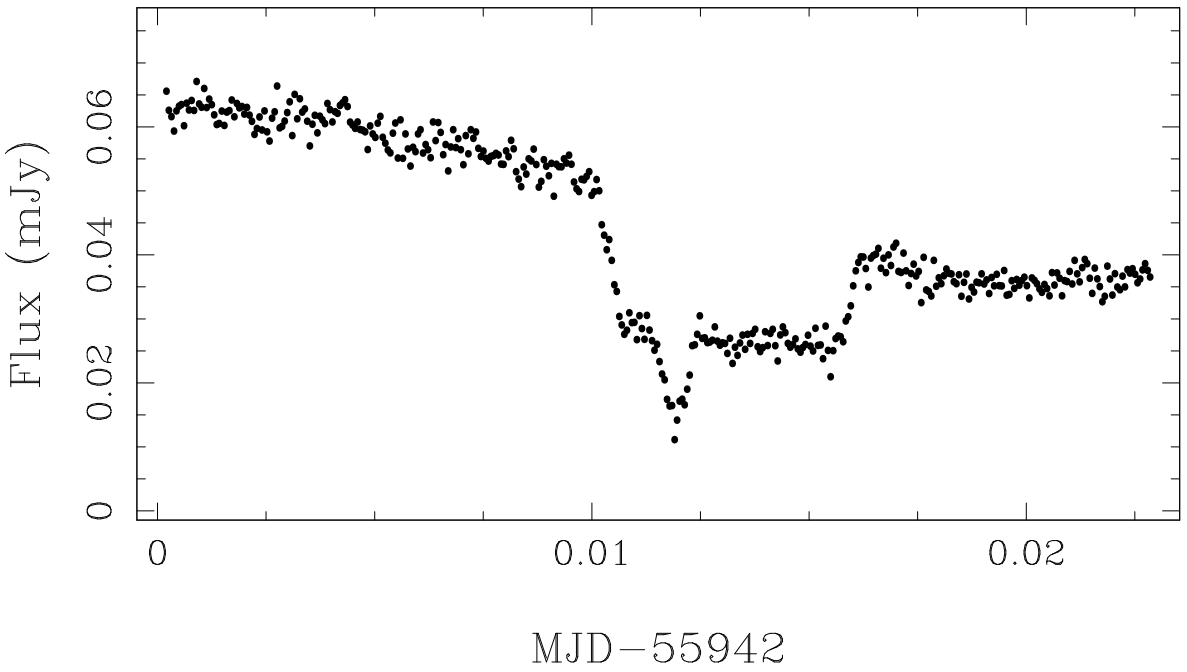

Figure 4.15: Eclipse of CSS110113:043112-031452 observed in g' with ULTRACAM on 2012-01-15. The WD and bright spot features are very clearly separated.

with the WD, but throughout 2012 the ingress and egress is clearly separable. In this system it appears that the WD begins exiting eclipse before the bright spot has finished entering eclipse. The light curves are suitable for modelling. Figure 4.15 shows our clearest light curve.

CSS111003:054558+022106

This object, also known as Te 11, was initially discovered as a planetary nebula (Jacoby et al., 2010). It was shown to be eclipsing, with an orbital period of 0.12 days, before its reclassification as a CV (Miszalski et al., 2016). It was flagged as a dwarf nova candidate by CRTS with numerous outbursts of amplitudes around 1-2 magnitudes from its quiescent state of $g = 19.4$ (SDSS). Drake et al. (2014) took spectra of the object, which showed very strong double-peaked Balmer emission lines, as well as [OIII] emission. Recently, Miszalski et al. (2016) presented a detailed study of the eclipses, using high cadence observations to extract some approximate system parameters using the light

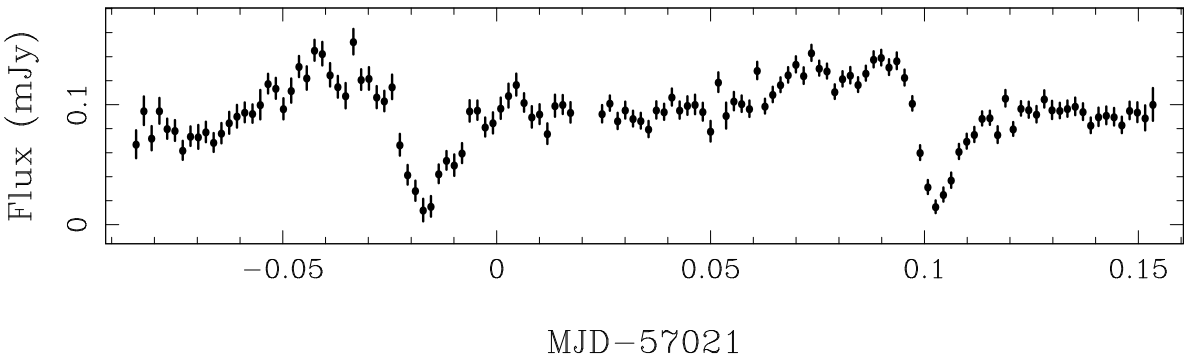

Figure 4.16: CSS111003:054558+022106 light curve observed by *pt5m* on 2014-12-29, showing two eclipses.

curve modelling technique. They also proposed that the nebula surrounding the system may in fact be the remnants of a nova eruption 1500 years ago.

We studied this system prior to the publication of Miszalski et al. (2016), observing 10 eclipses with *pt5m* and ULTRASPEC. Figures 4.16 and 4.17 show examples of these eclipse observations. Our ULTRASPEC light curves are noisy due to the faintness of this system, but a few eclipses showed well-separated bright spot and WD features (e.g. Figure 4.17). Further observations from a larger telescope should allow more precise measurement of the system parameters than those published by Miszalski et al. (2016).

CSS111101:233003+303301

This object was discovered by CRTS with an outburst of around 2 magnitudes from the quiescent level observed in SDSS (g=18.9-19.5). The long-term CSS light curve shows variability over 3 magnitudes, with some short-term variability of around one magnitude. We see this short-term variability in our own observations (see Table 4.2). There are also numerous faint dips, leading CRTS to suspect this system might be eclipsing. Our *pt5m* observations confirm these eclipses, as shown in Figure 4.18.

We also conducted high speed observations with ULTRASPEC (see Figure 4.19). The light curve is dominated by strong flickering, and the eclipse is shallow, suggesting a possible grazing eclipse. This would make the system a poor candidate for light curve

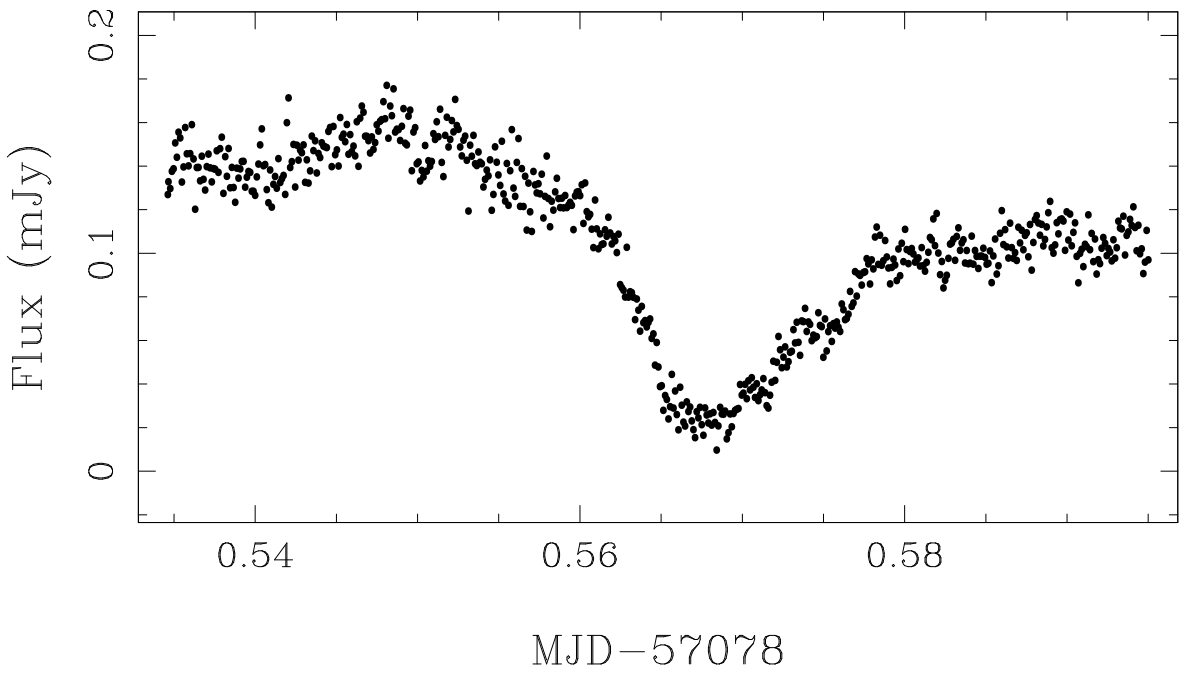

Figure 4.17: Eclipse of CSS111003:054558+022106 observed with ULTRASPEC on 2015-02-25. In this observation the ingress of the bright spot is separated from that of the WD, but in other observations this separation was not so clear.

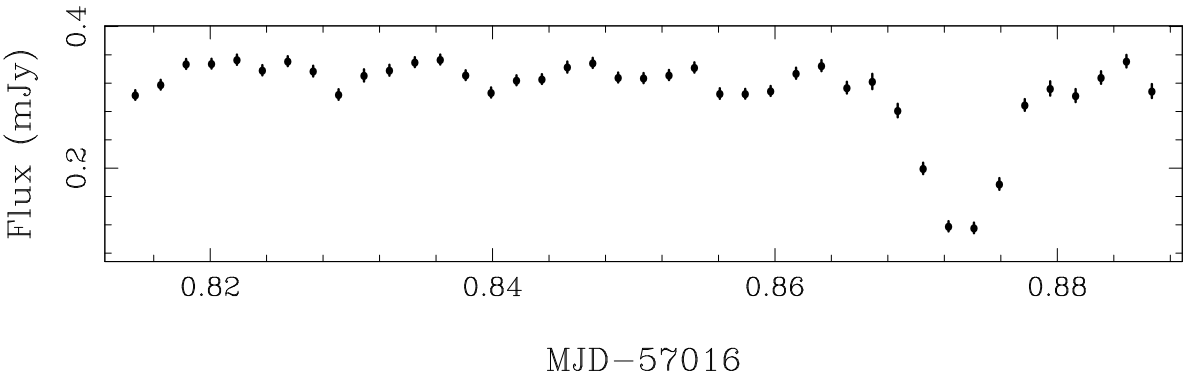

Figure 4.18: CSS111101:233003+303301 light curve showing one eclipse, observed on 2014-12-25 with *pt5m*.

modelling.

CSS131106:052412+004148

This object was discovered as a dwarf nova by CRTS, having a quiescent magnitude in SDSS of $g = 18.3$. The CSS light curve shows multiple outbursts in recent years, as well as occasional faint measurements (see Figure 4.6), prompting us to search for and

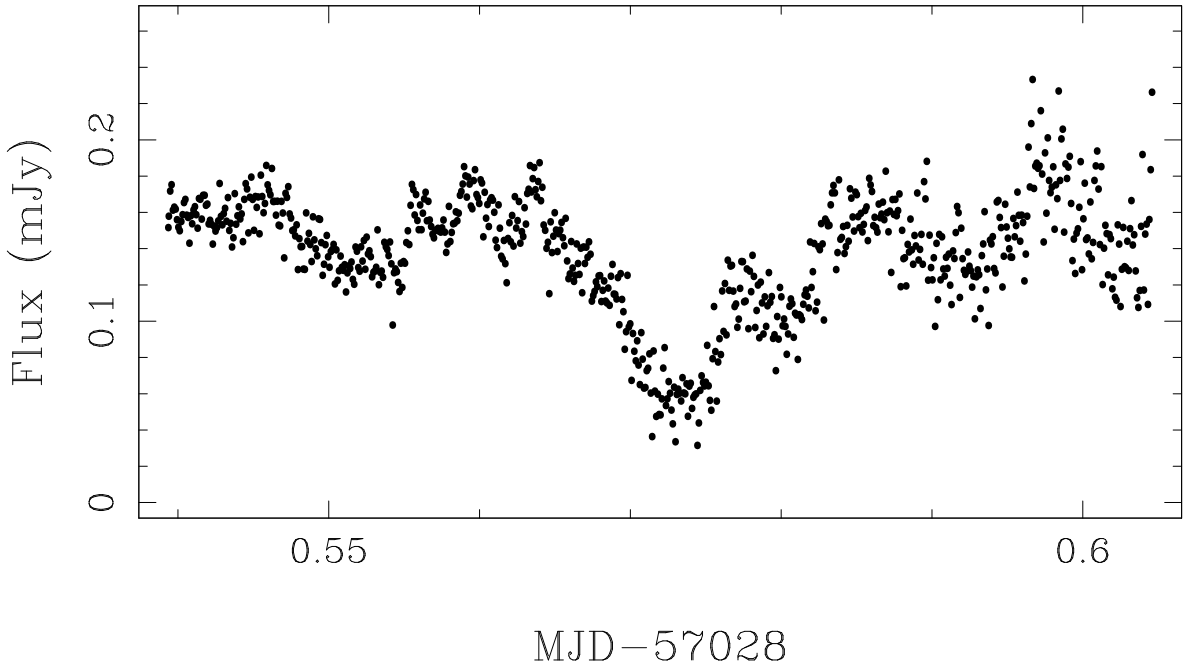

Figure 4.19: Eclipse of CSS111101:233003+303301 observed with ULTRASPEC on 2015-01-06. The light curve shows possible bright spot features but is dominated by flickering.

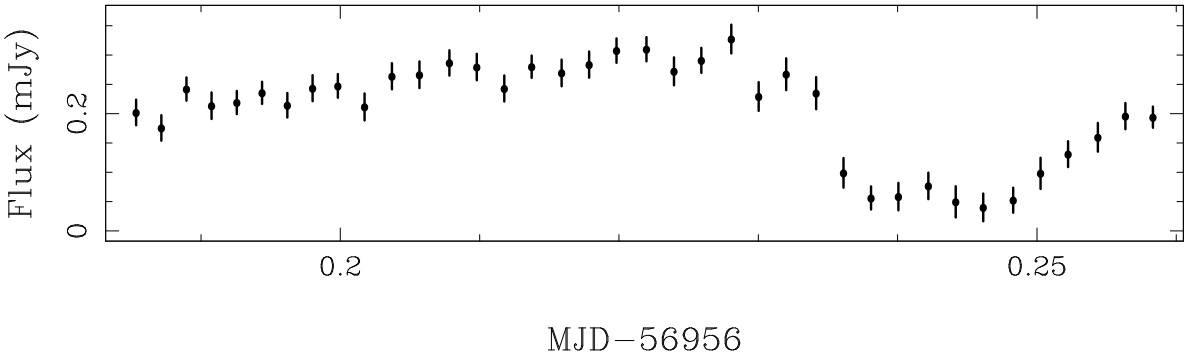

Figure 4.20: Quiescent eclipse of CSS131106:052412+004148 observed by *pt5m* on 2014-10-25.

subsequently discover eclipses with *pt5m*. An example light curve from *pt5m* is shown in Figure 4.20.

We observed CSS131106:052412+004148 at high speed with ULTRASPEC to investigate the structure of the eclipses. Figure 4.21 shows an example light curve. The eclipses show clear separation between the WD egress and bright spot egress, although separation of the ingresses is less clear. We hope that with more data from a larger telescope

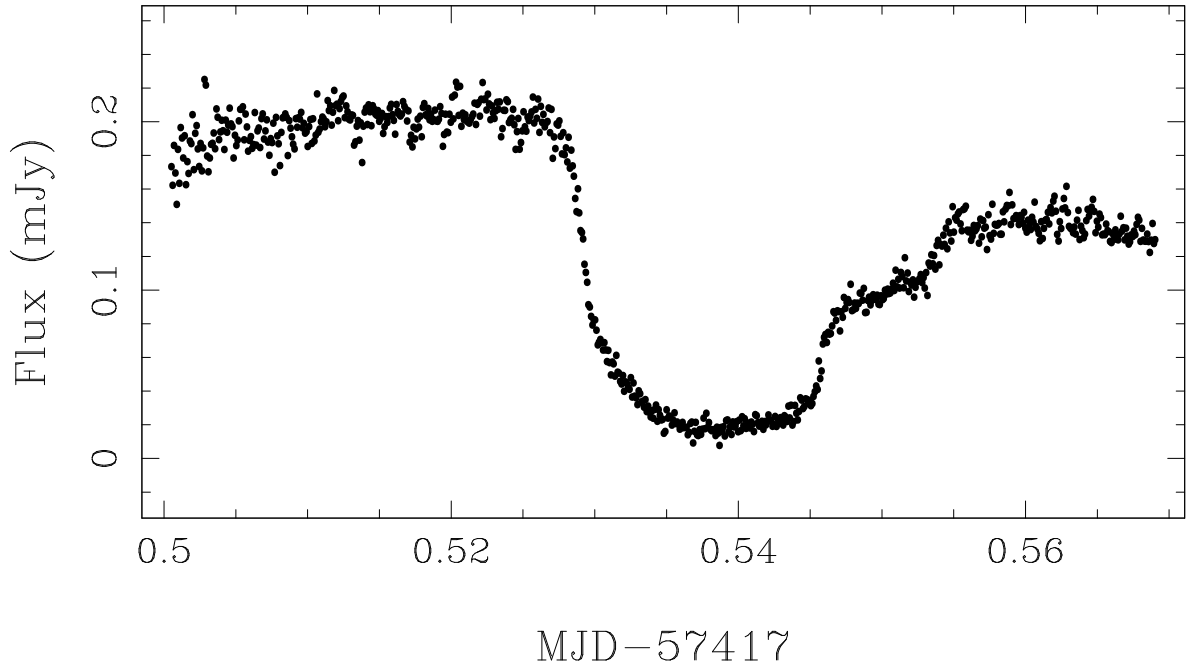

Figure 4.21: Eclipse of CSS131106:052412+004148 observed with ULTRASPEC on 2016-01-30.

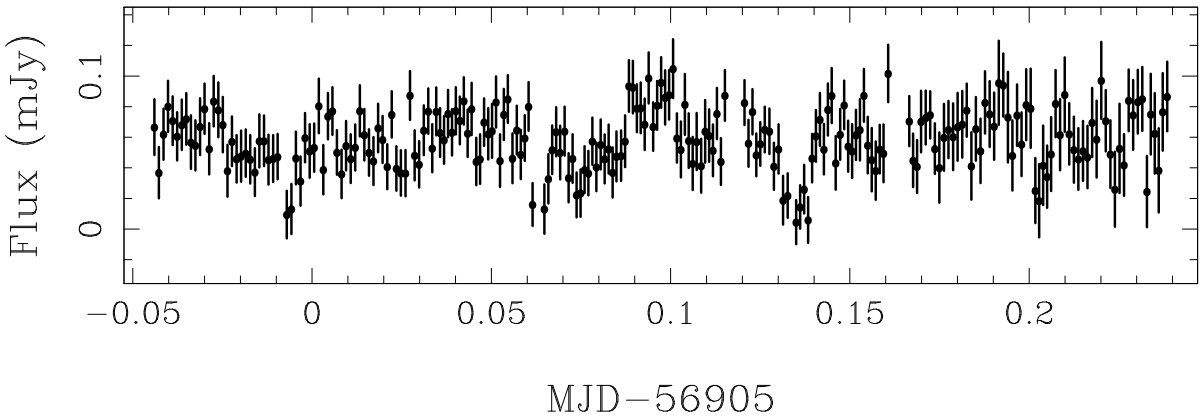

Figure 4.22: MASTER OT J003059.39+301634.3 light curve observed with *pt5m* on 2014-09-04, showing four eclipses separated by 0.07 days.

this system will be a good candidate for modelling.

MASTER OT J003059.39+301634.3

This object was discovered by MASTER (Shurpakov et al., 2014a) with an outburst to 16.1 magnitudes (unfiltered). It is listed as $g = 18.1$ in SDSS, though we suspect

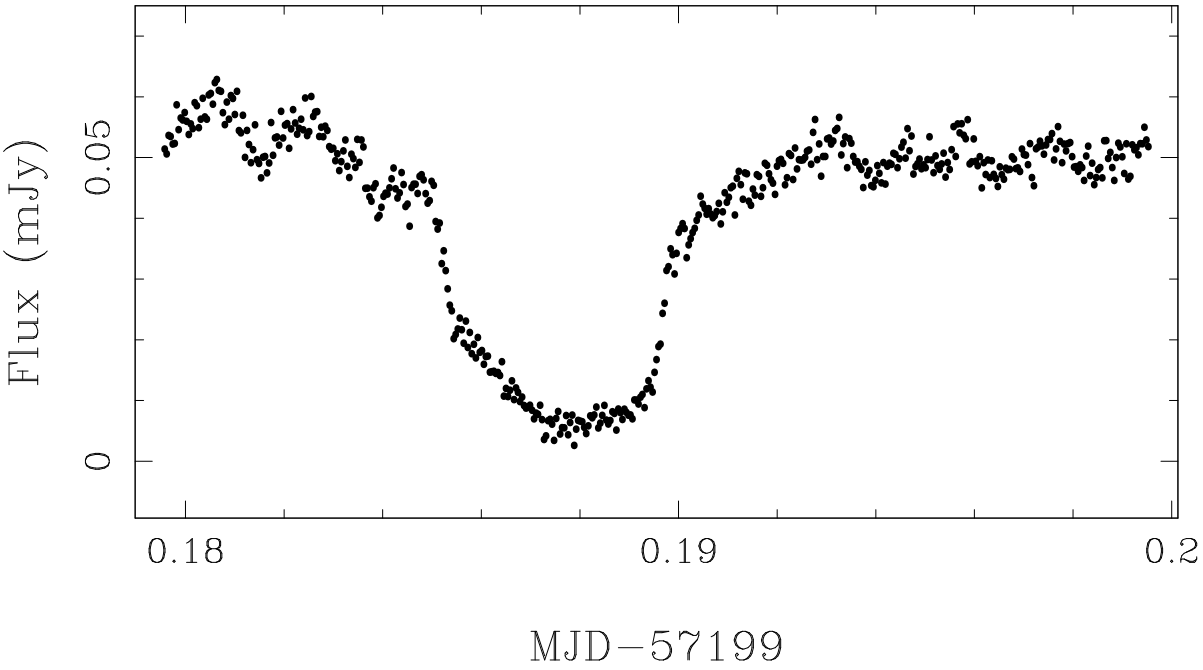

Figure 4.23: MASTER OT J003059.39+301634.3 eclipse observed in g' with ULTRACAM on 2015-06-25. It shows a clear WD ingress and egress, and a gradual bright spot eclipse. Flickering is also present.

this magnitude may have been observed during a brightened state, as the USNO-B1.0 catalogue (Monet et al., 2003) and our own observations suggest it is fainter than 19th magnitude in B and V. We first spotted eclipses in a 7-hour observing run with *pt5m*, shown in Figure 4.22.

We later observed MASTER OT J003059.39+301634.3 with ULTRASPEC and ULTRACAM, obtaining several high cadence light curves. The eclipse light curve with the highest signal-to-noise ratio is shown in Figure 4.23, and displays a clear WD ingress and egress. There is also a suspected bright spot component, but one which eclipses slowly and follows directly after the WD eclipse. It is also unclear exactly where the transition from bright spot egress to full out-of-eclipse flux occurs due to flickering. Unless the bright spot becomes more prominent in the future, this system presents a challenging light curve to model.

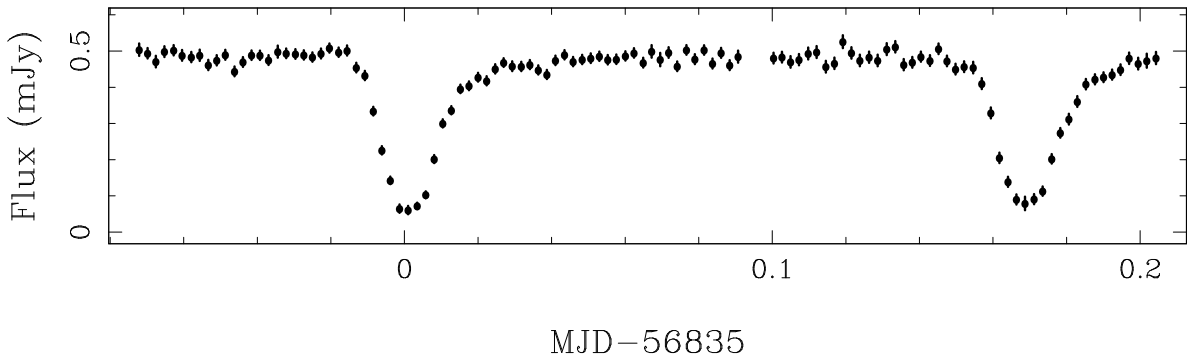

Figure 4.24: MASTER OT J192328.22+612413.5 light curve observed with *pt5m* on 2014-06-26 during a dwarf nova outburst.

MASTER OT J192328.22+612413.5

This object was discovered in outburst at 17.5 magnitudes by MASTER (unfiltered, Balanutsa et al. 2014a), with a historical Palomar light curve showing significant variability. It was suspected to be eclipsing, and this was confirmed by observations at the Vatican Observatory (Garnavich & Kennedy, 2014). We independently confirmed its eclipsing nature with *pt5m* observations of 5 eclipses. A light curve showing two eclipses observed with *pt5m* during outburst is shown in Figure 4.24. Very recently, Kennedy et al. (2016) published further photometric and spectroscopic studies of this system, showing that it shares similarities with the SW Sex class of CVs during outburst.

We also studied MASTER OT J192328.22+612413.5 at high time resolution with ULTRACAM, in both a quiescent and outburst state. The outburst eclipse mimics the sharp, featureless profile seen with *pt5m*, whilst the quiescent eclipse (see Figure 4.25) shows clearly the WD ingress and egress. There is no obvious bright spot component. Our observations are similar to those of Kennedy et al. (2016) in both outburst and quiescence, and confirm that the system is not suitable for light curve modelling.

MASTER OT J232100.42+494614.0

This object was discovered by MASTER (Shumkov et al., 2014b) with an outburst to 15.5 magnitudes. In quiescence, it can be as faint as $V = 19$ magnitudes. We first discovered

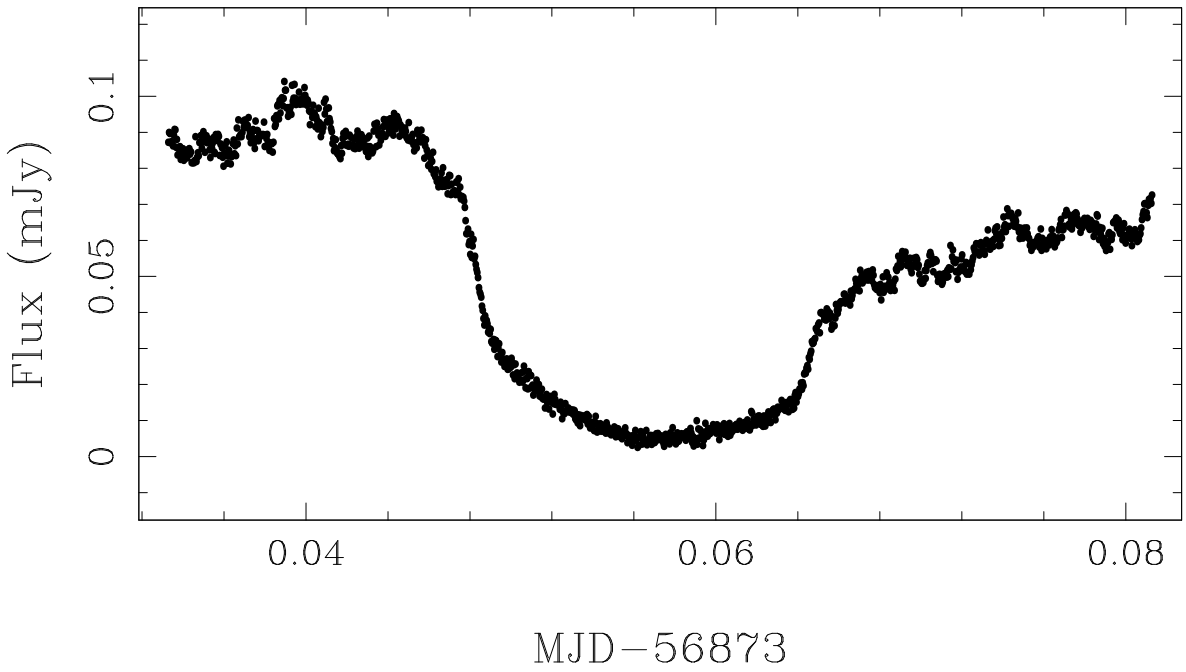

Figure 4.25: MASTER OT J192328.22+612413.5 g'-band light curve observed on 2014-08-03 with ULTRACAM. There is some flickering and a lack of clear bright spot features.

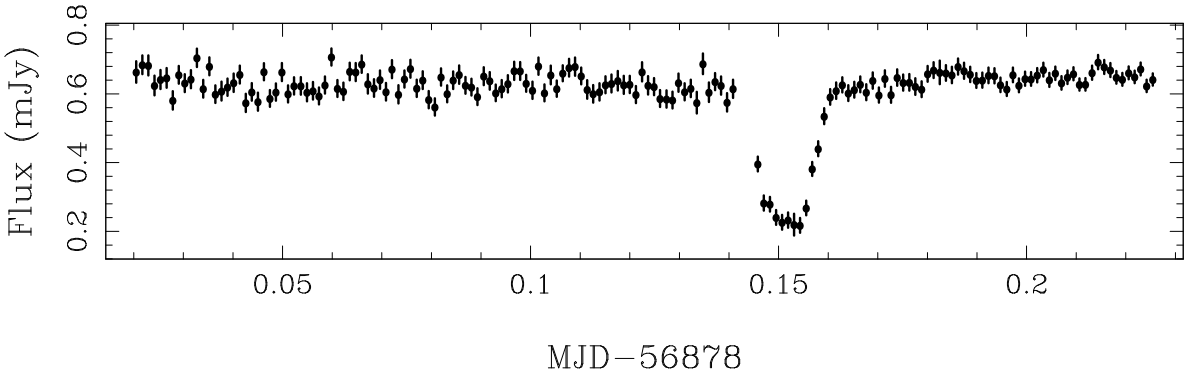

Figure 4.26: Eclipse of MASTER OT J232100.42+494614.0 during outburst, observed by *pt5m* on 2014-08-08.

eclipses when the system was in outburst, and found MASTER OT J232100.42+494614.0 to have a long period of over 5 hours. Observations are detailed in Table 4.2 and an example light curve from *pt5m* is shown in Figure 4.26.

We observed the system at high speed with ULTRASPEC (see Figure 4.27). The eclipse is noisy, and although there are signs of bright spot features, better data is needed before we can judge conclusively the potential for modelling this system. This

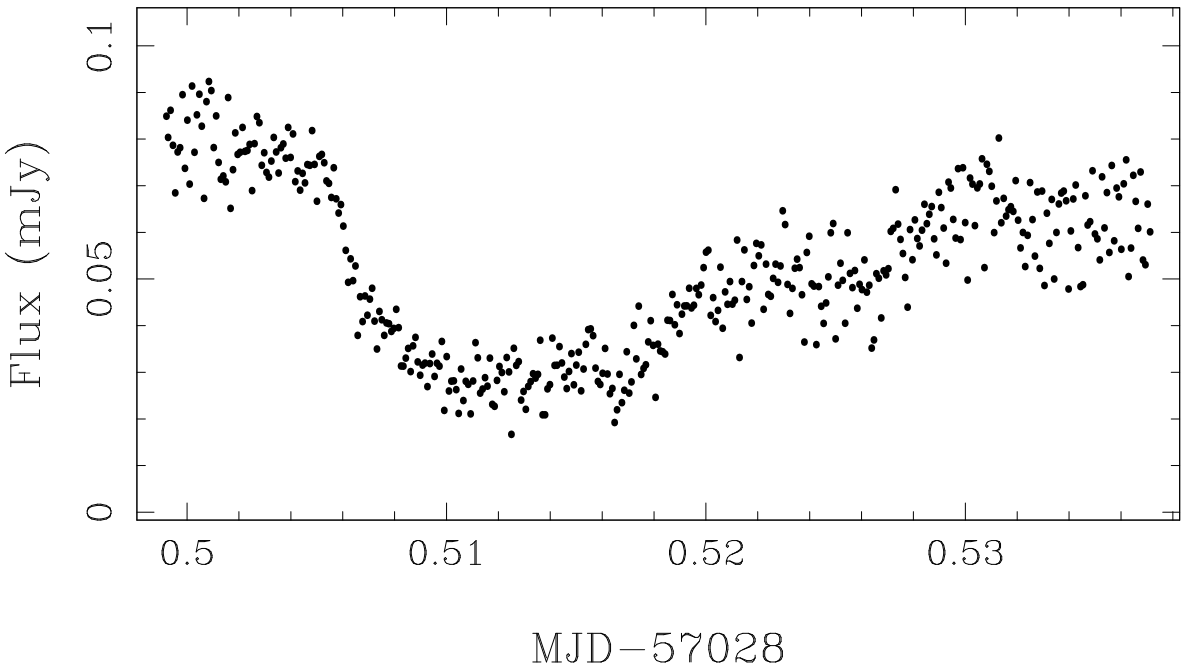

Figure 4.27: MASTER OT J232100.42+494614.0 eclipse observed with ULTRASPEC on 2015-01-06. The eclipse is too noisy to determine any definite features.

is a priority, as a system with such a long period would be particularly useful to the CV evolution studies, since measurements of system parameters at long periods remain scarce (Zorotovic, Schreiber & Gänsicke, 2011).

MLS101226:072033+172437

This object was discovered by CRTS in the MLS (Drake et al., 2009). The MLS long-term light curve shows obvious variability, which was interpreted as eclipses (Drake et al., 2014). However, no detailed studies have, to our knowledge, been published. We observed the system on 8 occasions with *pt5m* and confirmed its eclipsing nature with 7 eclipses. An example of one of these eclipses is shown in Figure 4.28.

We observed MLS101226:072033+172437 with ULTRASPEC on 2015-02-26. Unfortunately we began observations a little late, missing the beginning of eclipse ingress. The eclipse (shown in Figure 4.29) is smooth and dominated by the disk feature, showing no strong WD or bright spot components. If the eclipse structure remains like this, the

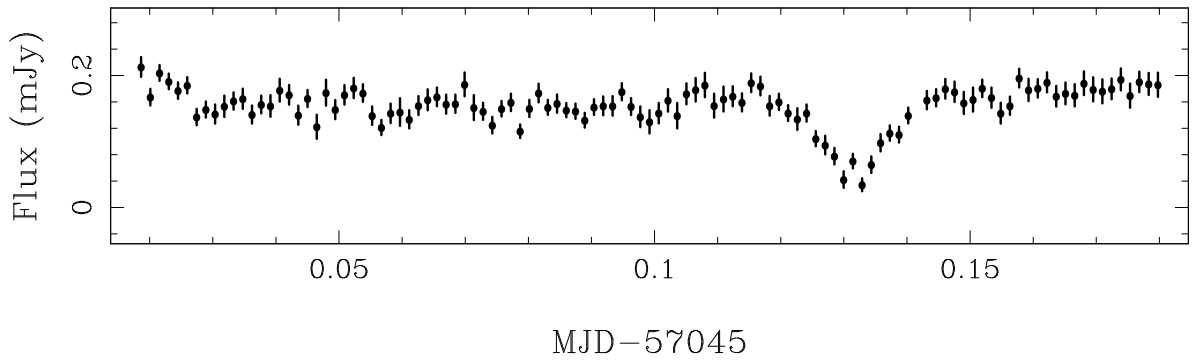

Figure 4.28: MLS101226:072033+172437 light curve observed with *pt5m* on 2015-01-22 confirming the eclipsing nature of the system.

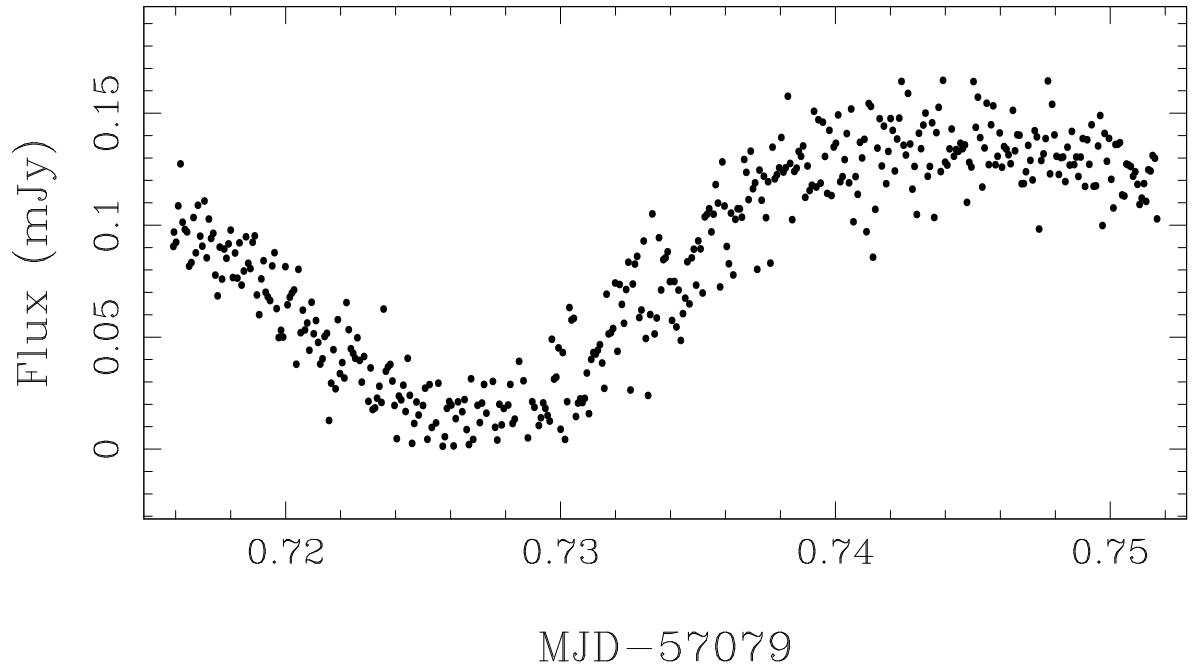

Figure 4.29: MLS101226:072033+172437 eclipse observed on 2015-02-26 with ULTRASPEC. The eclipse is fairly smooth and featureless.

system will not be suitable for light curve modelling.

MLS120517:152507-032655

This object was flagged as a possible MLS CV candidate in 2012, after also showing substantial variability in the CSS data as far back as 2007. The source has a quiescent magnitude of around V=19, but appears to have been brightening for several years and

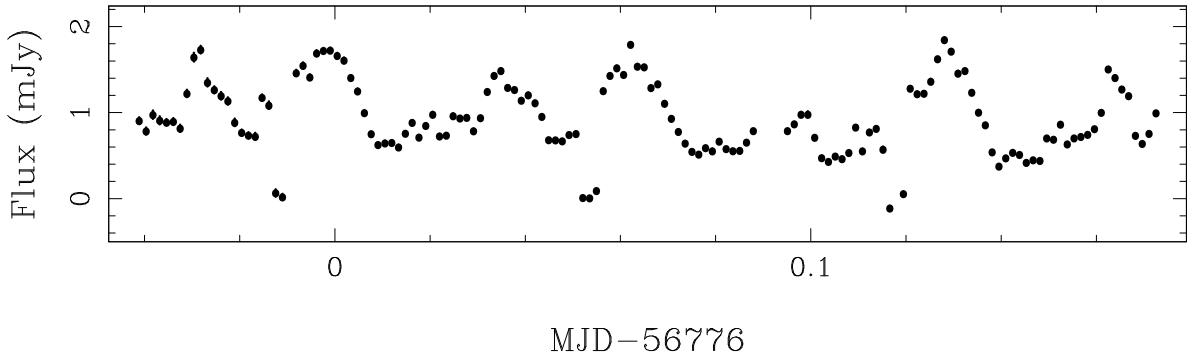

Figure 4.30: MLS120517:152507-032655 light curve observed on 2014-04-28 with *pt5m* confirming the eclipsing nature of the system. The high amplitude (>1 magnitude) and unstable variability in the light curve suggests that the system may be a polar.

is now seen at $V\sim16$. The object, also known as 1RXS J152506.9-032647, is detected in X-rays and the UV (Voges et al., 1999). It was proposed to show deep eclipses by Drake et al. (2014), and although no further studies have been published on MLS120517:152507-032655, it has been observed by amateur astronomers[3] and does indeed show eclipses.

We made four observations of the system with *pt5m* and one observation with ULTRACAM, all of which show the system to be brighter after eclipse than before. This behaviour is typical for a polar CV (Hellier, 2001), when the eclipse signal is primarily due to the obscuration of a hot spot on the surface of the WD. Most of the light in the system comes from this hot spot where the accretion stream impacts the surface, and from the stream itself that trails between this spot and the L1 point on the donor star. The projection effect of this stream causes the system to appear bright just after the eclipse, when the stream is being viewed more side-on than before the eclipse.

As a polar, the light curve is not suitable for modelling to determine the system parameters. Figures 4.30 and 4.31 show the *pt5m* and ULTRACAM light curves.

[3] `http://var2.astro.cz/EN/obslog.php`

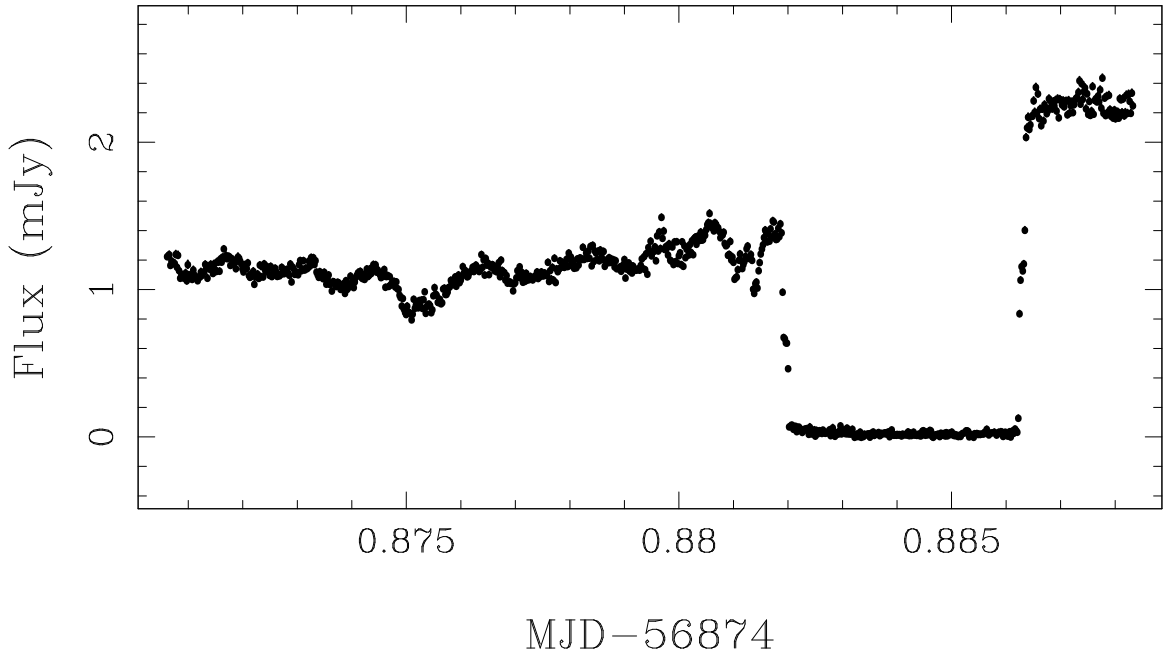

Figure 4.31: MLS120517:152507-032655 eclipse light curve observed on 2014-08-05 with ULTRACAM in i'. The sharp eclipse and step-up after eclipse is characteristic of polars.

SSS130413:094551-194402

Long suspected of being a variable star, this system was first known as NSV4618 (Kukarkin et al., 1981). It was flagged as a CV by CRTS when it went into outburst, and the Siding Springs Survey (SSS, part of CRTS) light curve also shows numerous faint measurements associated with eclipses. The *vsnet* collaboration reported eclipses and estimated an orbital period of 1.6 hours (alert 15615).

We confirmed the eclipsing nature with observations of 15 eclipses. Figure 4.32 shows an example of the eclipse structure of SSS130413:094551-194402, which displays clear WD and bright spot eclipses that should be suitable for light curve modelling.

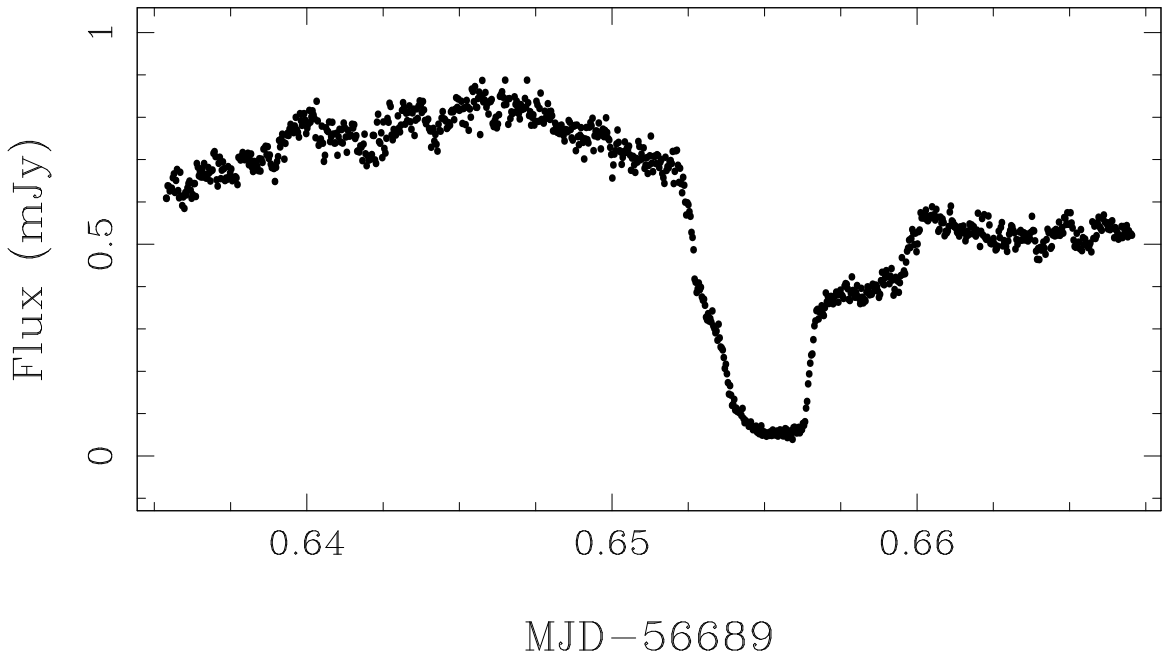

Figure 4.32: SSS130413:094551-194402 eclipse observed with ULTRASPEC on 2014-02-01, showing clearly separated WD and bright spot features.

4.5 Results: Systems Not Showing Eclipses

All data in this section were reduced and analysed by me.

The vast majority (>80%, Gänsicke et al., 2009) of CV systems have orbital periods less than 5 hours (Ritter & Kolb, 2003). In addition, eclipse lengths tend to scale with the orbital period, such that longer period systems have longer eclipses. This means that a 5-hour observation showing no sign of eclipses would suggest the system either does not eclipse, or has a period of at least 5.5 hours. We can therefore say that any systems we have observed continuously for 5 or more hours that show no signs of eclipses are unlikely to be eclipsing.

Sometimes shorter observations can also help to rule out eclipses, if for example the system shows periodic variability. This could be due to superhumps, the projection effect of the orbiting bright spot, or the oblation of the secondary star as it fills its Roche lobe. These modulations occur on or close to the orbital period, thus any observation covering at least one cycle of such modulations without showing an eclipse is sufficient to rule out

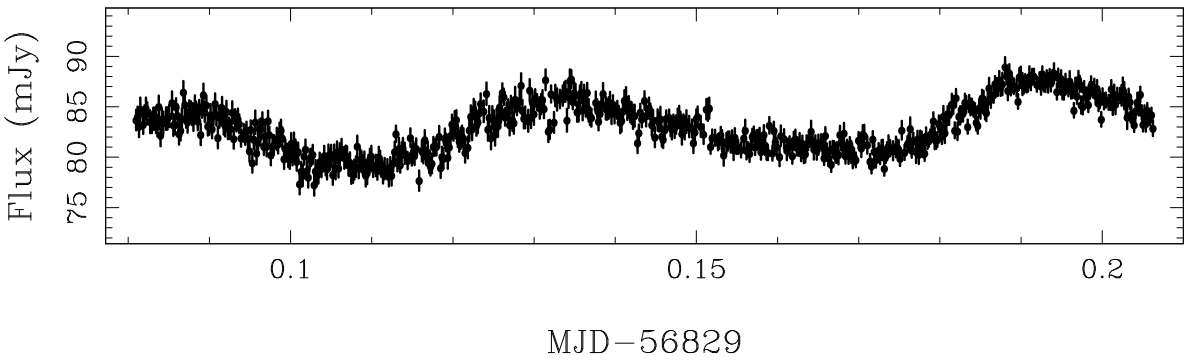

Figure 4.33: ASASSN-14cl light curve observed on 2014-06-20 with *pt5m*, showing superhumps with a period of approximately 1.4 hours.

the presence of eclipses in the system.

In this section we discuss those systems which did not show eclipses, and present light curves of those which show clear or possible periodic variability. Those systems which show only flickering or no variation are not shown here. Instead these light curves are available in Appendix B. Details of the observations are given in Table 4.2.

ASASSN-14cl

This object was discovered by ASAS-SN with a super-outburst to $V = 10.7$ in June 2014 (Stanek et al., 2014a). It has a quiescent magnitude of $g = 18.8$ in SDSS, and was studied extensively by the *vsnet* collaboration (Kato et al., 2015). We observed the object on several occasions, with some observations showing superhump behaviour with an estimated period of 1.4 hours. The light curve with the clearest superhumps is shown in Figure 4.33. No eclipses were seen.

ASASSN-14ds

ASASSN-14ds was discovered in outburst at $V = 15.7$ (Holoien et al., 2014). It is a known X-ray source (1RXS J204455.9-115151) and was studied spectroscopically (Masetti et al., 2014), confirming its nature as a CV. In quiescence this object is known to be variable over long timescales, as is seen in the CSS light curve, with a mean magnitude of around

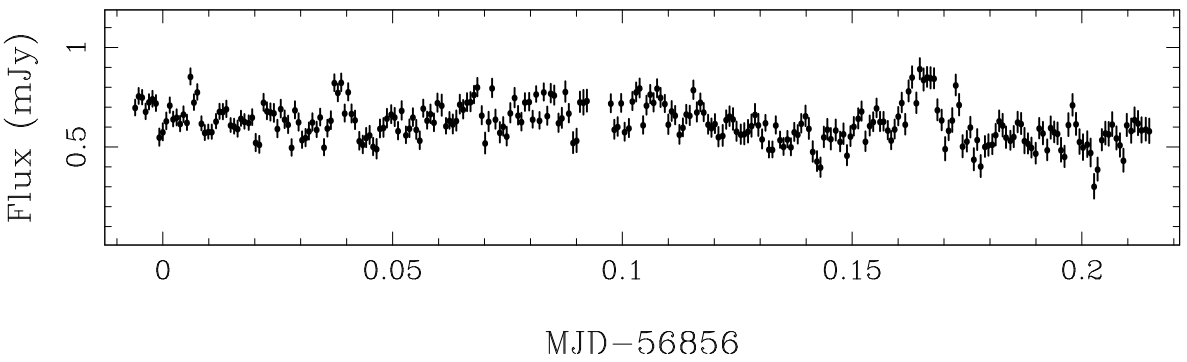

Figure 4.34: ASASSN-14ds light curve observed during outburst on 2014-07-17 with *pt5m*. There is obvious variability, probably associated with flickering, as well as a possible sinusoidal modulation, but no eclipses.

16.8. The object was observed on several nights following its discovery, with the longest, spanning 5.3 hours, being shown in Figure 4.34. All observations showed flickering, but no eclipses were observed.

ASASSN-14gl

This object was discovered by ASAS-SN during an outburst to $V = 15.7$, in a region outside the CSS footprint. The quiescent magnitude is listed as $g = 18.8$ in SDSS. The object was observed with *pt5m* on 5 occasions during quiescence, with the longest observation spanning 5.7 hours. The system shows signs of flickering with an amplitude of 0.8 magnitudes, but we observed no eclipses.

ASASSN-14gu

This object was discovered by ASAS-SN with an outburst to $V = 15.2$ in September 2014. It is listed with a quiescent magnitude of $V = 17.4$ in the NOMAD catalogue. We observed the object on three occasions, during outburst and quiescence, all of which showed a slow variability with a period of approximately 2.8 hours. Since this variability persists even in quiescence, it cannot be attributed to superhumps. It may be due to the projection effect of the bright spot, or of the oblate secondary star. The longest

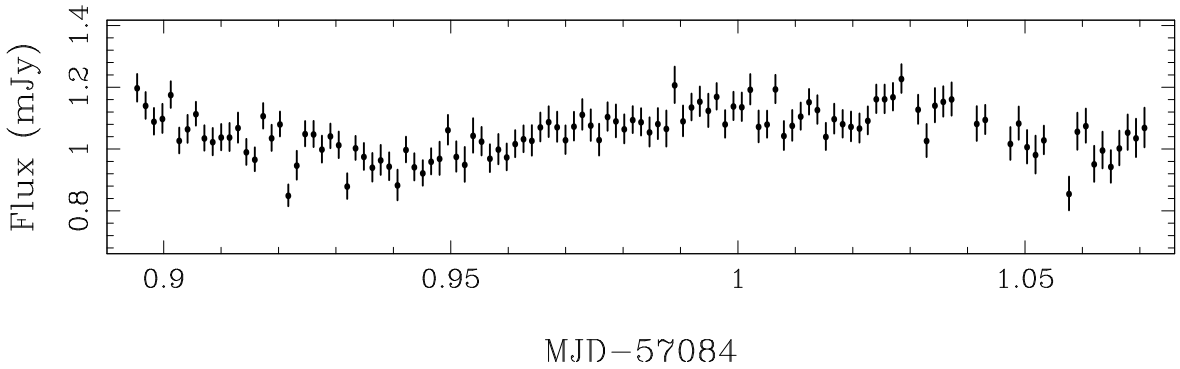

Figure 4.35: ASASSN-14gu light curve observed on 2015-03-03 with *pt5m*. The periodic modulation is present in outburst and in quiescence, meaning that it cannot be attributed to superhumps.

observation, shown in Figure 4.35, covers 4.2 hours and at least one period of this modulation. No eclipses were seen.

ASASSN-14hk

ASASSN-14hk was discovered by ASAS-SN during a bright outburst to $V = 14.7$ (Stanek et al., 2014b). The object is suspected to be an SU UMa system (defined as showing large super-outburst amplitudes), as the outburst amplitude is ~ 6 magnitudes, with the quiescent magnitude listed as $B = 20.8$ in the USNO-B1.0 catalogue. We were only able to observe this system for 3 hours during outburst, the light curve for which is shown in Figure 4.36. The light curve shows evidence for superhumps, with an approximate period of 90 minutes. This observation agrees with that reported by Kato et al. (2015) and *vsnet* (alert 17773). We conclude that the system is not eclipsing, and is likely to have an orbital period close to 90 minutes.

ASASSN-14mv

This object was discovered by ASAS-SN with a super-outburst to $V = 12$. There are also recorded outbursts in historical data, by MASTER, CSS, and archival photographic plates (Denisenko et al., 2014a). It has a quiescent magnitude of $B = 17.3$ in the

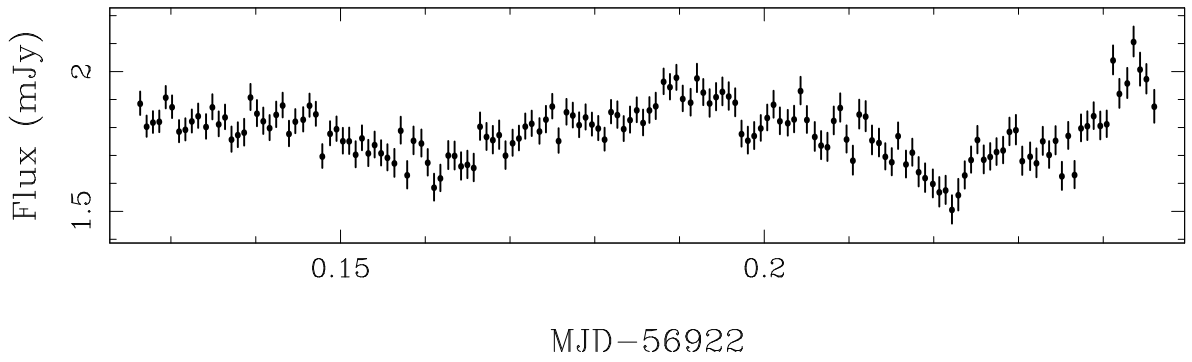

Figure 4.36: ASASSN-14hk light curve showing superhumps, observed during super-outburst on 2014-09-21 with *pt5m*.

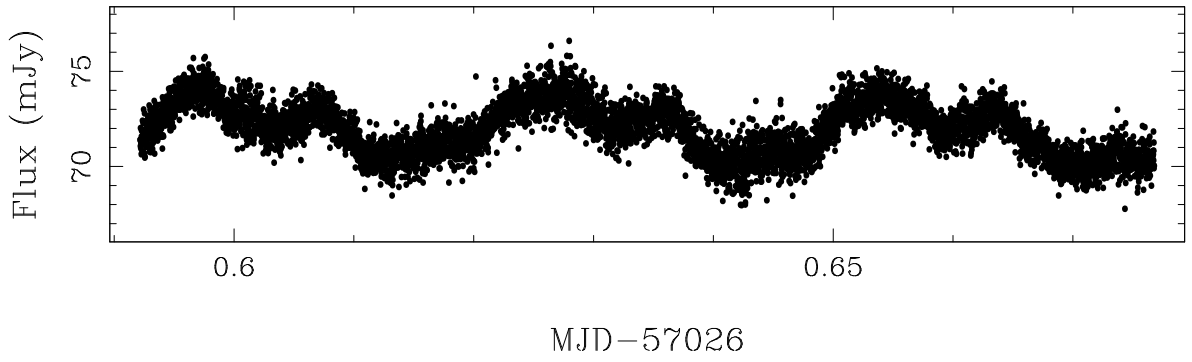

Figure 4.37: ASASSN-14mv light curve observed on 2015-01-04 with ULTRASPEC. The periodic modulation is associated with superhumps, and has a period of approximately 40 minutes. Error bars have been omitted for clarity.

USNO-B1.0 catalogue. We observed ASASSN-14mv for two hours during outburst with ULTRASPEC. The light curve is shown in Figure 4.37, showing strong periodic variability associated with superhumps, suggesting an orbital period of the order of 40 minutes. This, along with other observations (e.g. *vsnet* alert 18124 and 18160) suggests that ASASSN-14mv may be an AM CVn system. The system is not eclipsing.

ASASSN-15ni

This object was discovered by ASAS-SN (Dong et al., 2015) with a high amplitude outburst to $V = 12.9$ magnitudes, from an unknown quiescent magnitude (undetected in all surveys). Spectroscopy confirmed its nature as a dwarf nova outburst (Berardi, 2015), and we observed the target on four nights with *pt5m*, the longest spanning 5.2

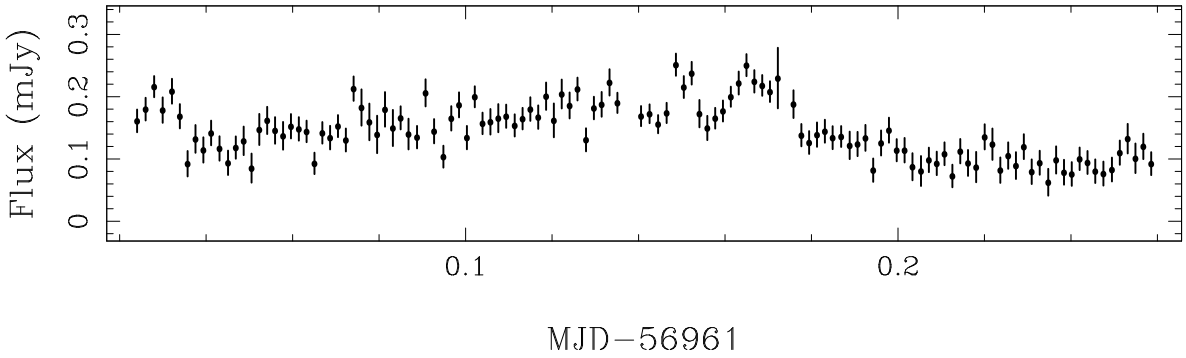

Figure 4.38: CSS090219:044027+023301 light curve observed during quiescence on 2014-10-30 with *pt5m*.

hours. We saw evidence for flickering, with an amplitude of up to 0.1 magnitudes, but no eclipses.

CSS090219:044027+023301

This object showed several outbursts in CSS, to around 16th or 17th magnitude in V, from a quiescent level varying between 18 and 19 magnitudes. It was studied spectroscopically by Thorstensen & Skinner (2012), who noted double-peaked emission lines, indicative of an edge-on accretion disk. This made it a good candidate eclipsing system. CSS090219:044027+023301 was observed on four occasions with *pt5m*, the longest spanning 5.7 hours on 2014-10-30. The light curve, shown in Figure 4.38, shows variability but no eclipses. The variability may be periodic, but is more likely to be flickering.

CSS091116:232551-014024

This object was discovered by CRTS with an outburst to $V = 15.9$, and shows several other outbursts to similar magnitudes in the CSS light curve (Djorgovski et al., 2011). The quiescent magnitude shows long-term variability, with the CSS light curve ranging between 18 and 19 magnitudes, as well as occasional fainter measurements. This behaviour often indicates eclipses, but we did not observe any. Our longest observation, over 6 hours, shows some flickering with an amplitude of 0.5 magnitudes.

CSS100508:085604+322109

This object has shown several outbursts since its discovery by CRTS in 2010, most of them reaching magnitudes of around 16 or 17 in V. The quiescent brightness is listed as $g = 19.6$ in SDSS. Kato, Maehara & Uemura (2012) predicted an orbital period of around 95 minutes based on the SDSS colours, though this is uncertain. A 4.8 hour light curve observed during quiescence on 2015-03-16 with *pt5m* shows some flickering of amplitude 0.8 magnitudes, but no sign of eclipses.

CSS100520:214426+222024

This object was discovered by CRTS with an outburst to $V = 14.7$. The quiescent magnitude is listed in SDSS as $g = 17.6$, and the average CSS magnitude is $V = 17$. We observed this object for 5.9 hours on 2014-10-08 in windy, non-photometric conditions. The light curve shows no sign of eclipses or any variability above the noise level of 0.4 magnitudes.

CSS110114:091246-034916

This object was discovered by CRTS with an outburst to around $V = 16$. The quiescent brightness is listed as $R = 18$ in USNO-B1.0. CSS110114:091246-034916 was studied spectroscopically by Thorstensen & Skinner (2012), who estimated an orbital period of approximately 0.06 days. The target was observed on several occasions, with two observations spanning almost 5 hours each. We found no eclipse signals in either of these 5-hour observations, but see some possible flickering with amplitude 0.5 magnitudes.

CSS110226:112510+231036

This object was discovered by CRTS with an outburst to $V = 16$. The quiescent brightness is listed as $R = 19$ in USNO-B1.0 and $g = 20$ in SDSS. Based on the SDSS colours,

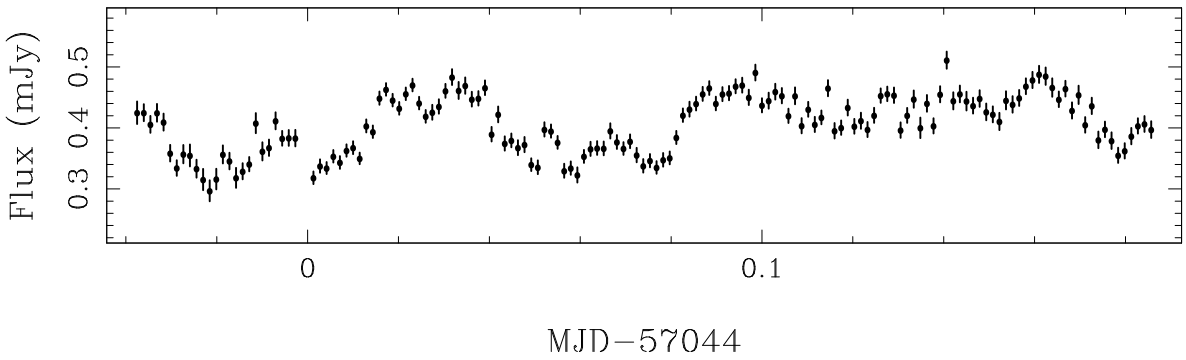

Figure 4.39: CSS130906:064726+491542 light curve observed during quiescence on 2015-01-21 with *pt5m*. Variability could be attributed to flickering, or movement through hot pixels on the CCD detector (see text for details).

Kato, Maehara & Uemura (2012) estimate the orbital period to be 1.8-1.9 hours, although again this is uncertain. The discovery of double-peaked emission lines (Breedt et al., 2014) suggested that the system might eclipse. The target was observed on three occasions, with the longest run covering 5.3 hours. We found no eclipses and no variability above the noise level (0.5 magnitudes).

CSS130906:064726+491542/MASTER OT J064725.70+491543.9

This object's 2013 outburst was recorded by CRTS and MASTER (Tiurina et al., 2013), reaching $V = 14$. The quiescent brightness is around $V = 17$. We observed CSS130906:064726+491542 once during outburst, just after discovery, and again over a year later, during quiescence. The longer, 5.3 hour observation is shown in Figure 4.39 and shows variability but no eclipses. The *pt5m* auto-guiding system failed during these observations, and so the variability could be attributed to the target and comparison stars moving across the field, through hot pixels or bad pixels, or it could be attributed to flickering. In addition, *vsnet* (alert 15476) reported signs of superhumps during outburst with a period of 1.6 hours, although superhump behaviour was not observed in our observations taken during outburst. They did not mention any sign of eclipses.

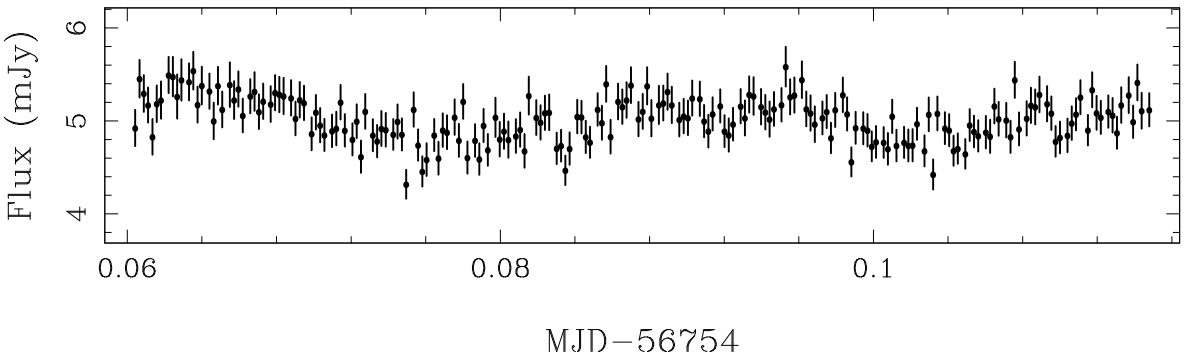

Figure 4.40: CSS140402:173048+554518 light curve showing superhumps, observed during outburst on 2014-04-06 with *pt5m*.

CSS140402:173048+554518

This object, also known as SDSS J173047.59+554518.5, was seen by CRTS exhibiting a bright outburst to $V = 14.5$. The quiescent brightness is listed as $g = 20.1$ in SDSS. The object is a confirmed AM CVn system, with an orbital period of 35 minutes (Carter et al., 2014). Although the object was not likely to be eclipsing due to it being a relatively face-on system, we observed the target in outburst during an automated follow-up from the initial CRTS transient trigger (Hardy et al., 2015b). A 1.3 hour light curve is shown in Figure 4.40, showing superhump behaviour with a period of 39 minutes, close to the orbital period. This is also similar to the 35 minutes measured by Kato et al. (2015). Unfortunately, because of the large uncertainty in the spectroscopic orbital period, and since the accretion disk is made predominantly of helium and not hydrogen, the superhump excess cannot be reliably used to measure a mass ratio (Pearson, 2007; Kato et al., 2014b), as can be done for classical SU UMa systems (Patterson et al., 2005).

CSS140901:013309+133234/ASASSN-14gk

This object was discovered by both CRTS and ASAS-SN during an outburst to $V = 15.3$. It was confirmed as a CV spectroscopically by Kaur et al. (2014). Its quiescent magnitude is listed as $g = 21.3$ in SDSS, placing it outside the usual capabilities of *pt5m*.

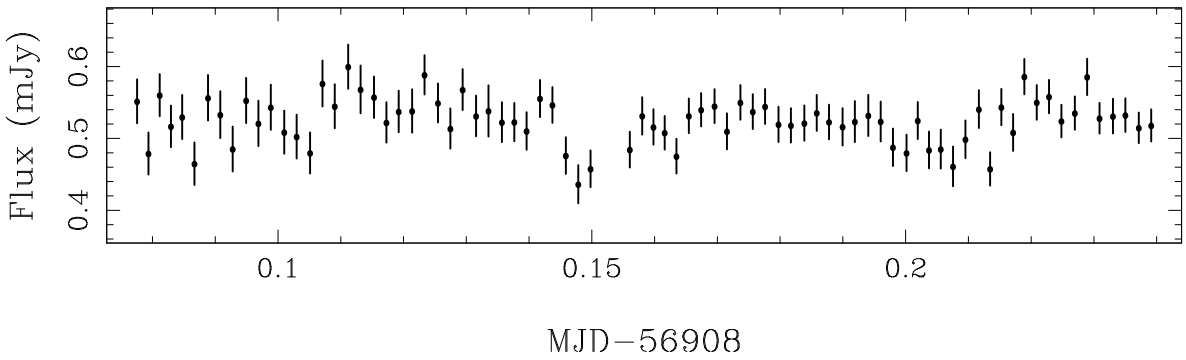

Figure 4.41: CSS140901:013309+133234 light curve observed on 2014-09-07 with *pt5m*. Here the system is fading from outburst but some superhump signals are still present.

However, the object was observed with *pt5m* on several nights during its outburst, with the longest observation spanning 3.9 hours. There were some signs of superhumps in the first observation (not shown), and these persisted until the next observations 4 days later (see Figure 4.41). No eclipse signals were observed, and although the light curves were not sufficient to derive a conclusive period, the superhumps appeared to cycle on a timescale of approximately 2.5 hours.

CSS141005:023428-045431/ASASSN-14dx

This object was discovered by both CRTS and ASAS-SN during an outburst to $V = 14.5$. It was confirmed as a CV spectroscopically by Kaur et al. (2014). The quiescent magnitude is listed as $g = 16.3$ in SDSS, making it a good target for *pt5m*. The object was observed on 5 occasions during outburst and quiescence, with the longest observation spanning 4.6 hours, though this suffered from some cloud cover. A 3-hour long observation is shown in Figure 4.42, and shows possible periodic modulation. We attribute this to superhumps, in agreement with the interpretation found by *vsnet* (alert 18017), although the signal-to-noise ratio of the light curve is insufficient to completely rule out flickering as the cause of the variability. If the signal is indeed superhumps, we would estimate a period of approximately 83 minutes. No eclipses were seen.

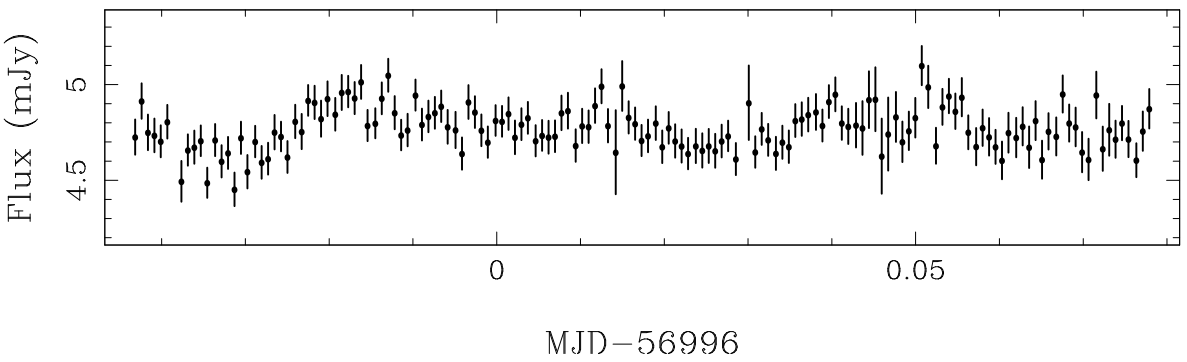

Figure 4.42: CSS141005:023428-045431 light curve observed whilst fading from outburst on 2014-12-04 with *pt5m*. The light curve is one of several observations on different nights that display periodic features resembling superhumps, although stochastic flickering is not ruled out as the cause for this variability.

CSS141117:030930+263804

This object was discovered by CRTS during a bright outburst reaching around $V = 12$. In quiescence, its SDSS g-band magnitude is 18.9. It showed no signs of eclipses in several *pt5m* observations, the longest stretching 3.75 hours, although this run was interrupted by passing cloud cover. In data taken on clearer nights, it showed no variability down to the noise limit of 0.4 magnitudes. According to Kato et al. (2015), in outburst it shows superhumps, and no eclipses.

Gaia15aan/ASASSN-14mo

This object was first discovered by Gaia in December 2014, but was not published until January 2015. By this time, ASAS-SN had also detected it as ASASSN-14mo. The outburst was recorded to reach 13 magnitudes in the Gaia G-band. It has a quiescent magnitude of $g = 19.7$ in SDSS. Although it was originally proposed to be an AM CVn system based on its colours (Carter et al., 2013), it was later seen to display a normal, hydrogen-dominated CV spectrum (Campbell et al., 2015b). We observed the object on four occasions with *pt5m*, seeing no signs of eclipses or superhumps. The longest observation, of 107 minutes, showed possible flickering, but no periodic variability above

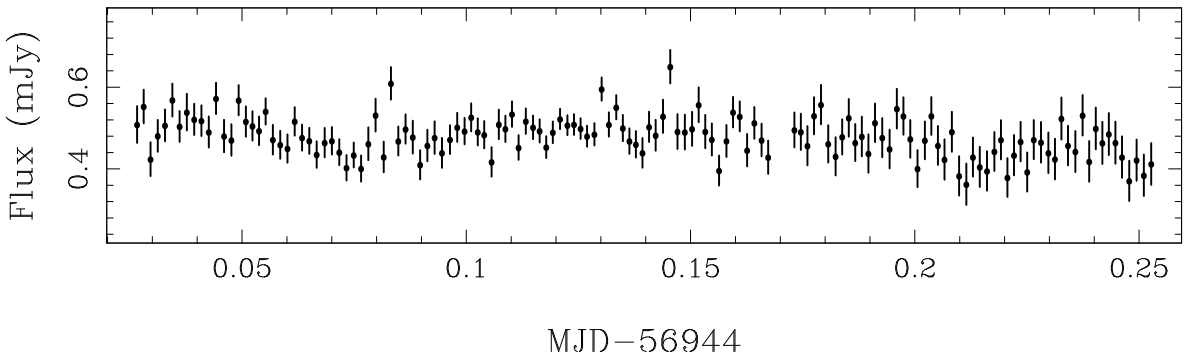

Figure 4.43: MASTER OT J041923.57+653004.3 light curve observed on 2014-09-21 with *pt5m*. The variability may be periodic, or attributed to flickering.

the noise level, which was ∼0.5 magnitudes.

MASTER OT J034045.31+471632.2

This object was discovered by MASTER during an outburst to 15.6 magnitudes (unfiltered, Denisenko et al. 2013). The quiescent magnitude is listed as $g = 18$ in SDSS, and $R = 18.5$ in USNO-B1.0. The object was observed with *pt5m* for 6.7 hours. We observed no eclipses and no variability above the noise level of 0.6 magnitudes.

MASTER OT J041923.57+653004.3

This object was identified as a potential CV candidate from its X-ray counterpart 1RXS J041924.8+653006 in 2006 (see *vsnet* alert 16729), but was not confirmed as a dwarf nova until an outburst was seen by MASTER in 2013 (Balanutsa et al., 2013). The object was measured at 13.9 magnitudes (unfiltered) during outburst, and has a quiescent magnitude of $V = 17.9$ in the NOMAD catalogue. Our longest observation (5.4 hours, Figure 4.43) showed no eclipses, but does include some variable behaviour resembling a periodic modulation, though this could also be attributed to flickering. If the modulation is periodic, it could have a period of either 1.8 or 3.6 hours.

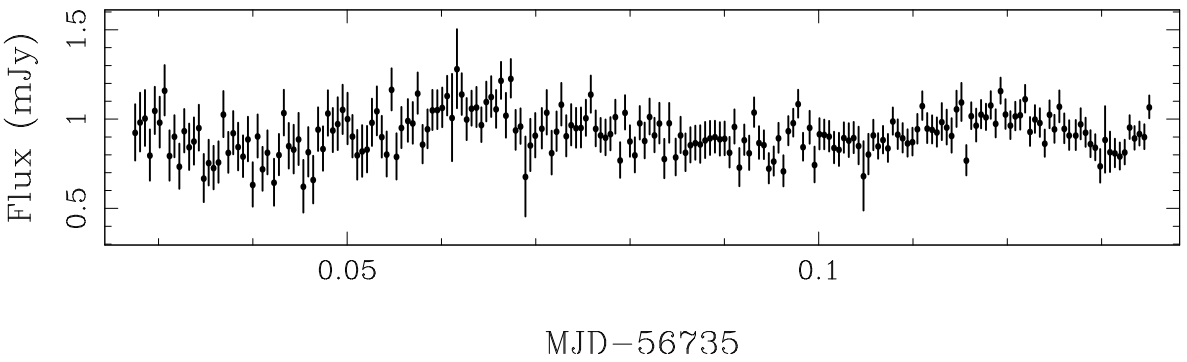

Figure 4.44: MASTER OT J171921.40+640309.8 light curve showing superhump-like variability, observed during outburst on 2014-03-18 with *pt5m*.

MASTER OT J171921.40+640309.8

This object was discovered by MASTER with an outburst to 15.6 magnitudes (unfiltered, Balanutsa et al. 2014b). The quiescent brightness is listed as $g = 21.3$ in SDSS, making it too faint for study with *pt5m*, though we were able to observe it briefly during outburst. In both observations we saw signs of periodic variability, and associate this behaviour with superhumps. A 2.5 hour light curve is shown in Figure 4.44. We estimate the superhump period to be approximately 80 minutes, and see no sign of eclipses.

MASTER OT J194955.17+455349.6

This object, also known as KIC 9358280, was discovered by MASTER (Denisenko et al., 2014c) and is located in the Kepler field. The Kepler-INT Survey (Greiss et al., 2012) gives g-band magnitudes of 19.7 and 19.4 for two different epochs, placing it at the limit of observability with *pt5m*. It showed no significant signs of eclipses or any periodic modulation above the 0.5-magnitude noise level in a 4.9 hour observing run, nor in a supplementary 100 minutes of observation.

MASTER OT J201121.95+565531.1

This object was found in outburst at 16.5 magnitudes (unfiltered) by MASTER (Shurpakov et al., 2014b). Its quiescent magnitude is unknown, but it was seen very close

to two faint stars measuring approximately $R = 20$ in USNO-B1.0. We were only able to observe the system on one occasion before it faded to quiescence, with a 3.3-hour observation on 2014-06-27. There are no signs of any variability above the noise level of 0.3 magnitudes.

MASTER OT J202157.69+212919.4

This object was discovered in outburst at 16.3 magnitudes (unfiltered) by MASTER (Denisenko et al., 2014b), and is associated with the known X-ray source RXP J202157.8+212919 (Voges et al., 2000). In quiescence it is very difficult to study with *pt5m*, listed with a V-band magnitude of 20 in the NOMAD catalogue. However, we attempted one observation of the object, collecting a faint 5 hour light curve. There were no signs of any variability, although the noise was $\sim$1 magnitude.

MASTER OT J203421.90+120656.9

This object was discovered in outburst at 17.4 magnitudes (unfiltered) by MASTER (Yecheistov et al., 2014). With a quiescent brightness of $B = 20.4$ in USNO-B1.0, this was a challenging target for *pt5m*. We observed the object on two nights, shortly after its discovery outburst. A 4.4 hour light curve showed no eclipses and no variability above the noise level of 0.3 magnitudes.

MASTER OT J210316.39+314913.6

This object was discovered in a bright outburst to 13.9 magnitudes (unfiltered) by MASTER (Shumkov et al., 2014a). The quiescent magnitude is listed as $B = 19.4$ in the USNO-B1.0 catalogue. We observed MASTER OT J210316.39+314913.6 on several occasions, the two longest of which are shown in Table 4.2. The longest (7 hour) light curve was taken on 2014-08-10 and showed no variability above the noise level of 0.5 magnitudes.

4.6 Results: Previously Known Eclipsing Systems

Approximately half the data in this section were reduced by me, while the other half were reduced by Martin McAllister. All analysis and text is my own.

We present high cadence light curves of CV systems which are already known to be eclipsing. Studies of such high time resolution do not appear elsewhere in the literature, and therefore many of these systems have never been studied before in such detail. Details of the observations are given in Table 4.2, and mid-eclipse times are provided whenever an eclipse was observed. In those systems with resolved eclipse structure, we measure the mid-eclipse time as the mid point between the WD mid-ingress and mid-egress, when these features are visible, rather than using a Gaussian fitting procedure. If the WD ingress or egress is not clear but there is visible structure, we use the ingress and egress of the most consistent features, and account for this in the uncertainties. We derived updated ephemerides for most systems, and these are given in Table 4.2.

We discuss the eclipse structure of each system in the subsections below, and comment on whether or not each system is suitable for light curve modelling to determine the system parameters. A summary of these judgements is also included in Table 4.2.

1RXS J180834.7+101041

This object was originally classified as a polar (Denisenko, Kryachko & Satovskiy, 2008), but spectroscopic observations soon showed that it had an accretion disk (Bikmaev & Sakhibullin, 2008). Its eclipses have been studied previously (Yakin et al., 2010; Southworth & Copperwheat, 2011), but attempts to measure the system parameters have had limited success because the inclination of the system is relatively low, and the WD itself might not be eclipsed.

We observed three eclipses of this system, finding strong flickering and only weak eclipse structure, possibly associated soley with the accretion disk and bright spot but

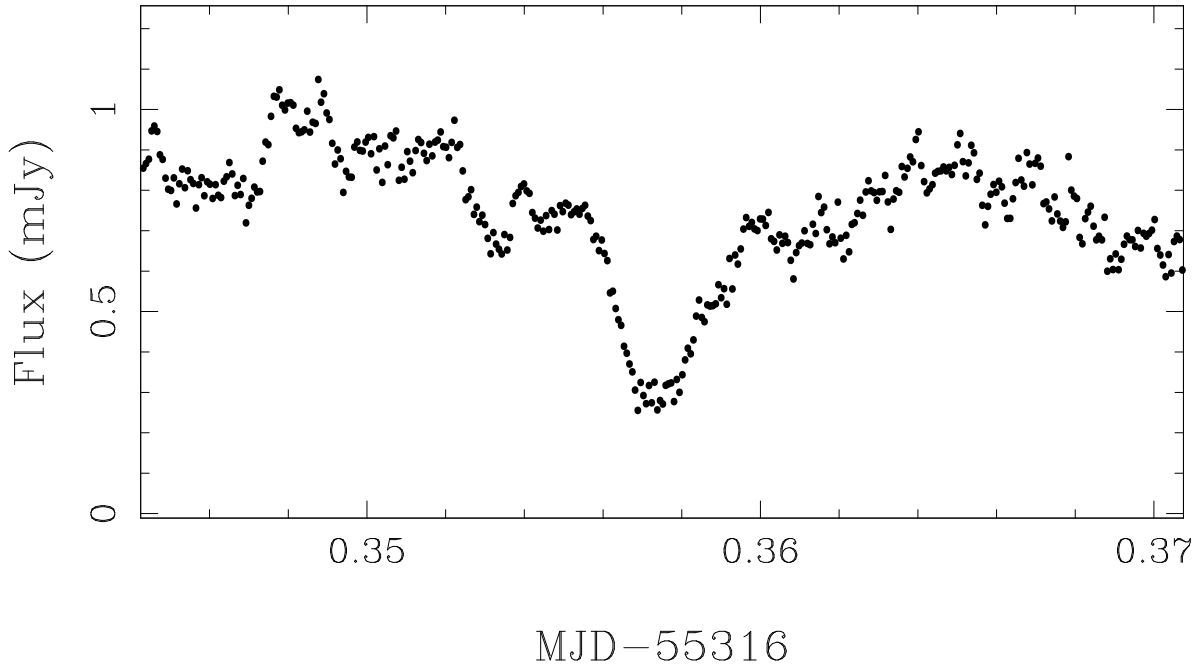

Figure 4.45: 1RXS J180834.7+101041 eclipse observed with ULTRACAM in g' on 2010-04-29. The system shows high amplitude flickering, and suspected WD and bright spot eclipses. However, it is unclear if this is an eclipse of the WD or simply the accretion disk.

not the WD. Figure 4.45 shows one of these eclipses. The light curve is a poor candidate for modelling.

CSS080227:112634-100210

This system was discovered via a dwarf nova outburst by CRTS (Drake et al., 2008) and confirmed as a CV from its spectrum (Djorgovski et al., 2008a; Drake et al., 2009). The CSS light curve shows numerous dips due to eclipses, and these were later observed with time-resolved photometry (Woudt & Warner, 2010). An ULTRASPEC eclipse of CSS080227:112634-100210 is shown in Figure 4.46, which displays a separated bright spot ingress and egress, although the egress is less clear due to flickering. Other eclipses show similar features, and with additional observations this system may be suitable for light curve modelling.

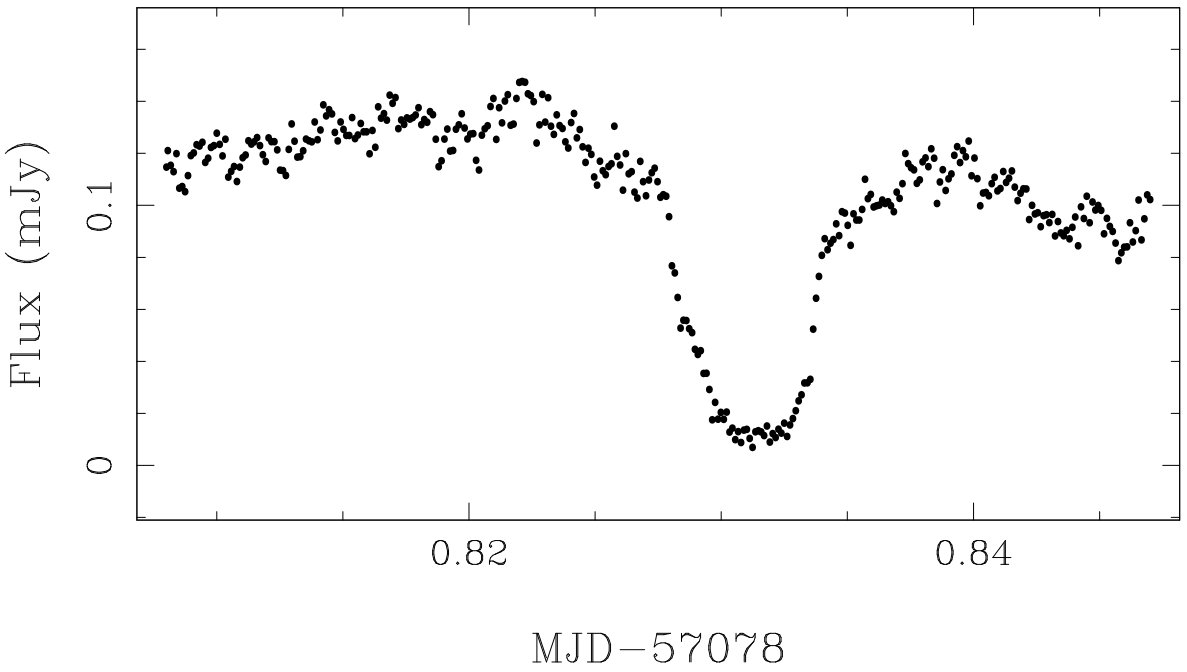

Figure 4.46: ULTRASPEC eclipse of CSS080227:112634-100210 observed on 2015-02-25. The bright spot ingress is visible and separated from the WD, but the egress is difficult to trace due to flickering.

CSS080306:082655-000733

This CV was discovered in dwarf nova outburst by CRTS. Photometry by Woudt et al. (2012) found eclipses and persistent negative superhumps, even in quiescence. This is evident in the CSS light curve, which shows short-term variability with amplitude of ∼1 magnitude, as well as eclipse-like dips and several outbursts. We have observed 8 eclipses. The eclipse structure is suitable for light curve modelling, with clear bright spot features, as shown in Figure 4.47.

CSS080623:140454-102702

This system was found by CRTS, with archival data available from all three CRTS surveys (CSS, MLS and SSS). The light curves show occasional outbursts from a $V = 19.5$ quiescent state, along with eclipse-like dips. Time-resolved photometry revealed eclipses (Woudt et al., 2012) and spectroscopy found both the WD and the secondary star to be

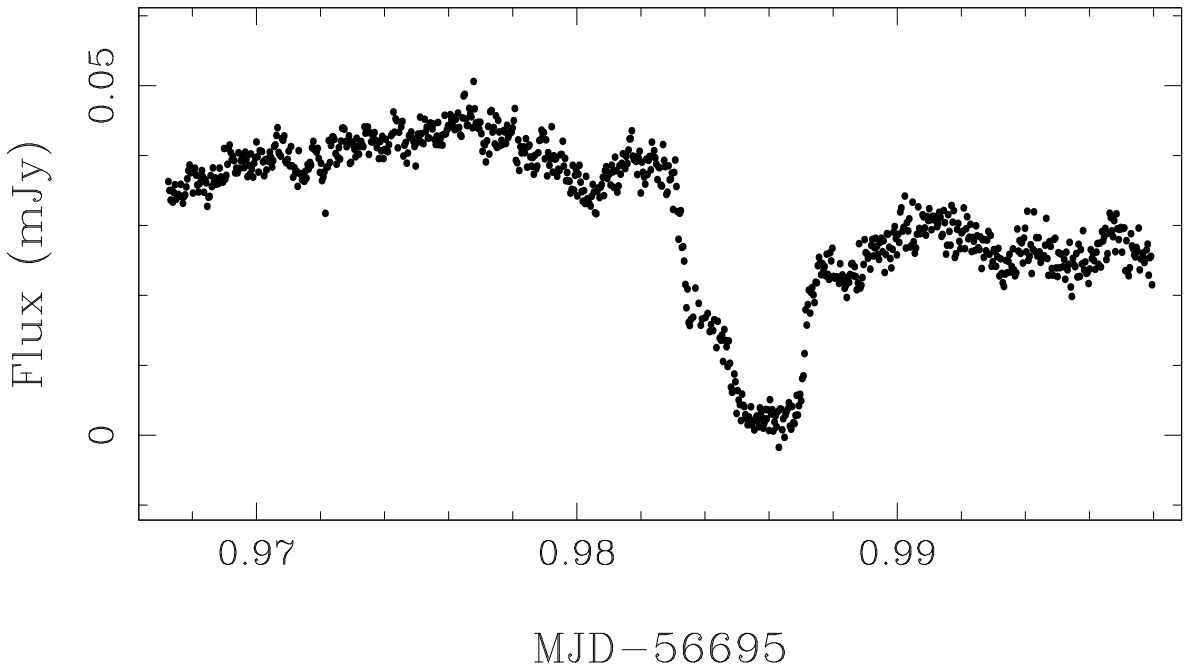

Figure 4.47: CSS080306:082655-000733 eclipse observed with SALTICAM on 2014-02-07. The bright spot eclipse is clearly visible.

present in the spectrum (Breedt et al., 2014). We observed 10 eclipses with ULTRACAM, which all show classic CV eclipse structure that should be suitable for modelling. An example light curve is shown in Figure 4.48.

CSS081220:011614+092216

Discovered by CRTS, this CV shows several short outbursts in its CSS light curve, as well as faint measurements signalling eclipses. Time-resolved photometry confirmed deep eclipses (Coppejans et al., 2014) and found an orbital period of around 1.6 hours. We have observed a total of 15 eclipses, most of which find either the bright spot ingress blended with that of the WD, or the bright spot egress lost in flickering. Occasionally the system was found in a low-flickering state, with the bright spot features visible. These eclipses may be sufficiently clear for light curve modelling. An example is shown in Figure 4.49.

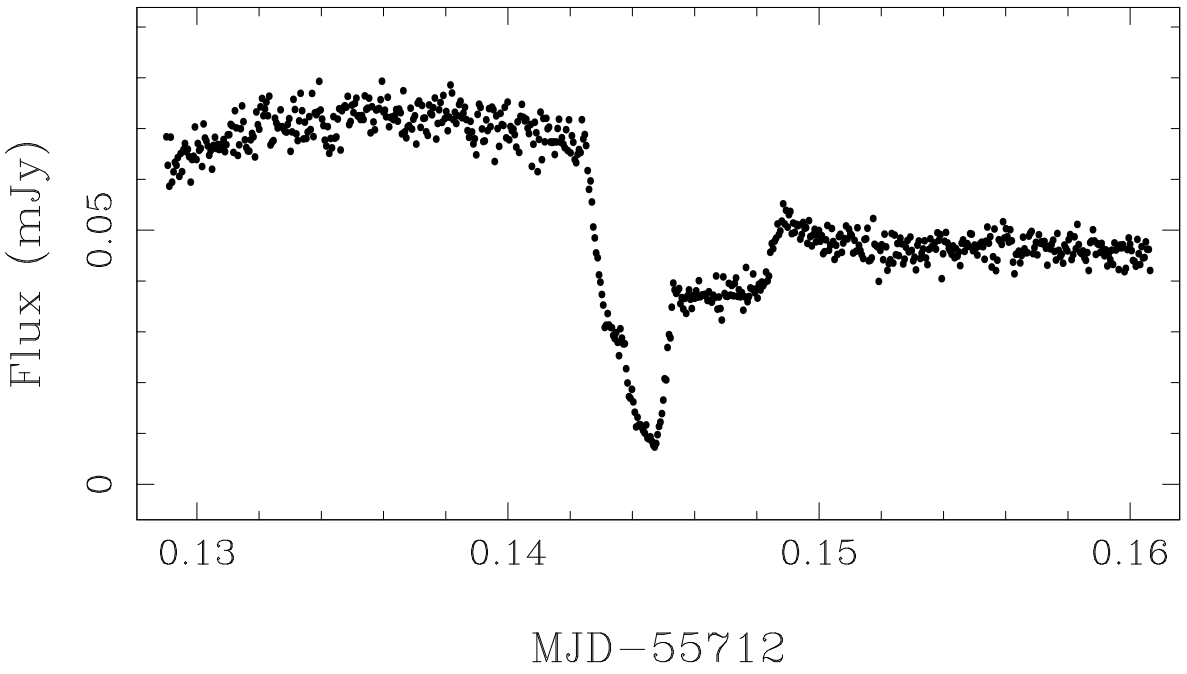

Figure 4.48: Eclipse light curve of CSS080623:140454-102702 observed in g' with ULTRACAM on 2011-05-30. The WD and bright spot features are easily resolved.

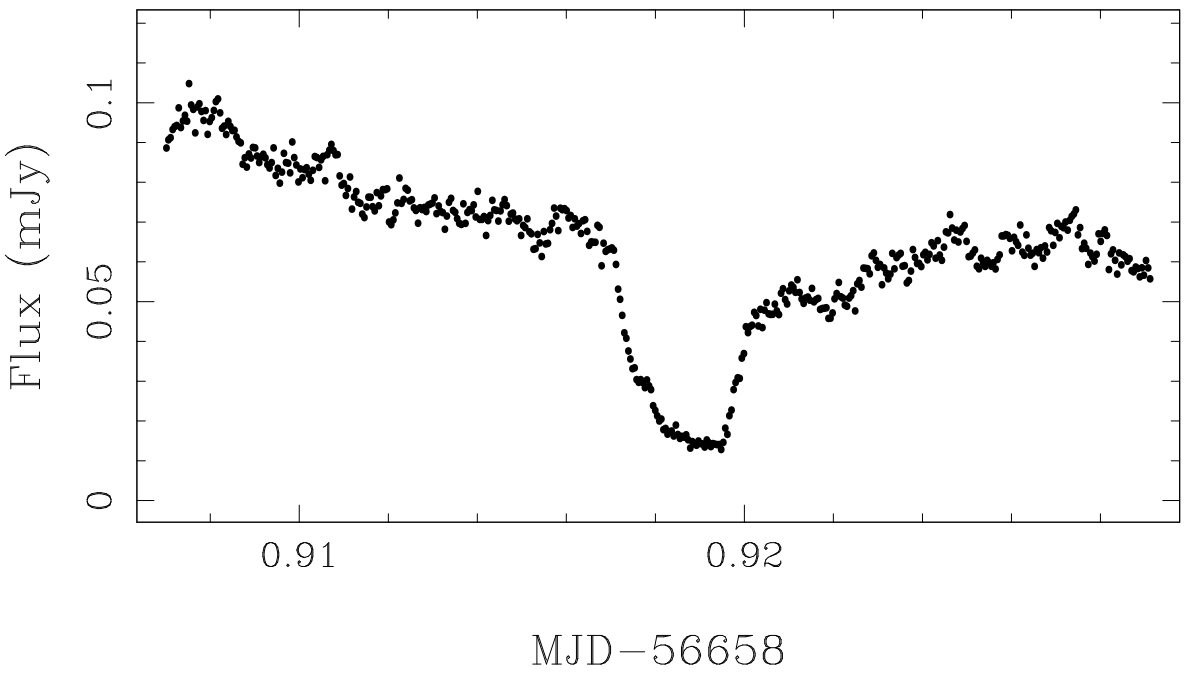

Figure 4.49: CSS081220:011614+092216 eclipse observed in g' with ULTRACAM on 2014-01-01. This system usually suffers from strong flickering, and often the bright spot features are not visible. However in this instance the ingress is clearly separated from that of the WD, and the egress is apparent despite the flickering.

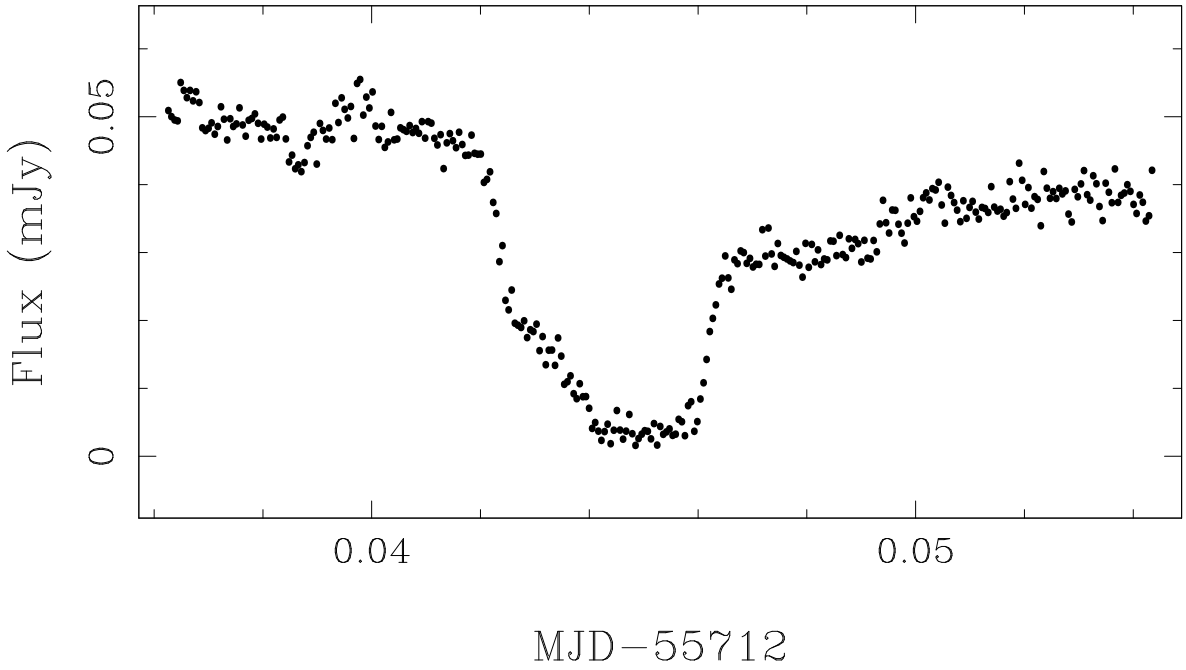

Figure 4.50: Eclipse of CSS090102:132536+210037 observed in g' with ULTRACAM on 2011-05-30. The WD and bright spot eclipses are clearly visible.

CSS090102:132536+210037

This system was discovered via outbursts in CRTS, and shows numerous eclipse-like dips in its CSS light curve. Also known as SDSS J132536.05+210036.7, it was flagged as a dwarf nova in retrospective data mining of SDSS objects by Wils et al. (2010) because of its strong blue colour and very odd g'-r' colour, which must have been affected by an eclipse. Southworth et al. (2015) studied the system in detail and confirmed its eclipsing nature. We have observed 8 eclipses with ULTRACAM, most of which show resolved WD and bright spot features, ideal for light curve modelling. An example light curve is shown in Figure 4.50.

CSS090419:162620-125557

This system was discovered via dwarf nova outbursts in CRTS. It is so faint (r'=20.4) that in quiescence it is not detected by CSS at all. Follow-up photometry identified eclipses (Woudt et al., 2012), and we have observed 7 eclipses at high speed. Figure 4.51

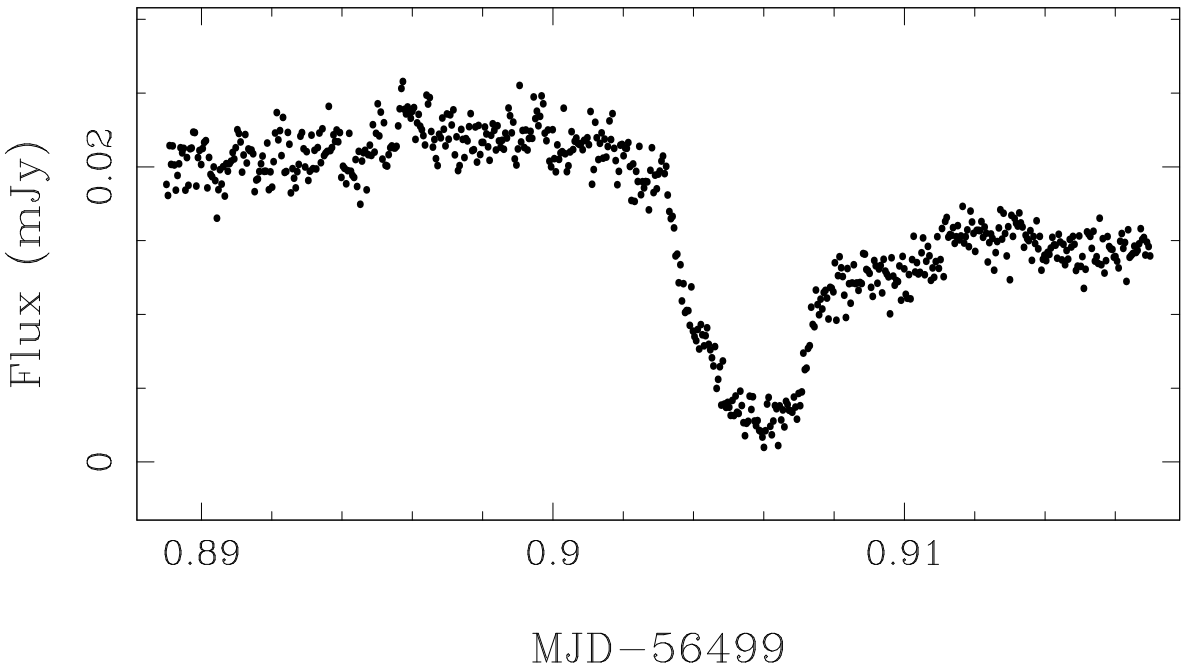

Figure 4.51: CSS090419:162620-125557 g'-band eclipse observed with ULTRACAM on 2013-07-26. Because this system is faint, the data are quite noisy, but the eclipse structure is visible and the bright spot is separated from the WD.

shows an example eclipse light curve, which displays clear WD and bright spot eclipses. This system should be suitable for light curve modelling.

CSS090622:215636+193242

This CV was discovered by CRTS via dwarf nova outbursts from $V = 19$ to $V = 16$-17. The long-term light curve also shows many dips suggestive of eclipses, and the eclipsing nature was confirmed by Drake et al. (2014). We observed a total of 9 eclipses. We show an example eclipse light curve in Figure 4.52, with clear WD and bright spot features. This system should be suitable for light curve modelling.

CSS100218:043829+004016

This object showed dwarf nova outbursts in CSS, and was subsequently found to be eclipsing by Coppejans et al. (2014), who derived a period of 1.5 hours. We observed

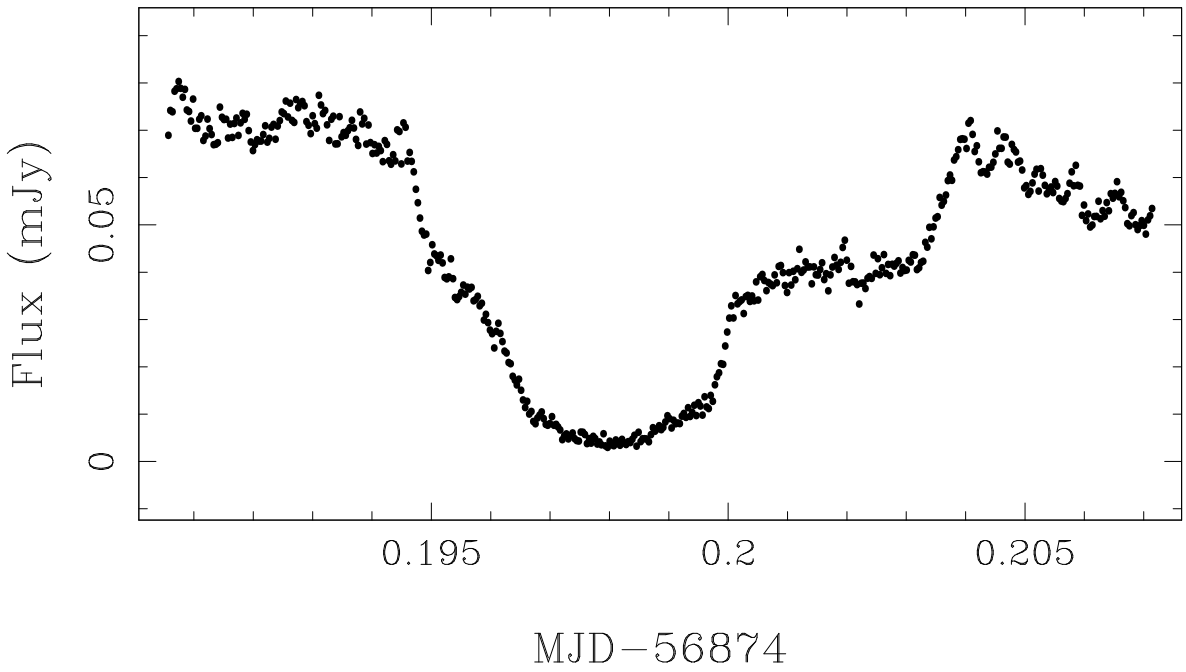

Figure 4.52: Eclipse light curve of CSS090622:215636+193242 observed in g' with ULTRACAM on 2014-08-04. The WD and bright spot eclipses are clear.

three eclipses over four nights. The system is faint, and our best eclipse observations are shown in Figure 4.53. Strong flickering and a weak bright spot make this system a poor candidate for light curve modelling.

CSS110513:210846-035031

This CV was discovered by CRTS via a dwarf nova outburst from $V = 18$ to $V = 15$. It shows faint measurements in the long-term light curve indicating eclipses, and Coppejans et al. (2014) confirmed eclipses with an orbital period of 3.8 hours. We observed 4 eclipses with *pt5m* to derive a reliable ephemeris, and then 4 eclipses with ULTRACAM to examine the eclipse structure. Unfortunately most observations found the system in outburst, showing an eclipse dominated by the disk. However, on 2013-07-15 a high quality quiescent eclipse was observed (see Figure 4.54). This eclipse is suitable for light curve modelling.

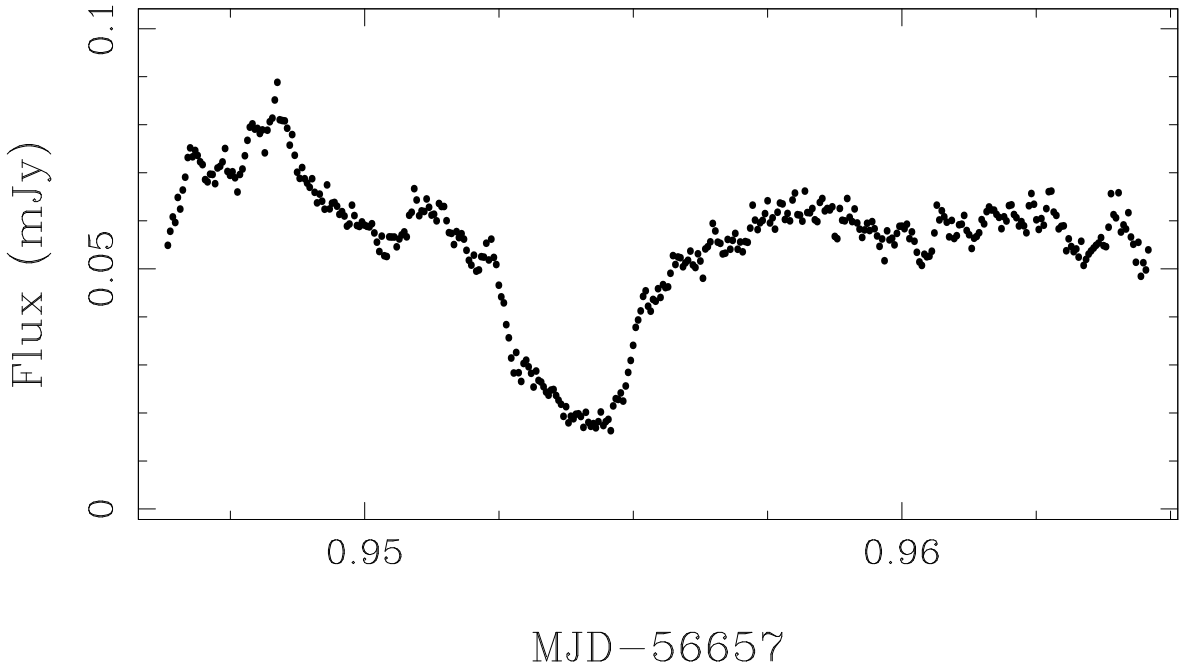

Figure 4.53: ULTRACAM g'-band observation of the eclipse of CSS100218:043829+004016 on 2013-12-31. The bright spot ingress is visible, but the egress is lost in flickering.

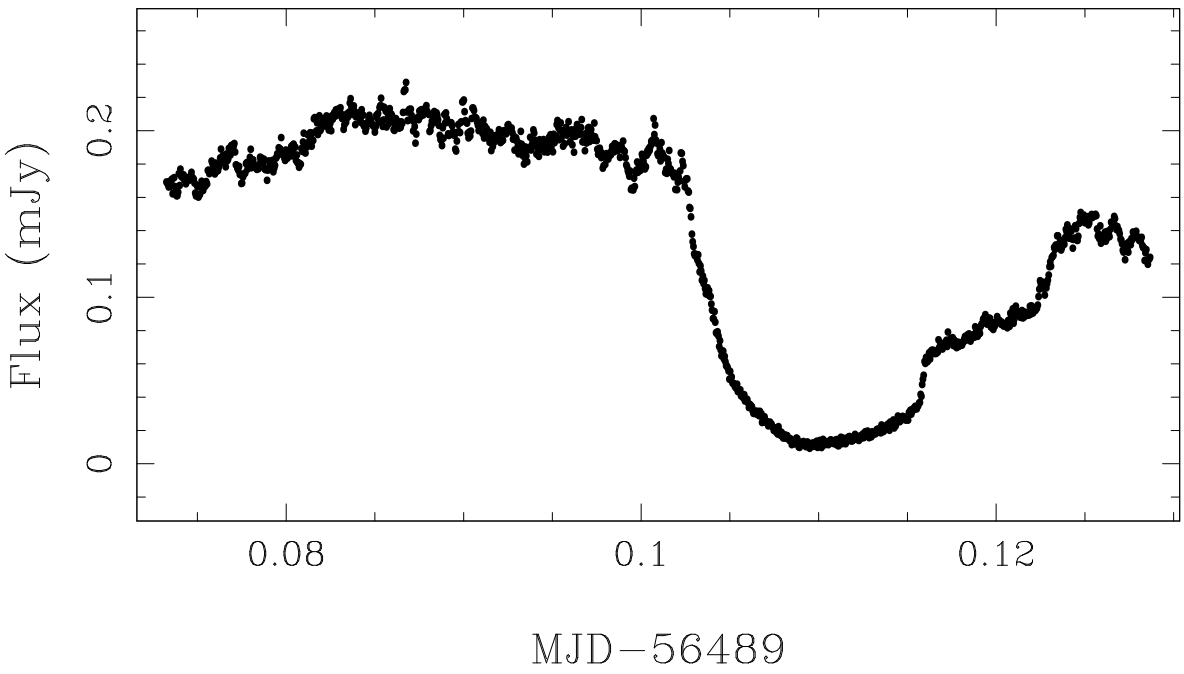

Figure 4.54: Eclipse light curve of CSS110513:210846-035031 observed with ULTRACAM on 2013-07-15 in g'. The bright spot and WD components are visibly separated.

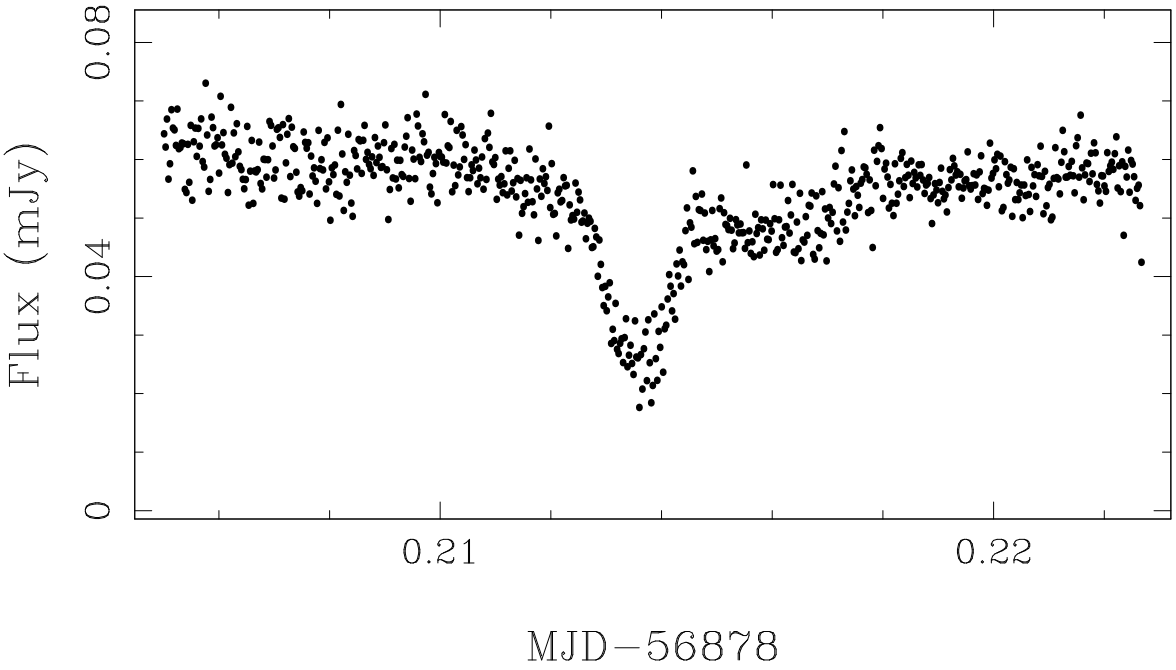

Figure 4.55: CSS111019:233313-155744 g'-band eclipse observed with ULTRACAM on 2014-08-08. The bright spot eclipse appears to occur after the WD eclipse, suggesting that this may be an AM CVn system.

CSS111019:233313-155744

This object was discovered during a super-outburst by CRTS, and seen to have a short period of around 62 minutes (Woudt & Warner, 2011). We observed 6 eclipses with ULTRACAM. The eclipses are shallow, and the bright spot eclipse appears to arrive very late - almost after the WD has come out of eclipse (see Figure 4.55). This behaviour is similar to that seen in the AM CVn system SDSS J092638.71+362402.4 (Marsh et al., 2007; Copperwheat et al., 2011), where the secondary star is so bright and small in size (e.g. a WD), that the eclipses are shallow and the bright spot ingress occurs after the primary WD egress. This would also explain the short period. CSS111019:233313-155744 would be an interesting object for further study.

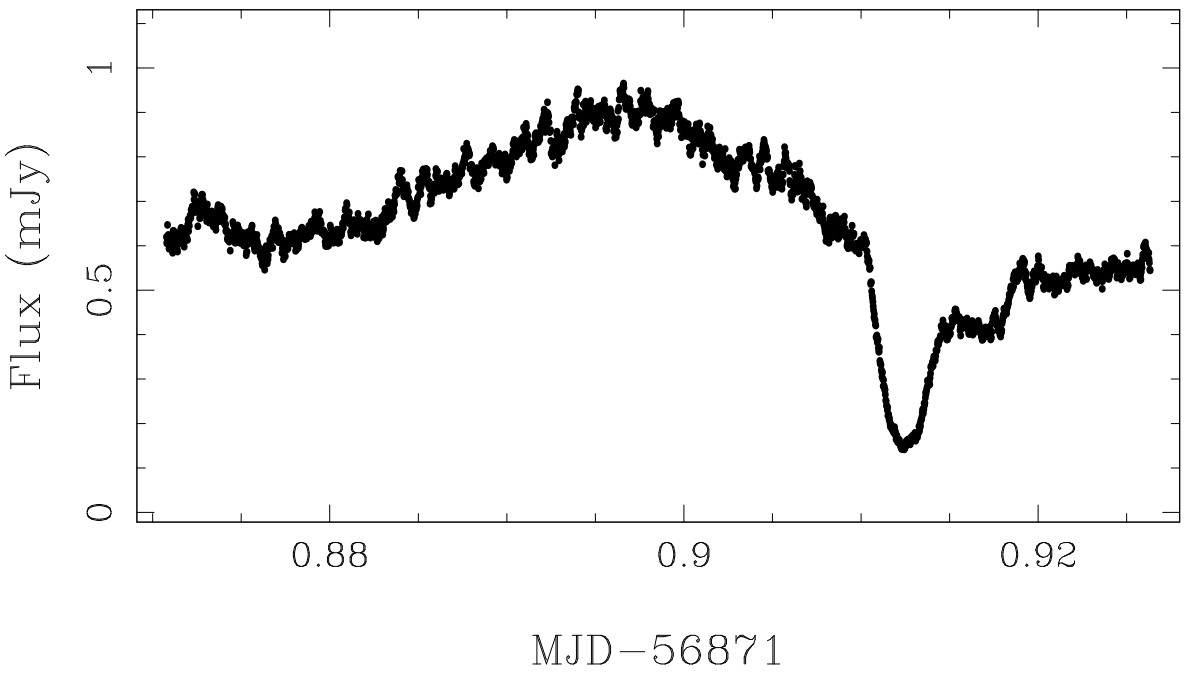

Figure 4.56: Eclipse of CzeV404 Her observed with ULTRACAM in the g' filter on 2014-08-02. The separation between the ingresses of the WD and bright spot are not clear.

CzeV404 Her

This CV was discovered in outburst by Cagas & Cagas (2014), and further studied by Bąkowska et al. (2014), who observed multiple super-outbursts and estimated an orbital period within the period gap. We have observed 7 additional eclipses with *pt5m*, ULTRASPEC and ULTRACAM. Figure 4.56 shows an eclipse observed in quiescence. The bright spot and WD ingresses appear to be blended, although there may be a hint of separation. In this state it will be very challenging to model the light curve of this system.

GALEX J003535.7+462353

This object showed long-term variability in SuperWASP data (Butters et al., 2010), prompting time-resolved photometry which discovered dwarf novae outbursts and eclipses (Wils et al., 2011). We observed 16 eclipses, with most showing little or no sign of bright spot features. Some distinct structure is visible in the ULTRASPEC light curve shown in

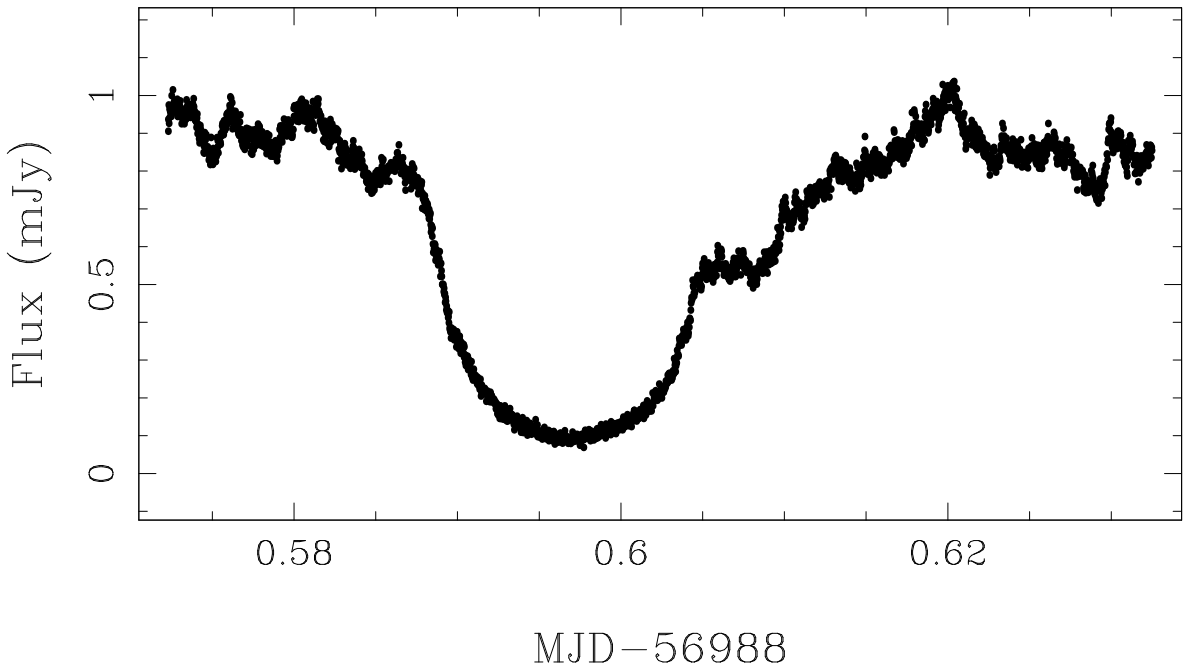

Figure 4.57: ULTRASPEC light curve of GALEX J003535.7+462353 observed on 2014-11-27. The bright spot egress is clearly separated from the WD, but the ingresses are difficult to separate. In other observations the bright spot component is less visible.

Figure 4.57, but more observations are needed before modelling can be conducted.

GY Cnc

A well-studied system, GY Cnc already has spectroscopically derived masses (Thorstensen, 2000) and published ULTRACAM data (Feline et al., 2005). Nevertheless, with additional data we expect that the system parameter measurements could be improved. We observed 9 additional eclipses. There are signs of a varying orbital period for this system, although since the system appears to have a variable disk radius, we cannot be certain we are always measuring the exact mid-eclipse time of the same features. Figure 4.58 shows an example ULTRACAM eclipse with clear bright spot features.

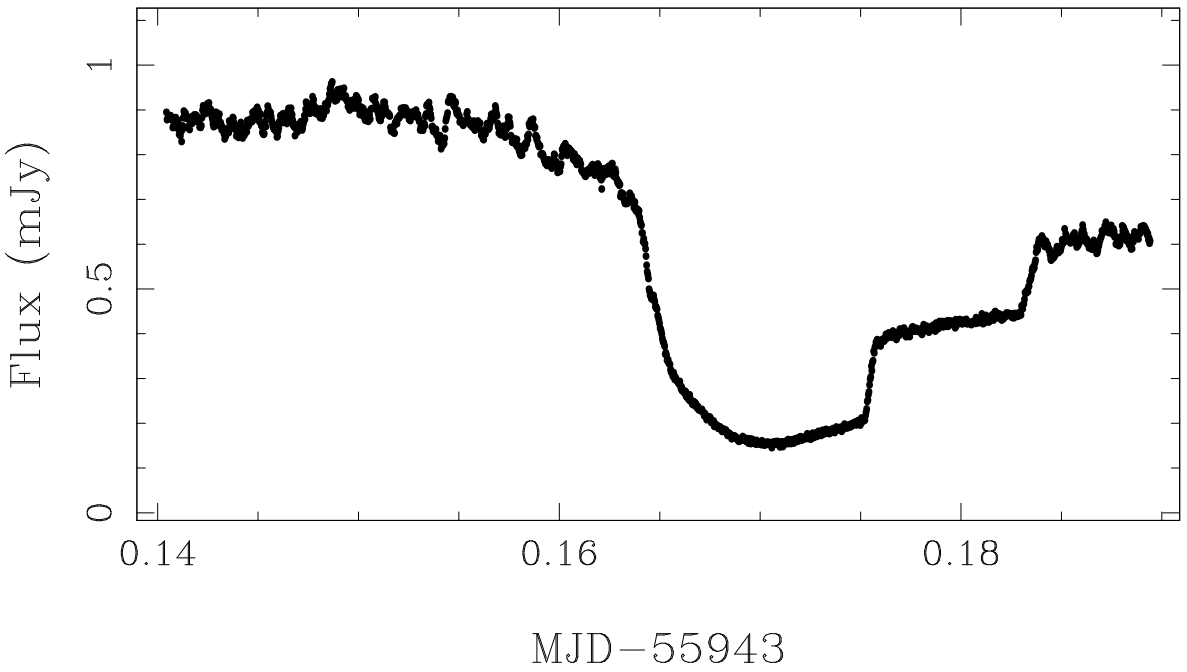

Figure 4.58: ULTRACAM g'-band eclipse of GY Cnc observed on 2012-01-16. The ingress and egress of the WD and bright spot appear separately in this instance, though other observations of this system often find these components blended.

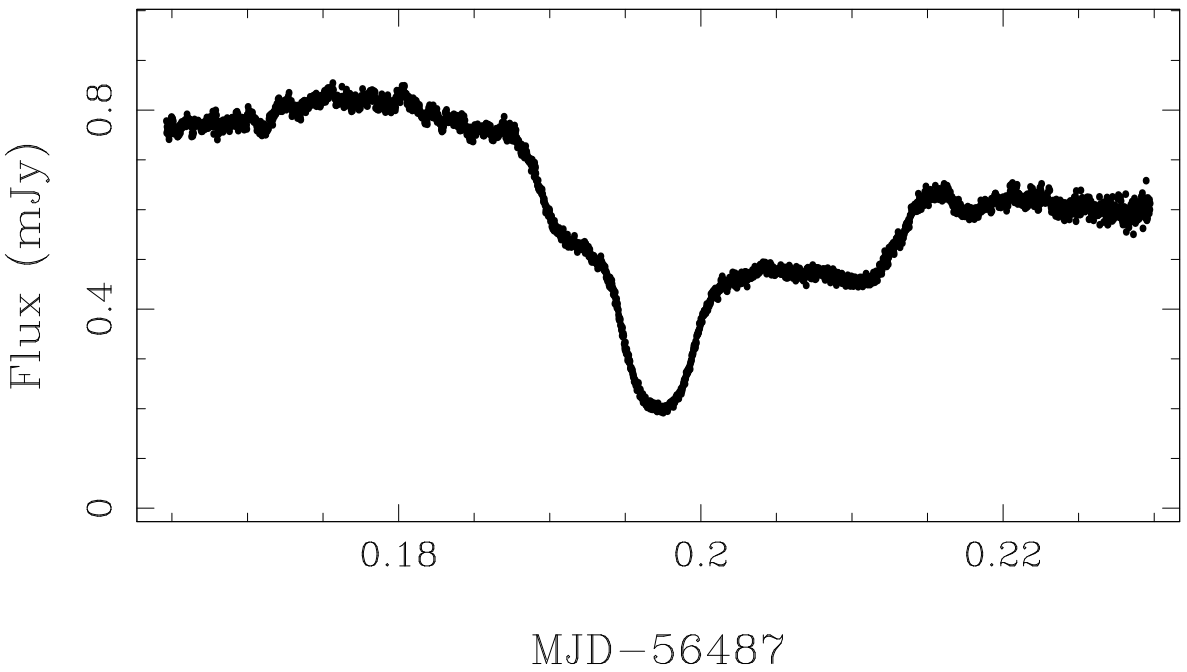

Figure 4.59: Eclipse of HS 2325+8205 observed with ULTRACAM in the g'-band on 2013-07-13. The WD and bright spot components are clearly visible.

HS 2325+8205

This object has already been studied extensively (Pyrzas et al., 2012), and was found to be a candidate eclipsing Z Cam type CV, with a long orbital period. Here we present an updated ephemeris from three eclipses observed with ULTRACAM, two of which were during outburst (see Table 4.2). The single quiescent eclipse observation is shown in Figure 4.59, and appears to be well suited for modelling.

HT Cas

Another well-studied CV (Horne, Wood & Stiening, 1991), HT Cas has been observed by ULTRACAM before (Feline et al., 2005). Here we present additional observations and an updated ephemeris based on 8 eclipses and the 4 previously published eclipse times (see Table 4.2). There are signs of an evolving orbital period for this system according to the linear fit to the mid-eclipse times. This is discussed in detail in Bours et al. (2016). Normally HT Cas shows lots of flickering, and eclipses with no visible bright spot. However, for a short time in 2014, the bright spot eclipse was visible, but with the egress masked by flickering, meaning the light curve is still not suitable for modelling. An example of this is shown in Figure 4.60.

IY UMa

Known to be eclipsing for many years (Uemura et al., 2000), IY UMa has been studied extensively (Steeghs et al., 2003). System parameter determinations already exist, but have large uncertainties. We obtained 7 additional eclipse observations which should allow for improved light curve modelling and refinement of the parameters. An example of an ULTRASPEC eclipse light curve is shown in Figure 4.61.

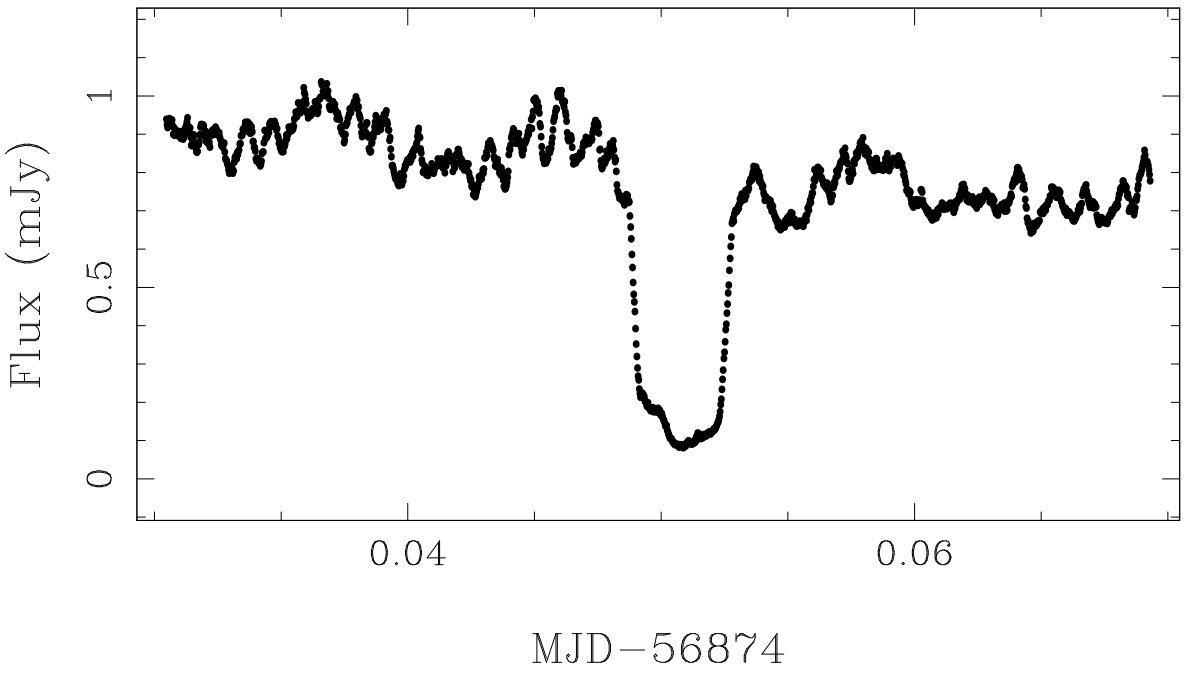

Figure 4.60: HT Cas eclipse observed with ULTRACAM in the g'-band on 2014-08-04. A notable bright spot feature was seen during these observations which is normally not present, although the egress is lost in the flickering.

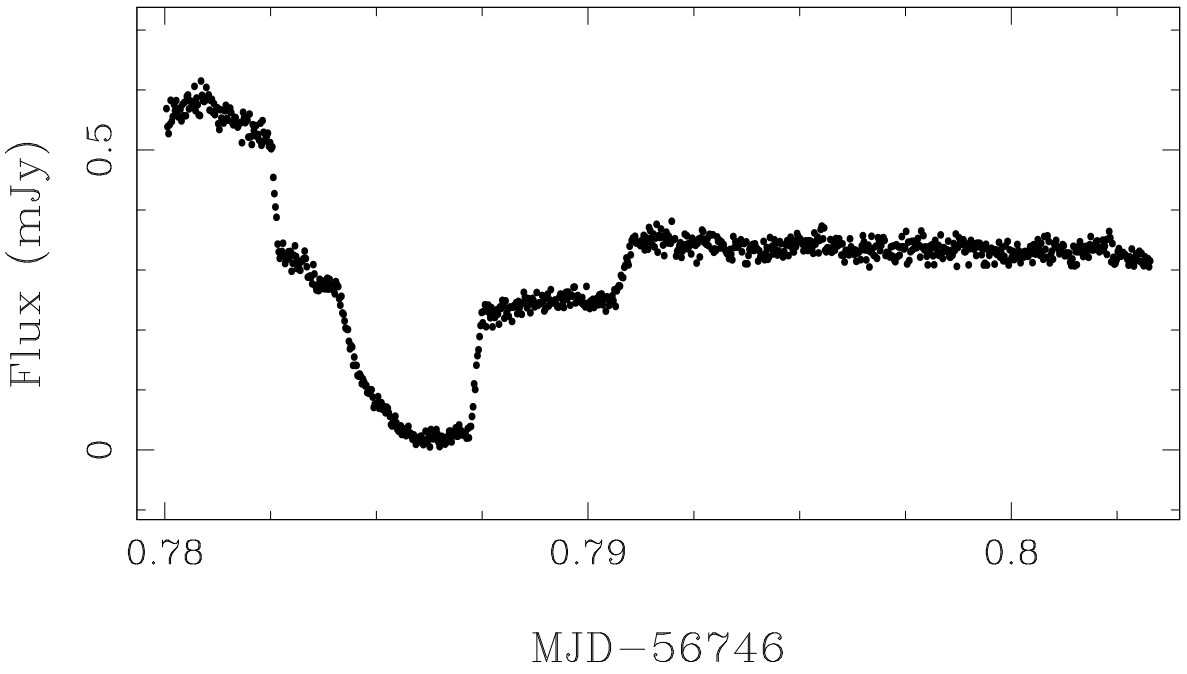

Figure 4.61: Eclipse of IY UMa observed with ULTRASPEC on 2014-03-30. The light curve shows the bright spot and WD components clearly.

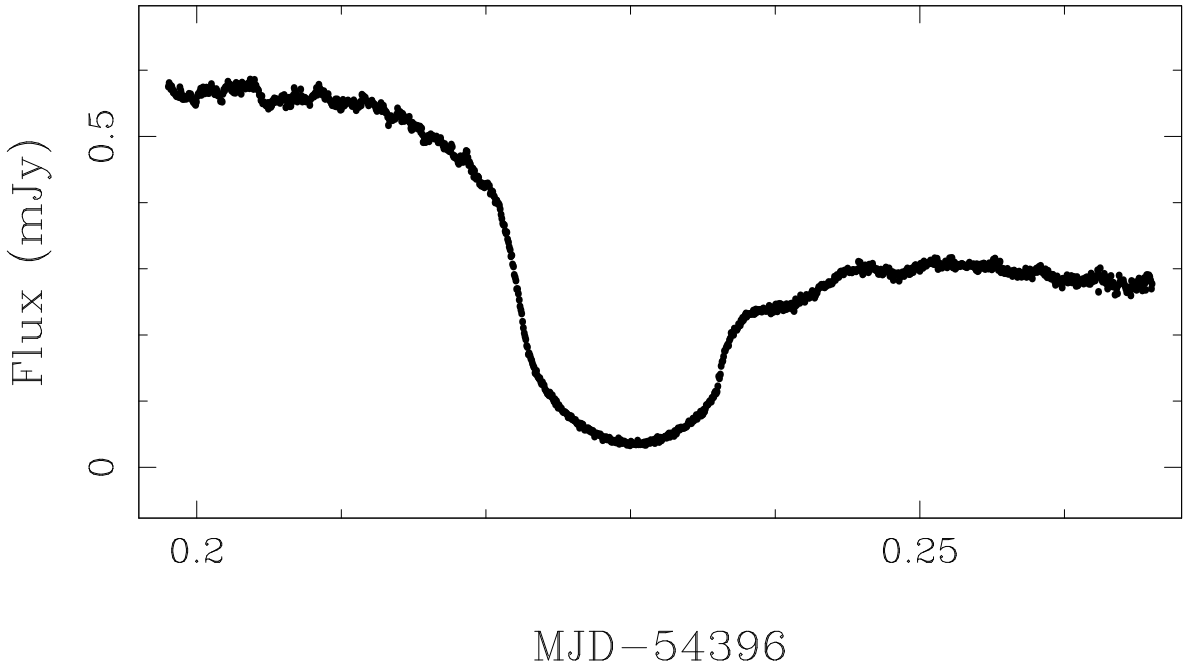

Figure 4.62: Eclipse of SDSS J040714.78-064425.1/LT Eri observed with ULTRACAM in the g'-band on 2007-10-22, with the system in a mildly high state. Whilst the bright spot egress is clearly separated from the WD, the ingresses are blended.

SDSS J040714.78-064425.1

Also known as LT Eri, this system was discovered as an eclipsing CV in SDSS follow up (Szkody et al., 2003). Further studies confirmed a period of around 4 hours (Ak et al., 2005), placing it well above the period gap. The CSS light curve shows many rapid ∼2 magnitude outbursts.

We have observed SDSS J040714.78-064425.1 with ULTRACAM, ULTRASPEC and *pt5m* on multiple occasions over the past 8 years (see Table 4.2). The system is often in outburst, but even in quiescence it usually shows blended WD and bright spot ingresses. An example is shown in Figure 4.62, when the system was in a slightly higher state than in quiescence. Unless the disk radius changes significantly, separating the times of ingress for the WD and bright spot, this system will not be suitable for light curve modelling.

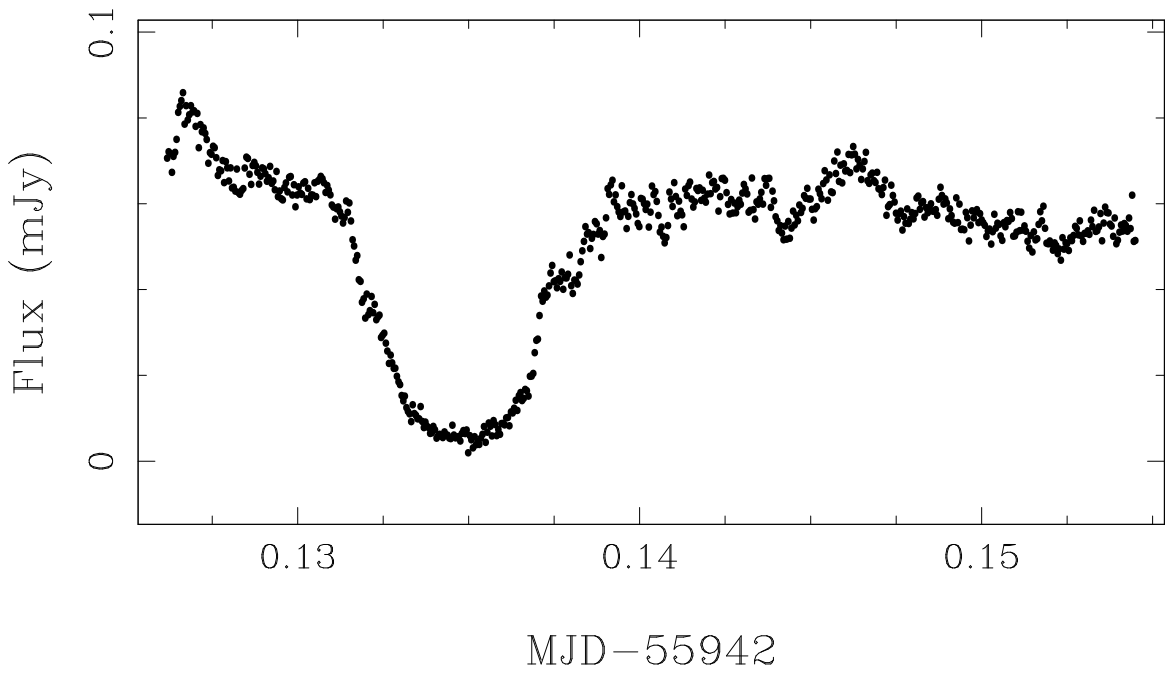

Figure 4.63: Eclipse of SDSS J075059.97+141150.1 in g' observed with ULTRACAM on 2012-01-15. The bright spot and WD components are clearly separated.

SDSS J075059.97+141150.1

Flagged as a CV in SDSS spectra (Szkody et al., 2007), SDSS J075059.97+141150.1 showed frequent outbursts in its CSS light curve. Southworth et al. (2010) presented time-resolved photometry and were able to measure the system parameters through light curve modelling. However, these measurements suffered from high uncertainties due to the low cadence of the observations.

We observed 10 eclipses at high speed, with Figure 4.63 showing our clearest example of the eclipse structure. The bright spot is visible and our additional observations should allow for refinement of the system parameters found in Southworth et al. (2010).

SDSS J090103.93+480911.1

Another CV found in SDSS spectra (Szkody et al., 2003), SDSS J090103.93+480911.1 has been studied photometrically already (Dillon et al., 2008; Shears et al., 2012; Kato et al., 2013). We observed 10 eclipses with ULTRACAM over 6 years, which showed an

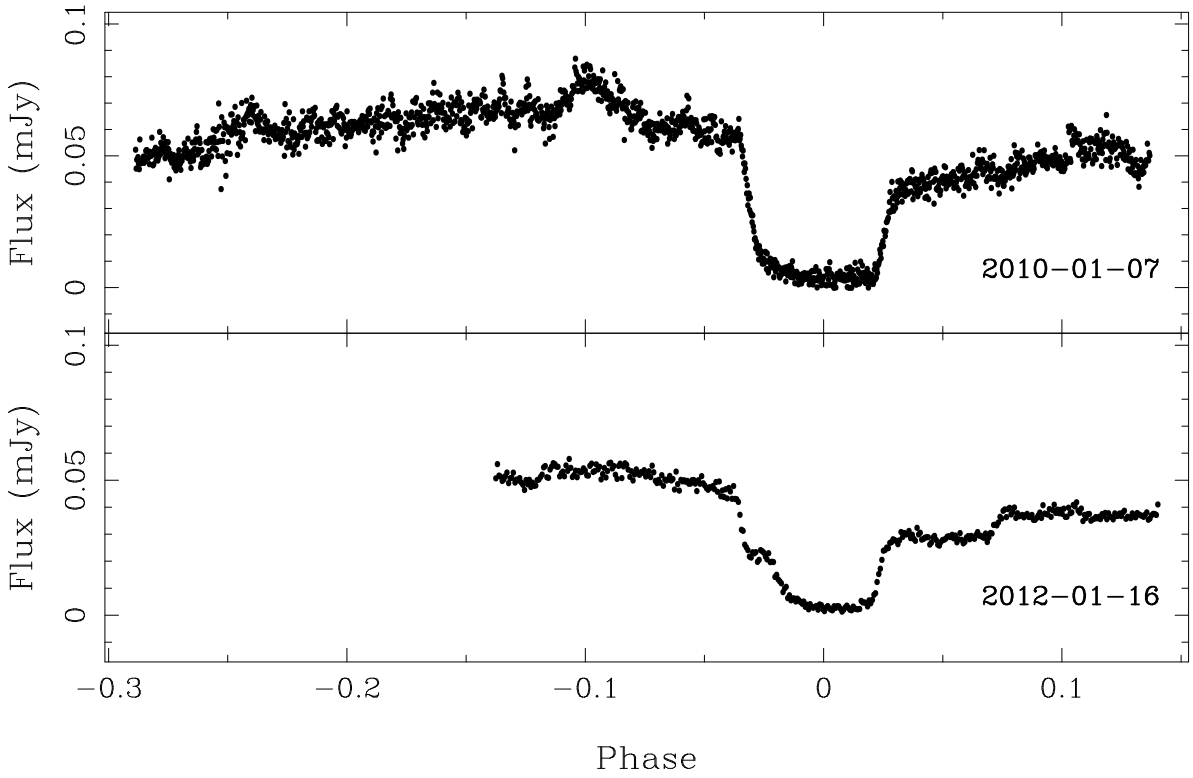

Figure 4.64: Phase-folded eclipses of SDSS J090103.93+480911.1 observed in g' with ULTRACAM on 2010-01-07 (top) and 2012-01-16 (bottom). The bright spot seems to appear (bottom) and disappear (top) over time.

evolution of bright spot features. At first the bright spot was clearly visible, then it disappeared and the flickering intensified. Later the flickering decreased and the bright spot returned. Two phase-folded light curves are shown in Figure 4.64 to demonstrate this behaviour. These later observations should be suitable for light curve modelling.

SDSS J090403.49+035501.2

This object was identified as a CV from its spectrum (Szkody et al., 2004), and has already been studied photometrically (Woudt et al., 2012). It shows no outbursts in the CSS archival light curve. We observed a single eclipse of this system confirming the grazing nature of its eclipses (see Figure 4.65). The system is not suitable for light curve modelling and we cannot improve on the precision of the ephemeris in Woudt et al. (2012).

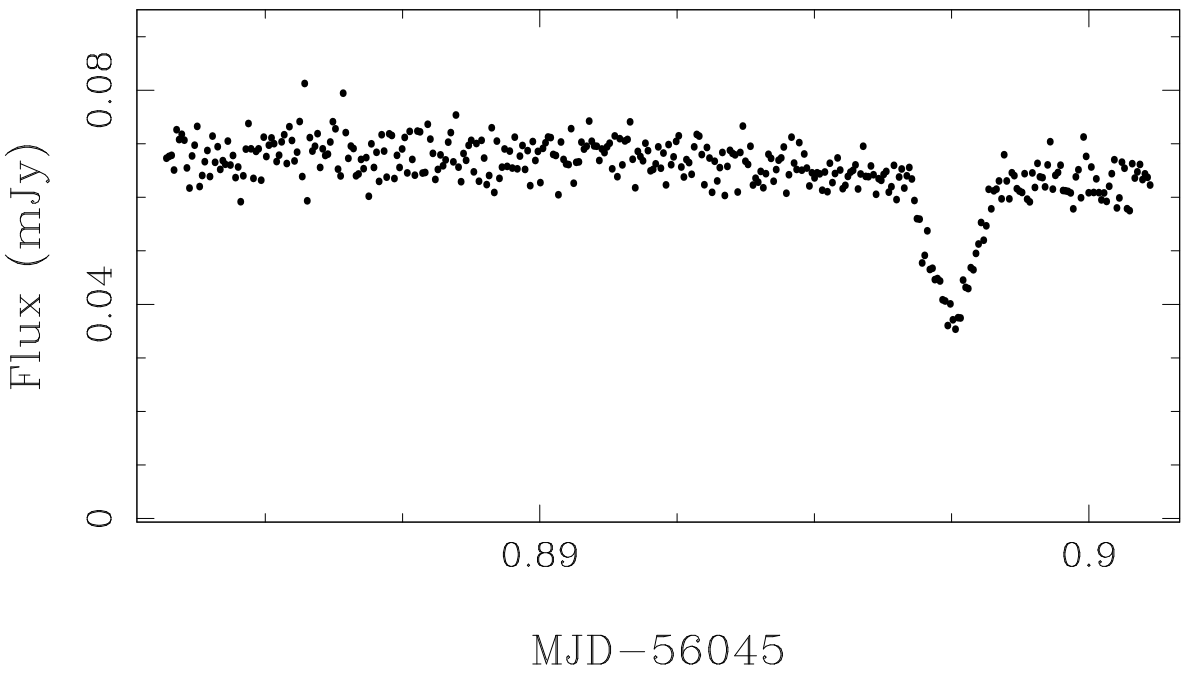

Figure 4.65: Shallow eclipse of SDSS J090403.49+035501.2 observed in g' with ULTRACAM on 2012-04-28. The large orbital hump seen in other observations (Woudt et al., 2012) is not present.

SDSS J092009.54+004245.0

This object was discovered as an eclipsing CV in the spectroscopic and photometric follow-up of the SDSS (Szkody et al., 2003). Further studies revealed a period of 3.5 hours (Gänsicke et al., 2009; Zengin Çamurdan, İbanoğlu & Çamurdan, 2010), and the CSS light curve shows numerous eclipse dips, as well as occasional measurements which are ∼1 magnitude brighter than usual. It has since been suggested to be a SW Sex system (Schmidtobreick, Rodríguez-Gil & Gänsicke, 2012), a nova-like with a strong, flared disk.

We observed SDSS J092009.54+004245.0 on three occasions. We found that the eclipses have some structure, but without any clear sign of a WD component (see Figure 4.66). This supports the SW Sex classification as a nova-like variable, and confirms that this system is not suitable for light curve modelling.

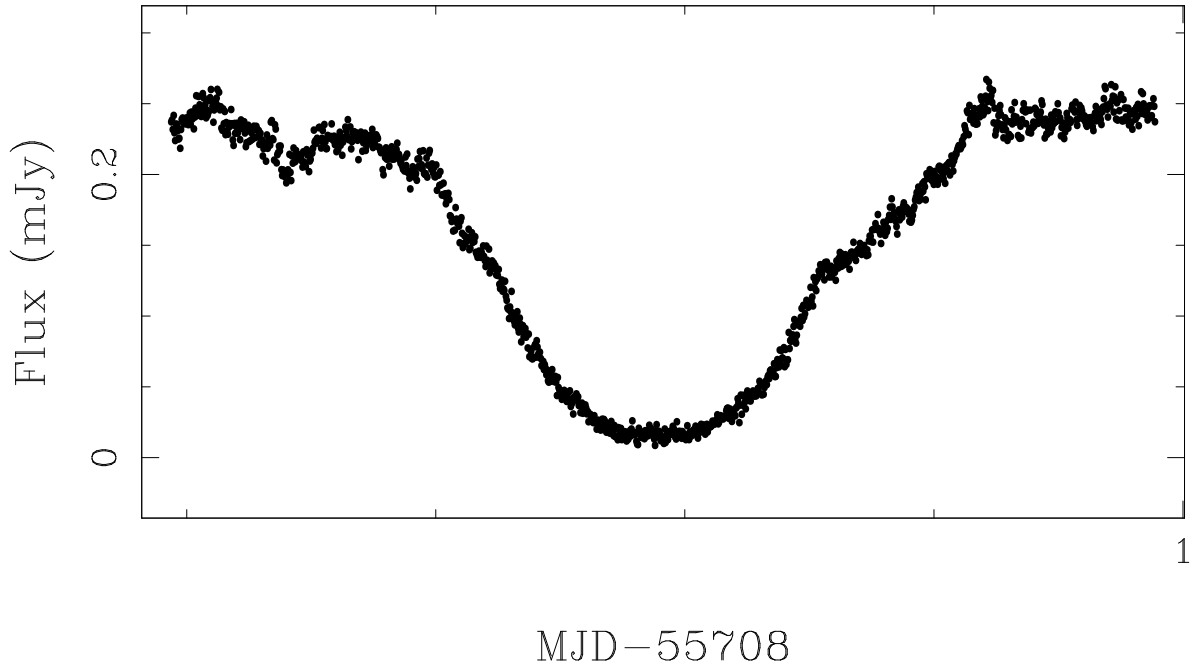

Figure 4.66: ULTRACAM g'-band eclipse of SDSS J092009.54+004245.0 observed on 2011-05-27. The eclipse is disk dominated, supporting the SW Sex classification.

SDSS J092444.48+080150.9

Also known as HU Leo, this object was discovered in SDSS spectra (Szkody et al., 2005). Numerous eclipses have been observed (Southworth et al., 2010, 2015), and we present one additional eclipse (Figure 4.67). The system shows no signs of an accretion disk, and has the possible characteristics of a polar light curve, showing a bright post-eclipse hump due to the projection effect of the accretion stream. The light curve is not suitable for modelling. We are unable to refine the ephemeris further.

SDSS J093537.46+161950.8

This system was found in SDSS spectroscopic follow-up (Szkody et al., 2009), with a spectrum that could have been classified as a polar or an old nova. It has also been studied photometrically (Southworth et al., 2015). We observed a single eclipse with ULTRACAM, finding a typical polar light curve which is unsuitable for modelling (see Figure 4.68).

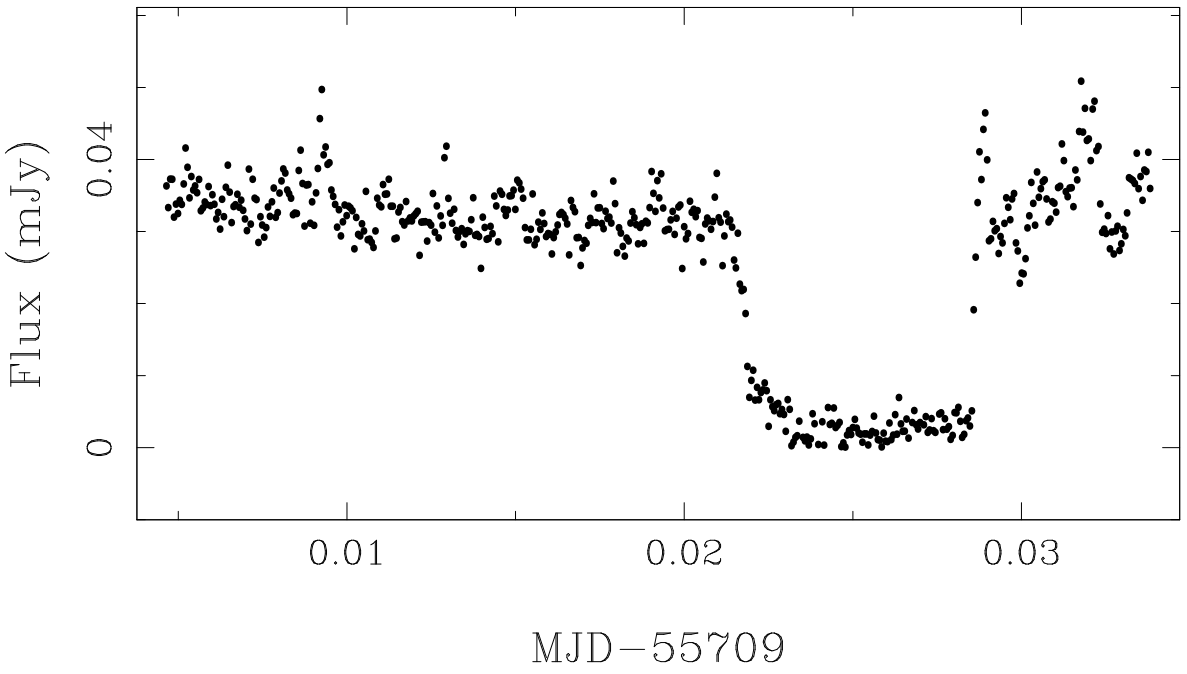

Figure 4.67: Eclipse of SDSS J092444.48+080150.9/HU Leo observed in g' with ULTRACAM on 2011-05-27. The eclipse shows no sign of an accretion disk, and the brighter flux after the eclipse may indicate that this is a polar.

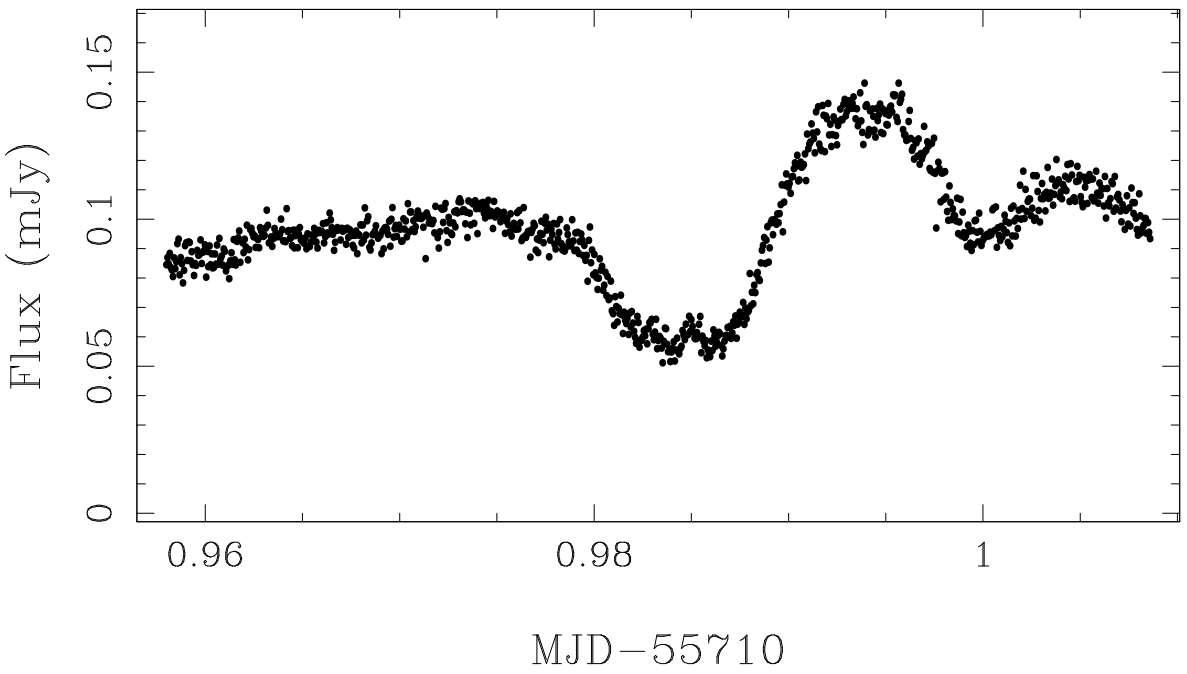

Figure 4.68: Eclipse of SDSS J093537.46+161950.8 in g' observed with ULTRACAM on 2011-05-29. The bright bump after the eclipse is similar to that seen in HU Aqr (Harrop-Allin et al., 1999) and therefore this system is likely to be a polar.

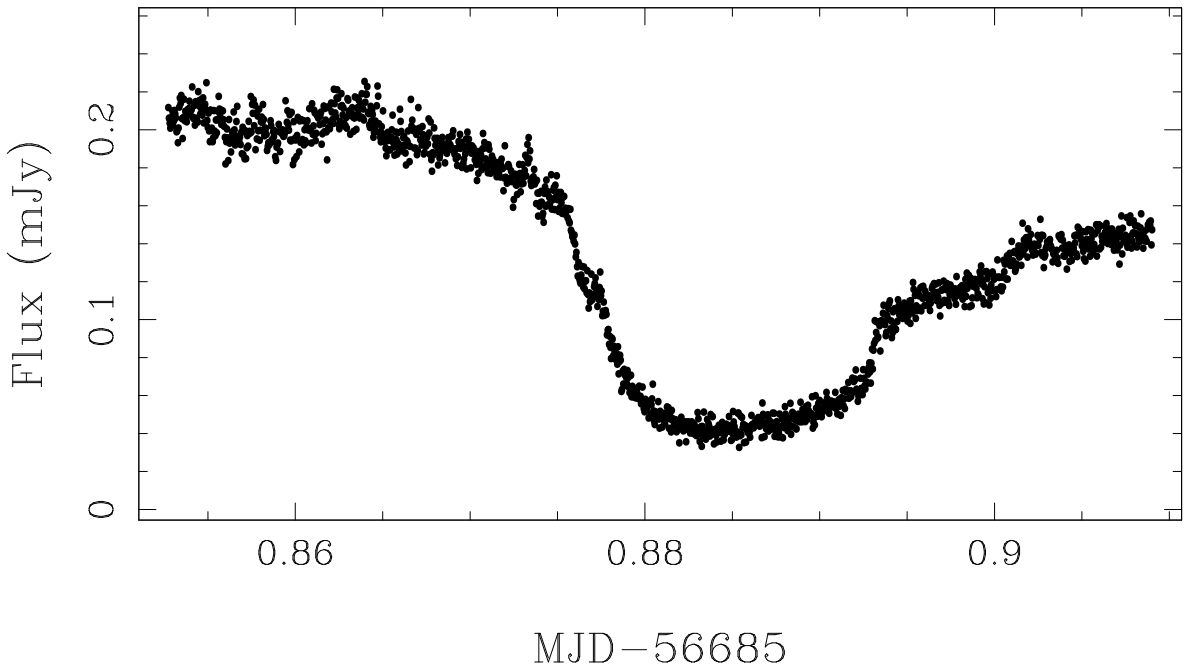

Figure 4.69: Eclipse of SDSS J100658.40+233724.4 observed with ULTRASPEC on 2014-01-28. The bright spot and WD components are clearly separated.

SDSS J100658.40+233724.4

Discovered in SDSS spectra (Szkody et al., 2007), this system has already been studied in detail. Southworth et al. (2009) published some system parameters based on time-resolved photometry and spectroscopy. We continued to observe SDSS J100658.40+233724.4 in an attempt to improve on the previous measurements. Figure 4.69 shows an example eclipse light curve observed with ULTRASPEC, displaying the clear bright spot structure. Other observations suffer from flickering, but there should be enough eclipses that can be modelled together to refine the system parameters.

SDSS J115207.00+404947.8

This system already has some published ULTRACAM observations, with its parameters determined through light curve modelling (Savoury et al., 2011). We later discovered that the ephemeris presented in Savoury et al. (2011) is incorrect. We have since observed 4 more eclipses of this CV, and a new ephemeris is given in Table 4.2. The new eclipse

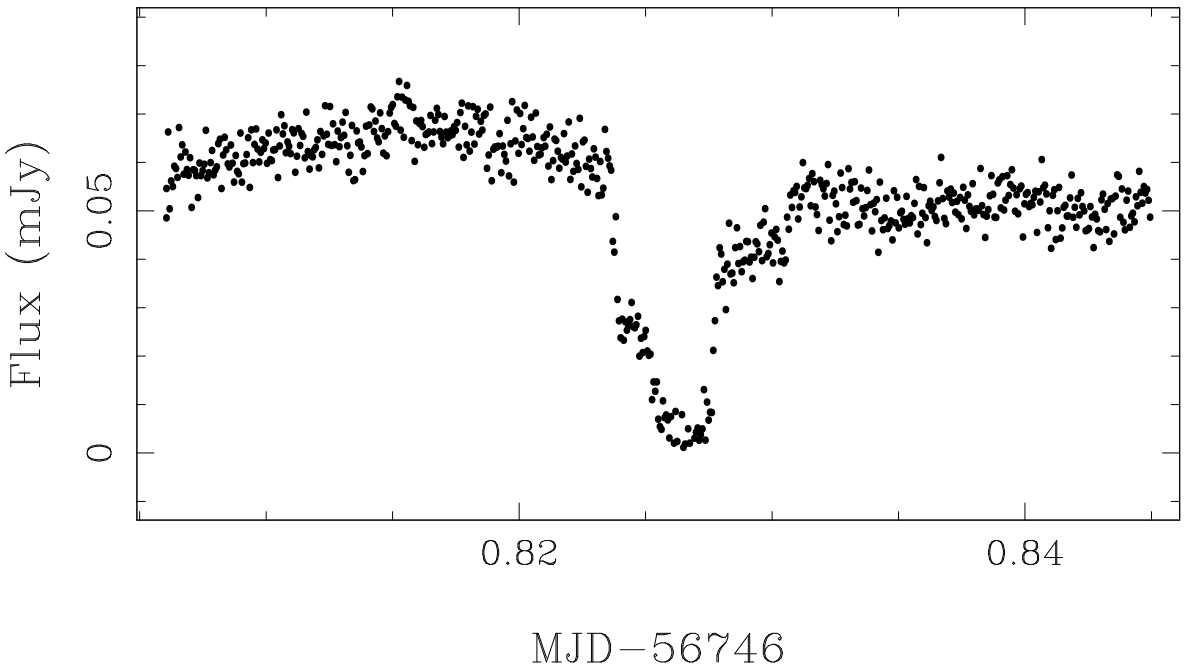

Figure 4.70: ULTRASPEC eclipse of SDSS J115207.00+404947.8 observed on 2014-03-30. The bright spot and WD components are clear and separated.

observations (see Figure 4.70) have stronger bright spot features than previously seen. This should allow for improved measurement of the system parameters.

SDSS J125023.84+665525.4

Also known as OV Dra, this CV was found with a high inclination in spectroscopic SDSS studies (Szkody et al., 2003). The archival CSS light curve shows long-term variability, eclipse-like dips, and several outbursts. Time-resolved photometry confirmed eclipses (Dillon et al., 2008; Kato et al., 2014a), but no high-speed studies have been published so far. We have observed 14 eclipses of SDSS J125023.84+665525.4, and although the bright spot eclipse is visible, unfortunately it is often interrupted by flickering, and further studies will be needed before modelling can be conducted. We show an example eclipse light curve in Figure 4.71.

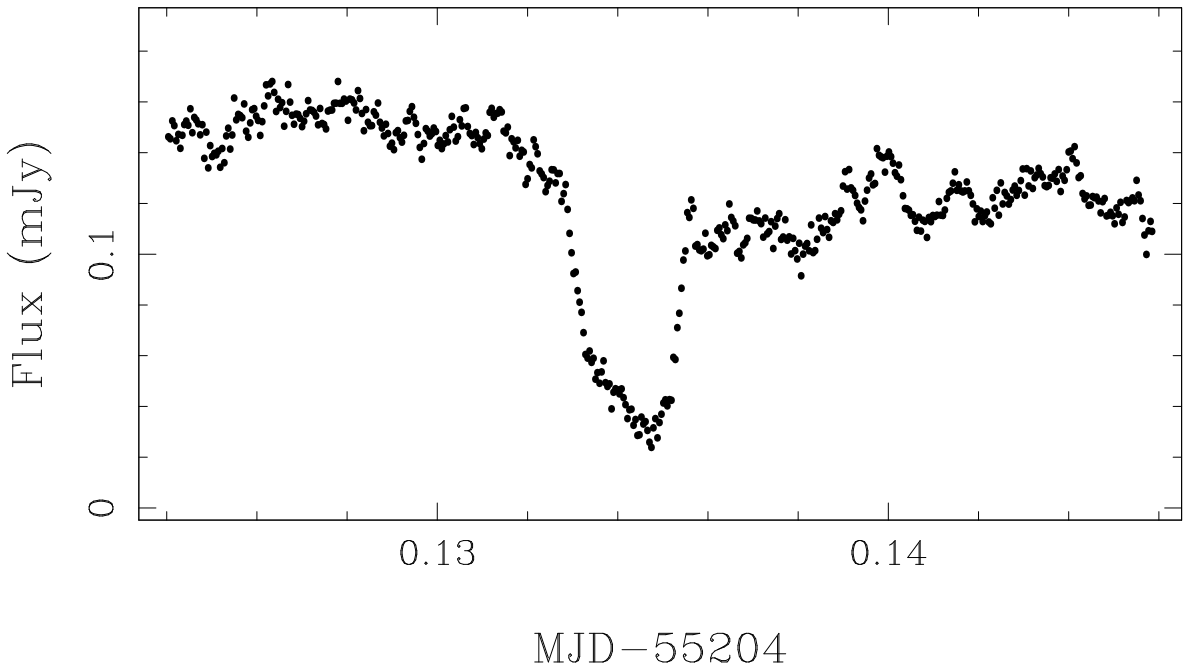

Figure 4.71: Eclipse of SDSS J125023.84+665525.4/OV Dra observed in g' with ULTRACAM on 2010-01-07. This is the clearest eclipse observed, with visible bright spot features. Other observations show stronger flickering.

SDSS J152419.33+220920.1

This CV was found in SDSS spectra (Szkody et al., 2009), which revealed double-peaked emission lines suggesting a high inclination. It was also independently discovered as a dwarf nova by CRTS and is therefore also known as CSS090329:152419+220920. Time-resolved photometry later revealed eclipses (Southworth et al., 2010), which were modelled to estimate the inclination and mass ratio of the system, but the data were not of sufficient quality to measure the individual component masses. Further photometry of 27 eclipses improved the ephemeris greatly (Michel, Echevarría & Hernández Santisteban, 2013), and we derived only a slightly improved ephemeris (see Table 4.2).

In our first observations SDSS J152419.33+220920.1 showed minimal bright spot features, and often its eclipse was completely blended with that of the WD. However, the disk radius and/or bright spot flux changed in later observations, showing a clear bright spot component which is separate from the WD. Figure 4.72 shows two phase-folded

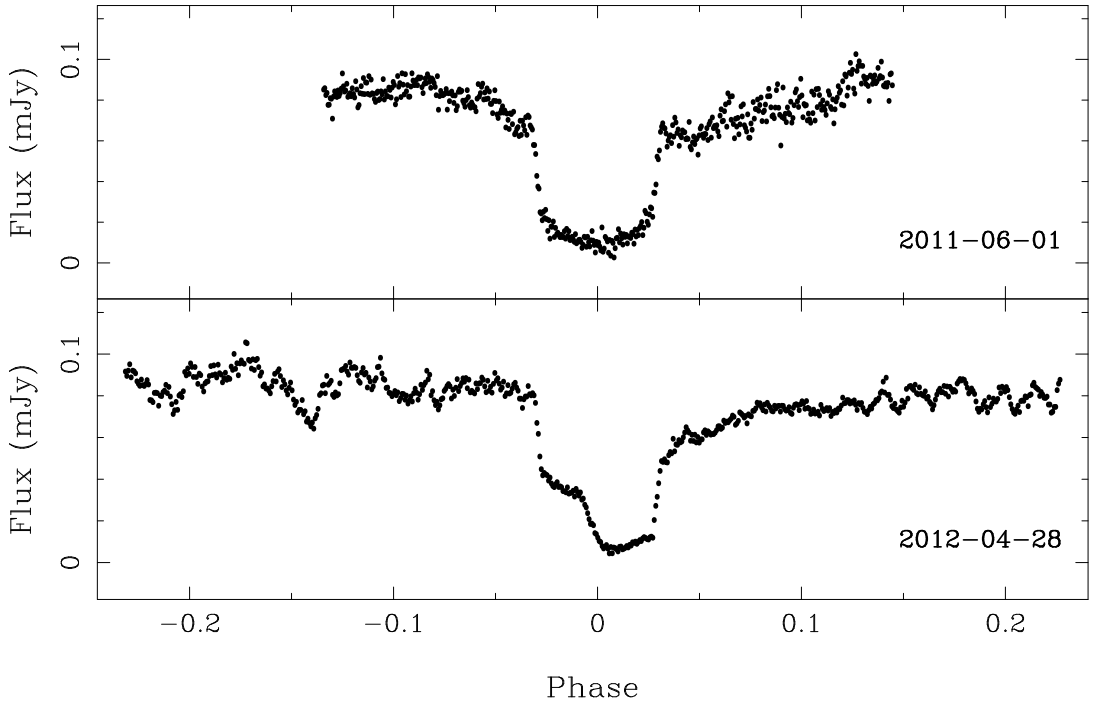

Figure 4.72: Phase-folded eclipses of SDSS J152419.33+220920.1 observed in g' with ULTRACAM on 2011-06-01 (top) and 2012-04-28 (bottom). In 2011 the bright spot was very weak, or perhaps blended with the WD, but in 2012 it was clearly visible.

eclipses which demonstrate this change in eclipse structure. The light curves showing a visibly separated bright spot ingress should be suitable for light curve modelling, although the system does show flickering which may interfere with model fitting.

SDSS J155531.99-001055.0

This system was found in the first spectroscopic search for CVs in the SDSS (Szkody et al., 2002), with double-peaked emission lines indicating a high inclination. The CSS light curve showed no outbursts, but does hint at eclipses with occasional faint measurements. Photometric follow-up confirmed the eclipses and an orbital period of 113 minutes (Southworth et al., 2007). We have observed 14 eclipses since 2007 with ULTRACAM and ULTRASPEC. The eclipse structure clearly changes over time, as shown in Figure 4.73. In 2007 there were no bright spot features, but in 2013 the bright spot ingress, at least, is clear. Further observations may provide sufficient data to enable light

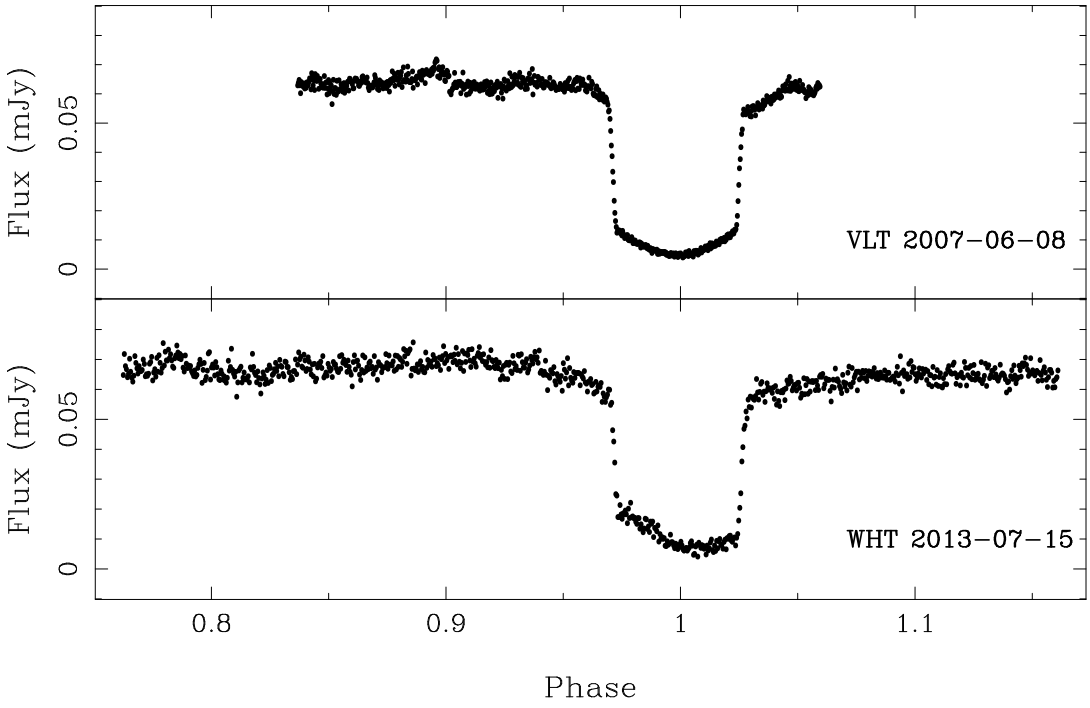

Figure 4.73: Phase-folded eclipses of SDSS J155531.99-001055.0 observed in g' with ULTRACAM on 2007-06-08 (top) and 2013-07-15 (bottom). The strength of the bright spot feature evolves with time.

curve modelling.

SDSS J155656.92+352336.6

This object, also known as BT CrB, was discovered in follow-up studies of the SDSS (Szkody et al., 2006). Like LT Eri, it shows multiple outbursts in the CSS archival light curve, from around $V = 18$ to around $V = 17$. An initial orbital period estimate of approximately 2 hours was found by Szkody et al. (2006), which we confirmed and refined (see Table 4.2).

We observed SDSS J155656.92+352336.6 with ULTRACAM, ULTRASPEC and *pt5m*. It was seen in outburst on several occasions. High cadence observations of the eclipse often show the WD and bright spot egresses well separated, but with blended ingresses. On one occasion (shown in Figure 4.74), the ingresses are marginally separated. Clearly this system has a variable disk radius, and at times could present an eclipse structure that

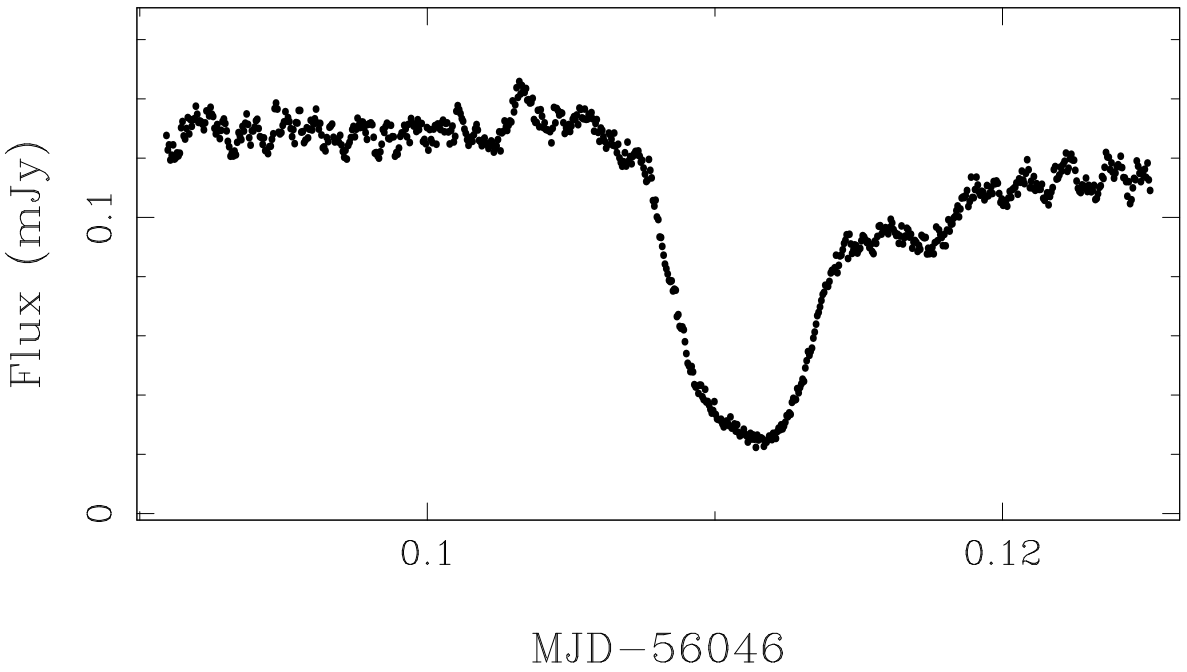

Figure 4.74: Eclipse of SDSS J155656.92+352336.6/BT CrB observed with ULTRACAM in the g'-band on 2012-04-28. The WD ingress is only just separated from the bright spot ingress, but the egresses are well separated.

can be modelled, though more data is needed.

SSS120402:134015-350512

This object, found by CRTS, shows multiple dwarf nova outbursts in its archival light curve, as well as fainter observations associated with eclipses. Coppejans et al. (2014) observed two eclipses, and derived a period of 0.059 days. We observed a single eclipse with ULTRASPEC (Figure 4.75) and confirm the eclipse is only partial in nature, with the eclipse of the WD suspected to be grazing. This makes it unsuitable for light curve modelling. We cannot improve on the ephemeris due to cycle count ambiguities.

SSS100615:200331-284941

This system was discovered via an outburst in CRTS, and later observed by Coppejans et al. (2014). They found deep eclipses and an orbital period of 1.4 hours. We observed

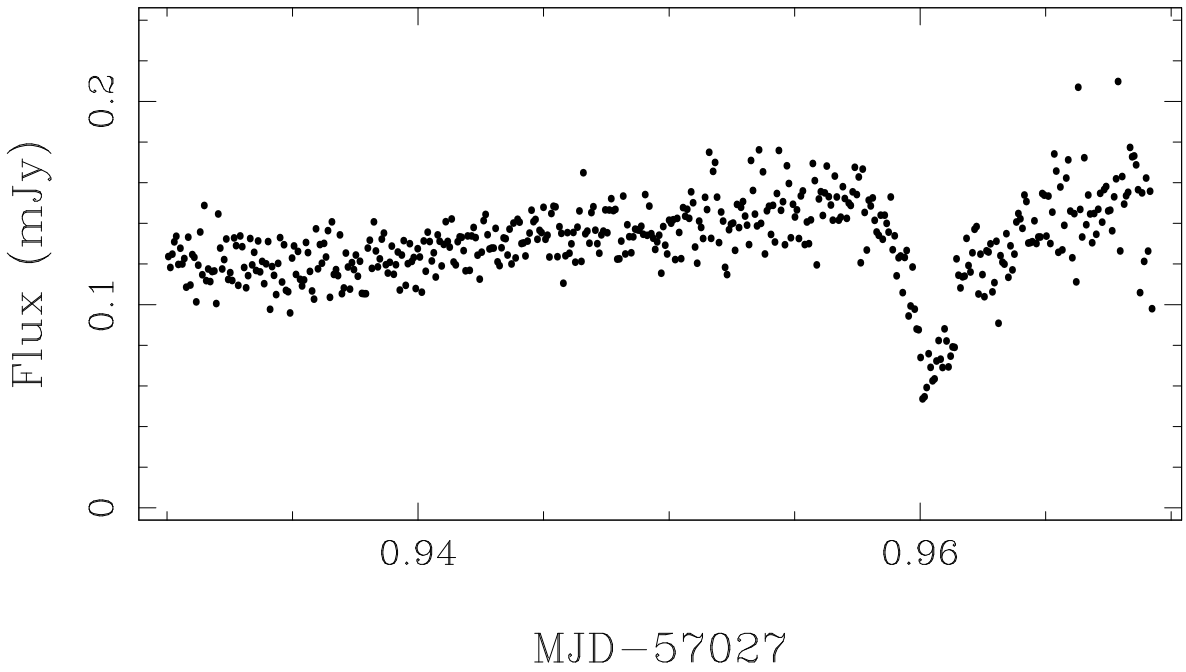

Figure 4.75: SSS120402:134015-350512 suspected grazing eclipse observed with ULTRASPEC on 2015-01-05.

three eclipses with ULTRACAM, enough to clearly see the eclipse structure and to constrain the ephemeris further (see Table 4.2). Figure 4.76 shows an example light curve. This system should be suitable for modelling.

V2051 Oph

This well-known eclipsing CV was originally discovered on objective prism plates (Sanduleak, 1972). It was studied photometrically by Warner & Cropper (1983), who observed eclipses at a time resolution of 5 seconds. These observations found bright spot features, making this system a good target for eclipse modelling. V2051 Oph has been the subject of numerous studies, via eclipse mapping and Doppler tomography (Vrielmann, Stiening & Offutt, 2002; Baptista & Bortoletto, 2004; Saito & Baptista, 2006; Longa-Peña, Steeghs & Marsh, 2015), as well as eclipse-timing observations which show it has cyclical period changes and a larger than expected orbital decay rate (Baptista et al., 2003; Qian et al., 2015) - this may be due to the presence of a third body and/or magnetic braking.

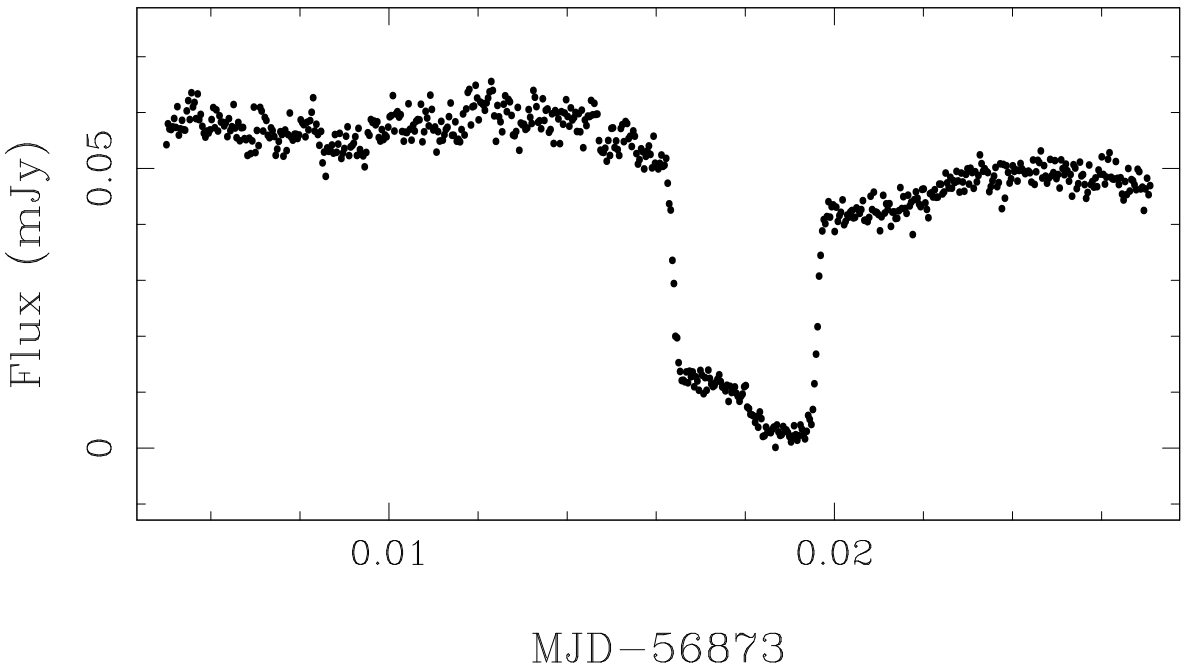

Figure 4.76: Eclipse light curve of SSS100615:200331-284941 observed in g' with ULTRACAM on 2014-08-03. Although the data are noisy, the bright spot ingress and egress are visible.

In the hope of being able to model the eclipse structure of V2051 Oph, we have observed a total of 26 high-speed eclipses over 8 years (Table 4.2). On many occasions the light curves are not useful, due to either cloudy conditions, the system being in outburst, the presence of strong flickering, or the bright spot eclipse not being visible. However, several observations appear to be useful, with visible bright spot features which should be suitable for light curve modelling. Figure 4.77 shows two phase-folded eclipses showing the differences in eclipse structure at different epochs. We confirm the period changes discussed by Baptista et al. (2003) and Qian et al. (2015), though we cannot confirm that these changes are cyclical in nature due to the relatively short baseline of our observations.

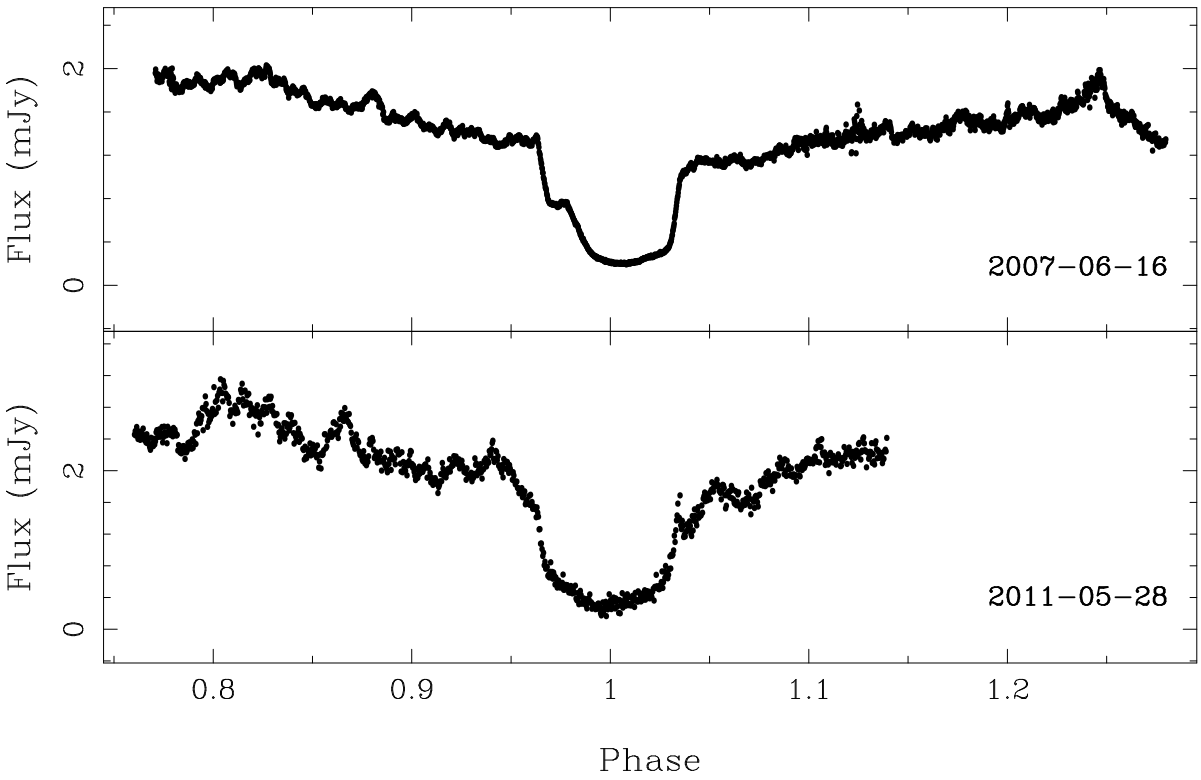

Figure 4.77: Phase-folded eclipses of V2051 Oph observed in g' with ULTRACAM on 2007-06-16 (top) and 2011-05-28 (bottom). The bright spot ingress is much clearer in 2007, but the egress is less clear. The opposite seems to be the case in 2011, although flickering may be affecting this.

V713 Cep

Originally discovered as variable by inspection of archival photometric plates (Antipin & Kroll, 2003), eclipses were soon discovered in this CV (*vsnet* alert 9516). More detailed photometry showed that it had the classic CV eclipse structure that lends itself to light curve modelling (Boyd et al., 2011). We have observed 15 eclipses. There are signs of a varying orbital period for this system according to the linear fit to the mid-eclipse times. This is discussed in detail in Bours et al. (2016).

Interestingly, the bright spot in V713 Cep changes significantly over time. In some observations it is clearly seen, and in others there is no bright spot feature at all. The top two panels in Figure 4.78 demonstrate this.

In addition to this, we also observed an apparent complete switch-off of accretion

on 2015-06-24. This eclipse, shown in the bottom panel of Figure 4.78, is typical of detached binaries in which the WD is not accreting gas from the secondary star. A short while later, on 2015-09-17, the system appeared to be accreting again. This low-state behaviour is unexpected and very rare among non-magnetic dwarf novae. The only other example of this was seen in IR Com, with an extended low state in which accretion was confirmed to have switched off (Manser & Gänsicke, 2014). Low states are however found in the VY Scl class of nova-like variables, in both magnetic and non-magnetic systems (Hellier, 2001).

One theory for temporary interruptions to mass transfer is that they occur when a starspot crosses the L1 point on the donor star, lowering the stellar surface and disconnecting it from the L1 point. In CVs with strong magnetic fields (i.e. polars), accretion onto the WD ceases immediately, and the system enters a low state until mass transfer resumes. However, in non-magnetic dwarf novae the system remains bright, as the accretion disk is thought to outlive the short break in mass transfer, only draining slowly onto the WD (Hessman, 2000; Manser & Gänsicke, 2014).

Our observations of V713 Cep show that this break in mass transfer is not constrained to nova-likes, polars and the single instance of IR Com, but may be more common than first thought. Unfortunately V713 Cep is not covered in the CSS footprint, and we cannot say exactly how long the low state lasted. Our observations of eclipses before and after the non-accreting observation on 2015-06-24 both show normal eclipse structure including an accretion disk. These are separated by 403 days and thus we constrain the low state of V713 Cep to a maximum duration of 403 days.

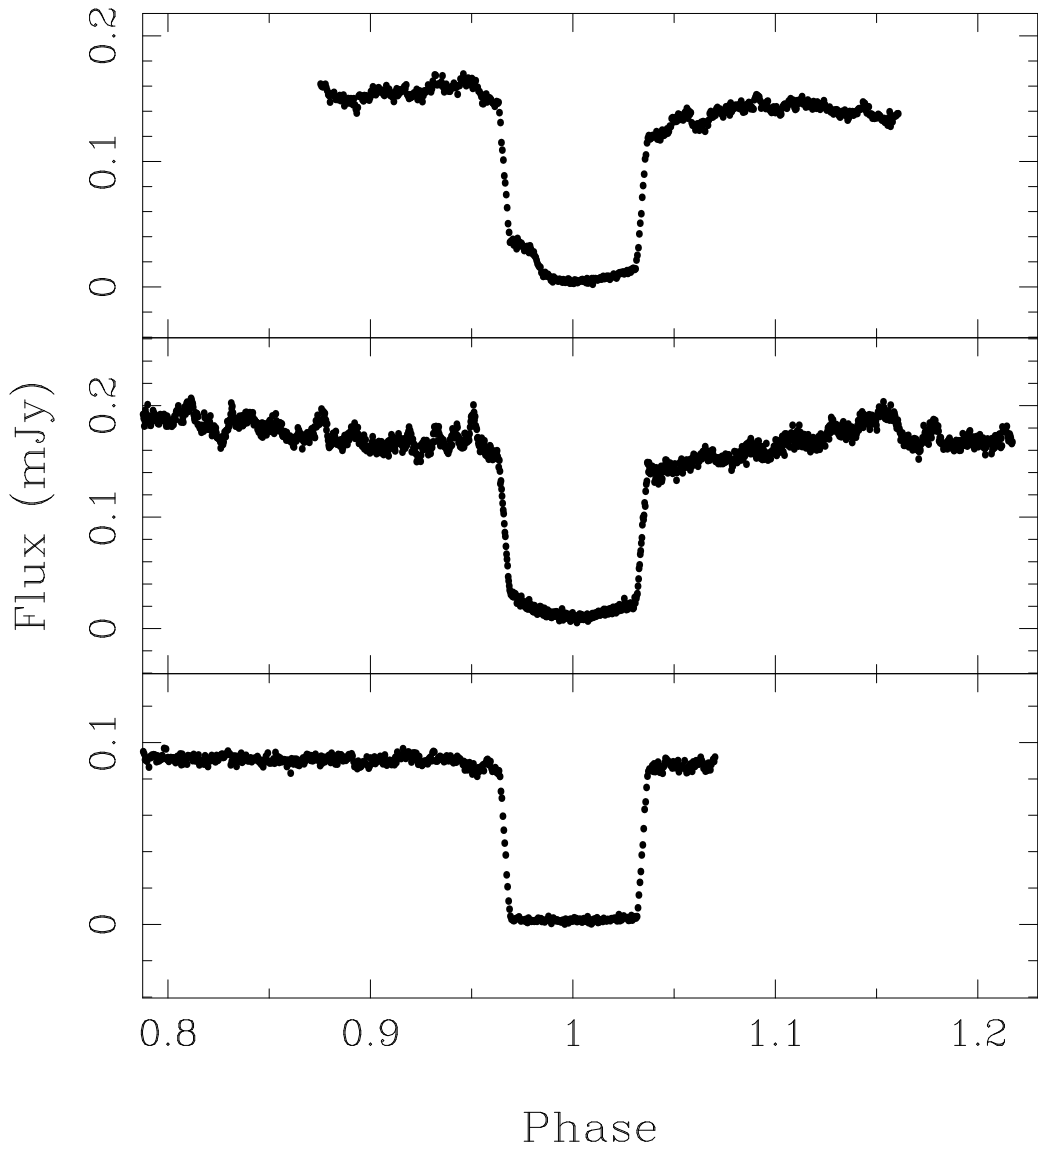

Figure 4.78: Phase-folded eclipses of V713 Cep observed in g' with ULTRACAM on 2013-07-15 (top), 2014-08-02 (middle) and 2015-06-24 (bottom). The bright spot feature appears and disappears over time, and in June 2015 we saw accretion switching off altogether, with an eclipse that resembles a bare WD.

4.7 Conclusions

We have presented time-resolved photometry of 74 CVs, most of which have never before been observed in detail. We have discovered or confirmed 13 new eclipsing systems with orbital periods ranging from 1.5 to 5.1 hours. We studied 27 new systems which did not show any eclipses. The remaining 34 systems presented here were already known to be

eclipsing, but we provide the highest cadence observations of these objects to date.

Systems of particular interest include V713 Cep, in which we observed a temporary switching-off of accretion; and ASASSN-14mv and CSS111019:233313-155744, which both have orbital periods well below the CV period minimum. The low orbital periods and light curve shapes suggests that they may be double degenerate (AM CVn) systems or evolved CVs.

For all eclipsing systems we provide the mid-eclipse times in Table 4.2 and the ephemerides in Table 4.2. We have also discussed the feasibility of modelling the eclipse structure for all eclipsing systems, to determine their system parameters. We found that approximately 20 of the objects discussed above should be suitable for system parameter studies using light curve modelling, while a further 15 systems may be feasibly modelled if additional observations are acquired. The modelling of these systems is outside the scope of this thesis, and will be presented in a forthcoming paper by our group.

We found 13 new eclipsing systems in a study of approximately 40 systems. This cannot be directly compared to the expected rate of eclipses in a general population of CVs (approximately 10% according to Ritter & Kolb, 2003), because we deliberately targeted systems which had a greater chance of showing eclipses, for example systems with CSS light curves showing occasional faint states, and high inclination systems identified as such by their double-peaked spectral lines.

To broaden our collection of CVs with precisely measured parameters, we would like to study more systems with longer orbital periods, and also more systems close to the period minimum. Our search for new eclipsing systems focused on new CVs discovered by their transient behaviour in dwarf nova outburst. The frequency of dwarf nova outbursts increases with mass accretion rate, which itself increases with orbital period. This means that longer period systems are easier to detect in wide-field optical surveys, and there is a bias towards these kinds of systems. This is good for our study, since we hoped to find more long period eclipsing systems. On the other hand, longer orbital periods need

longer observations to detect eclipses, and it is possible that we may have missed some exceptionally long period systems (longer than 5 hours for example).

Our search for eclipsing CVs is poorly suited to discovered systems close to the period minimum. This is because at such short orbital periods, CVs tend to have very low mass accretion rates. This in turn means dwarf novae outbursts are very rare, and thus selecting CVs based on their transient detections in wide-field surveys means we are biased against finding the shortest period systems. To avoid this bias, we would need to search for eclipses in CVs which were discovered spectroscopically (e.g. Gänsicke et al., 2009).

Finally, we note that there are a number of systems which might be suitable for light curve modelling, but that our data is insufficient to allow us a full judgement. Some of these systems appear to have variable rates of accretion and variable disk radii, meaning that they will sometimes be in a state which is conducive to modelling. These systems require continued monitoring and then special and urgent attention when they reach such states. Other systems are simply too faint to provide sufficient signal-to-noise ratios for investigating the light curve structure or for modelling. In many cases, such as MASTER OT 232100.42+494614.0, observations with a larger telescope and in dark time are required in order to attain high signal-to-noise ratios. This will be a priority for our group going forward, especially for those systems with longer periods.

Chapter 5

Optical follow-up of FRBs with the Thai National Telescope

This chapter discusses the relatively new and poorly-understood phenomena known as Fast Radio Bursts (FRBs). I review the literature on FRBs and our observational attempts to detect their optical counterparts. Some of the observations were conducted by others (as stated below), but all data reduction and analysis was carried out by myself. I did most of the coordination of the follow-up efforts and liaising with collaborators, including the radio astronomers searching for FRBs, led by Evan Keane (SKA, Jodrell Bank Centre, previously Swinburne, Australia), Emily Petroff (ASTRON, Netherlands, previously Swinburne, Australia) and Sarah Burke Spolaor (NRAO, New Mexico, previously CalTech). Several papers are currently in preparation which will include my follow-up efforts. The following data reduction, analysis and text is all my own.

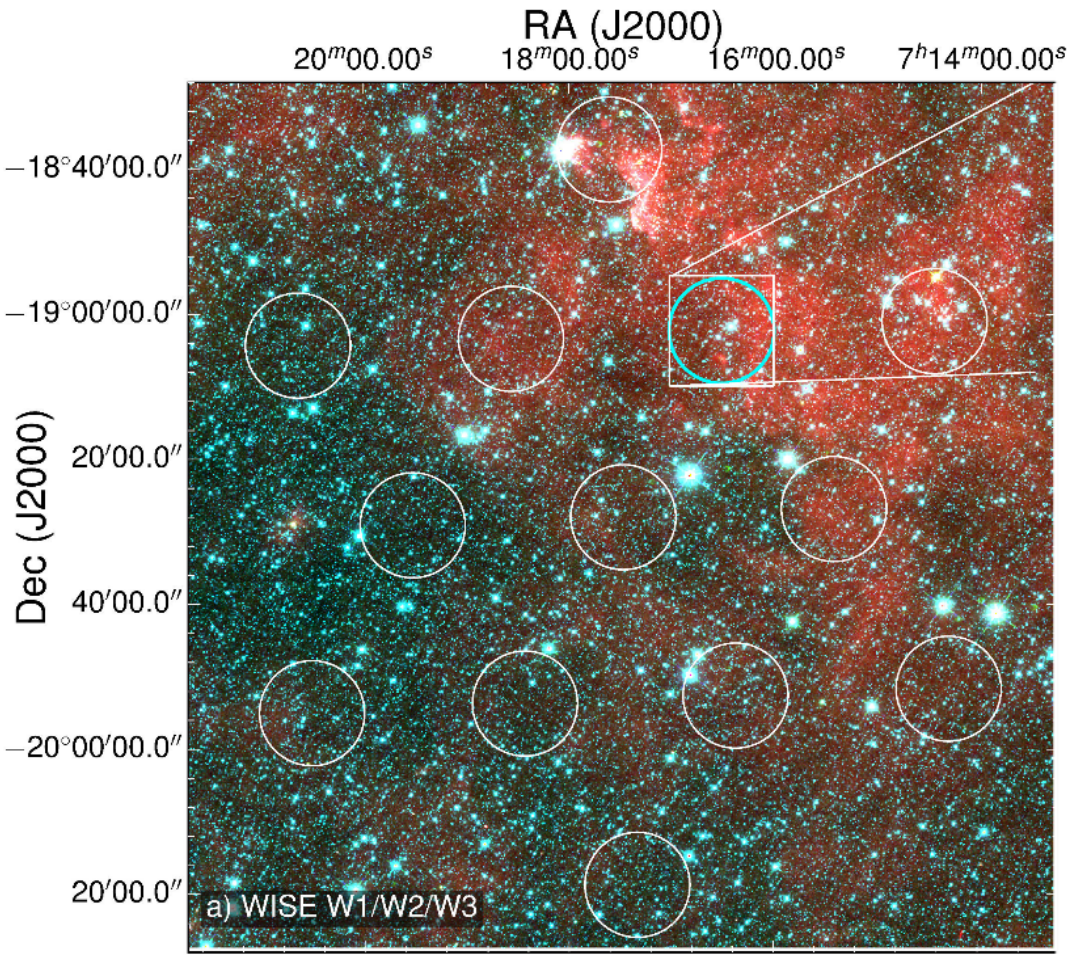

Figure 5.1: Example field-of-view of the Parkes multibeam receiver, superimposed on a WISE infrared image. Reproduced directly from Keane et al. (2016).

5.1 Introduction and Theory

5.1.1 What do FRBs look like?

Fast radio bursts (FRBs) are short (few millisecond), highly dispersed pulses of coherent radio emission, only seen so far at 0.8 and 1.4 GHz. With measured fluxes of 0.02-30 Jy, they can be extremely bright. Due to the nature of their discovery, none have so far been conclusively associated with a quiescent host source, be it galactic or extra-galactic. This is because of their poorly constrained localisation, due to the wide field of view of directional radio telescopes. For example the Parkes telescope, where most

FRBs have been observed, has a 13-beam receiver where the half-power width of each beam is around 14′. The beams are arranged in a double hexagon around a central beam, covering a total area of approximately $2^\circ \times 2^\circ$, with some gaps between the beams having lower sensitivity. Figure 5.1 shows the scale of the Parkes multiple beams on a background infrared image.

Some bursts have shown double-peaked structure (Champion et al., 2015), whilst most appear to be single pulses. Until very recently, no events had definitely been seen to repeat, but Spitler et al. (2016) and Scholz et al. (2016) have now reliably reported repeated bursts from FRB 121102. We discuss these recent developments in more detail, as well as the claimed host galaxy localisation of FRB 150418 (Keane et al., 2016) at the end of this Chapter in Section 5.8.

The first prototype burst was discovered by Lorimer et al. (2007) in archival data from a survey of the Magellanic Clouds (Manchester et al., 2006). This burst, dubbed 'the Lorimer burst', showed an enormous peak flux of 30 Jy, over ten times brighter than any other FRB discovered since. Over the last 9 years, 16 further events have been detected and published. A complete catalogue of published FRBs is available and being maintained at `http://www.astronomy.swin.edu.au/pulsar/frbcat/` (Petroff et al., 2016). Figure 5.2 shows a sky distribution plot of all FRBs observed so far, with most being below the celestial equator due to their discovery at Parkes in Australia.

FRBs are deeply interesting because, if proven to be from extra-galactic sources, they could provide new information about the intergalactic medium through their dispersion measures. In turn this could effect our understanding of the baryonic content of the Universe (McQuinn, 2014). This is discussed further in Section 5.1.3.

Dispersion Measure

FRBs are the first transient radio events with properties suggesting an extra-galactic origin. The key parameter here is the dispersion measure (DM). Radio emission from

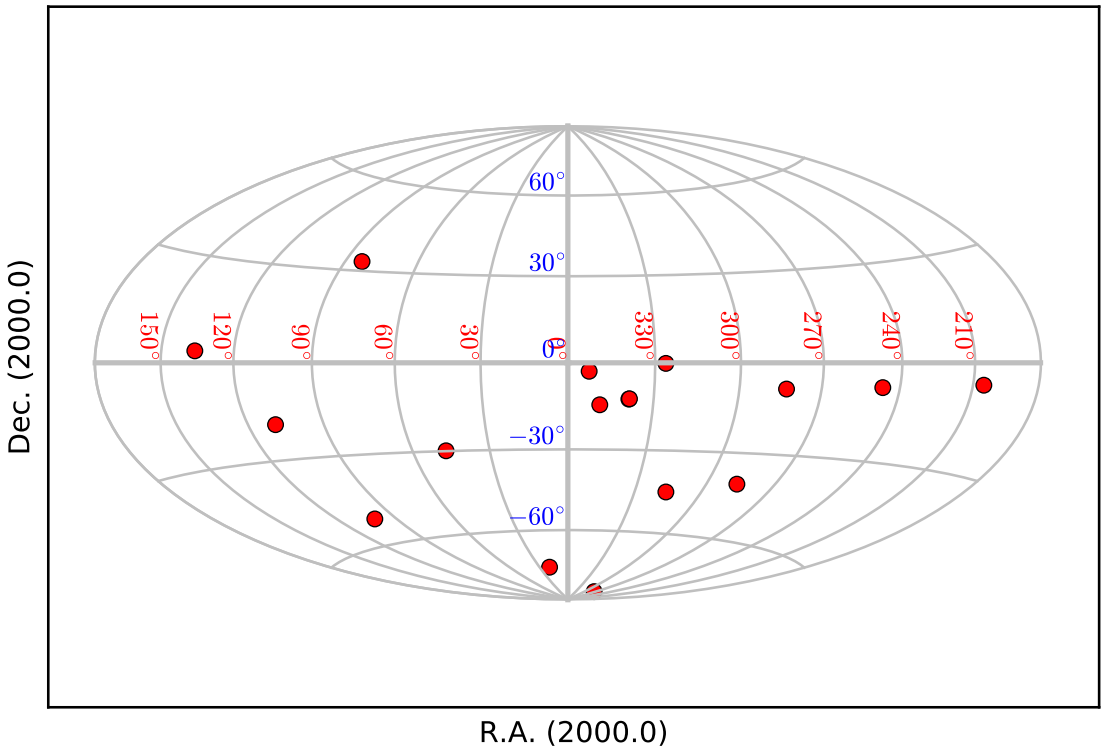

Figure 5.2: Sky distribution plot of all FRBs observed so far.

celestial sources has to pass through a cold ionised plasma of free electrons, which permeates the Galaxy and also the space beyond. The interstellar medium (ISM) has three main components: the cold ($\sim$ 100K), dense, neutral dust and gas clouds which are thought to be the birth places of stars; the warm, photo-ionised gas, with temperatures around 8000K, which is considered to be the largest component of the ISM by mass; and the hot, highly ionised gas which permeates most of the Galaxy, with temperatures approaching 10^6K (Spitzer, 1998). It is the warm photo-ionised gas which interacts

most with radio signals. The low density prevents cooling and recombination, even at distances far away from strong sources of photons (e.g. O-type stars). Due to the ionisation level being relatively low, it is often called a 'cold' plasma, despite being considered the 'warm' component of the ISM.

The ionised plasma has a frequency-dependent refractive index, and thus the arrival time of a celestial radio signal is delayed by an amount proportional to its frequency (Lorimer & Kramer, 2005; Hassall, Keane & Fender, 2013). A single pulse emitted across a range of frequencies thus arrives with a spread of time delays, with lower frequencies arriving later. Figure 5.3 helps explain this effect. The total column density along a particular line of sight (known as the dispersion measure, or DM, measured in pc cm^{-3}) denotes the amount of dispersion due to this effect. The time delay, t_{DM} (in seconds), at any given frequency is given by Equation 5.1.

$$\Delta t_{DM} = \frac{DM}{C_d \nu^2}, \tag{5.1}$$

where C_d is the dispersion constant, equal to 2.41×10^{-4} MHz^{-2} pc cm^{-3} s^{-1}, and ν is the frequency in MHz (Manchester & Taylor, 1972; Lorimer & Kramer, 2005).

Searches for single pulses from FRBs and from rapidly rotating radio transients (RRATs, galactic neutron stars which show bright, non-periodic pulses, see e.g. McLaughlin et al. 2006; Lorimer et al. 2007) exploit this dispersion property of radio signals. Raw radio data are de-dispersed for a large range of DM values, with the resulting data then searched for single pulses with significant signal-to-noise (Rane et al., 2016). Radio frequency interference is masked along the way, and filters select only those pulses which match certain criteria, such as only being observed in a single beam (terrestrial signals are often seen simultaneously in many/all of the multiple receivers utilised on directional radio telescopes). This process produces a definitive pulse signal at a given DM, shown by diagnostic plots which must be inspected visually for final confirmation. An example

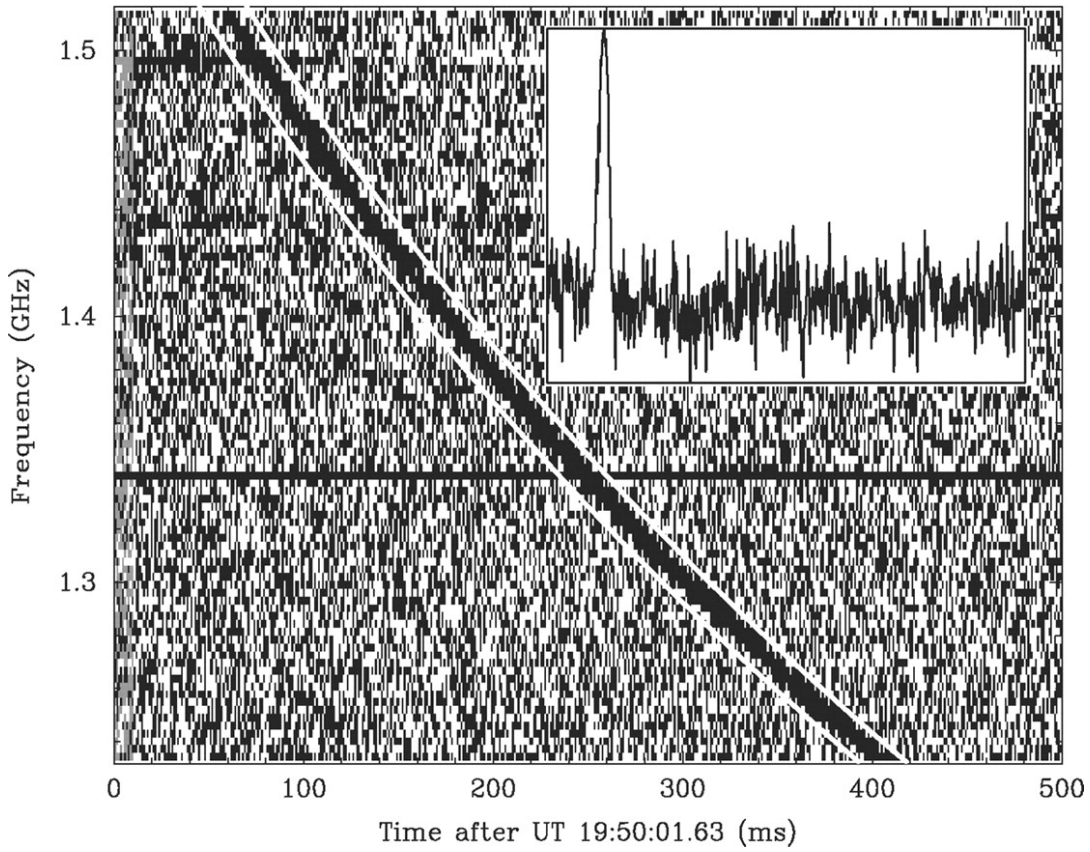

Figure 5.3: Dispersion of radio emission as it propagates through a cold plasma, as seen in FRB010724. Intensity (darkness) is shown on a plot of frequency versus time. The dark horizontal band at 1.34 GHz represents a malfunctioning channel, and the white lines either side of the pulse show the theoretical behaviour for the cold plasma dispersion law (Equation 5.1) with a DM of 375 pc cm^{-3}. The inset shows the average pulse signal after de-dispersion, for a reference frequency of 1.52 GHz. Reproduced directly from Lorimer et al. (2007).

of such a plot is given in Figure 5.3.

With information on the density of the interstellar plasma, the DM can be equated to a distance. Using DM measurements and models of the galactic distribution of free electrons (Cordes & Lazio, 2002), the distance to a burst can be estimated. One of the most interesting properties of the FRBs reported so far is that they show DMs which are significantly larger than expected for their Galactic lines of sight. In other words, the pulses are dispersed by more than can be accounted for by the free electrons in the ISM of the Milky Way. This is the main evidence we have that these bursts are extra-galactic,

and therefore highly luminous. The excess DM above the expected values could be explained by ionised plasma in host galaxies, and/or in the intergalactic medium (IGM). The strength of both of these contributions are uncertain, but in the case of FRB 110523 (Masui et al., 2015) the direct detection of scintillation and polarisation suggests that the emission was scattered primarily in a dense, compact, magnetised nebula or within the dense, magnetised, inner regions of a host galaxy, rather than whilst traversing the IGM or the Milky Way ISM.

Temporal Scattering

Another effect seen in celestial radio pulses is the broadening of the pulse width at any given frequency. This is caused by scattering due to fluctuations in the electron density along the line of sight. The multipath propagation of the emission broadens a narrow pulse in time, reducing the peak intensity. This means that whilst FRB pulses are already short (few ms), their origins may be even shorter, with some claiming that FRBs are consistent with a delta-function broadened by electron scattering (Lyutikov, Burzawa & Popov, 2016). This also implies that even those FRBs which appear to be single pulses, may in fact have multiple original peaks that have been scatter-broadened into a single signal (Champion et al., 2015).

Only some FRBs have signals strong enough to measure the amount of scatter broadening. These show less broadening than expected for their assumed distances, and this is thought to be due to the positioning of the scattering screen between the observer and source (Hassall, Keane & Fender, 2013). Scattering models rely on the concept of a thin screen of electrons which cause the scatter, as they are unable to model the total scattering effect along the entire line of sight. Whether this screen is placed close to the observer, the source, or at locations between the two, affects the results. The measured scattering is lower than expected for a screen between the source and observer, and matches more closely a screen located closer to the observer. This is somewhat

contrary to the findings of Masui et al. (2015), who detected signals of magnetisation and scattering with an origin close to the source of FRB 110523.

Interference

Radio frequency interference (RFI) is a common plague in radio surveys, causing a reduced signal-to-noise ratio and creating false sources. Interference can come from many physical origins, including satellites, solar radiation, and even slowly varying gain changes in the detector systems.

There are several methods which aim to correct for RFI (Eatough, Keane & Lyne, 2009; Kocz et al., 2012). In general, terrestrial interference can be removed using the natural cold plasma dispersion law described above. When the pulse is de-dispersed (when trialling DMs in search of bursts), any RFI signals will be dispersed. Unfortunately, sometimes the RFI signal is strong enough to still be picked up by automatic detection routines even after dispersion, and the vast numbers of RFI signals present make manual analysis unmanageable. Often RFI signals are detected simultaneously in multiple beams of the telescope, so this can also be used to help guard against RFI (Kocz et al., 2012).

An interesting development soon after the first FRBs were discovered was the observation of numerous bright signals with similar characteristics to the original Lorimer burst, but with properties clearly demonstrating a terrestrial origin. These so called 'perytons' (the name comes from the mythological winged stag which casts the shadow of a man) showed a frequency-dependent delay which closely matches those of FRBs (Equation 5.1). An example peryton detection is shown in Figure 5.4.

It was initially unclear as to whether these events were the same as FRBs, related to them, or totally different. Some events exhibited sharp deviations from the frequency dispersion law. These 'kinks' were strong enough to suggest a non-astrophysical origin. In addition, they were usually seen in many telescope receivers simultaneously, and

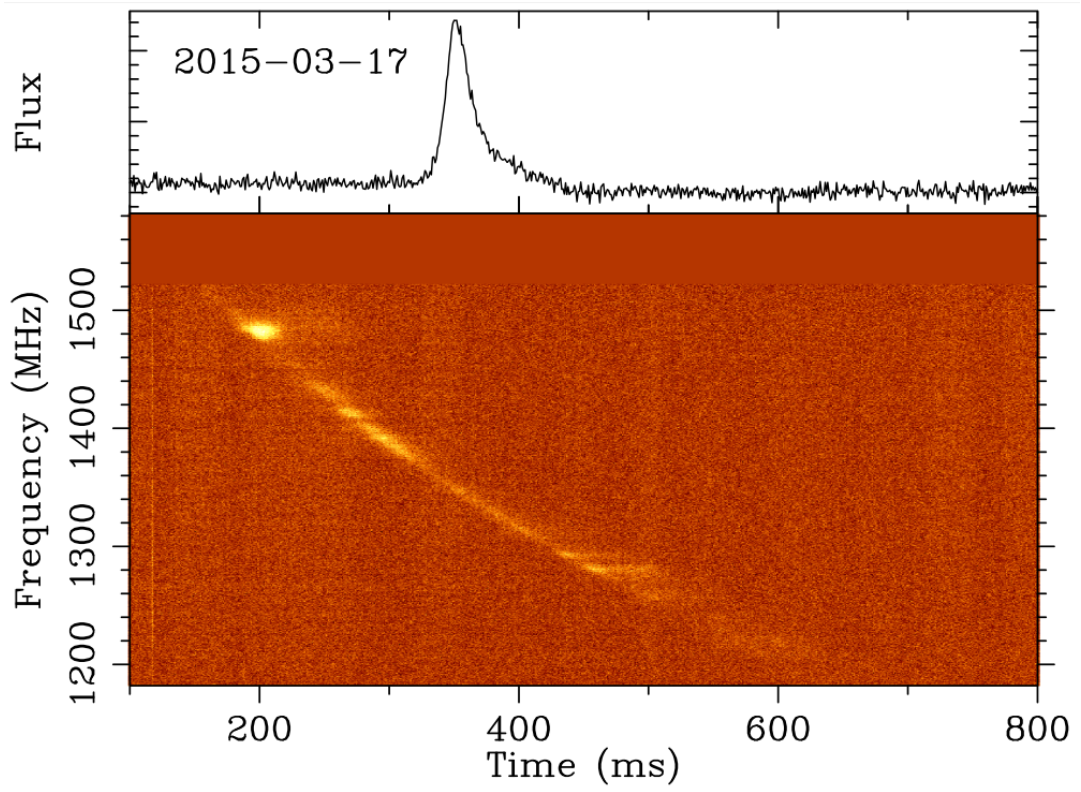

Figure 5.4: Pulse shape (top) and frequency dispersion curve (bottom) of a peryton detected by Parkes, with blips and kinks in the dispersion curve. Reproduced directly from Petroff et al. (2015c).

seemed to occur more often at certain times of the civil day, which confirms their terrestrial origin (Burke-Spolaor et al., 2011; Kocz et al., 2012; Bagchi, Nieves & McLaughlin, 2012). Many questioned the astrophysical origin of FRBs, since they looked so similar to perytons, especially the original 'Lorimer burst' (Kulkarni et al., 2014).

Eventually Petroff et al. (2015c) discovered the true source of perytons at Parkes - microwave ovens in a nearby kitchen. The powering down of the magnetrons inside the microwave ovens, combined with premature opening of the door (and thus breaking the Faraday cage), allow for the production of swept frequency signals which produce perytons, when the Parkes receiver was located in certain positions. This discovery reinvigorated the study of FRBs, since the clarified understanding of perytons has dissociated them from FRBs. This lends more credence to the idea that FRBs have truly

astrophysical origins.

5.1.2 Where could FRBs come from?

For a time there were more theories for FRB progenitors than there were observed bursts, and some exotic explanations still remain as plausible candidates. The question remains open, but certain observational characteristics allow us to constrain possible progenitor theories and favour or disfavour them accordingly.

Observational Constraints

The observed properties of the FRB population give certain constraints on progenitor theories. The following parameters are all important:

- Dispersion Measure: This dictates inferred distances - if a large fraction of the DM contribution can be attributed to phenomena very close to the FRB source, as has been suggested for FRB 110523 (Masui et al., 2015), then lesser distances are inferred.

- Pulse Width: This dictates the size of the emission region - widths are scatter broadened when passing through the free electron plasma, so true pulse widths are likely even smaller than the milliseconds that have been observed (Lyutikov, Burzawa & Popov, 2016). The speed of light (or speed of information travel) thus constrains the size of the emission region to less than $\sim$ 1000km.

- Polarisation: Some FRBs have measurable polarisation signatures, e.g. FRB 110523 (Masui et al., 2015), and FRB 140514 (Petroff et al., 2015a), which can be used to infer information about the magnetic field conditions between the source and the observer, or the nature of the progenitor itself (Cordes & Wasserman, 2016).

- Energy/Flux/Fluence: The scale of the energy involved in the production of FRBs can constrain progenitors - if sources are shown to be at cosmological distances, and/or with emission that is isotropic in nature rather than beamed, higher total energy emission is needed.

- Spectral Index, α: The shape of the flux as a function of frequency ($f_\nu \propto \nu^\alpha$) can suggest a great deal about the emission mechanism. Indices (α) of 0 to 2 often indicate thermal emission, whereas negative indices (e.g. -1) tend to be related to non-thermal, synchrotron radiation (Burke & Graham-Smith, 2014). Bursts have been observed with differing spectral indices (Spitler et al., 2016), but some of this variation could be explained by free-free absorption dampening the spectral steepness (Lyutikov, Burzawa & Popov, 2016).

- Scatter Broadening: Some bursts show measurable scatter broadening of the pulse profiles (see Figure 5.5), e.g. FRB 110220 (Thornton et al., 2013), and FRB 131104 (Ravi, Shannon & Jameson, 2015). Sometimes this cannot be attributed solely to passing through the IGM (Luan & Goldreich, 2014), suggesting relatively dense source environments.

- Event Rate: The observed event rate of FRBs must be reconciled with the predicted event rate of any given progenitor. The current estimate for the event rate at 1.4 GHz, above a fluence of 3.8 Jy ms, is $\sim 3 \times 10^3$ events per sky per day (Crawford et al., 2016).

- Repetition: FRB 121102 has been seen to repeat (Spitler et al., 2016). If all FRBs transpire to be repeating, this might disprove any catastrophic (meaning violent and totally disruptive) progenitor theories.

Considering these constraints, there appear to be roughly three different classes for possible progenitor theories. One option is some nearby galactic source, which somehow

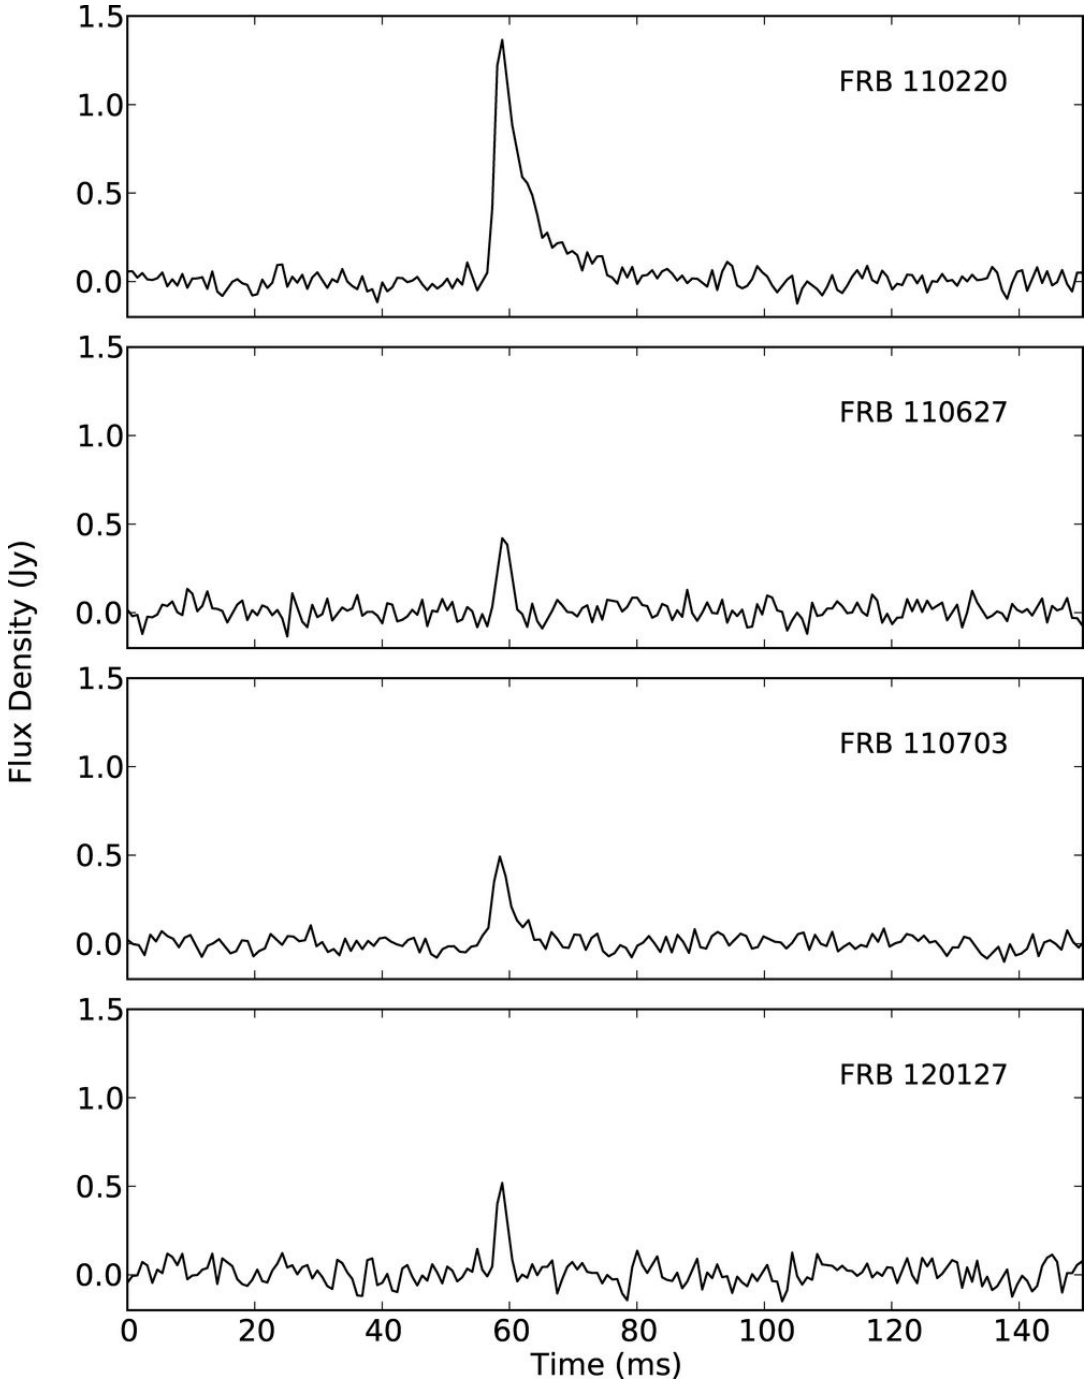

Figure 5.5: Four FRB flux density profiles, as measured and published by Thornton et al. (2013). The effect of scatter broadening is apparent in the top panel, which shows an exponential tail, whilst the other three bursts were too weak to show any clear broadening.

produces the necessary DM but does not require extreme energies. Another class might include extra-galactic events that are non-catastrophic - they may repeat on observable timescales and might be expected to be found within a few hundred Mpc. The final category of models includes catastrophic events at higher redshifts that would not repeat. They would need larger energies, but the corresponding volume from which the events could be generated is equally much larger.

Galactic theories

In general the DM values for observed FRBs exclude Galactic origins. In order to account for the excess DM, Galactic theories must provide another source of the dispersion. One suggested theory involves hidden HII regions in the line of sight which might contribute additional DM. This possibility is usually discussed and discredited by FRB observers based on deep observations along the lines of sight of known bursts, which discover no evidence for previously undiscovered nebulae (Lorimer et al., 2007; Luan & Goldreich, 2014; Spitler et al., 2016).

It has been proposed that FRB signals could have originated from nearby stellar flares (Loeb, Shvartzvald & Maoz, 2014; Maoz et al., 2015). Considering that radio bursts have long been observed from flare stars (Lang & Willson, 1986; Bastian et al., 1990), it seems logical that these could also produce FRBs. The additional DM could be generated in dense plasmas in the coronae or coronal mass ejections. Loeb, Shvartzvald & Maoz (2014) and Maoz et al. (2015) discovered active stars in the fields of several FRBs, which they claim are unlikely to be chance detections. One key prediction of this theory is that bursts could and should be repeated, but with different DMs due to variable coronal density and structure. The identical DM measured from FRB 121102 over several-year timescales obviously casts doubt on this theory (Spitler et al., 2016).

Luan & Goldreich (2014) proposed the possibility that extraterrestrial civilisations might be using dispersed radio signals in an effort to alert us to their presence. However,

given the growing number of FRB detections in completely different areas of the sky (and therefore the Galaxy), this scenario seems implausible as it would require multiple different extraterrestrial civilisations to have reached the same technological age and discovered our presence at the same time.

Extra-galactic but not catastrophic theories

The simplest and most obvious way to explain the observed dispersion measures is to attribute the excess (beyond what is expected from the ISM within the Galaxy) to the free electron density of the intergalactic medium (IGM) and a progenitor host galaxy. Depending on the density and structure of this predicted IGM component, and the true distance to FRB sources, a wide range of DM contributions is possible. Several extra-galactic but non-cosmological (e.g. within a few hundred Mpc, redshift less than z$\sim$0.1) models have been proposed. Such models invoke significant DM contributions close to the source of FRBs, allowing their distances to be smaller and therefore requiring less extreme energy demands. Often these models are not catastrophic in nature, and are sometimes repeatable on a variety of timescales. All of the following theories involve neutron stars in some way.

One of the most likely explanations for FRBs is that they are simply a high energy extension of the giant radio pulses already known to be emitted by pulsars. Giant pulse emission has been discovered originating from several Galactic pulsars (e.g. Knight 2006), displaying extremely high brightness temperatures (up to 10^{37}K, requiring non-thermal, coherent emission, Hankins et al., 2003) and sometimes being composed of multiple short 'nano-shots' (Hankins & Eilek, 2007; Majid et al., 2011). Cordes & Wasserman (2016) and Popov & Pshirkov (2016) argue that such pulses could very occasionally reach much higher energy levels. They show that even rare pulses (<0.01 per year), in a relatively small, non-cosmological volume of $d < 100$ Mpc, could match the observed FRB event rate. Superposition of many incoherent nano-shot pulses is required to explain the

required high energy level, though this superposition is widely accepted as the source of the giant pulses in the Crab pulsar and others (Löhmer et al., 2008; Jessner et al., 2010). The existence of such short pulses (few nanoseconds), requires very small emission regions of a few metres (Hankins et al., 2003).

A potential downside to this theory is that the measured DMs are expected to place FRB sources at larger, often cosmological distances. Connor, Sievers & Pen (2016) discuss the possibility that such bursts may come from young pulsars which are still shrouded in their supernova remnants (see also Lyutikov, Burzawa & Popov 2016, who focus on young, rapidly rotating pulsars). This would provide additional material to contribute to the high DMs, and would explain the observational evidence that some of the scatter broadening occurs close to the source (Masui et al., 2015). Another possible stumbling block for the giant radio pulse theory is the need for a viable trigger or emission mechanism, which might generate such high energy bursts. It is not clear what triggers giant pulses in Galactic pulsars, and what would trigger equivalent supergiant pulses in FRBs. Several options have been discussed, including polar cap discharges after cascading pair creation (Jessner et al., 2010) and spin glitches. Another explanation involves the pulsar colliding with, or accreting, small orbiting bodies, such as comets or asteroids (Cordes & Shannon, 2008; Geng & Huang, 2015). In addition, the supergiant pulse theory neatly explains the measured polarisation signals from some FRBs, since giant pulses from Galactic pulsars are known to be highly polarised (Connor, Sievers & Pen, 2016).

Another possible model for FRBs related to neutron stars is that they are produced by hyperflares of soft gamma repeater (SGR) magnetars. Popov & Postnov (2010) and later Lyubarsky (2014) showed that expected rates of such hyperflares are consistent with FRB detection rates. Magnetic pulses are proposed to originate from the gamma-ray flares, triggered for example by a spin glitch inside the magnetar. When the propagating magnetic pulses reach the termination shock boundary between the magnetised,

relativistic magnetar winds and the surrounding nebulae, they could induce powerful synchrotron maser emission which could explain FRBs (Lyubarsky, 2014). Since magnetars tend to be young, they will typically trace regions of massive star formation, which can also help explain some of the excess DM if FRBs are found in nearby galaxies. In principle this theory should be testable with simultaneous detection or non-detection of both a radio and gamma ray component, although only nearby events ($d < 100$Mpc) would produce detectable gamma-ray signals. Some evidence for simultaneous detection of gamma-ray signals and radio bursts has been published (Bannister et al., 2012), though this was inconclusive. Tendulkar, Kaspi & Patel (2016) placed further limits on simultaneous detection using archival Parkes data taken at the same time as a galactic SGR flare, but could not fully reject the magnetar hyperflare model for FRBs.

Lastly, another neutron star source for FRBs, involving orbiting companions, was proposed by Mottez & Zarka (2014). They discussed the possibility that highly collimated radio emission from Alfven waves attached to bodies orbiting extra-galactic pulsars might provide the necessary supply of FRBs. Pulsars with highly magnetised winds and small orbiting companions (e.g. planets, asteroids or white dwarfs) may generate bursts from cyclotron maser instabilities in the Alfven wings - two stationary Alfven waves attached to the body are known as 'wings', driven by the electric currents produced as the body passes through the Poynting-flux dominated wind. Relativistic beaming can explain the observed properties of FRBs, and this theory also expects circular polarisation, as was seen in FRB 140514 (Petroff et al., 2015a). Mottez & Zarka (2014) also expect periodic repetition of bursts matching the companion orbital period. Such repetition was searched for in some FRBs (Petroff et al., 2015b) but no such periodicities were found on timescales <8.6hr. In addition, the repeating bursts of FRB 121102 found by Spitler et al. (2016) were separated by as little as 23 secs in some cases, and were not seen to be periodic. This is too short to be an orbital period, and thus the pulsar with orbiting companion model cannot easily explain the source of FRB 121102.

Cosmological and catastrophic theories

The final class of FRB progenitor models assume that the majority of the DM contribution comes from the radio signals traversing large distances through the IGM. These require extremely high luminosities if they originate so far away, and the models producing FRBs in this case are often catastrophic in nature. These would obviously not be seen to repeat in the same way as FRB 121102 (Spitler et al., 2016), but certain FRB host galaxies could perhaps be 'FRB-loud', producing more than one catastrophic FRB signal.

The merger of two compact objects is a popular theory for FRB progenitors, possibly due in part to the potentially exciting prospect of linking FRBs with gravitational wave signals, which have recently been observed by LIGO (Abbott et al., 2016b). White dwarfs are too large to produce signals as short as those seen in FRBs, according to their dynamic timescales (defined as $t_D = (2r^3/GM)^{0.5}$, with values around 10s). However, Kashiyama, Ioka & Mészáros (2013) proposed that coherent emission from the polar regions of a large, strongly magnetised white dwarf shortly after its birth via the merging of two smaller white dwarfs, could explain the FRB phenomenon. Such systems appear to be a subset of the white dwarf population (Liebert, Bergeron & Holberg, 2003; Kashiyama, Ioka & Mészáros, 2013). The coherent emission, released shortly after the merger, would need to be heavily amplified by the magnetic fields (caused by the differential rotation throughout the body) to supply the required luminosity. This complex mechanism can be made to be consistent with the FRB event rate, total energy, and pulse duration, but would not be expected to repeat due to the rapid spin down of the white dwarf.

After white dwarfs, another logical option for energetic mergers are neutron star (NS) mergers. Wang et al. (2016) argue that an electromotive force induced by the inspiral of two neutron stars could accelerate electrons, emitting coherent radiation in the style of

an FRB. Lipunov & Pruzhinskaya (2014) found consistency between the predicted rates of NS+NS mergers and the observed FRB event rate. In this scenario, the radiation for FRBs would be produced during the collapse of a highly magnetised, rapidly rotating, short-lived object, immediately after the merging of the neutron stars (Lipunova, 1997; Pshirkov & Postnov, 2010). However, this process would naturally produce radio signals at lower frequencies, which have not been observed so far (Karastergiou et al., 2015). Additionally, Totani (2013) proposed that FRBs could result from the magnetic braking induced when the different magnetic fields in a NS-NS merger coalesce. Conversely, it has been argued that the radio emission from such a process would not be able to escape the expected dense ejecta surrounding the system, or the dense baryon wind at the poles (Zhang, 2016b).

An extension of the NS-NS merger scenario, or perhaps formed directly after a supernova, is the possible existence of a supra-massive rotating neutron star. Such a star has a mass larger than theoretically allowed, but is sustained from gravitational collapse into a black hole by its rapid rotation. Magnetic braking would eventually spin down the star until it collapses, on a timescale which could range from seconds (Zhang, 2014) to millions of years (Falcke & Rezzolla, 2014). During the collapse, magnetic field lines could snap energetically, accelerating electrons in bunches to generate the required radio emission (Ravi & Lasky, 2014; Zhang, 2016b). An FRB produced in this way might be expected to show a short GRB at some time preceding the radio signal, i.e. during the initial merger, from seconds up to several million years earlier. The delay between the formation of the supra-massive star and the collapse, if long enough, could allow the surrounding ejecta to disperse, thus allowing the radio pulse to propagate beyond the ejecta.

The merging of two black holes (BHs) is not directly expected to produce any electromagnetic radiation (Abbott et al., 2016a). However, Zhang (2016a) argued that if one of the merging BHs is highly charged, a magnetosphere would form and produce an FRB.

Such emission is expected to be isotropic rather than beamed (Zhang, 2016b), requiring large amounts of energy (since beaming lowers the energy requirements, by reducing the volume in which the radiation must propagate). However, we know such large amounts of energy can be emitted from merging black holes, at least in the form of gravitational waves (Abbott et al., 2016b). Such a detection might again be associated with a GRB, strengthening the case for collaborative searches in radio, gamma-rays and gravitational waves.

Almost 40 years ago, Rees (1977) predicted that the annihilation of miniature black holes could produce short, bright radio pulses. If a BH drops below a certain critical mass, it will evaporate into relativistic electron-positron pairs, which interact with ISM magnetic fields and produce a burst. However, the predicted energies involved in this scenario are not sufficient, at the assumed distances, to provide the observed high fluxes (Keane et al., 2012).

Whilst supernovae themselves are not considered possible progenitors for FRBs, the interaction of a SN shockwave with a companion neutron star has been proposed as a possible source (Egorov & Postnov, 2009). The shock impact on the magnetosphere of a binary neutron star could produce an extensive magnetospheric tail, which eventually reconnects violently, accelerating charged particles and forming a pulse of radio emission. However, a progenitor process like this would link FRBs to the detection of supernovae, and no such observational evidence for this exists thus far (Petroff et al., 2016). In addition, Kulkarni et al. (2014) showed that the strong winds emitted by massive stars before core-collapse supernovae would create significant nebulous material around the system such that the radio signals proposed by Egorov & Postnov (2009) are unlikely to propagate beyond the local region.

Finally, an exotic source of FRBs has been proposed in the form of superconducting cosmic strings (Vachaspati, 2008). Apparently, kinks in these current-carrying strings could generate strong bursts of radio emission during oscillation. However, this source

of FRBs predicts linearly-polarised emission (Cai et al., 2012), which is rarely observed (e.g. Petroff et al. 2015a).

Since only 17 independent FRB sources have so far been published (Petroff et al., 2016), it is obvious that further observations of FRBs and detailed follow-up at other wavelengths are needed to confirm the true origin(s) of FRBs. Nevertheless, some progenitor theories currently seem more likely than others, namely magnetar hyperflares and giant pulses from neutron stars, since they match the most observable properties. It may in fact transpire that multiple different classes of FRB exist (Li et al., 2016b).

5.1.3 How might FRBs be useful?

Understanding and utilising FRBs could lead to a much greater understanding of numerous cosmological unknowns. Examples include the ingredients of the IGM, its turbulence and magnetic properties (Macquart & Koay, 2013), the cosmic reionisation model (if FRBs can be studied at both low and high redshifts; Inoue 2004 and Deng & Zhang 2014), locating the 'missing baryons' (Ioka, 2003; McQuinn, 2014) and even the universal equation-of-state parameters related to dark energy (Zhou et al., 2014). Most of these breakthroughs require an independently measured distance to an FRB. This, along with some information about the location of the substantial scattering screen relative to the burst source and the observer (Masui et al., 2015), should allow us to infer the free electron density in the intergalactic medium.

A good example of this is the recent discovery by Keane et al. (2016), who claimed to have associated FRB 150418 with a host galaxy, and thus determine a redshift and distance. This allowed them to measure the density of ionised baryons in the IGM, and therefore confirm the presence of the 'missing baryons'. Unfortunately this discovery has been solidly refuted (Williams & Berger, 2016). This is discussed further in Section 5.8.

Even if FRBs turn out to be unsuitable as standard candles, cosmic rulers or probes

of the IGM, understanding their origin will also help develop our understanding of stellar and galactic evolution. More information about the FRB population can be gained from further detections, but great progress will surely be made if we are able to confidently discover a host source for FRBs.

5.1.4 Using ULTRASPEC to further FRB science

In order to investigate the interesting cosmological impacts and progenitor origins of FRBs, reliable redshifts must be measured. This itself requires precise localisation of FRB host galaxies. As mentioned above, recently Keane et al. (2016) claimed to have successfully associated an FRB with a visible host galaxy at a redshift of z=0.492 via a slow radio transient coincident with the original FRB direction. However, this claim has already been refuted quite convincingly (Williams & Berger, 2016), and we discuss this in more detail in Section 5.8. Thus we are yet to see a conclusive detection of an FRB host source, and the cosmological impacts of FRB studies therefore remain unanswered.

The simplest and most accurate way to measure a redshift is through optical spectroscopy. Given the large positional uncertainty in FRB detections, the association of an FRB with an optical counterpart or host galaxy is vital. It is still highly unclear whether or not we should expect emission at optical frequencies from FRBs. This very much depends on the progenitors of the bursts, as well as their total energies and distances. Yi, Gao & Zhang (2014) considered the high energy, optical and radio counterparts that might be expected to be associated with an FRB generated by the shock synchrotron emission from a standard fireball of a given total energy and Lorentz factor. They considered both forward shock and reverse shock scenarios, following the standard GRB afterglow model, which might be driven by the collapse of a supra-massive neutron star, for example. They found that only in the most nearby and energetic FRBs might we expect to see optical emission from forward shocks, and even then only at faint magni-

tudes (>18th magnitude). Reverse shocks might produce brighter emission (up to 13th magnitude in the most extreme cases), but this might only live for a few minutes after the FRB. The emission mechanism considered by Yi, Gao & Zhang (2014) may not be appropriate for FRBs, but it gives a flavour of the kinds of emission we might expect to observe at other wavelengths, in the most extreme cases.

An optical counterpart to an FRB could therefore be highly elusive, if it should exist at all. Nevertheless, the detection of such counterparts would help enormously with the precise localisation of the burst, and for this alone they are worthy of 'target of opportunity' status. The rest of this Chapter describes our attempts at optical follow-up of FRBs.

5.2 Parkes Optical Followup Collaborations

In September 2014, we entered the SUPERB collaboration for FRB follow up, led by Evan Keane. SUPERB (SUrvery for Pulsars and Extragalactic Radio Bursts) is an ongoing observing project using the 13-beam receiver at the 64m Parkes radio telescope, centred on 1.4 GHz (Keane et al., 2016). It aims to discover pulsars and FRBs in real time using recently optimised data processing techniques, and has so far missed no burst detections that are found in slower, more detailed analyses of the collected data. We joined the collaboration offering observations with ULTRASPEC at the TNT (Dhillon et al., 2014). In return for immediate notification of new FRB detections and their best measured positions, we would endeavour to follow up the FRB field in the optical as quickly and as deeply as possible, and supply any detections of afterglows, or detection limits, to the collaboration for publication as a group.

In addition to SUPERB, we also joined a very similar collaboration with the same procedures and end goals in February 2015, coordinated by Sarah Burke-Spolaor and named P879. This project is essentially the same as SUPERB, except that P879 deliber-

ately searches for FRBs a little further from the Galactic plane than SUPERB (which is also searching for pulsars, and thus points closer to the plane). An additional advantage of the P879 project is that they were able to supply us with precise times and rough sky coordinates for the FRB search observations, such that we could plan in advance when we might expect to receive an observable real-time trigger.

At one point the idea of directly and simultaneously shadowing the Parkes pointing was discussed with the SUPERB team. This might allow simultaneous detection (or non-detection) of an optical counterpart to an FRB. Considering the very short nature of the radio bursts, it is conceivable that the optical counterpart to these events, if any such component exists, might be equally short lived (Yi, Gao & Zhang, 2014). FRBs might offer an optical afterglow much like those of GRBs, which are often short-lived (seconds to hours) and fast-fading. There is also the possibility that the optical signals from whatever events generate FRBs could actually arrive at Earth before the radio signals (since the radio signals are dispersed and delayed as a function of frequency, with lower frequency photons arriving later than higher frequency photons). We decided not to attempt simultaneous shadowing of Parkes fields with ULTRASPEC (or with *pt5m*, which was also considered as a potential FRB follow-up facility), purely because the field of view of all 13 beams at Parkes is around $2^\circ \times 2^\circ$, meaning we would have needed to conduct 225 tiles of the ULTRASPEC field of view ($207 \times$ *pt5m* tiles) to cover this area.

Nevertheless, we have considered other possible options for optical shadowing of Parkes. At one stage, we discussed the possibility of mounting a small, wide-field reflecting telescope on the Parkes dish itself, such that it always points in the same direction. This turned out to be unfeasible since Parkes doesn't smoothly track the diurnal motion of the sky, but instead moves in quite large steps which would render optical images useless. We also considered the construction of a small but dedicated wide-field robotic telescope, somewhat similar to *pt5m*, at the Parkes site or nearby, which would be slaved to observe the same field as Parkes. We discuss this idea further in Chapter 6.

In general, most FRB detections from Parkes come with a positional uncertainty of $14'$, although if the burst is detected two beams, the positional uncertainty could be larger. This means that even in the best (and most common) case of a single beam detection, multiple tiles are needed in order to cover the entire error region. With an uncertainty of $14'$, and an ULTRASPEC FoV of $7.7' \times 7.7'$, four slightly overlapping tiles would be needed. If the uncertainty were $28'$, 16 tiles would be needed. To collect the most light and observe as faint sources as possible, we select wide-band filters with high throughputs. To begin with, this was r', though later we opted to use the $KG5$ filter, which has a wider throughput, although the calibration of fluxes in this filter is more difficult (see Chapter 2). If the delay time between the burst detection and our optical follow-up were to be short (few minutes), it would be worth attempting high cadence observations (e.g. few seconds up to a minute), however usually the delay is several hours, and so we opt for the longer exposures of up to 10 minutes, with multiple frames (at least 3) taken of each tile.

After the initial observations, which are always requested immediately after receiving a trigger, and conducted as soon as possible (weather and visibility permitting), it is important to conduct another observation of the field some time later. Re-visiting the field and observing the same sources then allows for a comparison study to be conducted, which should show any variability or new or disappearing sources.

So far we have received a total of 5 near-real-time triggers of FRBs. Two came from the SUPERB and P879 collaborators but outside of official operation of those projects. Three came from SUPERB. Of these five, we were able to obtain useful follow-up observations of two of them. The other three were received during periods of bad weather, and where observing was conducted, the data are not usable. We give details on these five triggers below.

5.3 FRB 140515

We received an FRB alert email just 1.5 hours after detection of an FRB on 2014-05-14. Unfortunately by this time the TNT had been closed for several weeks already due to the monsoon season. This FRB has been published with follow-up at many wavelengths with other facilities in Petroff et al. (2015a). No significant detections of possible counterparts were made, leading to the rejection of certain progenitor models for that particular FRB at least, such as supernovae or long GRBs.

5.4 FRB 150215

At 20:41:41 UT on 2015-02-15, the Parkes radio telescope observed a new FRB at RA: 18:17:27.0, DEC -04:54:15.3 (J2000) with a DM of 1105 pc cm^{-3}. The event was discovered almost in real-time, and we received an email alert requesting follow-up observations two hours after the burst. Incidentally, this burst was not observed during allocated P879 or SUPERB time, but was found during other observations. Our collaborators in Thailand chose to trigger optical follow up, because they had a convenient scheduling gap. Unfortunately ULTRASPEC was not in use at the time, and the observers instead used the Thai ARC 4K camera.

5.4.1 Observations

Because the FRB trigger came through at the end of the night in Thailand, observations of the field could not begin until the following night. In addition, the field was only visible for the last hour before morning twilight, meaning the observations were conducted more than 25 hours after the burst was first detected.

The ARC camera has a single square CCD with 4096 pixels on each side, cooled to -110°C. The pixel scale is $0.13''$/pixel, giving a field of view of $8.8' \times 8.8'$. The

observations were conducted in the Bessell R filter, as this was considered to have the best throughput. The observations were conducted by Utane Sawangwit and Grittiya Pongsupa.

Because the error circle of this FRB was $30'$ (it was detected in two of the Parkes beams rather than one), nine tiled observations were made in order to cover as much of the field as possible. Each tile was observed at least three times, to reach a deeper limit and to allow the removal of artefacts such as cosmic rays. Each image had an exposure time of 60 seconds, meaning a total observation time of at least 3 minutes for each tile.

On 2015-04-14, 57 days later, the same field was re-imaged using the same camera and set up, in order to allow a search for any transient or variable objects in the field. Table 5.1 shows the tiling information for the two observing sessions. Note that the final tiles (7, 8, 9) observed in the second session were unusable due to the brightening morning twilight sky, meaning the total area covered by this study is considerably less than the full $30'$ error circle.

As can be seen in Table 5.1, the seeing for the first visit was much worse than the second. Unfortunately the poor conditions during the first observations severely limited the depth to which we could search for an optical FRB counterpart.

5.4.2 Data Reduction and Analysis

For each pointing, the individual frames were reduced as follows:

1. Subtract master bias made from a median combination of 31 individual bias frames
2. Divide by master twilight flat made from the normalised median of 6 flat frames
3. Align science frames using 5 selected bright stars
4. Stack science frames by averaging each pixel with sigma clipping
5. Calibrate astrometry using `astrometry.net`

Table 5.1: Useful tiles observed with the 4K camera at the TNT, covering the field of FRB 150215 on the first and second visits. The coordinates listed correspond to the centre pixel of each tile and are relative to J2000. The FWHM is the mean measured FWHM across the image.

Tile No.	**RA**	**DEC**	**FWHM (″)**	**No. of frames**
	First	**Visit**		
1	18:17:27.9	-04:54:14	3.0	6
2	18:17:52.0	-04:54:15	3.0	5
3	18:17:52.0	-04:50:15	2.9	3
4	18:17:32.0	-04:50:15	2.8	3
5	18:17:11.9	-04:50:15	2.8	3
6	18:17:11.9	-04:55:15	2.7	3
7	18:17:11.8	-05:01:15	2.5	3
8	18:17:31.9	-05:01:16	2.6	3
9	18:17:52.0	-05:01:16	2.5	3
	Second	**Visit**		
1	18:17:27.3	-04:54:16	1.1	5
2	18:17:59.4	-04:54:19	1.2	5
3	18:17:59.4	-04:46:19	1.2	5
4	18:17:27.3	-04:46:17	1.1	5
5	18:16:55.2	-04:46:16	1.3	5
6	18:16:55.1	-04:54:16	1.0	5

Images of all 9 tiles from the first visit are shown in Figure 5.6, whilst the 6 usable tiles from the second visit are shown in Figure 5.7. The difference in seeing conditions between the two observations is clearly seen in the images.

To search for an FRB counterpart after stacking the individual frames, we needed to compare the objects detected in the first visit with those seen in the second visit. One method to do this, and which can show transient or variable sources clearly, is difference imaging, in which one image is subtracted from another. This method is useful in crowded fields, or for automated transient detection pipelines. We used difference imaging in a later analysis (see Section 5.6.2), but it was unsuitable for these observations because the tiles in the first and second visits were not directly overlapping. Instead we employed another technique. We used SourceExtractor (SE, Bertin & Arnouts

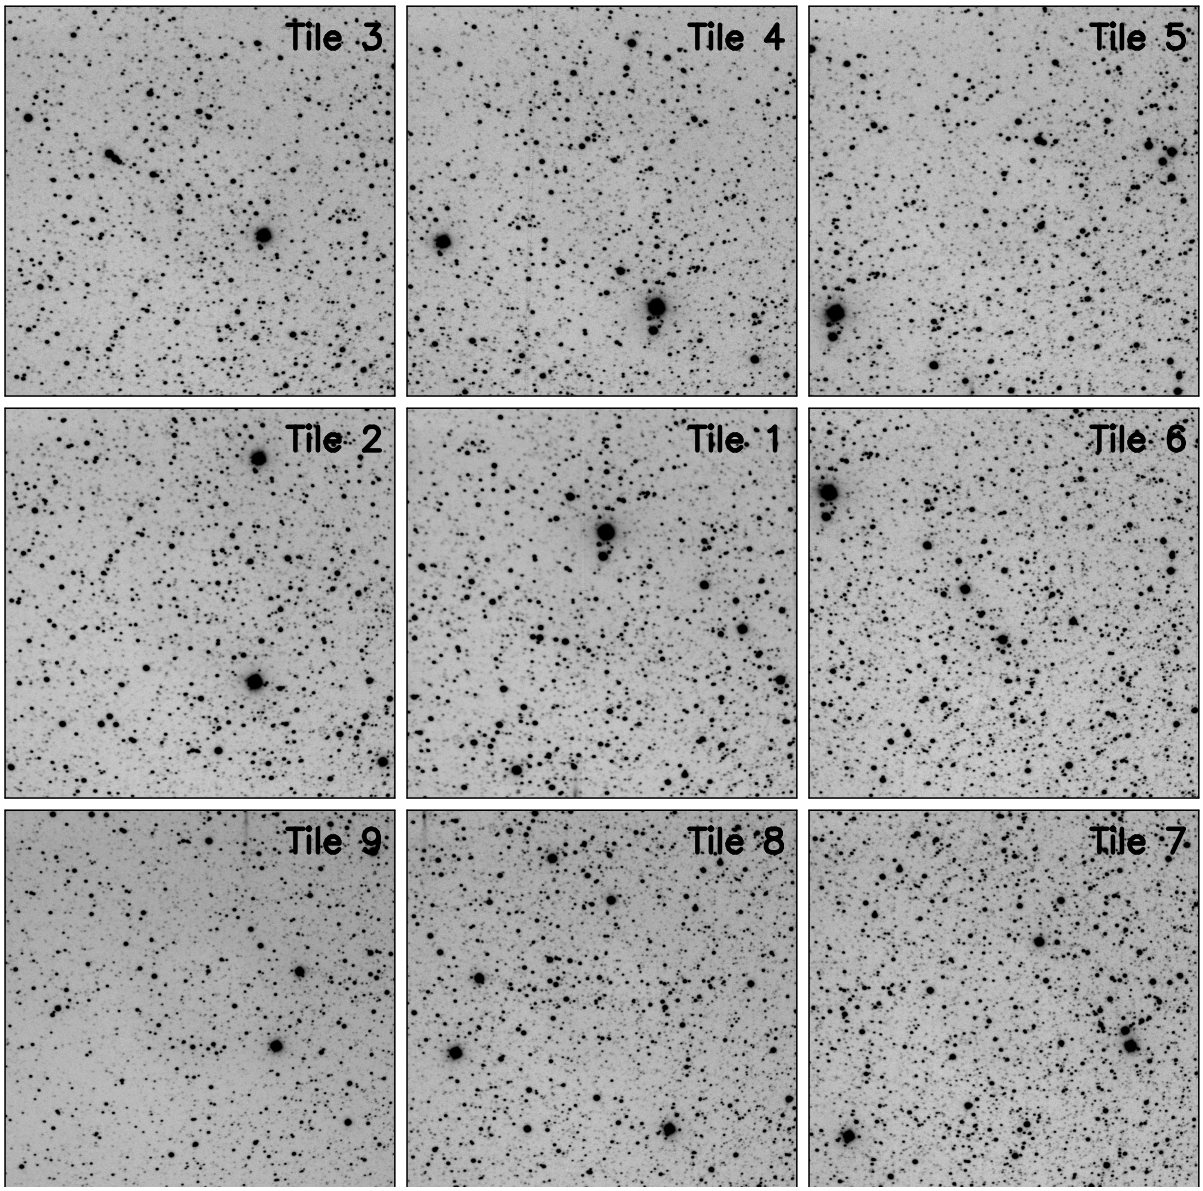

Figure 5.6: The 9 tiled images from the first observation of the FRB 150215 field. The 3′-4′ overlap is evident in some tiles. North is up and East to the left.

1996) to create a catalogue of all sources in all the tiles. These catalogues can then be compared to each other for variable or transient sources.

SE first analyses the background levels of the image by splitting the image into many boxes of a given size, defined by the parameter BACK_SIZE. For this analysis SE was found to accurately extract the most sources when this parameter was set to 32 pixels. For each box, SE measures the mean pixel value (or median when an image is crowded),

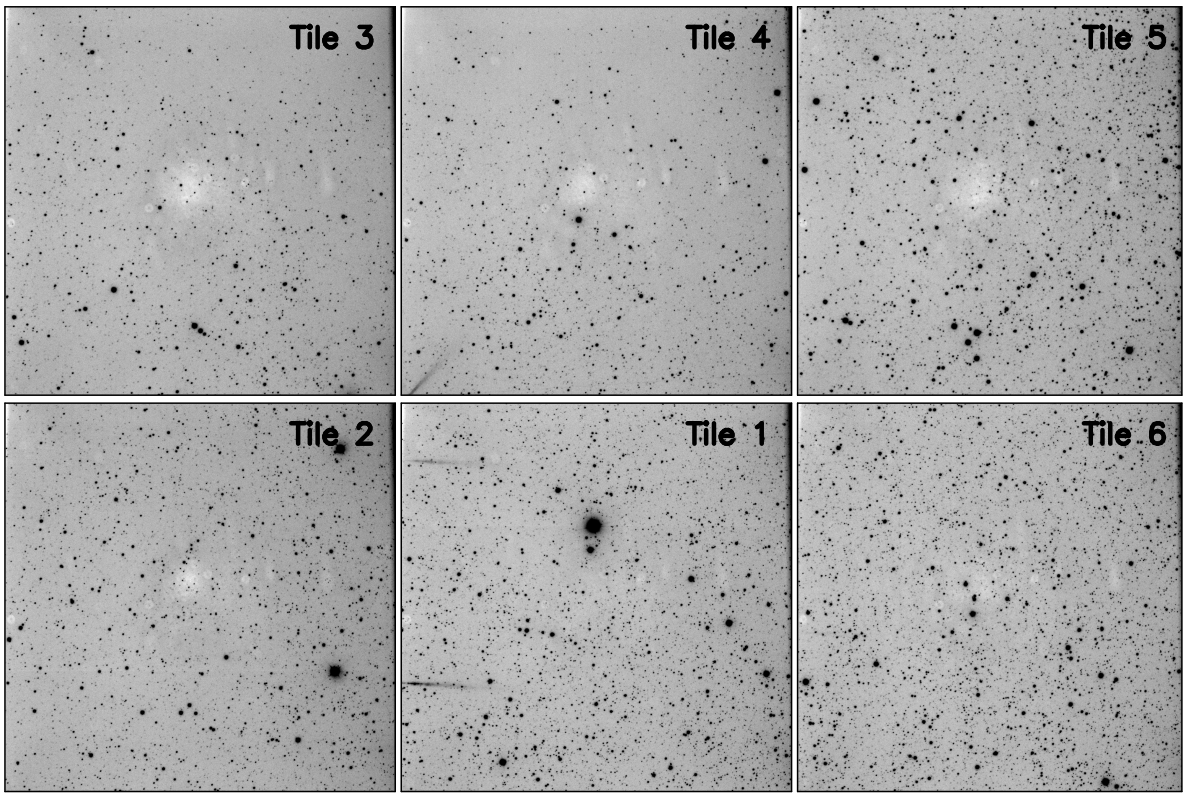

Figure 5.7: The 6 usable tiled images from the second observation of the FRB 150215 field. There is less overlap between these tiles than in the first visit shown in Figure 5.6. The lighter patches in the centre of some tiles is an artefact of the camera, presumably highlighted more clearly by the bright twilight sky background. The order of observations was 1, 6, 5, 4, 3, 2, and the lighter patch seems to grow in strength with time, as the sky brightens. North is up and East to the left.

as well as the RMS noise (σ) across the box. With this information for every box in the image, SE then calculates a smoothed function for the background level and σ for every pixel in the image. If a box contains a bright star, the background level may be somewhat skewed, meaning that unfortunately SE is less efficient at detecting faint sources close to bright ones.

To extract candidate sources, SE then examines every pixel and highlights any that have values higher than a certain threshold above the mean, dictated by a multiple of the measured σ for that pixel. The multiple factor is dictated by the DETECT_THRESH parameter. The lower this parameter, the more sources are extracted, though of course

after a certain point SE begins to extract unreal sources which are simply noise. By analysing the extracted sources and comparing them to the image for a selection of different values, we found that the lowest value for DETECT_THRESH which extracted the most sources and the least unreal sources (classified as such by eye) was 0.85σ. I believe this value was found to be so low because of the poor image quality. Many genuine sources appear very faint and blurry in the images, meaning that the value of σ is likely already quite large, and that most faint sources may otherwise be missed if the threshold was set higher.

Next, for SE to classify something as a genuine source, candidate sources must display multiple adjacent pixels which are all above the detection threshold. The minimum number of adjacent pixels needed to constitute a source is defined by the DETECT_MINAREA parameter. If this parameter is set too low, SE will pick up many unreal sources such as hot pixels and cosmic rays. If it is set too high, SE will ignore some genuine sources which are faint and compact. We found that the best value for DETECT_MINAREA in extracting the most real sources and least unreal sources was 11 pixels. This parameter is expected to depend on the seeing, but 11 pixels appeared to also be suitable for the second visit images, presumably because the majority of obvious artefacts (due to cosmic rays or hot pixels) are smaller than 11 pixels.

At this point, SE has compiled a preliminary catalogue of sources, but many of them will be blends of two or more stars which are close to each other. This is especially relevant with our images of the FRB 150215 field, because the field is close to the Galactic plane and very crowded.

To deblend sources in the preliminary catalogue, SE takes each source and tests for blending by searching for multiple 'branches' in a 'tree'. First a number of levels are defined between the threshold level and the maximum pixel value observed in the source. The number of levels is set by the DEBLEND_NTHRESH parameter, which is hard-wired to be 64 levels. SE then builds a 'tree' of objects within the source, creating a new

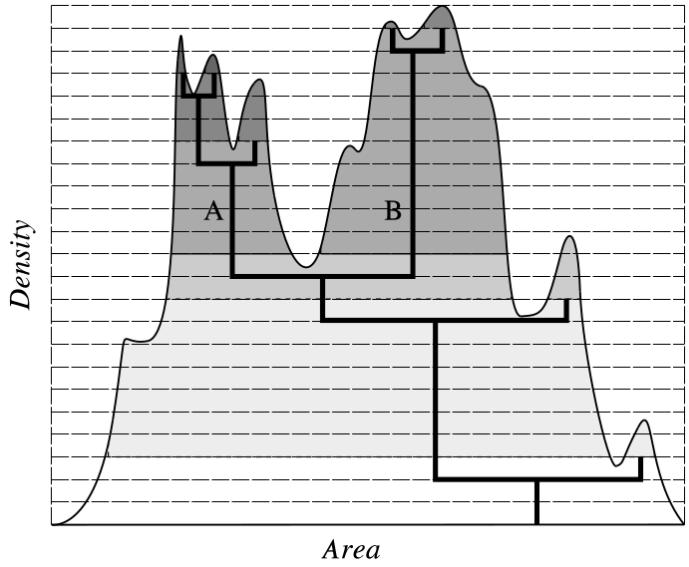

Figure 5.8: Schematic of the 'branches of a tree' deblending procedure for SE, reproduced directly from Bertin & Arnouts (1996). In this hypothetical case, two sources A and B are extracted, but other shallower peaks do not contain a large enough fraction of the total counts to be considered separate sources. The counts in these 'false peaks' are then treated as contributions to their nearest progenitor source.

'branch' every time there are pixels above a threshold level that are separated by pixels at a lower level. To be considered a genuine separate source, a branch must contain a certain fraction (defined by DEBLEND_MINCONT) of the total counts of the original source. We found an appropriate value for this parameter to be 0.005 (i.e. 0.5%). A schematic example of this process is shown in Figure 5.8 for ease of understanding.

Finally, SE attempts to clean the list of sources of any artefacts or unreal sources, due to e.g. diffraction spikes from bright stars, by checking to see if sources would still have been detected if their neighbours did not exist. The contributing counts from all neighbours are removed by fitting Moffat profiles to them and subtracting the modelled counts. In this way artefacts related to particularly bright sources are often removed. The resulting catalogues, extracted for each pointing, then contain as many real, faint sources, and as few unreal sources as possible. Each extracted source comes with a measured FWHM, as well as a flux, calculated using the pixel counts contained in the

source, with the background levels subtracted.

With catalogues of sources extracted from each pointing by SourceExtractor, we then merged together all of the source catalogues from the tiles observed on the first visit to the field, and we did the same for the tiles observed on the second visit. Where the tiles overlap (by about $3'$ in the tiles from the first visit, and by about $1'$ for the second visit), we removed any duplicate sources in the catalogues which are located within a given distance from each other. For the first visit, when the seeing was worse, we removed duplicates within $3''$, and for the second visit we removed duplicates within $1.8''$.

Next we made cuts to the catalogues in RA and DEC, so as to only consider the extracted objects which lie within the overlapping region between the two observations (because not all tiles were observed in the second visit). We selected only objects with RA between 18:16:54.60 and 18:18:08.16, and with DEC between -04:58:27.3 and -04:46:03.4. This represents an area of $18.4'$ by $12.4'$ centred on 18:17:40, -04:51:55 (J2000). Unfortunately this only covers about one quarter of the error circle of the original FRB position.

We made some final cuts which removed more artefacts and blends. We selected only those sources which have a FWHM above 8 pixels in the first observation (we assumed that objects with FWHM less than this were not true sources). We also rejected any sources that were flagged directly by SE as being likely to be blended. This still left 1,500 sources in the first visit catalogue, and 5,500 in the second.

Finally, by comparing the extracted fluxes of several identifiable stars in the field to the known R-band magnitudes of these stars given by the USNO-B1.0 catalogue, we estimated the local zeropoint for the observations on the first visit to the field. We then used this zeropoint to convert the extracted fluxes to R-band magnitudes for sources from both visits.

An estimate of the limiting magnitude for detection of sources in this study can be found using a histogram of the source detections as a function of magnitude. The

turnover in the number counts of sources for a given magnitude in the first observations (which were less deep than the second, because of the worse seeing) had an R-band magnitude of 21.3. We consider this a rough limiting magnitude for detection of an FRB counterpart in these observations.

5.4.3 Results

To search for transient or variable sources, we compared the catalogues of sources from the two visits in two ways. Firstly, we looked for objects which appeared in the first visit but not the second. Any fast-fading transients visible in the first observation are likely to have disappeared entirely by the time the second visit was made. Secondly, we merged both catalogues and selected only the objects which appeared in the same place (to within $2''$) in both observations. We then used this merged catalogue, which has measured fluxes for both epochs, to search for highly variable objects.

Initially our cross-matching revealed 123 sources which appeared in the first catalogue but were not accurately matched to a source in the second catalogue. These 123 sources were inspected visually in the images, and were all found to be either artefacts in the first images, or present as sources in the second images but had not been picked up by SE (or, more likely, had been rejected as blends). Thus we found no genuine objects in the first observations that were not present in the second observations conducted 57 days later (i.e. we detected no transient sources).

In the cross-matched catalogue showing sources detected on both visits, we found 1788 sources. For the second observations, there was a general trend for sources to appear fainter than in the first, which we attribute to a changing local zeropoint, probably related to differing airmass and extinction. The change in average magnitude from one epoch to the next is not important - what we need to look for are sources which have changed in brightness by an amount significantly different to the average. A subtlety here

is that a faint star (e.g. 21st mag) changing by 0.1 magnitudes would be unsurprising, due to the inherent Poisson noise in the measurement. Contrastingly, a change of 0.1 magnitudes in a bright star (e.g. 17th mag) would likely be above the 1-σ level for that source, and would suggest an intrinsic change in brightness.

To account for this dependence on brightness, we split the sources into 5 bins by magnitude; <18, 18, 19, 20, ≥21, and calculate a differential magnitude (*Visit1Mag* - *Visit2Mag*) for each source. There was large scatter in this measurement, with standard deviations ranging from 0.18 in the <18th magnitude bin, to 0.425 in the ≥21st magnitude bin. The scatter is likely due to crowding in the field and the considerable difference in image quality between the two visits.

A histogram of the magnitude difference (*Visit1Mag* - *Visit2Mag*) for all sources binned by magnitude is shown in Figure 5.9. It is interesting to note that the distributions in each bin are centred at different values, with brighter objects appearing to fade from the first visit to the second. We suspect that this is a result of the poor image quality in the first visit, whereby many bright sources are observed with blended additional contributions from nearby faint sources. In the second visit they are more cleanly separated with less contamination from other sources, and therefore these sources appear fainter.

Defining the standard deviation of the scatter in magnitude difference as σ for each bin, we then investigate sources which have differential magnitudes greater than 3σ above or below the mean. We find between 2 and 10 potential sources in each bin that vary by more than 3σ. A total of 29 sources were selected in this way, and we inspected each of them by eye. All of these potential varying sources could be explained by blending or edge of CCD effects, and none appear to be due to genuine variability.

We therefore made no detection of a potential optical FRB counterpart, in observations to a limiting depth of 21 magnitudes in the Bessell R-band, 25.3 hours after the first detection of the burst. Our contributions to the optical follow-up of this burst will

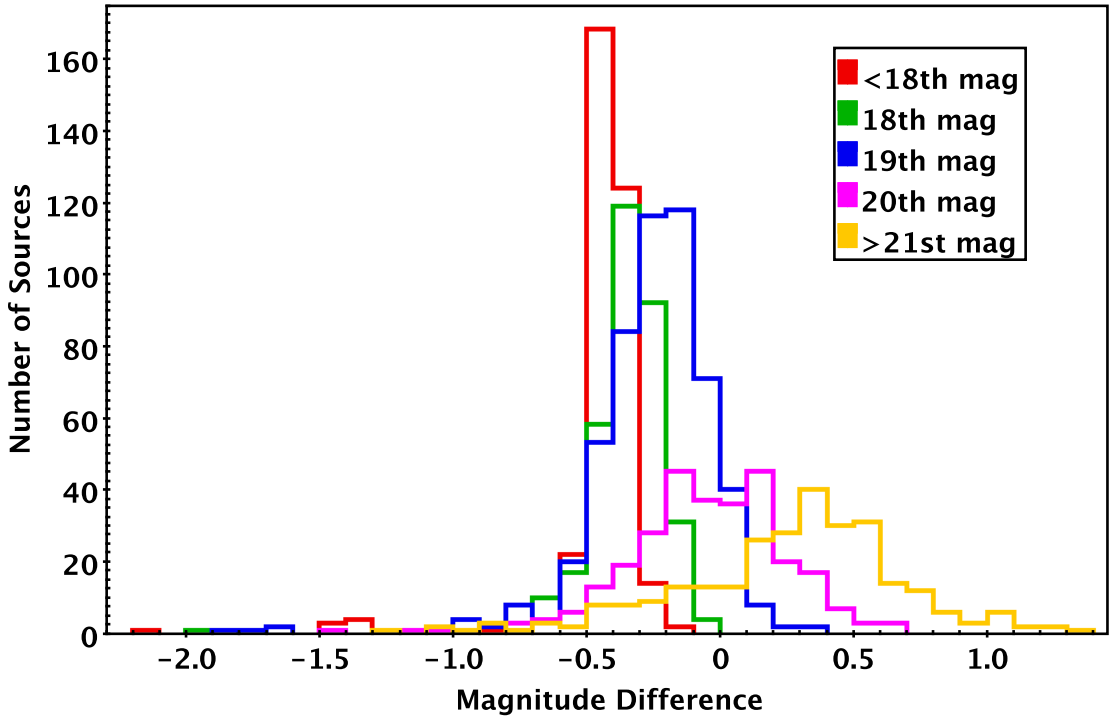

Figure 5.9: Histogram showing the distribution of magnitude differences between the two observations of the FRB 150215 field. The sources are split into 5 separate magnitude bins to investigate variability in each case.

be included in Petroff et al. 2016 (in prep.).

5.5 FRB 150418

Another follow-up campaign was triggered on an FRB detected by Parkes on 2015-04-18 at 04:30 UT. Observations were requested immediately, and observed that night with ULTRASPEC by the NARIT team, around 7.5 hours after the burst. Due to scheduling constraints, the observations were conducted in evening twilight. Unfortunately there was significant cloud coverage present, and the resulting images had a 3-σ limiting magnitude of only $g' = 16.5$ mags.

A second visit to the field was also conducted on 2015-05-11, but sadly these images

were even worse than the first visit, with the weather in Thailand deteriorating rapidly throughout the Spring. Variability analyses could not be conducted, and we were unable to contribute any data to the follow-up effort for this FRB.

It was particularly unfortunate that we were unable to make deep observations of this field, because this particular FRB was later published and claimed to be the first localised FRB detection with an optically confirmed host galaxy. Keane et al. (2016) used early radio follow-up observations to localise a long (6-day) transient source at 5.5 and 7.5 GHz with the Australia Telescope Compact Array (ATCA). Deep optical imaging with the SUBARU telescope found a faint galaxy ($r' = 23.45$) within $1''$ of the ATCA transient source. Later, spectroscopic observations were able to measure a redshift of the galaxy of $z = 0.492$, which ultimately allows a measurement of the density of ionised baryons in the intergalactic medium. Keane et al. (2016) found a value for the cosmic density fraction of the IGM of $\Omega_{IGM} = 0.049$, which independently and directly verifies the ΛCDM cosmological model when applied to WMAP observations of the ionised baryonic matter (Hinshaw et al., 2013), and thus also represents a direct detection of the 'missing baryons' (Bregman, 2007). If the claims by Keane et al. (2016) are correct, this is a momentous breakthrough in the study of FRBs and in astrophysics. However, the association of the FRB with the long radio transient and thus the host galaxy may be incorrect (Williams & Berger, 2016), and further observations are needed to clarify this. This issue is discussed further in Section 5.8.

5.6 FRB 151206

SUPERB at Parkes made another FRB detection at 07:30 UT on 2015-12-06, with a dispersion measure of 1909 cm^{-3}pc. This corresponds to an excess of 1666 cm^{-3}pc above the expected DM for that particular line-of-sight through the Galaxy, according to the NE2001 model of Cordes & Lazio (2002). We received an email alert request-

ing optical follow up of the FRB field around 25 hours later. At the time, a scheduled Sheffield/Warwick ULTRASPEC observing run was underway, and the observers in Thailand were happy to attempt observations of the field. I prepared finding charts and had them ready in time for the start of the night of 2015-12-07.

5.6.1 Observations

Unfortunately, the field of this FRB was quite close to the Sun, and only visible for about an hour at the very beginning of the night. With coordinates 19:21:25.6, -04:07:54.5 (J2000) it was setting quickly during evening twilight, meaning that some of the data collected has a bright sky background, and the rest was observed at high airmass (3-5), with possible vignetting by nearby trees surrounding the TNT.

Nevertheless, several usable images of four tiles were taken with ULTRASPEC in the *r'* filter. The field of view is $7.8' \times 7.8'$, and the four tiles (which overlap by up to $10''$ on each edge) cover the $14'$ error circle associated with the Parkes coordinates. Observations began at 11:30 UT, but the first frames suffer from too much twilight sky background, and the useful frames were taken from 12:00 to 12:25 UT, 29 hours after the original burst. 6 viable images of each tile were observed, with exposure times of 60 seconds.

Unlike FRB 150215, the tiles were observed symmetrically around the centre of the Parkes error circle, covering a total area of $15' \times 15'$. The same tilings were repeated 3 and 4 days later in order to search for variability, but the images on 2015-12-10 were not of sufficient quality to be useful, so we only analysed the data from 2015-12-07 and 2015-12-11. The average seeing during observations was $3''$ on both nights. The observations were conducted by Vik Dhillon, David Sahman and Somsawat Rattanasoon (Tim).

5.6.2 Data Analysis

As with FRB 150215, the individual images were bias-subtracted and divided by a flat field before further processing. Next, the individual frames were astrometrically calibrated using `astrometry.net`, before being stacked in IRAF by averaging the pixel values, with sigma clipping. Finally, each tile from the second visit was aligned with its equivalent tile from the first visit, using the image alignment program AliPy. Figures 5.10 and 5.11 show the stacked and calibrated images for each of the four tiles, on the first and second visit. The overlap between each tile is just $20''$.

To analyse the images and search for objects of interest, we employed two complimentary methods of image analysis; difference imaging and source catalogue comparison. Recall that difference imaging was not feasible for FRB 150215 since the first and second visit tiles were not of the exact same fields. The two methods used for FRB 151206 are described below.

Difference Imaging

We conducted a difference imaging analysis (Alard & Lupton, 1998) by subtracting the second epoch tiles from their equivalent first epoch tile. This was done using the High Order Transform of PSF And Template Subtraction (HOTPANTS) software of Becker (2015). Difference imaging is a useful tool for analysing multi-epoch data and searching for variable or transient objects. It is particularly useful in crowded fields (as is the case for FRB 151206), or for rapid transient detection programmes. For other situations it is less accurate in measuring variability, because the process includes subtraction of source counts, whilst increasing the overall sky background noise.

Difference imaging is essentially the subtraction of one image from another, after careful astrometric and photometric matching between the images has taken place. This is necessary to ensure stars which are intrinsically the same in both images are perfectly

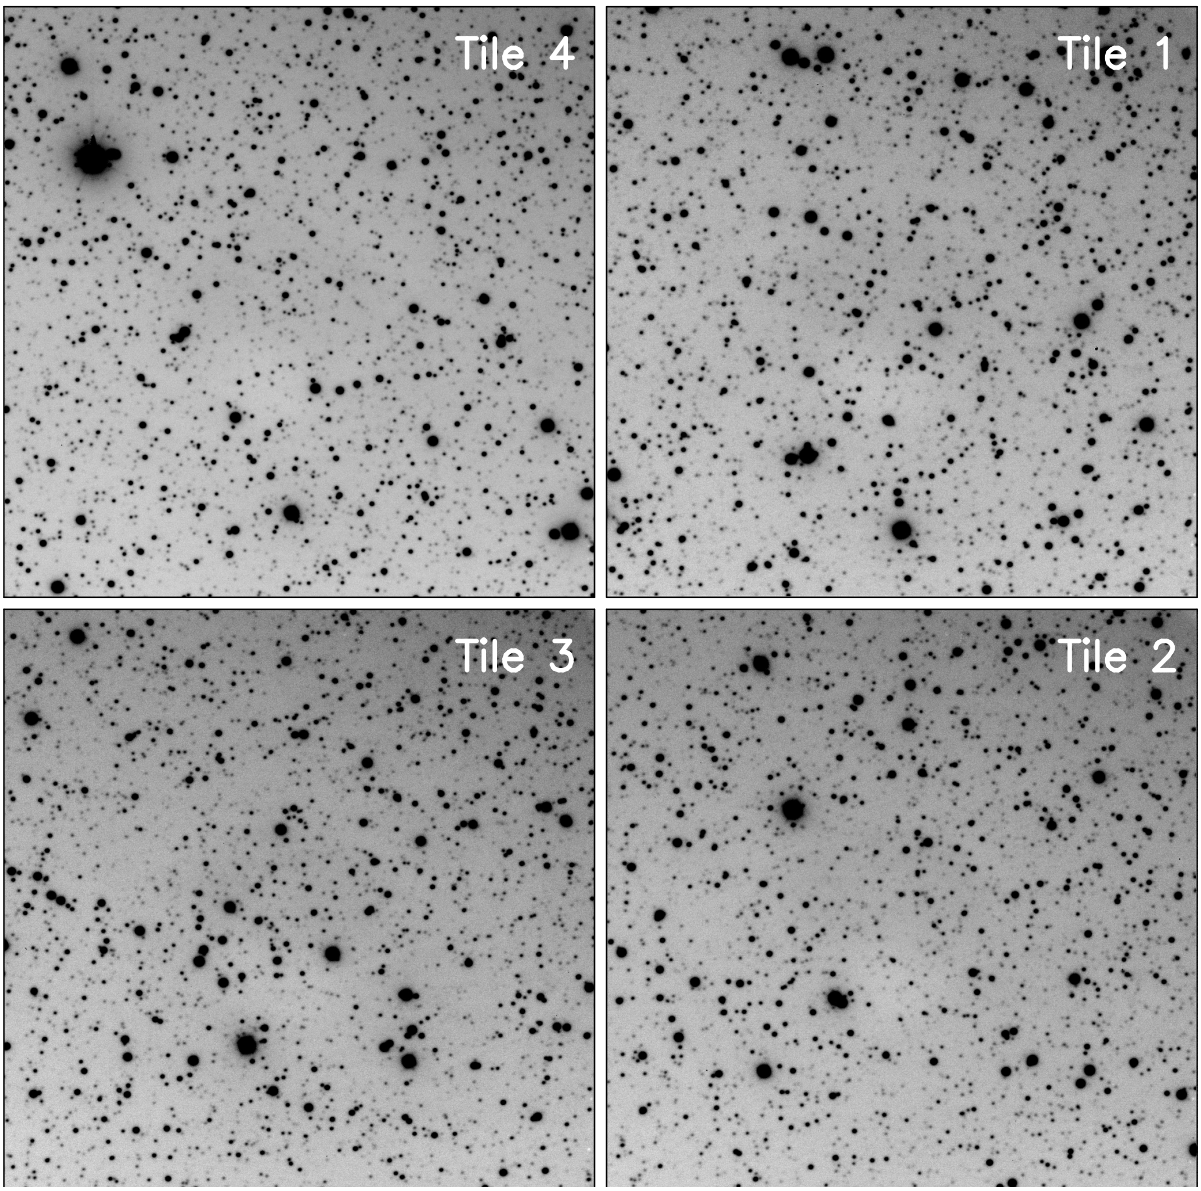

Figure 5.10: The 4 tiled images from the first observation of the FRB 151206 field. North is up and East to the right.

subtracted from one another. Otherwise the effects of varying image quality, seeing, atmospheric extinction, airmass, exposure time or detector efficiency, can skew this subtraction. The method works by building a kernel that minimises the difference between the PSFs and photometric intensity of sources in the two images. In general, one image will have worse seeing than another, or may have small coma-like trails due to tracking errors. The crisper image is then smoothed with a kernel such that the PSFs across the

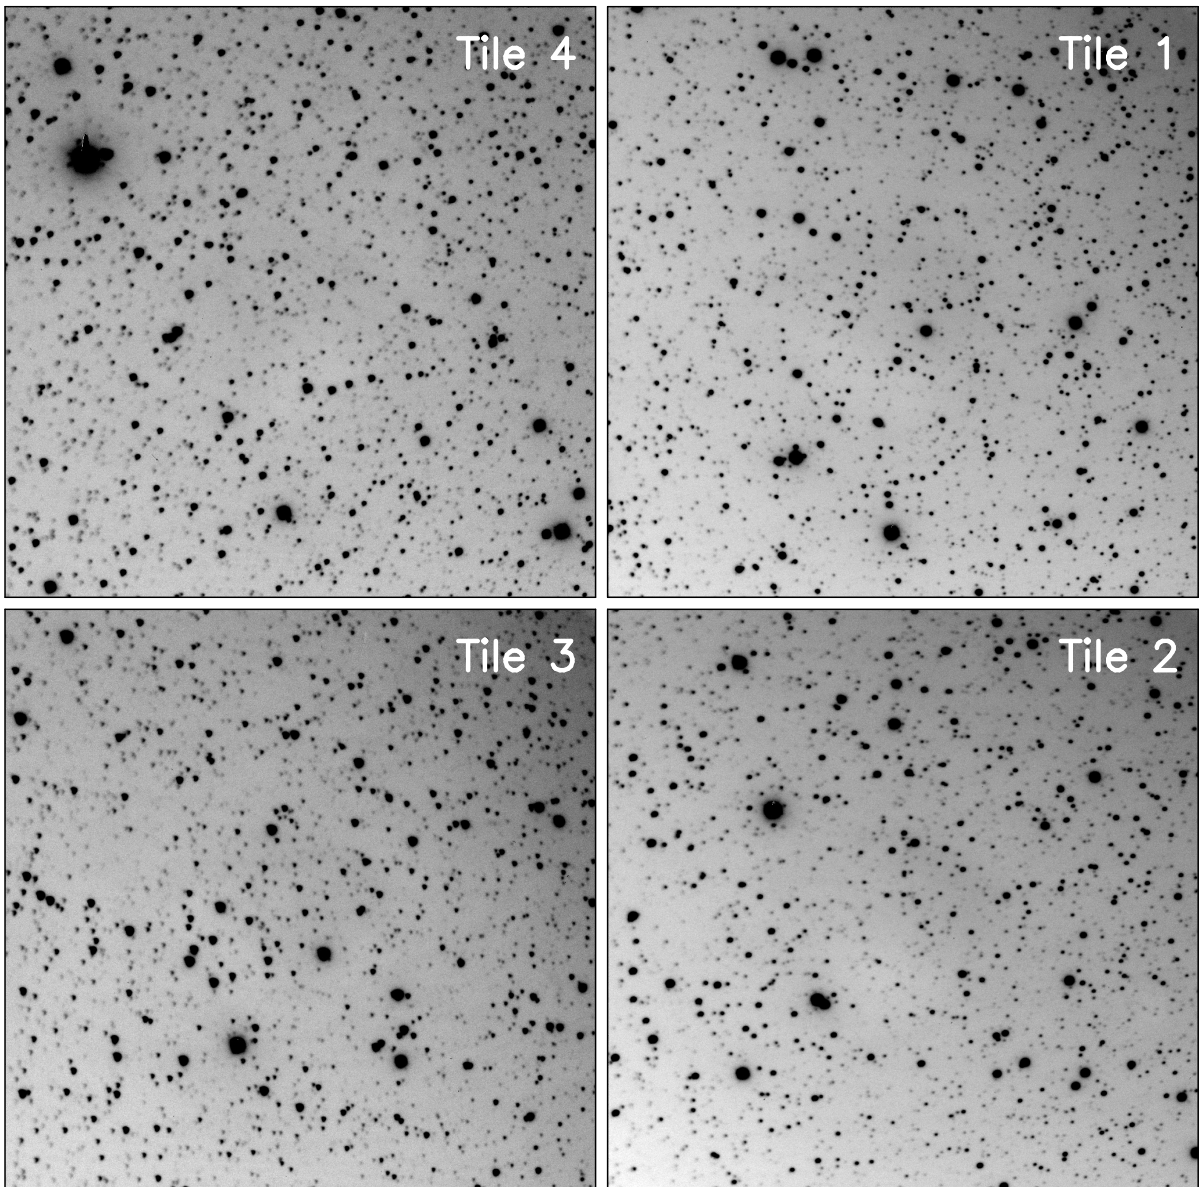

Figure 5.11: The 4 tiled images from the second observation of the FRB 151206 field. North is up and East to the right.

image match the other image with worse image quality. This allows neat subtraction of all non-variable sources, leaving behind only those sources which have genuinely changed in brightness or PSF shape. A section of a completed difference image is displayed in Figure 5.12.

The difference images were examined visually, and with a simple SourceExtractor analysis. We found 5 clear point-like sources in the difference images, all of which

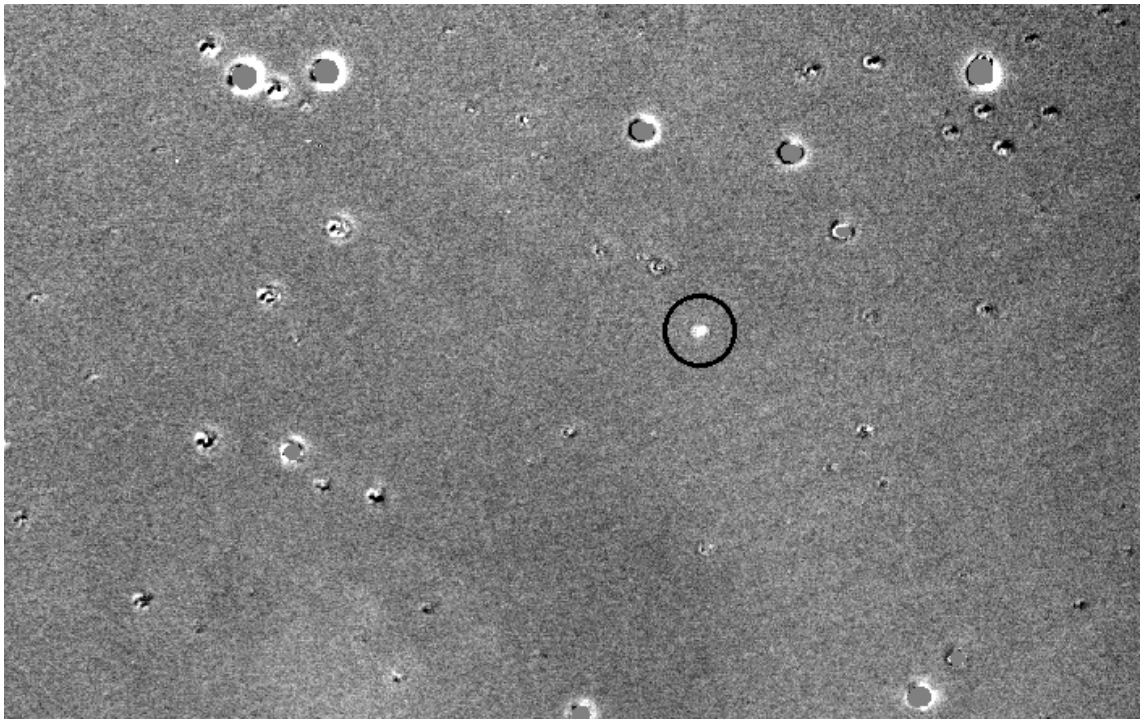

Figure 5.12: Cut-out showing $6.3' \times 4.0'$ of one of the ULTRASPEC difference images of FRB151206. Some sources have not been subtracted successfully, due to saturation or other linearity effects. One clear residual source is marked with a circle. This source was significantly brighter in the first image than in the second, and corresponds to the star UCAC4 430-099010.

correspond to known stars (i.e. there were no transient source detections). One was found to be brighter in the second observations than the first, while the other four are seen to fade from 2015-12-07 to 2015-12-11. The sources are analysed in detail in Section 5.6.2 and the results are discussed in Section 5.6.3.

Catalogue Extraction

In addition to difference imaging, we also employed the SOURCEEXTRACTOR (SE) catalogue extraction and comparison technique discussed in Section 5.4.2. After extensive testing of different parameters and visual inspection, we found the following to be the best values for extraction parameters in order to investigate the most sources: BACK_SIZE = 16, DETECT_THRESH = 0.8, DETECT_MINAREA = 10, DEBLEND_MINCONT = 0.0005

(i.e. 0.05%).

We also derived an internal zeropoint for the images using SE and the UCAC4 catalogue. Using TOPCAT, we matched extracted sources with UCAC4 stars, and calculated an *r'*-band magnitude from the UCAC4 *B*-*V* colours and SDSS transformations (Jester et al., 2005). These magnitudes were then compared with the SE extracted magnitudes, and the SE zeropoint parameter was tweaked to minimise the average difference between the two values, across all sources. For the first epoch, the internal zeropoint (ignoring extinction and airmass) was found to be $r' = 24.6$ and for the second epoch it was found to be $r' = 24.3$, both with a standard deviation of 0.3 magnitudes.

With the zeropoint parameter set, we then extracted catalogues for all four tiles on both epochs, and concatenated them with removal of duplicates where the fields overlap. To remove unreliable sources, we cut out any that were flagged by SE for containing saturated or nearly saturated pixels, or being too close to the edge of the image. Then we removed all sources with unreliable or large full-width-half-maximum measurements, which could have been cosmic rays, galaxies, or moving objects. This meant cutting out any sources with FWHM less than 4 pixels or greater than 20 pixels. A histogram of FWHM value from the first epoch after this cut is shown in Figure 5.13.

Next we investigated all sources which were cleanly extracted from one observation and which were not seen in the other. We initially found around 1000 sources which were seen in the first epoch but not in the second, and similarly approximately 1000 sources were seen in the second epoch but not the first. To trim this down to a manageable number, we cut out anything with a magnitude error above 0.2 magnitudes (corresponding to a signal to noise ratio of less than 5) or with a magnitude fainter than the detection limits discussed below in Section 5.6.2. This left around 400 sources from each epoch which, according to the catalogue analysis, were suspected of being transients.

These sources were examined by eye in the images from both epochs, and all of them could be explained by blending with other nearby sources (the most common cause for

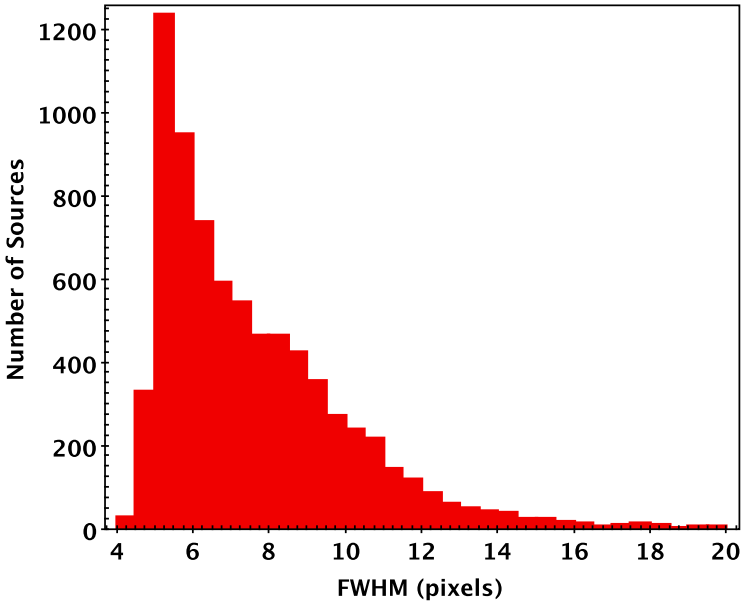

Figure 5.13: Histogram of FWHM measurements for sources extracted from images taken on 2015-12-07, after cutting out spurious sources below 4 pixels and above 20 pixels. The tail towards larger FWHM is likely due to the blending of multiple sources in the crowded field.

these anomalies), edge of CCD effects, hot pixels, or cosmic rays. This left no genuine sources showing transient behaviour.

Finally, we searched for sources which appeared in both catalogues and which varied in brightness between the two epochs. Again, we cut out sources with uncertainties above 0.2 magnitudes and with magnitudes fainter than the detection limits. The remaining sources (around 6000) were then split into 6 bins by magnitude; <17, 17, 18, 19, 20, ≥ 21. A differential magnitude was calculated, being simply the difference in magnitude from one epoch to another. As with FRB 150215, there was substantial scatter in this measurement, and more than would be expected from noise alone. The standard deviation ranged from 0.15 mags for objects brighter than 17 mag, to 0.4 mags for objects fainter than 21 mag. As well as the expected noise, the scatter is likely also due to crowding in the field, the varying bright sky background and the high airmass. A

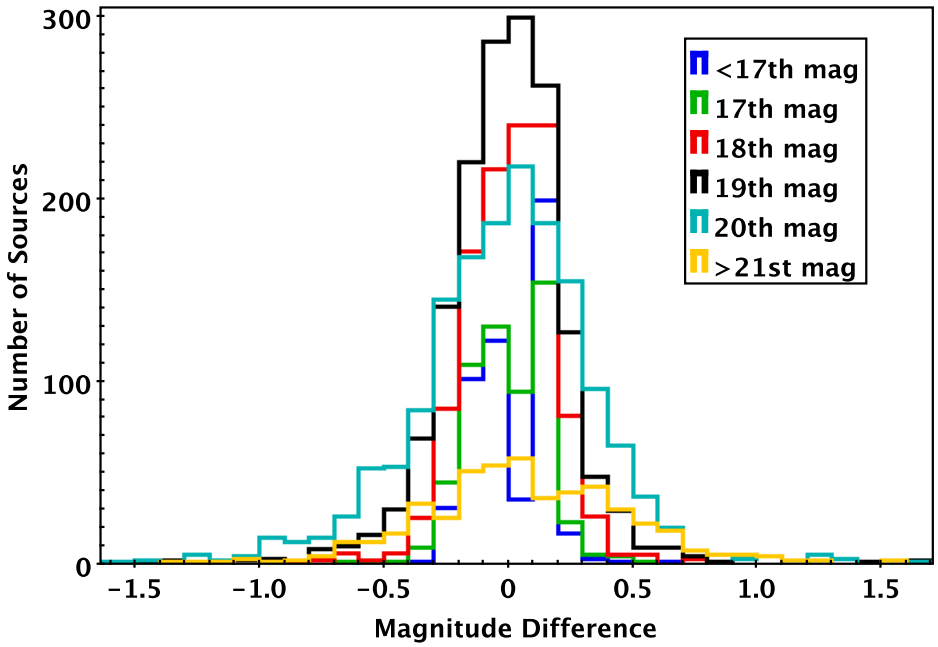

Figure 5.14: Histogram of differences in measured magnitudes from the observations on 2015-12-07 and 2015-12-11 of FRB 151206, as observed by ULTRASPEC on the TNT. The sources are split into 6 bins by initial magnitude.

histogram of differential magnitudes in each bin is shown in Figure 5.14.

We treated the standard deviation of this scatter as σ for each bin, and then selected sources which have differential magnitudes greater than 3σ above or below the mean. Table 5.2 shows the breakdown of the standard deviation and number of sources matching this criterion for each bin. A total of 67 sources were selected in this way, and all were inspected visually in the stacked images. Only two of these were considered to be associated with genuine variability, as the rest can easily be explained by the effects of blending or being close to the edge of the image. One of the two variable sources matches one of the five found in the difference imaging analysis described in Section 5.6.2.

Table 5.2: Survey of variability in the extracted source catalogues from 2015-12-07 and 2015-12-11 for FRB 151206. All sources are split into 6 bins based on their 2015-12-07 magnitudes, and those sources showing variability above the 3σ level are selected for examination by eye. N is the number of sources in each bin and σ is the standard deviation in differential magnitudes for the bin.

Mag. Bin	**N**	σ	$\mathbf{3\sigma}$	**N>3σ**
<17	513	0.154	0.44	3
17	580	0.199	0.60	4
18	1118	0.216	0.65	12
19	1587	0.259	0.78	20
20	1580	0.369	1.11	24
≥21	490	0.423	1.27	4

Differential Photometry of Variable Sources

The six stars which displayed variability strong enough to be detected by our two analysis methods (difference imaging and source catalogue comparison), were then studied carefully by differential photometry. Using the ULTRACAM data reduction pipeline, we conducted aperture photometry on each of the varying sources, as well as five bright comparison stars around each source, on the original unstacked images taken by ULTRASPEC. We took the mean of the flux from all five comparison stars in each case, to generate an artificial comparison star flux with minimal photon noise. We divided the target flux by this mean comparison flux to remove any variation in the transparency of the atmosphere. In addition to this, we also extracted the source counts using fixed, wide apertures, to ensure we collect all the counts from the comparison stars. We used these measured counts along with the measured local zeropoints (which already account for airmass and extinction) to flux calibrate the resulting light curves. The results of this differential photometry analysis are presented in Section 5.6.3.

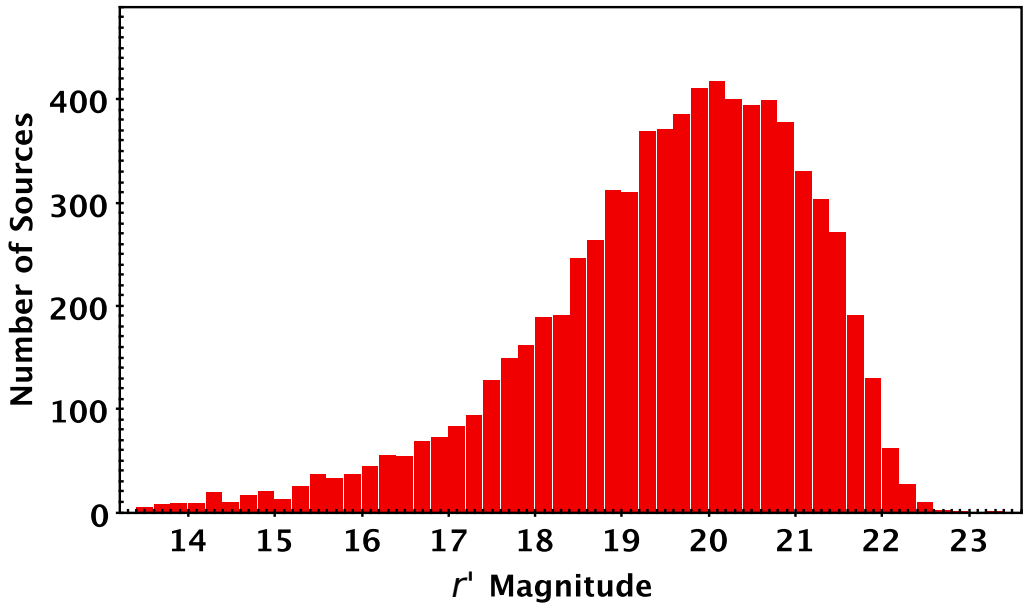

Figure 5.15: Histogram of sources extracted from images taken on 2015-12-07, binned by magnitude. The peak of the histogram represents a crude estimation of the detection limit for the analysis, with a value of $r' = 20.2$ in this case.

Limiting Magnitudes

Finally, it is useful to estimate the limiting magnitude for detection of genuine astrophysical sources in the above analysis of FRB 151206. Here we used three independent methods to measure a magnitude limit, the first two of which are used in the MegaPipe image stacking pipeline and described in Gwyn (2008), and use the extracted source catalogue directly. The third method uses a measure of the sky background noise on the stacked images and the measured zeropoint, and is independent of the source extraction process.

The first method is a simple and crude estimation, based solely on the number counts of sources detected by SE, binned by magnitude. In this instance, shown in Figure 5.15, the turnover of the histogram is a rough estimate of the limiting magnitude for detection of point sources. We measure a detection limit of 20.2 for the 2015-12-07 data and 20.4 for the 2015-12-11 data by this method.

The second method gives an estimate of the 5-σ limiting magnitude for detection by

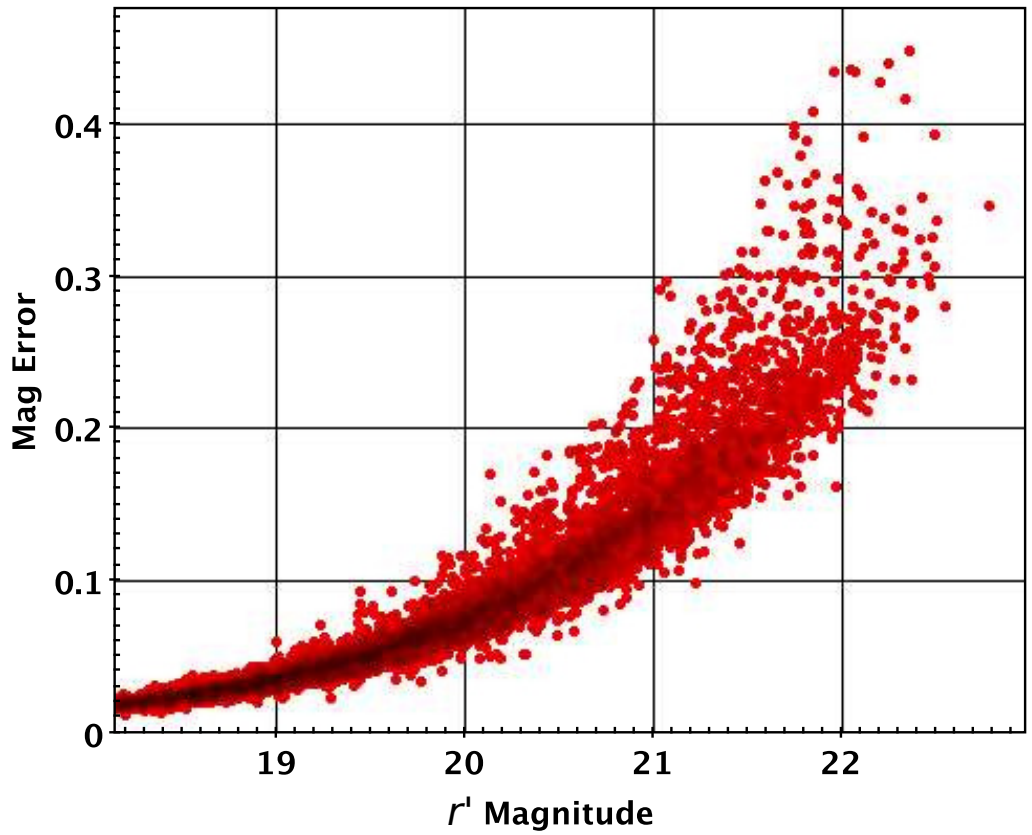

Figure 5.16: Plotting the magnitude error against magnitude allows us to estimate the 5-σ limiting magnitude of a source catalogue. In this case, for the data taken on 2015-12-07, the faintest source with a 0.2 magnitude error is found at $r' = 22$.

SE. A detection of 5-σ would mean a signal to noise ratio of 5 and an associated $\pm$ error of 20%. The magnitude error for such a detection would then be $2.5 \times \log_{10}(1.2) = 0.2$ magnitudes. The faintest source detected with an error of 0.2 magnitudes would then represent the 5-σ limiting magnitude for detection. To estimate this, we plotted the magnitude errors against the measured magnitudes for all sources extracted from each epoch, and find the limit to be $r' = 22.0$ for 2015-12-07 and $r' = 21.9$ for 2015-12-11. An example from the 2015-12-07 data is shown in Figure 5.16.

The final method gives a 5-σ limiting magnitude based on the noise in the sky background of the images. It is independent of the source extraction process, which makes it less relevant to the data being analysed here. Nevertheless, finding an agreement with the two values above would be reassuring. We measure the RMS error (σ) in the sky

background of a blank region of the image using IRAF. This was measured to be 17 counts on average, for all stacked images taken on 2015-12-07, and 18 counts on 2015-12-11. We also estimate the average number of pixels over which any given source is spread ($npix$), using the median FWHM measurement and assuming all sources produce circular PSFs. The 5-σ limiting magnitude is then given by Equation 5.2.

$$m_{lim} - m_{std} = -2.5 \log_{10}(\frac{5\sigma\sqrt{npix}}{F_{std}}) \tag{5.2}$$

where m_{std} and F_{std} are the magnitude and flux of a known standard star. Since in this case we have a measured zeropoint, we can use this instead of m_{std}, and replace F_{std} with 60, since the zeropoint corresponds to one count per second, and the images were exposed for 60 seconds. We measure the 5-σ limiting magnitudes to be $r' = 21.7$ on 2015-12-07 and $r' = 21.5$ on 2015-12-11 using this method. These values are in full agreement with the second method described above, especially when considering that the zeropoint estimations came with a scatter of 0.3 magnitudes.

An alternative and likely more accurate method for estimating the detection limits for a given image and extraction procedure involves the injection of fake sources. This method would see a large number of fake sources artificially introduced to the image at random locations, with a wide range of brightnesses. In this case, the limiting magnitude is that at which a certain threshold (say, 95%) of the injected sources are detected and extracted by the analysis software (SE in our case). This is usually considered to be a more reliable way of estimating a magnitude limit, but we did not employ this method due to time constraints, and because the methods described above are sufficient for this work.

Table 5.3: The six sources seen to be varying by difference imaging analysis or comparing source catalogues. The coordinates are relative to J2000. The final column represents the most accurate change in magnitude, measured by aperture photometry.

Designation	**GSC2.3 identifier**	**RA**	**DEC**	r'	**Cat. Mag. Diff.**	**Phot. Mag. Diff.**
Var1	SA4V041971	19:21:45.16	-04:00:15.2	19.8	0.79	0.20±0.15
Var2	SA4V039564	19:21:45.97	-04:02:02.9	15.7	0.07	0.07±0.01
Var3	SA4V029994	19:21:28.47	-04:08:50.5	17.8	0.74	0.45±0.05
Var4	SA4V022403	19:21:50.00	-04:13:38.2	17.9	0.21	0.22±0.03
Var5	SA4V025024	19:21:01.30	-04:12:00.4	18.3	0.12	0.20±0.04
Var6	SA4V025711	19:21:07.99	-04:11:38.7	15.2	-0.39	-0.095±0.005

5.6.3 Results

In Section 5.6.2 we discovered six stars which changed in brightness from 2015-12-07 to 2015-12-11. The majority were found via a difference imaging analysis, which is particularly useful in this instance due to the crowded fields being observed. The six stars are described in Table 5.3, along with their measured variability according to the extracted source catalogues, and as measured by aperture photometry. Figures 5.17 to 5.22 show the aperture photometry light curves for each star. Note that Var1 has measurement uncertainties large enough to render the observed variability insignificant.

None of the variables discovered here are listed as known variables in SIMBAD or the AAVSO Variable Star Index, or as transients in the CRTS (Drake et al., 2009), ASASSN (Shappee et al., 2014) or MASTER (Lipunov et al., 2010) surveys. The first five have unremarkable colours ($B-V$ <1), but the final source, Var6, has a $B-R$ colour of 3.12 according to the NOMAD1 catalogue. It is visibly red in DSS images, and is very bright in the infrared, according to 2MASS, with J, H, K magnitudes of 9.38, 8.31 and 7.93 respectively.

The optical colours of Var6 would suggest a late M-dwarf spectral classification, however its infrared colours point more to an L-dwarf interpretation. Such cool stars often exhibit violent stellar flares, which could be the explanation for the observed variabil-

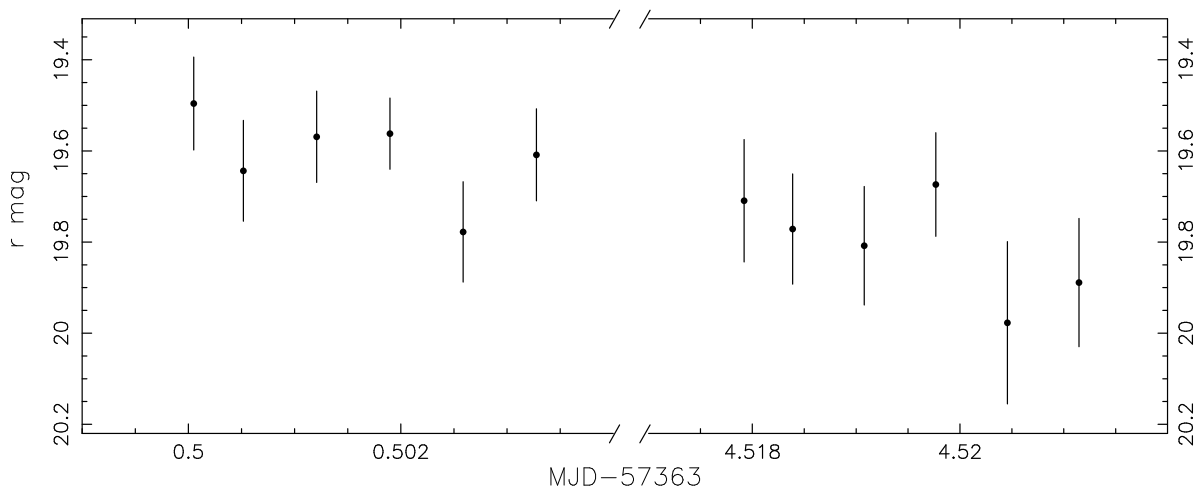

Figure 5.17: Light curve of Var1, measured by aperture photometry on the unstacked images, and calibrated using the local zeropoint. A break in the x-axis separates the two epochs on 2015-11-07 and 2015-11-12.

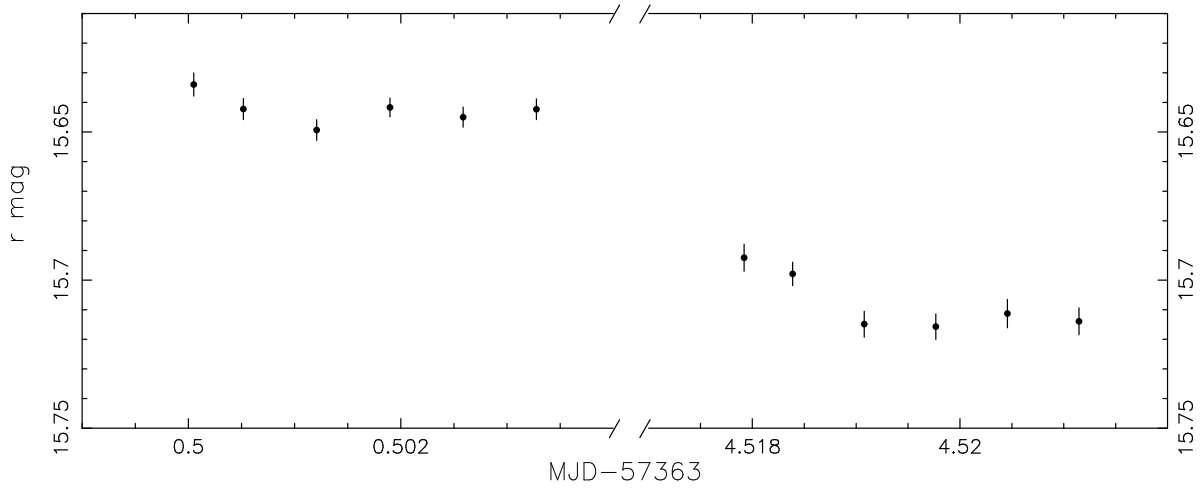

Figure 5.18: Same as Figure 5.17 but for Var2.

ity. This is of particular interest to FRBs, since there is ongoing discussion regarding whether or not flare stars could produce FRBs (Loeb, Shvartzvald & Maoz, 2014; Luan & Goldreich, 2014; Tuntsov, 2014; Maoz et al., 2015). That said, the $J - H$ and $H - K$ values themselves are not consistent with each other, meaning the source could perhaps be a brown dwarf in a binary system with another, earlier-type star. Recent results (Keane et al., 2016) claiming to confirm the extra-galactic nature of at least one FRB detection, and likely a whole class of detections, obviously casts doubt on the flare star

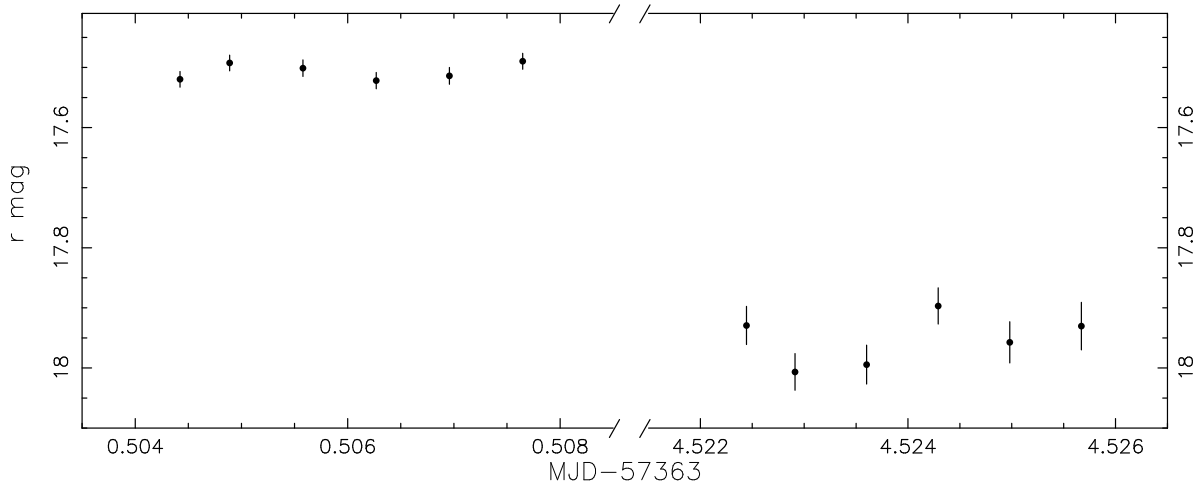

Figure 5.19: Same as Figure 5.17 but for Var3.

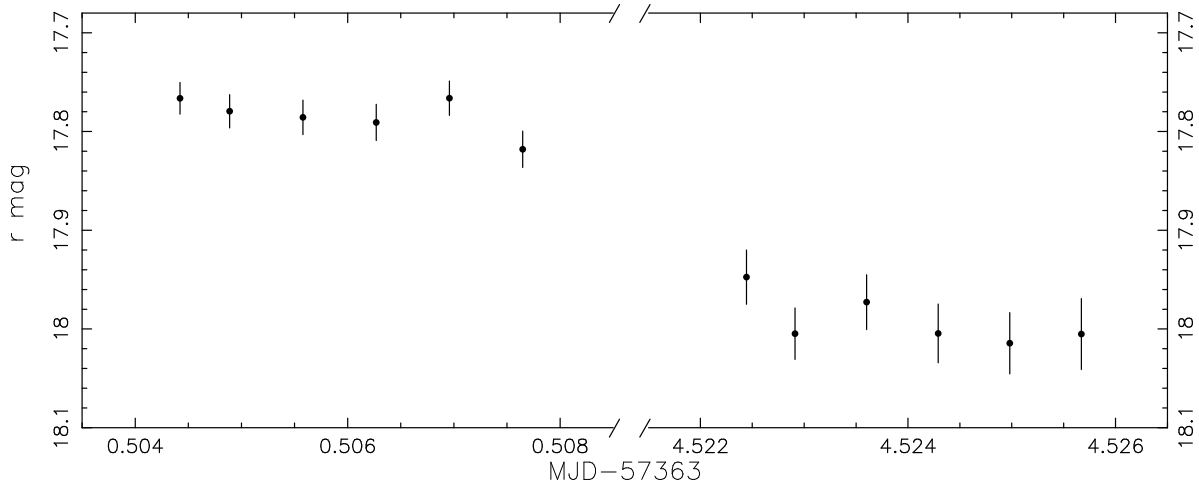

Figure 5.20: Same as Figure 5.17 but for Var4.

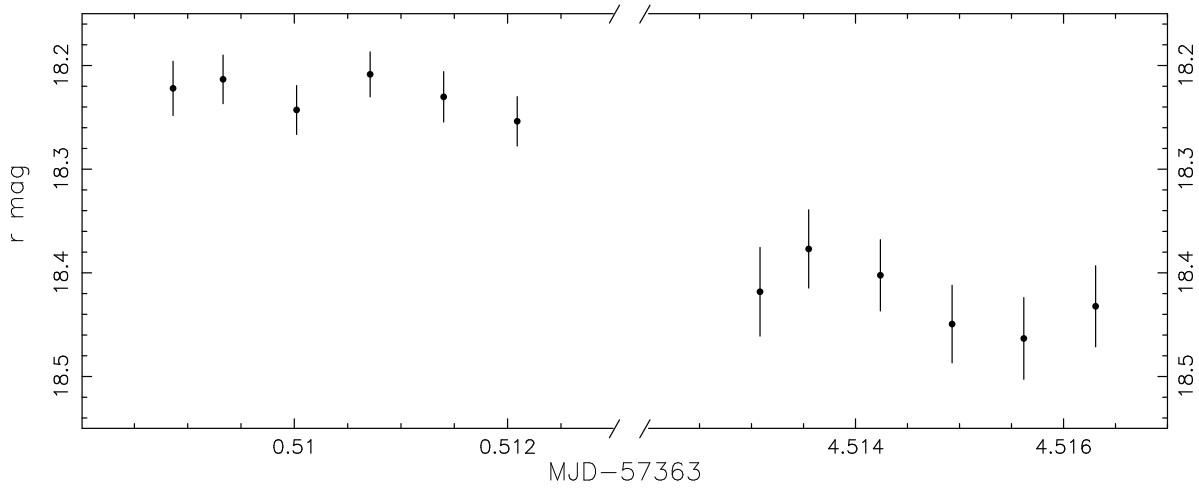

Figure 5.21: Same as Figure 5.17 but for Var5.

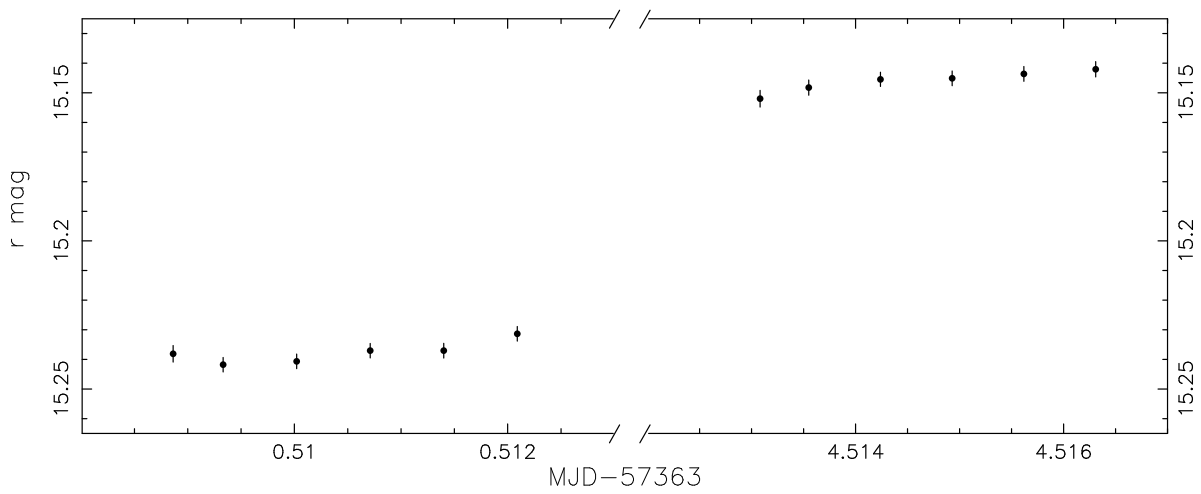

Figure 5.22: Same as Figure 5.17 but for Var6.

hypothesis. However, this claim is itself in doubt due to the possible coincidence of AGN radio activity being mistaken for an FRB afterglow (Williams & Berger, 2016; Williams, Berger & Chornock, 2016). See Section 5.8 for further discussion on this.

The variable sources we detected are likely to be of Galactic stellar nature, thus we can conclude that they are not associated with FRB 20151206, unless the excess dispersion measure (1666 $cm^{-3}pc$) can be assigned to local material around the star. This would place the progenitor within the Galaxy and exclude a large number of explanations for the nature of this FRB, but would go against the recent discovery of Keane et al. (2016). With substantial local material around a Galactic source, we would also expect significant optical extinction (reddening). This then requires any optical counterpart to be even more luminous than any variation we may have measured.

If the variability observed in one of these stars is somehow related to FRB 151206, we might expect to find the star to show repeatable flares, following the flare star explanation of Loeb, Shvartzvald & Maoz (2014). We therefore conducted further observations of these variable sources (excluding Var1, since this variability appears insignificant), at short (several minute) and long (several day) cadences, to test for continued variability.

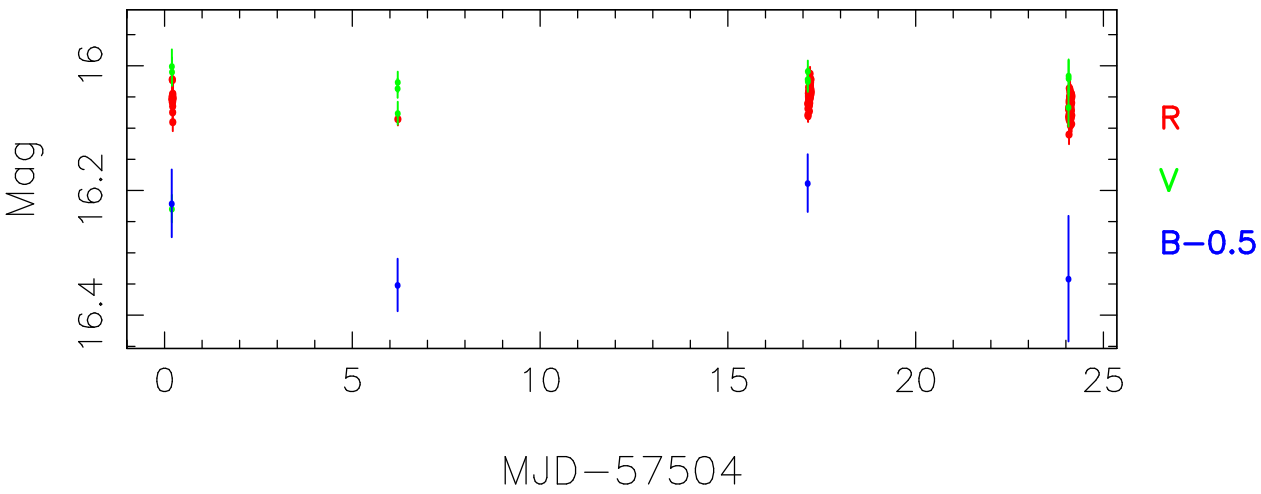

Figure 5.23: Long-term multi-colour light curve of Var 2, showing no significant variability above the photometric uncertainties of 0.05 magnitudes. The *B*-band points have been shifted by 0.5 magnitudes to improve clarity in the plot.

5.6.4 Further time-resolved observations of variable sources

During Spring 2016 we obtained further photometry of the five variable sources discovered in Section 5.6.3, using the robotic telescope *pt5m* (see Chapter 3). Given the small telescope collecting area (primary mirror diameter 0.5m), we would not expect to reach sufficient photometric precision to measure all the variability seen in Table 5.3.

Nightly observations in three colours were requested (*B, V, R*), as well as a short-term light curve with a cadence of two minutes, in the *R*-band. For each source, we reduced and analysed the images in the same way as in Section 5.6.2, except that instead of using a local zeropoint to flux-calibrate the extracted counts, we used the magnitudes of individual comparison stars as listed in the *Vizier* catalogues NOMAD, UCAC4 and USNO-B1.0. We present the long-term multi-colour light curves in Figures 5.23 to 5.27, and the most interesting or variable short-term *R*-band light curves in Figures 5.28 to 5.32.

Var2 showed no significant variability in either its long-term or short-term light curves. This is unsurprising, given that the initial variability observed by ULTRASPEC was only 0.07 magnitudes. A variation of 0.05 magnitudes is possible but too small to

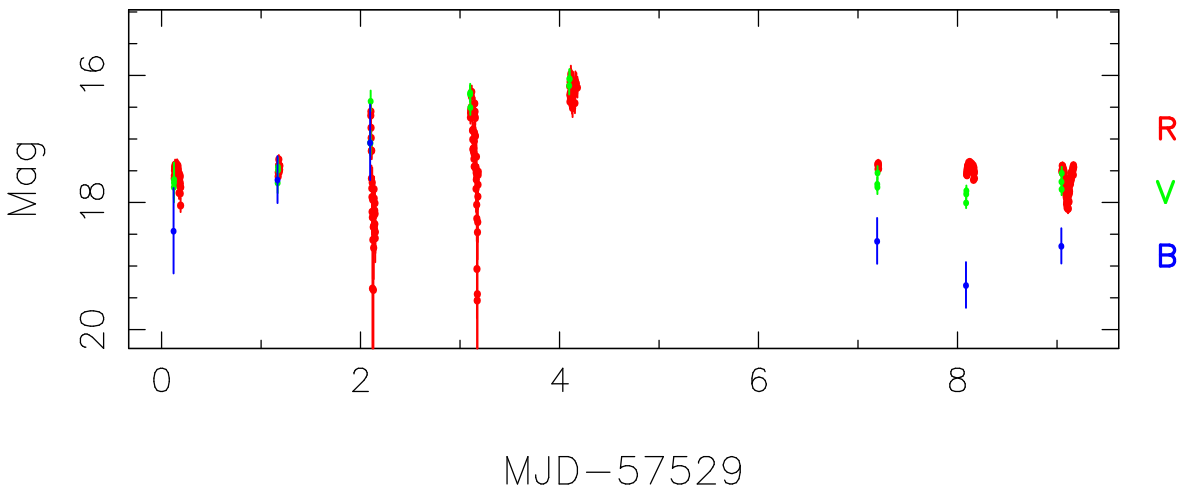

Figure 5.24: Long-term multi-colour light curve of Var 3, showing clear variability of up to 2 magnitudes.

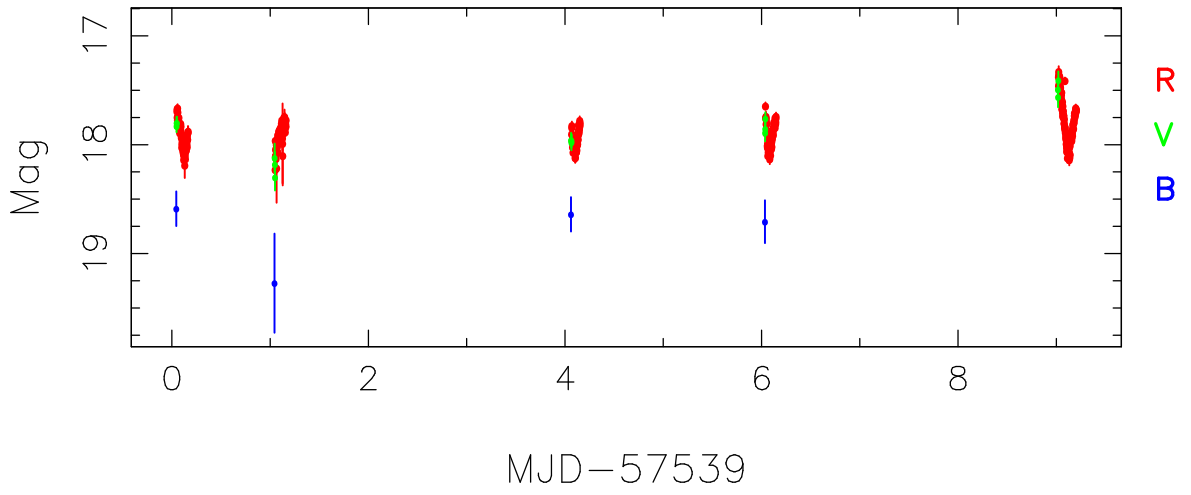

Figure 5.25: Long-term multi-colour light curve of Var 4. Short-term variability of up to 1 magnitude is visible, and this is likely to be the dominant cause of any long-term variability.

be detected by *pt5m*.

Var3 showed very strong variability of up to two magnitudes. The long-term light curve is similar to that of dwarf novae during short and mild outbursts. The short-term light curve shown in Figure 5.29 displays an eclipse-like signal, though this may also be some kind of periodic, sinusoidal variability.

Var4 showed a sharp 'V'-shaped sinusoidal light curve in the R-band short-term studies. The long-term variation is likely dominated by this. The cause may be stellar

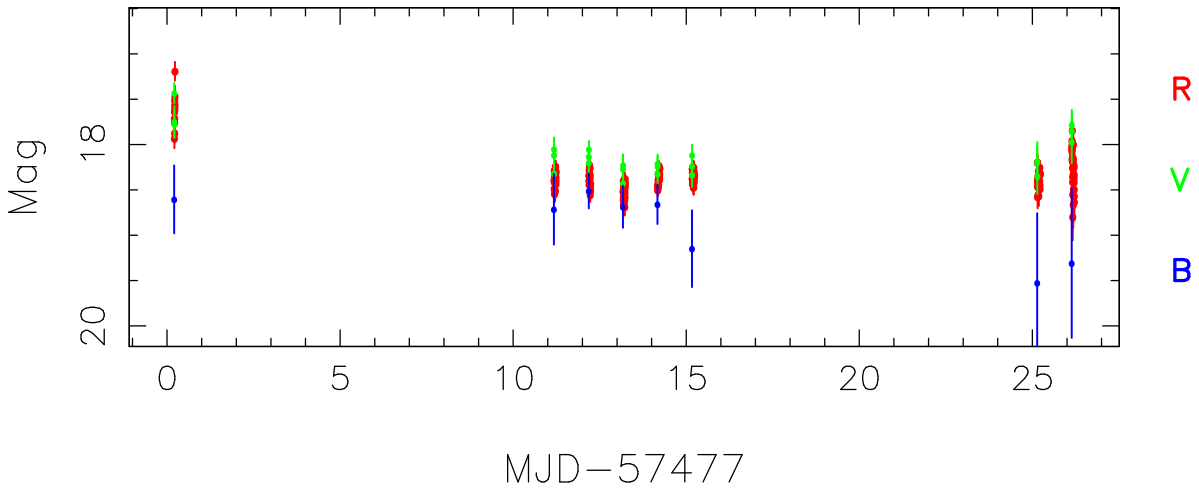

Figure 5.26: Long-term multi-colour light curve of Var 5, showing some variability of up to 1 magnitude.

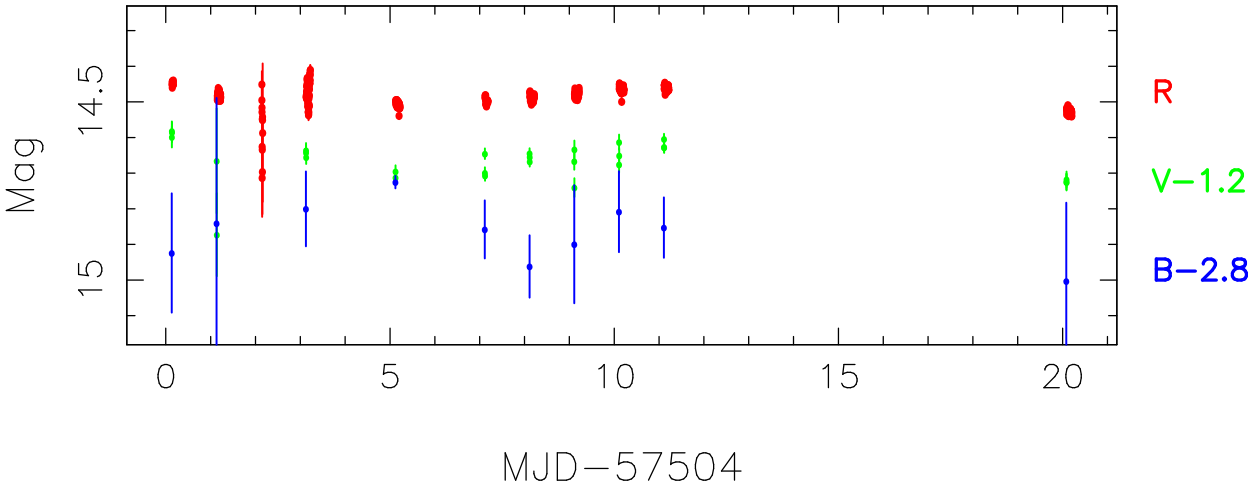

Figure 5.27: Long-term multi-colour light curve of Var 6, showing some minor variability of up to 0.1 magnitudes. The B-band points have been shifted by 2.8 magnitudes and the V-band points have been shifted by 1.2 magnitudes to improve clarity in the plot.

pulsation or ellipsoidal binary variability.

Var5 displays some variability of up to one magnitude in its long-term light curve (Figure 5.26). In the short-term light curve, a smaller variation is seen, which resembles stochastic variability such as that found in accretion-driven systems and flare stars.

Var6 shows weak (0.02-0.1 magnitude) variability at short and long timescales. The high cadence light curve has no obvious structure or periodicity, again suggesting stochastic variability. The source could be a flare star, AGN variability or a quasar.

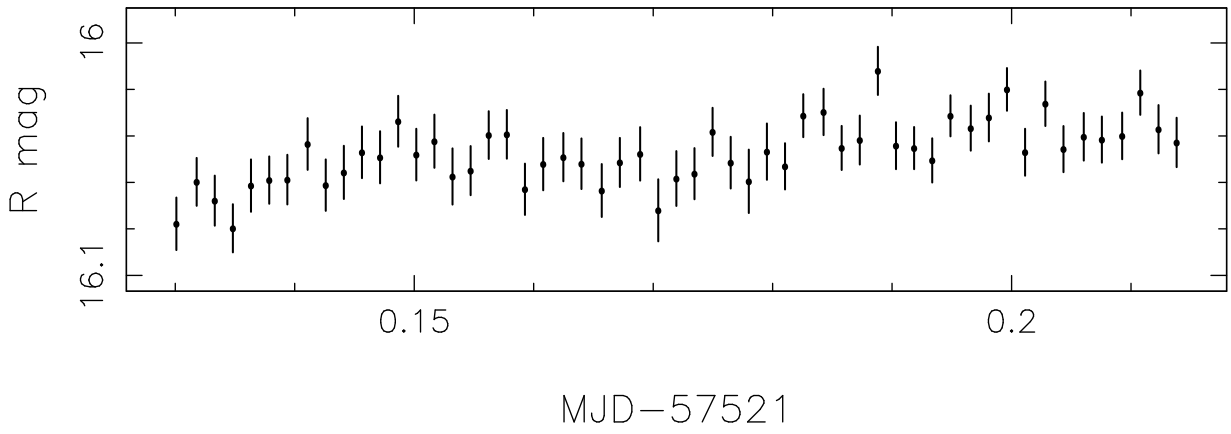

Figure 5.28: Short-term R-band light curve of Var 2 observed on the night beginning 2016-05-12. No significant variability is visible.

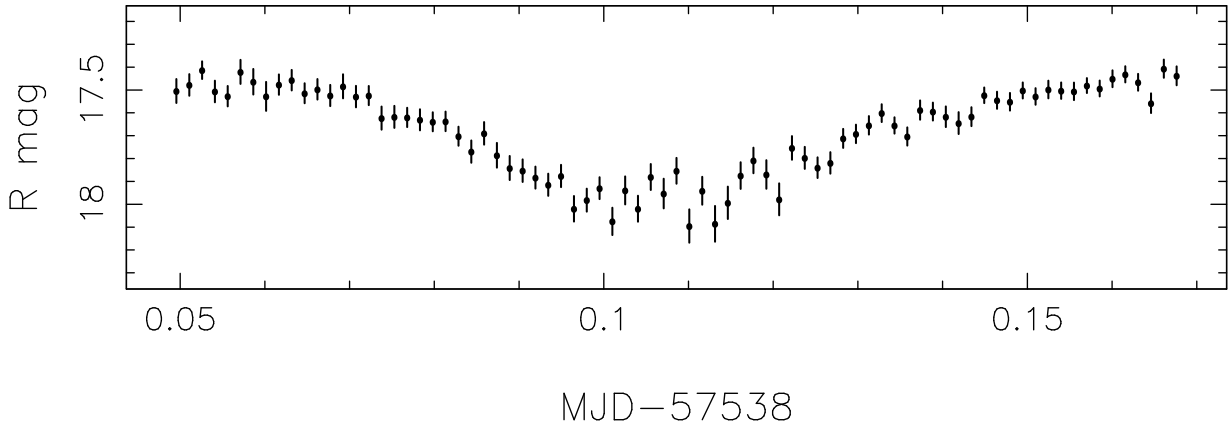

Figure 5.29: Short-term R-band light curve of Var 3 observed on the night beginning 2016-05-29. The variation seen could be attributed to an eclipsing event, or a periodic sinusoidal variation.

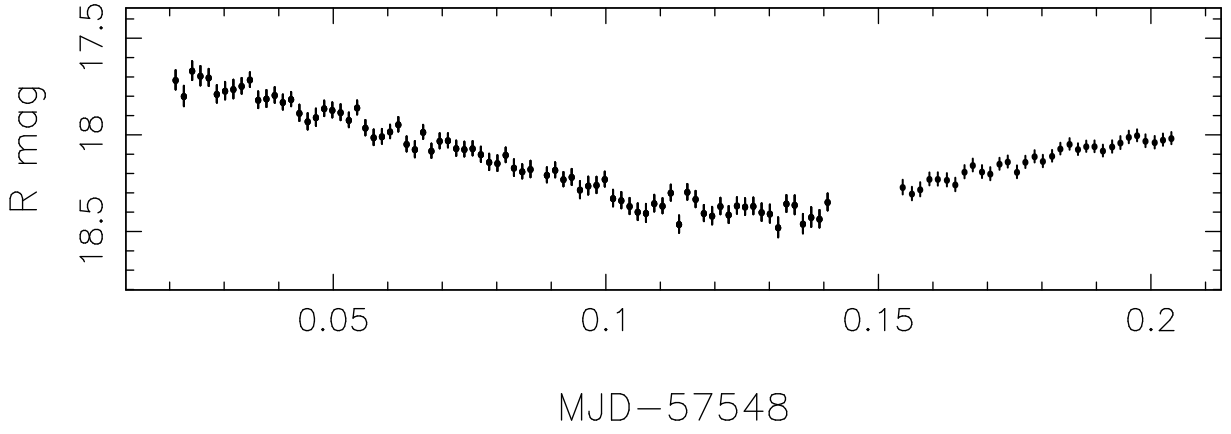

Figure 5.30: Short-term R-band light curve of Var 4 observed on the night beginning 2016-06-08. The 'V'-shaped variability is also observed on other nights.

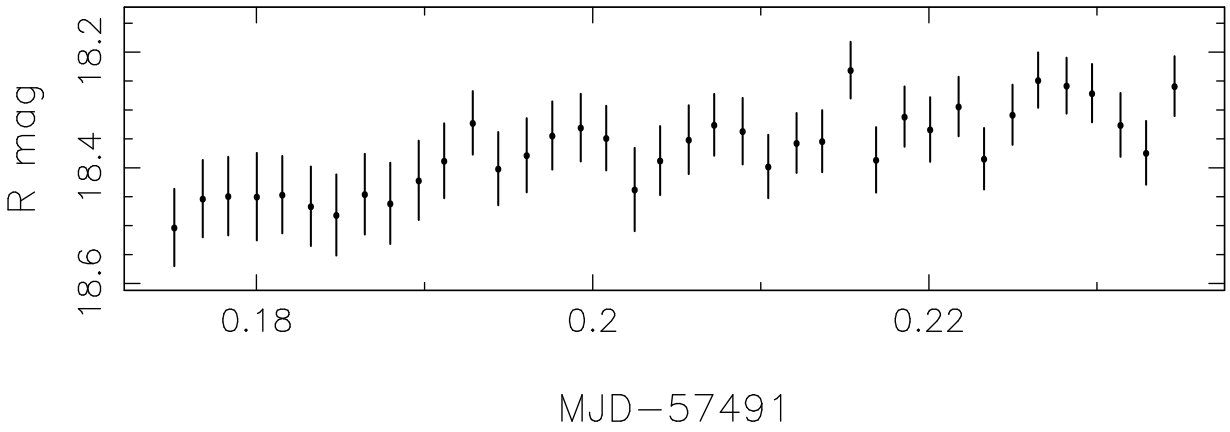

Figure 5.31: Short-term *R*-band light curve of Var 5 observed on the night beginning 2016-04-12. The variation of up to 0.2 magnitudes may be amplified by systematic effects such as the movement through hot pixels on the CCD.

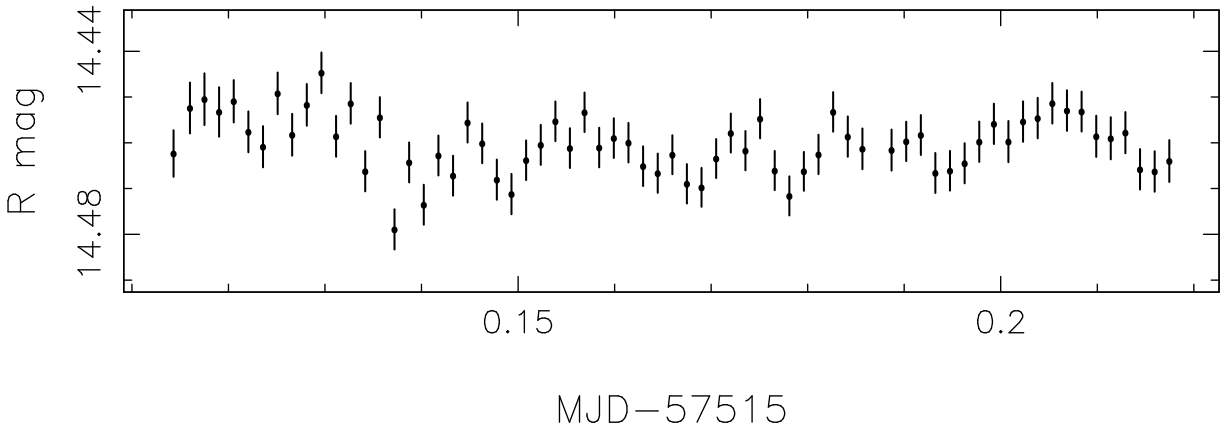

Figure 5.32: Short-term *R*-band light curve of Var 6 observed on the night beginning 2016-05-06. Fast, weak variability is visible.

For all of these 5 variable sources, low-resolution spectra would be particularly valuable in determining their nature. We hope to obtain spectra of them soon.

Our observations of FRB 151206 will be published as contributing data in Bhandari et al. (2016, in prep.).

5.7 FRB 151230

At 16:16:15 UTC on 2015-12-30, SUPERB at Parkes detected another FRB within 14 arcminutes of the coordinates 09:40:49.9, -03:27:05.1 (J2000). We received a trigger email

requesting followup observations about an hour later, and within another hour we had produced finding charts and contacted the telescope operators in Thailand and asked for observations to be conducted that night. The field would have become visible from around 23:00 UTC.

Unfortunately the TNT was completely clouded out, with no observations made that night, or on the two subsequent nights. We were unable to contribute any follow-up observations of FRB 151230.

5.8 Mystery Solved?

5.8.1 Host Galaxy Localisation

If the discovery of radio variability in a coincident location to FRB 150418 is genuinely associated with the burst (Keane et al., 2016), then this would represent an enormous breakthrough for the study of FRBs. As mentioned in Section 5.5, the identification of a possible host galaxy for FRB 150418 has allowed an accurate redshift to be determined which, if true, would provide definitive confirmation that FRBs are extra-galactic. This discovery could also open up a new era of observational cosmology tests, with the DM of FRBs being able to probe the intergalactic medium after being calibrated with a redshift.

Evidence for the 'missing baryons' (Bregman, 2007) follows directly from comparison of the cosmic density of ionised baryons (Ω_{IGM}) measured by the FRB 150418 detection, with that measured by fitting cosmological models to the WMAP observations (Hinshaw et al., 2013). Ω_{IGM} is derived from the measured DM, which arises from the free electrons in the IGM and is dependent on the density of those electrons, and the distance to the source, which is given by the measured redshift. Assumptions must be made about the ionisation factor of the IGM and the DM contributions from sources other than the IGM, but ultimately the baryon density can be estimated.

Keane et al. (2016) found that the values for Ω_{IGM} derived by the two methods (WMAP and FRB 150418) were in direct agreement. Alternatively, by adopting the value of Ω_{IGM} measured by WMAP, they predicted that the DM contribution from the host galaxy is negligible, which agrees with the observational constraints placed on the host DM contribution by the lack of observed scattering in the pulse profile and higher-order dispersion components (Dennison, 2014; Keane et al., 2016).

An alternative way to look at these data is to assume the cosmological parameters measured by WMAP are correct, and use the measured DM to constrain the rest mass of the photon via theories of Lorentz invariance violation. The DM is a time delay between the arrival of FRB photons as a function of frequency, with lower frequency photons arriving later. Assuming the photons are emitted at the same time from the same source, their delayed arrival could have two components: dispersion by electrons in the line of sight, and the effects of Lorentz invariance violation (if this exists). Wu et al. (2016) treat the WMAP measurements of intergalactic baryons as the true value, and use the corresponding lack of additional excess DM to constrain the strength of invariance violation, and thus set more stringent limits on the mass of the photon than previous studies. A photon mass limit of $\sim 10^{-50}$kg is derived by Wu et al. (2016), and a similar constraint is found by Bonetti et al. (2016). By similar comparison of the DM of FRB 150418 with the WMAP measurements, stronger constraints on the Einstein Equivalence Principle have also been made (Tingay & Kaplan, 2016).

Another exciting property that could be probed by further FRB localisation and redshift measurements is the optical depth of the cosmic microwave background (Fialkov & Loeb, 2016). If more FRBs can be observed at high redshifts, stronger constraints on this parameter can be made, thus deepening our understanding of the Epoch of Reionisation.

The measured radio afterglow potentially associated with FRB 150418 also helps constrain possible progenitor models (Keane et al., 2016). Zhang (2016b) analysed the

afterglow, constraining the energy output of the burst to being of the same order of magnitude as short GRBs ($\sim 10^{50}$erg). This supposedly rules out many FRB progenitor models, such as stellar flares, magnetar giant flares, blitzars and white dwarf merger events, since these are not expected to produce sufficient energy. Instead Zhang (2016b) discusses compact star merger events, such as BH-BH, NS-NS or BH-NS mergers, in the context of FRB progenitors. This is particularly interesting, because it raises the prospect of possible observational associations between GRBs, FRBs and gravitational wave events, which may all originate from such mergers.

However, despite all this excitement, there is reason to be cautious. The coincident radio afterglow which was associated with FRB 150418 by Keane et al. (2016) may in fact be a coincidence. Williams & Berger (2016) showed that the afterglow is more likely to be associated with simple AGN variability, rather than an FRB, on three counts. Firstly, the radio light curve does not show the expected behaviour for other highly energetic radio transients (synchrotron blast waves). Secondly, the quiescent radio luminosity predicts a much higher star formation rate than is inferred from inspection of $H\alpha$ in the source spectrum, and is therefore more consistent with the quiescent radio luminosity of an AGN. And thirdly, the predicted number of radio variables in the Parkes beam is of order unity, such that the discovery of one radio-variable source is not unexpected. This analysis lead Williams & Berger (2016) to reject the assertion by Keane et al. (2016) that FRB 150418 is associated with the faint galaxy WISE 0716-19.

Subsequent VLA observations of the proposed host galaxy found it to be brighter (Williams, Berger & Chornock, 2016), at more than 3-σ above the 'quiescent' level measured by Keane et al. (2016). This supports the AGN explanation for the source galaxy radio variability, and casts deeper doubt on the proposed localisation of the first FRB.

Finally, other authors have challenged the refuting claims of Williams & Berger (2016) with more detailed statistical analysis (Li & Zhang, 2016), finding that the probability of

seeing coincident AGN variability at that particular time is very low. Further monitoring of this source will assist in confirming the competing claims made by Keane et al. (2016) and Li & Zhang (2016), and Williams & Berger (2016). However, ultimately it is likely that at least one more indisputable FRB localisation will be needed in order to measure an FRB redshift for certain.

5.8.2 Repeating FRB

Another recent development that adds credence to the idea that there may be multiple populations of FRBs, is the discovery of a repeating FRB signal by the Arecibo telescope (Spitler et al., 2016). The field of FRB 121102, which was originally discovered by the Arecibo PALFA survey (Spitler et al., 2014), was observed extensively in May and June 2015 in a search for repeat signals. Spitler et al. (2016) found 10 additional burst signals during those observations, bringing the total observed to 11. All bursts were seen in the same direction ($1.3' \pm 3'$ apart from each other), and had dispersion measures that were consistent to within 2σ. All had dispersion indices of $\sim$2, as expected for radio waves travelling through a cold, ionised medium, although all had peak flux densities somewhat lower than is typical for other FRB detections. The burst spectra vary, with both positive and negative spectral indices, which were attributed to intrinsic variations of the source, and some bursts appeared to show multiple components, like FRB 121002 (Champion et al., 2015). No statistically significant periodicity was found in the arrival times of the bursts.

The repeating nature of FRB 121102 effectively rules out any progenitor models involving a catastrophic, disruptive event such as mergers or collapse of massive stars. Spitler et al. (2016) suggest a likely source of these repeated bursts could be an extra-galactic pulsar emitting giant pulses - a proposal supported by statistical analysis of other FRB properties (Katz, 2016b). This clearly contradicts the proposals of Keane et al.

(2016) and Zhang (2016b) which consider short GRB-like catastrophic progenitors, if the localisation of FRB 150418 is valid.

This discovery may emphasise the possibility that the FRB population constitutes several distinct classes, produced by different types of progenitors, in the same way that GRBs and supernovae are now understood (Li et al., 2016b). However, other explanations for the apparent lack of repetition in other FRBs have also been considered, for example the possibility that bursts arrive in groups as predicted by flicker noise (Connor, Pen & Oppermann, 2016). This would support the theory that FRBs are giant pulses from young neutron stars (Connor, Sievers & Pen, 2016; Lyutikov, Burzawa & Popov, 2016).

Simultaneous Radio and Optical observations of FRB 121102

Given the better localisation of FRB 121102 compared with those discovered at Parkes (error circle $\pm 3'$), continued optical follow-up of the field could provide valuable constraints, especially if coordinated with simultaneous radio observations which could pinpoint the times of arrival of bursts. This is exactly what we have attempted with FRB 121102, using both ULTRASPEC and *pt5m* to observe the field for several hours. Our ULTRASPEC observations were conducted partly during a time of simultaneous coverage in the radio with the Lovell Telescope at Jodrell Bank. The *pt5m* observations were analysed by Stuart Littlefair, who conducted a preliminary search for transient events in the images. No obvious astrophysical detections were made. We are still awaiting news from collaborators at Jodrell Bank as to whether they made any radio detections of repeated bursts of FRB 121102 during our simultaneous coverage.

5.9 Conclusion

The question of FRB progenitors, and their usefulness in studying cosmological and extra-galactic parameters, remains open. The efforts of Keane et al. (2016) and Williams & Berger (2016) are commendable, as are the numerous theoretical proposals for possible FRB progenitors. However it is clear that more observations are needed in order to confirm or reject every theory, and many more FRB observations will be needed if the corresponding cosmological questions are to be answered with any certainty.

We have conducted detailed optical follow-up of two new FRBs (150215 and 151206), and searched the resulting images for variable or transient sources. Five (possibly six) variable sources were discovered in the field of FRB 151206, and these were followed-up further with observations by *pt5m*.

As radio astronomers around the world continue to search for new FRB signals, and theorists continue to consider progenitor explanations, we will continue to follow up new real-time detections with ULTRASPEC whenever possible. Detecting a near-simultaneous optical afterglow of an FRB should allow definitive pinpointing of the location of the burst. Equally, additional limits on the absence of optical counterparts to FRBs will continue to help narrow the parameter space that progenitor theories may occupy.

Chapter 6

Conclusions and Future Work

6.1 Summary and Implications

In this thesis I have discussed the characterisation and commissioning of the ULTRASPEC camera on the 2.4m Thai National Telescope, the development of transient follow-up and data reduction software for the 0.5m robotic telescope *pt5m*, and the use of these new tools to undertake a search for new eclipsing CVs and optical afterglows of FRBs.

6.1.1 ULTRASPEC

In Chapter 2 I have explained the operation of ULTRASPEC, a high-speed imaging camera with electron-multiplying technology, now operational on the TNT. I conducted laboratory tests to characterise the CCD, vacuum system and liquid nitrogen cooling behaviour. I then helped with on-sky commissioning tests of the instrument in Thailand, working through initial problems and helping the local staff understand and utilise the instrument. I built a complete model for the instrument throughput, and later used this along with multiple on-sky observations of standard stars to produce a reliable method for flux calibrating the non-standard *KG5* filter.

Since November 2013 ULTRASPEC has performed well in 10 observing runs (totalling

77 nights) dedicated to Sheffield/Warwick observers, as well as being the work-horse instrument for most astronomers at NARIT. Technical downtime has been less than 1% of all allocated night time so far. The instrument has been used to study superluminous supernovae, eclipsing CVs (McAllister et al., 2015, Hardy et al., submitted)[1], eclipsing white dwarfs (Littlefair et al., 2014; Parsons et al., 2015)[1], other eclipsing binaries (Hemha et al., 2015), transit timing variations (Bours et al., 2014, 2016)[1], millisecond pulsars (Shahbaz et al., 2015), stellar pulsations (Bours et al., 2015), planetary debris transiting around white dwarfs (Irawati et al., 2016; Gänsicke et al., 2016)[1], lunar occultations (Richichi et al., 2014, 2016), and a mysterious white dwarf pulsar (Marsh et al. 2016, accepted in Nature)[1].

In addition to these studies with ULTRASPEC, I have also been a regular supporting observer for ULTRACAM on the 4.2m WHT, mounting and dismounting the instrument and observing a range of targets, including transiting exoplanets, eclipsing binaries, pulsating stars, and millisecond pulsars. Most of the ULTRASPEC studies referenced above also included some ULTRACAM data. One study of particular interest was of the outburst of the black hole X-ray transient V404 Cyg (Gandhi et al., 2016; Dallilar et al., 2016)[1], for which I helped with observing and data reduction. I have assisted with a total of 7 individual observing runs, leading two of them, and performing a total of 29 nights of observations. This experience was incredibly helpful in deepening my understanding of high-speed photometry and the operation of ULTRACAM and ULTRASPEC.

6.1.2 pt5m

I have described the components and operation of *pt5m* in Chapter 3, including the history of the telescope and the hardware and software architecture. I have also shown the results of the performance testing and detector characterisation which I conducted, and discussed in detail the Python scripts I developed to allow automated transient

[1]I am, or will be, a co-author on these papers

follow-up operations to take place.

The development of the automatic transient listening software, and the associated data reduction and analysis components, have allowed the robotic facility to become highly efficient. *pt5m* now regularly utilises over 80% of the available night time to observe new transients, as well as ongoing projects, and Chapter 3 also includes some recent scientific results. *pt5m* has been used to study transiting exoplanets (Hardy et al., 2015b; Garai et al., 2016)[1], variable stars in clusters (Fritzewski et al, in prep)[1], outbursting black hole transients (Gandhi et al., 2014; Hardy et al., 2015a, 2016a)[1], GRB optical afterglows (Hardy et al., 2015c,d)[1], an eclipsing AM CVn system (Campbell et al., 2015a)[1], hot subdwarf binaries (Kupfer et al., 2014), transit timing variations (Bours et al., 2016)[1], and eclipsing CVs (Littlefair et al., 2013; Hardy et al., 2016b)[1]. The telescope and automatic follow-up system have also contributed considerable data to the Gaia science alerts verification programme (Rixon et al., 2014; Wyrzykowski, 2016)[1]. *pt5m* has supplied more data points (over 7500) to the Gaia server than any other single facility, including data on supernovae, AGN, CVs and unknown objects. The first extensive results of this ongoing collaboration are to be published soon (Wyrzykowski et al., in prep)[1].

6.1.3 Eclipsing CVs

In Chapter 4 I presented a successful search for new eclipsing CVs, as well as a large catalogue of archival and recent observations of known eclipsing systems. New eclipsing CVs were discovered or confirmed with *pt5m*, and later observed at high cadence with ULTRASPEC or ULTRACAM to investigate the structure of their eclipses. In total we presented new photometry of 74 CVs. 47 of these are eclipsing, of which 13 are newly discovered. For all 47 eclipsing systems we showed high cadence (1-20 seconds) light curves. We provided new or refined ephemerides, and supply mid-eclipse times for all

observed eclipses. We assessed the potential for light curve modelling of all 47 eclipsing systems to determine their system parameters, finding 20 systems which appear to be suitable for future study.

Our efforts to discover new eclipsing CVs, and study their eclipse structure in fine detail, will help progress our knowledge of CV evolution. Only around 35 eclipsing systems have been studied at high enough cadence to allow precise measurement of the component masses, radii and other system parameters (Zorotovic, Schreiber & Gänsicke, 2011; Schreiber, Zorotovic & Wijnen, 2016). More mass and radius measurements are needed to refine models of CV evolution, and improve the empirical donor sequences already derived (Knigge, Baraffe & Patterson, 2011). As can be seen in Figure 8 of Savoury et al. (2011), the data on existing WD and donor star masses does match the theoretical models, but there are simply not enough studied systems, particularly near the period minimum (80 mins) and above the period gap (>160 mins), to say this with confidence. Our discovery of new eclipsing systems will help produce a substantial dataset from which precise masses, radii and other system parameters can be measured.

6.1.4 Fast Radio Bursts

In Chapter 5 I evaluated the extensive literature discussions on this exciting new class of transients, which produce bright (few mJy), short (few ms), highly-dispersed radio pulses through currently unknown emission mechanisms. I also presented our collaborative effort to conduct optical follow-up of the fields of FRBs detected in real-time, with the hope of detecting or constraining an optical counterpart.

We have made two attempts to observe optical afterglows of FRBs. FRB 150215 was observed 25 hours after the burst was detected. The full error circle of the radio coordinates could not be observed, and the data were not of optimal quality. Nevertheless, we were able to reach detection limits of around $R = 21$ magnitudes, finding no sign of

any obvious afterglow candidates. This, along with observations from other collaborating teams should help to constrain the emission mechanisms for the burst. Our optical follow-up efforts will be included in an upcoming paper discussing this burst (Petroff et al. 2016, in prep)[1].

FRB 151206 was observed with ULTRASPEC around 29 hours post-burst, this time in slightly better conditions. The full error circle of the burst was observed, and re-imaged four days later. Several variable objects were seen, but all are associated with existing sources, likely to be stars within the Milky Way. The discussion in Section 5.1.2 indicates that a Galactic progenitor for FRBs is unlikely, suggesting that the variable sources discovered in our observations are not expected to be associated with the radio burst itself. Further observations with *pt5m* showed continued variability in most of these sources, however, more data is needed in order to determine their intrinsic nature, and a spectroscopic study would be particularly valuable. Our follow-up efforts will be included in an upcoming paper discussing this burst (Bhandari et al. 2016, in prep)[1].

6.2 Future Work

This thesis has shown how small robotic telescopes with transient follow-up pipelines and high-speed cameras on 2m-class telescopes can contribute to variable star and transient astronomy. Nevertheless, there are a number of improvements that could be made, and ongoing projects that could be extended, if I were able to continue this work.

6.2.1 ULTRASPEC

To improve the accuracy of calibrations for ULTRASPEC, it would be beneficial to continue to regularly observe standard stars, also at a range of airmasses to help characterise and confirm the existing measurements of atmospheric extinction. This would allow more accurate tests of the throughput model by providing a clearer picture of the best

performance of the instrument, since it is possible that even the current best recorded performance included some unaccounted for transmission losses due to the hazy conditions (perhaps up to 10%, based on the variability of existing measurements). Continued standard star observations will also allow better understanding of those hazy conditions which are often prevalent at Doi Inthanon. A larger standard star dataset should enable us to search for patterns in atmospheric throughput with respect to local weather conditions or seasonal changes.

The flux calibration method for the *KG5* filter could also be improved with new, deeper observations of an open cluster in the five SDSS filters. The original study discussed in Chapter 2 detected few stars over a narrow range of colours using just 13-second exposures in each filter, but a deeper study (using hundreds of stacked frames) should improve on this. Observations have been conducted but not yet analysed. I would hope to be able to use these data along with the bolometric corrections to test the throughput model to better than 0.2 magnitudes (the current systematic uncertainty). If the model is found to be incorrect as a function of wavelength (i.e. found to show a non-flat trend in Figure 2.55), the model could be corrected using a polynomial modifier to improve its accuracy. This would also improve the theoretical zeropoint measurements in the *KG5* filter, although a systematic uncertainty of better than 0.1 magnitudes is probably not achievable due to the inherent uncertainties in the effective temperature and model spectra of any given standard star.

Related to ULTRASPEC, there are efforts underway at NARIT to begin a programme of automation at the TNT, with the end goal being a fully automated telescope and instrument suite, such as is operated by the Liverpool Telescope (Steele et al., 2004). The robotisation of ULTRASPEC will be complex and challenging, requiring careful interfacing of the telescope control system with a program which can automatically and efficiently operate ULTRASPEC. Unlike *pt5m*, ULTRASPEC observations have always needed a trained observer to be in control, monitoring signal-to-noise ratios, selecting appro-

priate windowing and binning settings, and producing a quick-look light curve on the fly. Automating this procedure presents a great challenge, but such an upgrade could provide great improvements in observing efficiency, as well as saving time and effort of staff and observers. I would certainly be interested in working on, or assisting with, this challenge.

6.2.2 pt5m

It is important that someone continues to manage the *pt5m* telescope, ensuring that any faults are spotted quickly and fixed accordingly, and making the most of the available observing time. The automated follow-up and data reduction software makes the management of the facility quite simple, but keeping track of and announcing exciting new observations, and removing uninteresting objects from the queue, are still necessary tasks.

If time were not limited, I would seek to develop the transient software with the following improvements. Firstly, I would like to rebuild the original `skyalert.org` listening software, such that it can listen to newer sources of alerts, e.g. VOEvents, now that the skyalert server is no longer operational. A listener that reads VOEvents, like those released by our AMI collaborators for GRBs (see Section 3.7.6), could be more directly combined with the AMI listening software, and could be more easily expanded to listen to additional streams of alerts. For example, it might be that in the future we would like to automatically trigger *pt5m* observations on all reported VOEvent transients from Gaia, ASAS-SN, MASTER, CRTS, Swift, LOFAR, LSST and others, possibly including FRBs from Parkes or elsewhere.

A possible architecture for such widespread follow-up could be that new transients are analysed by the listening software and sorted into two classes. If they are considered as urgent targets-of-opportunity (e.g. GRBs or FRBs), they would be submitted to

the *pt5m* queue immediately and as single, one-off observations. If instead they were less urgent targets requiring long-term monitoring (e.g. CVs, SNe, AGN), they would be added to the daily target list for job submission (Section 3.7.1), and remain in the nightly queue until a human user decides they are no longer of interest (e.g. because they have faded beyond detection, or were spurious to begin with). In both cases, the person managing the *pt5m* queue would be alerted to any new targets, and he/she would have the ongoing task of monitoring the target list and removing uninteresting objects.

For any objects with precise positions (positional uncertainty < 2", e.g. most transients discovered at optical wavelengths), the automatic reduction software is able to upload calibrated magnitude measurements to the Gaia Science Alerts server. However, for transients such as GRBs or FRBs, whose precise localisation (positional uncertainty > 2") is not supplied in the original alert trigger, these observations may need to be reduced manually, to search for FRB optical counterparts, or analyse the GRB afterglow once a more precise X-ray localisation is available. In the long term there may be options to improve this, such as using updated coordinates for GRBs when they become available to reduce and analyse the light curve.

Currently, staff effort is sometimes wasted downloading, reducing and analysing data which is taken in conditions which render the observations useless, such as thick cloud cover for example. A final improvement that could be made to the *pt5m* control architecture, and scheduling in particular, might be to implement a cloud limit for observations. For example, targets could be submitted with a maximum cloud okta value. This could be used to affect the priority of certain targets, allowing for more efficient scheduling, or in fact to reject some targets altogether whenever there is significant cloud present. We have not implemented a system to do this so far, as often we feel it would be better to collect some data, even through cloud, than no data at all. However, in the long term, if there is insufficient staff time to manage and analyse all data, it might prove useful to implement a system of this kind.

6.2.3 Eclipsing CVs

Both *pt5m* and ULTRASPEC could continue to be used in an ongoing search for eclipsing CVs, using the growing list of newly discovered dwarf novae from the Gaia, MASTER, CRTS and ASAS-SN surveys. ULTRASPEC and *pt5m*, along with ULTRACAM, are ideally suited to this work, and the search for eclipsing CVs could go on until sufficient systems have had their masses measured to refine the models of CV evolution and answer the ongoing question: do WDs in CVs grow in mass?

More mass measurements of medium (2-3 hours) and long (3+ hours) period systems would be especially useful in this regard, since these systems have the fewest measured masses to date. Even just another 5 systems with periods longer than 2 hours could be sufficient to detect a 0.2 $M_{\odot}$ discrepancy between short and long period systems at the 4σ confidence level (McAllister, priv. comm.). Our search for new eclipsing systems (Chapter 4) discovered or confirmed 6 new eclipsers with periods longer than 2.5 hours, although most of these are not suitable for light curve modelling or require further observations before they could be suitable. We may therefore need at least another 5 to 10 more long period discoveries before we can be sure of measuring another 5 precise WD masses.

More masses and system parameters from CVs close to the period minimum are also helpful, as these systems often contain sub-stellar donors (Littlefair et al., 2006). This will allow the CV donor sequences to be refined (Knigge, Baraffe & Patterson, 2011), which will add further constraints to models of CV evolution. Light curve modelling of all the suitable eclipsing CVs discussed in Chapter 4 will be performed by Martin McAllister and future PhD students, and will be published in a future paper.

6.2.4 Fast Radio Bursts

ULTRASPEC will continue to be used to conduct optical follow-up observations of FRBs detected in real time, via target-of-opportunity overrides. In addition, we would also like to investigate the possibility of building an optical telescope dedicated to shadowing the field of view of the Parkes radio telescope, perhaps even located at the same site. Simultaneous observations of the same field as Parkes could drastically constrain the properties of optical afterglows of FRBs and would be of immense value in this field. Vik Dhillon has had preliminary discussions with our SUPERB collaborators regarding this possibility.

The requirements for such a telescope would be that it is fully automated, requiring no human input, and that it has a wide field of view, ideally one that could observe the entire Parkes beam, which is $2^\circ \times 2^\circ$. The telescope could use long exposure times, since the Parkes telescope tends to observe the same field for many tens of minutes, meaning that the telescope primary mirror need not be particularly large to reach a faint limiting magnitude for detection. For example, a 0.5m diameter telescope could reach a 5σ limiting magnitude for detection of around $R = 20$ in a single 10-minute exposure. Nevertheless, if sufficient funds were available it would be beneficial to use a larger primary mirror, as this would allow the telescope to probe deeper for afterglows, or equivalently to reach a certain limiting magnitude at a higher cadence. If optical afterglows of FRBs are related to shock blast-wave phenomena, their durations are likely to be short (peaking within one minute of the radio burst Yi, Gao & Zhang, 2014), and higher cadence observations (few minutes or less) will be beneficial.

The costs associated with such a project would be approximately £150,000 for a high-quality 0.5m telescope, dome and mount, excluding staff effort. One option that we have considered, which would make substantial savings to the costs of such a project, would be to mount the optical telescope directly onto the Parkes radio dish, in such a

way that it always pointed to the exact same field of view. However, it transpires that this idea is not feasible, since the Parkes telescope does not track the sky in the smooth and linear fashion required for optical observations.

6.2.5 HiPERCAM

Aside from the instrumentation and observational campaigns discussed in this thesis, there is another project that is heavily related and of great interest to time-resolved astrophysics and the study of transients. HiPERCAM is the next generation high-speed multicolour imager which is currently in development by Vik Dhillon and collaborators. The camera will provide a significant upgrade from ULTRACAM and ULTRASPEC, having 5 colour channels enabling simultaneous observations in the u', g', r', i' and z' bands. HiPERCAM will use multi-stage Peltier cooling to achieve CCD temperatures of around 180K, resulting in no significant dark current, and frame transfer CCDs with multiple read-out amplifiers and state of the art controllers, reaching very high frame rates ($>$ 1000Hz). In addition, the camera will include a facility for scintillation noise correction (Osborn et al., 2011). The instrument is expected to be commissioned in 2017 and will be mounted on the 4.2m WHT and later on the 10.4m Gran Telescopio Canarias.

Bibliography

Abbott B. P. et al., 2016a, ArXiv e-prints, 1602.08492

Abbott B. P. et al., 2016b, Phys. Rev. Lett., 116, 061102

Ahn C. P. et al., 2012, ApJS, 203, 21

Ak T., Retter A., Liu A., Esenoğlu H. H., 2005, PASA, 22, 105

Alard C., Lupton R. H., 1998, ApJ, 503, 325

Andersen M. F. et al., 2014, in Revista Mexicana de Astronomia y Astrofisica Conference Series, Vol. 45, p. 83

Antipin S. V., Kroll P., 2003, Information Bulletin on Variable Stars, 5461

Arkharov A. A., Borman G. A., Di Paola A., Larionov V. M., Morozova D. A., 2015, The Astronomer's Telegram, 6942

Bagchi M., Nieves A. C., McLaughlin M., 2012, MNRAS, 425, 2501

Balanutsa P. et al., 2013, The Astronomer's Telegram, 5688, 1

Balanutsa P. et al., 2014a, The Astronomer's Telegram, 6097, 1

Balanutsa P. et al., 2014b, The Astronomer's Telegram, 5974, 1

Bannister K. W., Murphy T., Gaensler B. M., Reynolds J. E., 2012, ApJ, 757, 38

Baptista R., Borges B. W., Bond H. E., Jablonski F., Steiner J. E., Grauer A. D., 2003, MNRAS, 345, 889

Baptista R., Bortoletto A., 2004, AJ, 128, 411

Barbieri C., Naletto G., Zampieri L., Verroi E., Gradari S., Collins S., Shearer A., 2012, in IAU Symposium, Vol. 285, New Horizons in Time Domain Astronomy, Griffin E., Hanisch R., Seaman R., eds., p. 280

Barry M. A. T., Gault D., Pavlov H., Hanna W., McEwan A., Filipović M. D., 2015, PASA, 32, e031

Barthelmy S. D., D'Ai A., D'Avanzo P., Krimm H. A., Lien A. Y., Marshall F. E., Maselli A., Siegel M. H., 2015, GRB Coordinates Network, 17929

Barthelmy S. D., Page K. L., Palmer D. M., 2015, GRB Coordinates Network, 18716

Bastian T. S., Benz A. O., Gary D. E., 1998, ARA&A, 36, 131

Bastian T. S., Bookbinder J., Dulk G. A., Davis M., 1990, ApJ, 353, 265

Batygin K., Stevenson D. J., 2010, ApJL, 714, L238

Bąkowska K., Olech A., Pospieszyński R., Martinelli F., Marciniak A., 2014, Acta Astron., 64, 337

Becker A., 2015, HOTPANTS: High Order Transform of PSF ANd Template Subtraction. Astrophysics Source Code Library

Bell C. P. M., Naylor T., Mayne N. J., Jeffries R. D., Littlefair S. P., 2012, MNRAS, 424, 3178

Berardi P., 2015, The Astronomer's Telegram, 7854

Bertin E., Arnouts S., 1996, A&AS, 117, 393

Betoule M. et al., 2014, A&A, 568, A22

Bikmaev I. F., Sakhibullin N. A., 2008, The Astronomer's Telegram, 1648

Bildsten L., Shen K. J., Weinberg N. N., Nelemans G., 2007, ApJL, 662, L95

Blagorodnova N., Van Velzen S., Harrison D. L., Koposov S., Mattila S., Campbell H., Walton N. A., Wyrzykowski Ł., 2016, MNRAS, 455, 603

Bloom J. S. et al., 2012, PASP, 124, 1175

Bodenheimer P., Lin D. N. C., Mardling R. A., 2001, ApJ, 548, 466

Bonetti L., Ellis J., Mavromatos N. E., Sakharov A. S., Sarkisyan-Grinbaum E. K. G., Spallicci A. D. A. M., 2016, ArXiv e-prints, 1602.09135

Bours M. C. P. et al., 2014, MNRAS, 445, 1924

Bours M. C. P. et al., 2015, MNRAS, 450, 3966

Bours M. C. P. et al., 2016, ArXiv e-prints, 1606.00780

Boyd D., Denisenko D., Koff R., Miller I., Staels B., 2011, Journal of the British Astronomical Association, 121, 233

Boyd R. W., 1978, Journal of the Optical Society of America, 68, 877

Breedt E. et al., 2014, MNRAS, 443, 3174

Bregman J. N., 2007, ARA&A, 45, 221

Brooks J., Bildsten L., Marchant P., Paxton B., 2015, ApJ, 807, 74

Brown T. M. et al., 2013, PASP, 125, 1031

Buat-Ménard V., Hameury J.-M., Lasota J.-P., 2001, A&A, 366, 612

Burke B. F., Graham-Smith F., 2014, An Introduction to Radio Astronomy. Cambridge University Press

Burke-Spolaor S., Bailes M., Ekers R., Macquart J.-P., Crawford, III F., 2011, ApJ, 727, 18

Burrows D. N. et al., 2011, Nat, 476, 421

Butters O. W. et al., 2010, A&A, 520, L10

Cagas P., Cagas P., 2014, Information Bulletin on Variable Stars, 6097, 1

Cai Y.-F., Sabancilar E., Steer D. A., Vachaspati T., 2012, Phys. Rev. D, 86, 043521

Campbell H., Fraser M., Hodgkin S., Davis C., Kostrzewa-Rutkowska Z., Piascik A. S., Steele I. A., Wyrzykowski L., 2015a, The Astronomer's Telegram, 7152, 1

Campbell H. et al., 2015b, The Astronomer's Telegram, 7177, 1

Campbell H. C. et al., 2015c, MNRAS, 452, 1060

Carter P. J. et al., 2013, MNRAS, 429, 2143

Carter P. J., Steeghs D., Marsh T. R., Kupfer T., Copperwheat C. M., Groot P. J., Nelemans G., 2014, MNRAS, 437, 2894

Castelli F., Kurucz R. L., 2004, ArXiv e-prints, 0405087

Catala L. et al., 2013, MNRAS, 436, 590

Champion D. J. et al., 2015, ArXiv e-prints, 1511.07746

Chote P., Sullivan D. J., Brown R., Harrold S. T., Winget D. E., Chandler D. W., 2014, MNRAS, 440, 1490

Collier Cameron A. et al., 2010, MNRAS, 407, 507

Collins P. P., Shehan B., Redfern M., Shearer A., 2009, ArXiv e-prints, 0905.0084

Connor L., Pen U.-L., Oppermann N., 2016, MNRAS, slw026

Connor L., Sievers J., Pen U.-L., 2016, MNRAS, 458, L19

Coppejans D. L. et al., 2014, MNRAS, 437, 510

Copperwheat C. M. et al., 2013, The Astronomer's Telegram, 5195, 1

Copperwheat C. M. et al., 2011, MNRAS, 410, 1113

Copperwheat C. M. et al., 2012, MNRAS, 421, 149

Cordes J. M., Lazio T. J. W., 2002, ArXiv Astrophysics e-prints, 0207156

Cordes J. M., Shannon R. M., 2008, ApJ, 682, 1152

Cordes J. M., Wasserman I., 2016, MNRAS, 457, 232

Crawford F., Rane A., Tran L., Rolph K., Lorimer D. R., Ridley J. P., 2016, ArXiv e-prints, 1605.06074

Dallilar Y. et al., 2016, in American Astronomical Society Meeting Abstracts, Vol. 227, American Astronomical Society Meeting Abstracts, p. 218.08

Darnley M. J., Ribeiro V. A. R. M., Bode M. F., Hounsell R. A., Williams R. P., 2012, ApJ, 746, 61

Davenport J. R. A., Becker A. C., Kowalski A. F., Hawley S. L., Schmidt S. J., Hilton E. J., Sesar B., Cutri R., 2012, ApJ, 748, 58

de Ugarte Postigo A. et al., 2010, in SPIE Conference Series, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, p. 77351F

Deil C., Domainko W., Hermann G., 2008, in American Institute of Physics Conference Series, Vol. 984, High Time Resolution Astrophysics: The Universe at Sub-Second Timescales, Phelan D., Ryan O., Shearer A., eds., p. 140

della Valle M., Panagia N., 1992, AJ, 104, 696

Deming D. et al., 2012, ApJ, 754, 106

Deng W., Zhang B., 2014, ApJL, 783, L35

Denisenko D. et al., 2014a, The Astronomer's Telegram, 6857, 1

Denisenko D. et al., 2014b, The Astronomer's Telegram, 6328, 1

Denisenko D. et al., 2014c, The Astronomer's Telegram, 6325, 1

Denisenko D. et al., 2013, The Astronomer's Telegram, 5698, 1

Denisenko D. V., Kryachko T. V., Satovskiy B. L., 2008, The Astronomer's Telegram, 1640

Dennison B., 2014, MNRAS, 443, L11

Dermer C. D., 2007, ArXiv e-prints, 0711.2804

Désert J.-M., Lecavelier des Etangs A., Hébrard G., Sing D. K., Ehrenreich D., Ferlet R., Vidal-Madjar A., 2009, ApJ, 699, 478

Dhillon V. S., 1996, in Cataclysmic Variables and Related Objects, Evans A., Wood J. H., eds., Kluwer Academic Publishers, Dordrecht, p. 3

Dhillon V. S. et al., 2014, MNRAS, 444, 4009

Dhillon V. S. et al., 2016, ArXiv e-prints, 1606.09214

Dhillon V. S. et al., 2007, MNRAS, 378, 825

Dhillon V. S., Smith D. A., Marsh T. R., 2013, MNRAS, 428, 3559

Dillon M. et al., 2008, MNRAS, 386, 1568

Djorgovski S. G. . et al., 2008a, The Astronomer's Telegram, 1416

Djorgovski S. G. et al., 2008b, Astronomische Nachrichten, 329, 263

Djorgovski S. G. et al., 2011, ArXiv e-prints, 1102.5004

Dong S. et al., 2015, The Astronomer's Telegram, 7850

Dotter A., Chaboyer B., Jevremović D., Kostov V., Baron E., Ferguson J. W., 2008, ApJS, 178, 89

Drake A. J. et al., 2009, ApJ, 696, 870

Drake A. J. et al., 2016, The Astronomer's Telegram, 9112

Drake A. J. et al., 2014, MNRAS, 441, 1186

Drake A. J., Williams R., Mahabal A., Djorgovski S. G., Graham M. J., Beshore E. C., Larson S. M., Christensen E., 2008, The Astronomer's Telegram, 1399

e2v, 2004, Low-light technical note 4. Tech. rep., e2v technologies

Eatough R. P., Keane E. F., Lyne A. G., 2009, MNRAS, 395, 410

Egorov A. E., Postnov K. A., 2009, Astronomy Letters, 35, 241

Falcke H., Rezzolla L., 2014, A&A, 562, 137

Feline W. J., Dhillon V. S., Marsh T. R., Watson C. A., Littlefair S. P., 2005, MNRAS, 364, 1158

Fender R. P. et al., 2006, in VI Microquasar Workshop: Microquasars and Beyond, p. 104.1

Fialkov A., Loeb A., 2016, ArXiv e-prints, 1602.08130

Fischer R., Tadic-Galeb B., Yoder P. R., 2008, Optical System Design. McGraw Hill Professional

Fitch W. S., Pacholczyk A. G., Weymann R. J., 1967, ApJL, 150, L67

Foley R. J. et al., 2013, ApJ, 767, 57

Fornasa M., Sánchez-Conde M. A., 2015, Physics Reports, 598, 1

Fortney J. J., Marley M. S., Barnes J. W., 2007, ApJ, 659, 1661

Fortney J. J., Nettelmann N., 2010, Space Sci. Rev., 152, 423

Frail D. A. et al., 2001, ApJL, 562, L55

French J. et al., 2004, in American Institute of Physics Conference Series, Vol. 727, Gamma-Ray Bursts: 30 Years of Discovery, Fenimore E., Galassi M., eds., p. 741

Frieman J. A., Turner M. S., Huterer D., 2008, ARA&A, 46, 385

Fritzewski D. J. et al., 2016, MNRAS, 462, 2396

Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, AJ, 111, 1748

Fukugita M., Yasuda N., Doi M., Gunn J. E., York D. G., 2011, AJ, 141, 47

Gal-Yam A., 2012, Science, 337, 927

Gandhi P., Dhillon V. S., Tomsick J. A., Butterley T., Littlefair S. M., Wilson R. W., Kennea J. A., 2014, The Astronomer's Telegram, 6447, 1

Gandhi P., Littlefair S., Hardy L., Dhillon V., Marsh T. R., Shaw A. W., 2015, The Astronomer's Telegram, 7686

Gandhi P. et al., 2016, MNRAS, 459, 554

Gänsicke B. T. et al., 2016, ApJL, 818, L7

Gänsicke B. T. et al., 2009, MNRAS, 397, 2170

Garai Z. et al., 2016, Astronomische Nachrichten, 337, 261

Garnavich P., Kennedy M., 2014, The Astronomer's Telegram, 6116, 1

Gasparrini D., 2013, The Astronomer's Telegram, 5232, 1

Gehrels N. et al., 2004, ApJ, 611, 1005

Gehrz R. D., Truran J. W., Williams R. E., Starrfield S., 1998, PASP, 110, 3

Geng J. J., Huang Y. F., 2015, ApJ, 809, 24

Gershberg R. E., Shakhovskaia N. I., 1983, Ap&SS, 95, 235

Gillon M., Jehin E., Magain P., Chantry V., Hutsemékers D., Manfroid J., Queloz D., Udry S., 2011, in European Physical Journal Web of Conferences, Vol. 11, p. 6002

Gorbovskoy E. S. et al., 2013, Astronomy Reports, 57, 233

Greenstein J. L., Schmidt M., 1964, ApJ, 140, 1

Greiss S. et al., 2012, AJ, 144, 24

Guillot T., Showman A. P., 2002, A&A, 385, 156

Gunn J. E., Stryker L. L., 1983, ApJS, 52, 121

Gwyn S. D. J., 2008, PASP, 120, 212

Haisch B., Strong K. T., Rodono M., 1991, ARA&A, 29, 275

Hankins T. H., Eilek J. A., 2007, ApJ, 670, 693

Hankins T. H., Kern J. S., Weatherall J. C., Eilek J. A., 2003, Nat, 422, 141

Harding L. K. et al., 2016, MNRAS, 457, 3036

Hardy L., Gandhi P., Dhillon V., Littlefair S., Butterley T., Wilson R., 2016a, The Astronomer's Telegram, 8501, 1

Hardy L., Littlefair S., Dhillon V., Butterley T., Wilson R., 2015a, The Astronomer's Telegram, 7681, 1

Hardy L. K., Butterley T., Dhillon V. S., Littlefair S. P., Wilson R. W., 2015b, MNRAS, 454, 4316

Hardy L. K., Dhillon V. S., Littlefair S. P., Butterley T., Wilson R. W., 2015c, GRB Coordinates Network, 18619

Hardy L. K., Littlefair S. P., Dhillon V. S., Butterley T., Wilson R. W., 2015d, GRB Coordinates Network, 18703

Hardy L. K. et al., 2016b, ArXiv e-prints, 1611.07885

Harrop-Allin M. K., Cropper M., Hakala P. J., Hellier C., Ramseyer T., 1999, MNRAS, 308, 807

Hassall T. E., Keane E. F., Fender R. P., 2013, MNRAS, 436, 371

Heinis S. et al., 2016, ArXiv e-prints, 1605.02194

Hellier C., 2001, Cataclysmic Variable Stars. Springer-Praxis, New York

Hemha N., Sanguansak N., Irawati P., Dhillon V., Marsh T. R., 2015, Publication of Korean Astronomical Society, 30, 201

Herrero E., Morales J. C., Ribas I., Naves R., 2011, A&A, 526, L10

Hessman F. V., 2000, New Astron. Rev., 44, 155

Hillman Y., Prialnik D., Kovetz A., Shara M. M., 2015, MNRAS, 446, 1924

Hillman Y., Prialnik D., Kovetz A., Shara M. M., 2016, ApJ, 819, 168

Hinshaw G. et al., 2013, ApJS, 208, 19

Holmes S. et al., 2011, PASP, 123, 1177

Holoien T. W.-S. et al., 2014, The Astronomer's Telegram, 6306, 1

Horne K., Wood J. H., Stiening R. F., 1991, ApJ, 378, 271

Howell D. A. et al., 2009, ArXiv e-prints, 0903.1086

Howell S. B., 2006, Handbook of CCD Astronomy. Cambridge University Press, Cambridge

Hynes R. I., Robinson E. L., Morales J., 2015, The Astronomer's Telegram, 7677

Inoue S., 2004, MNRAS, 348, 999

Inserra C. et al., 2015, ApJL, 799, L2

Ioka K., 2003, ApJL, 598, L79

Irawati P. et al., 2016, MNRAS, 456, 2446

Ives D., Bezawada N., Dhillon V., Marsh T., 2008, in SPIE Conference Series, Vol. 7021, High Energy, Optical, and Infrared Detectors for Astronomy III, p. 70210B

Ivezic Z. et al., 2008, ArXiv e-prints, 0805.2366

Jacoby G. H., Hunter D. A., Christian C. A., 1984, ApJS, 56, 257

Jacoby G. H. et al., 2010, PASA, 27, 156

Janesick J. R., 2001, Scientific Charge Coupled Devices. SPIE Press, Washington

Jerram P. et al., 2001, in SPIE Conference Series, Vol. 4306, SPIE Conference Series, Blouke M. M., Canosa J., Sampat N., eds., p. 178

Jessner A. et al., 2010, A&A, 524, A60

Jester S. et al., 2005, AJ, 130, 873

Kaiser N. et al., 2002, in SPIE Conference Series, Vol. 4836, Survey and Other Telescope Technologies and Discoveries, Tyson J. A., Wolff S., eds., p. 154

Kanbach G., Stefanescu A., Duscha S., Mühlegger M., Schrey F., Steinle H., Slowikowska A., Spruit H., 2008, in Astrophysics and Space Science Library, Vol. 351, Astrophysics and Space Science Library, Phelan D., Ryan O., Shearer A., eds., p. 153

Karastergiou A. et al., 2015, MNRAS, 452, 1254

Kashiyama K., Ioka K., Mészáros P., 2013, ApJL, 776, L39

Kasliwal M. M. et al., 2010, ApJL, 723, L98

Kato T. et al., 2015, PASJ, 67, 105

Kato T. et al., 2013, PASJ, 65, 23

Kato T. et al., 2014a, PASJ, 66, 30

Kato T., Ishioka R., Uemura M., 2002, PASJ, 54, 1023

Kato T., Maehara H., Uemura M., 2012, PASJ, 64, 63

Kato T. et al., 2014b, PASJ, 66, L7

Kato T., Uemura M., Ishioka R., Nogami D., Kunjaya C., Baba H., Yamaoka H., 2004, PASJ, 56, S1

Katz J. I., 2016a, Modern Physics Letters A, 31, 1630013

Katz J. I., 2016b, ApJ, 818, 19

Kaur A., Porter A., Wilber A., Wagner R. M., Woodward C. E., Starrfield S. G., Hartmann D. H., 2014, The Astronomer's Telegram, 6624, 1

Keane E. F. et al., 2016, Nature, 532, 453

Keane E. F., Stappers B. W., Kramer M., Lyne A. G., 2012, MNRAS, 425, L71

Keller S. C. et al., 2007, PASA, 24, 1

Kennedy M. R. et al., 2016, AJ, 152, 27

Kepler S. O., Kleinman S. J., Nitta A., Koester D., Castanheira B. G., Giovannini O., Althaus L., 2007, in Astronomical Society of the Pacific Conference Series, Vol. 372, 15th European Workshop on White Dwarfs, Napiwotzki R., Burleigh M. R., eds., p. 35

Khargharia J., Froning C. S., Robinson E. L., 2010, ApJ, 716, 1105

King A. R., Kolb U., Szuszkiewicz E., 1997, ApJ, 488, 89

King D. L., 1985, Atmospheric extinction at the roque de los muchachos observatory, la palma. Technical Note 31, Isaac Newton Group, La Palma

Klebesadel R. W., Strong I. B., Olson R. A., 1973, ApJL, 182, L85

Knigge C., 2006, MNRAS, 373, 484

Knigge C., Baraffe I., Patterson J., 2011, ApJS, 194, 28

Knight H. S., 2006, Chinese Journal of Astronomy and Astrophysics Supplement, 6, 41

Kocz J., Bailes M., Barnes D., Burke-Spolaor S., Levin L., 2012, MNRAS, 420, 271

Kovács G. et al., 2013, A&A, 553, A44

Krimm H. A., D'Elia V., Gronwall C., Izzo L., Page K. L., Palmer D. M., Ukwatta T. N., 2014, GRB Coordinates Network, 16788

Kukarkin B. V. et al., 1981, in Moscow, Acad. of Sciences USSR Sternberg, Vol. 1

Kulkarni S. R. et al., 1998, Nat, 395, 663

Kulkarni S. R., Ofek E. O., Neill J. D., Zheng Z., Juric M., 2014, ApJ, 797, 70

Kulkarni S. R., Rau A., 2006, ApJL, 644, L63

Kupfer T. et al., 2014, in Astronomical Society of the Pacific Conference Series, Vol. 481, 6th Meeting on Hot Subdwarf Stars and Related Objects, van Grootel V., Green E., Fontaine G., Charpinet S., eds., p. 293

Kuroda D. et al., 2014, GRB Coordinates Network, 16794

Kurtz D. W., Shibahashi H., Dhillon V. S., Marsh T. R., Littlefair S. P., Copperwheat C. M., Gänsicke B. T., Parsons S. G., 2013, MNRAS, 432, 1632

La Dous C., 1993, NASA Special Publication, 507, 95

Lang K. R., Willson R. F., 1986, ApJ, 305, 363

Larson S., Beshore E., Hill R., Christensen E., McLean D., Kolar S., McNaught R., Garradd G., 2003, in Bulletin of the American Astronomical Society, Vol. 35, AAS/Division for Planetary Sciences Meeting Abstracts #35, p. 982

Larson S., Brownlee J., Hergenrother C., Spahr T., 1998, in Bulletin of the American Astronomical Society, Vol. 30, Bulletin of the American Astronomical Society, p. 1037

Lasota J.-P., 2001, New Astron. Rev., 45, 449

Law N. M. et al., 2009, PASP, 121, 1395

Lehmann H., Guenther E., Sebastian D., Döllinger M., Hartmann M., Mkrtichian D. E., 2015, A&A, 578, L4

Levan A. J. et al., 2016, ApJ, 819, 51

Li F., Zhu C., Lü G., Wang Z., 2016a, PASJ, 68, 39

Li K. L., Kong A. K. H., 2016, The Astronomer's Telegram, 8523

Li L., Huang Y., Zhang Z., Li D., Li B., 2016b, ArXiv e-prints, 1602.06099

Li Y., Zhang B., 2016, ArXiv e-prints, 1603.04825

Liebert J., Bergeron P., Holberg J. B., 2003, AJ, 125, 348

Lipunov V. et al., 2010, Advances in Astronomy, 30, 349171

Lipunov V. M., Pruzhinskaya M. V., 2014, MNRAS, 440, 1193

Lipunova G. V., 1997, Astronomy Letters, 23, 84

Littlefair S. P. et al., 2014, MNRAS, 445, 2106

Littlefair S. P., Dhillon V. S., Marsh T. R., Gänsicke B. T., Southworth J., Baraffe I., Watson C. A., Copperwheat C., 2008, MNRAS, 388, 1582

Littlefair S. P., Dhillon V. S., Marsh T. R., Gänsicke B. T., Southworth J., Watson C. A., 2006, Sci, 314, 1578

Littlefair S. P. et al., 2013, MNRAS, 431, 2820

Loeb A., Shvartzvald Y., Maoz D., 2014, MNRAS, 439, L46

Löhmer O., Jessner A., Kramer M., Wielebinski R., Maron O., 2008, A&A, 480, 623

Lohr M. E., Norton A. J., Kolb U. C., Maxted P. F. L., Todd I., West R. G., 2013, A&A, 549, A86

Lomb N. R., 1976, Ap&SS, 39, 447

Longa-Peña P., Steeghs D., Marsh T., 2015, MNRAS, 447, 149

Lorimer D., Kramer M., 2005, Handbook of Pulsar Astronomy. Cambridge University Press, Cambridge

Lorimer D. R., Bailes M., McLaughlin M. A., Narkevic D. J., Crawford F., 2007, Science, 318, 777

Lovegrove E., Woosley S. E., 2013, ApJ, 769, 109

Luan J., Goldreich P., 2014, ApJL, 785, 26

Lyubarsky Y., 2014, MNRAS, 442, L9

Lyutikov M., Burzawa L., Popov S. B., 2016, ArXiv e-prints, 1603.02891

Mackay C. D., Tubbs R. N., Bell R., Burt D. J., Jerram P., Moody I., 2001, in SPIE Conference Series, Vol. 4306, Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications II, Blouke M. M., Canosa J., Sampat N., eds., p. 289

Macquart J.-P., Koay J. Y., 2013, ApJ, 776, 125

Mahabal A. A. et al., 2011, Bulletin of the Astronomical Society of India, 39, 387

Majid W. A., Naudet C. J., Lowe S. T., Kuiper T. B. H., 2011, ApJ, 741, 53

Manchester R. N., Fan G., Lyne A. G., Kaspi V. M., Crawford F., 2006, ApJ, 649, 235

Manchester R. N., Taylor J. H., 1972, Astrophys. Lett., 10, 67

Manser C. J., Gänsicke B. T., 2014, MNRAS, 442, L23

Mao S., Paczynski B., 1991, ApJL, 374, L37

Mao Y.-N., Lu X.-M., Wang J.-F., Jiang X.-J., 2013, Research in Astronomy and Astrophysics, 13, 239

Maoz D. et al., 2015, MNRAS, 454, 2183

Maoz D., Mannucci F., 2012, Publications of the Astronomical Society of Australia, 29, 447

Marsh T. R. et al., 2007, in Astronomical Society of the Pacific Conference Series, Vol. 372, 15th European Workshop on White Dwarfs, Napiwotzki R., Burleigh M. R., eds., p. 431

Martí J., Luque-Escamilla P. L., García-Hernández M. T., 2016, A&A, 586, A58

Martin D. C. et al., 2005, ApJL, 619, L1

Masetti N., Chavushyan V., Torrealba J., Jimenez-Bailon E., 2014, The Astronomer's Telegram, 6588, 1

Masui K. et al., 2015, Nat, 528, 523

Maxted P. F. L. et al., 2013, Nat, 498, 463

Mazin B. A. et al., 2013, PASP, 125, 1348

McAllister M. J. et al., 2015, MNRAS, 451, 114

McLaughlin M. A. et al., 2006, Nat, 439, 817

McQuinn M., 2014, ApJL, 780, L33

Meegan C. A., Fishman G. J., Wilson R. B., Horack J. M., Brock M. N., Paciesas W. S., Pendleton G. N., Kouveliotou C., 1992, Nat, 355, 143

Metzger B. D. et al., 2010, MNRAS, 406, 2650

Metzger M. R., Djorgovski S. G., Kulkarni S. R., Steidel C. C., Adelberger K. L., Frail D. A., Costa E., Frontera F., 1997, Nat, 387, 878

Michel R., Echevarría J., Hernández Santisteban J. V., 2013, A&A, 554, A25

Miszalski B. et al., 2016, MNRAS, 456, 633

Moffett T. J., 1974, ApJS, 29, 1

Monet D. G. et al., 2003, AJ, 125, 984

Morris T. et al., 2009, in Frontiers in Optics 2009/Laser Science XXV/Fall 2009 OSA Optics & Photonics Technical Digest, Optical Society of America, p. AOTuA5

Motta S. E., Sanchez-Fernandez C., Kuulkers E., Kajava J., Bozzo E., 2016, The Astronomer's Telegram, 8500, 1

Mottez F., Zarka P., 2014, A&A, 569, A86

Moya A., Bouy H., Marchis F., Vicente B., Barrado D., 2011, A&A, 535, A110

Munari U. et al., 2002, A&A, 389, L51

Mushotzky R. F., Done C., Pounds K. A., 1993, ARA&A, 31, 717

Nakar E., 2007, Physics Reports, 442, 166

Nather R. E., Mukadam A. S., 2004, ApJ, 605, 846

Naylor T., 1998, MNRAS, 296, 339

Negoro H. et al., 2015, The Astronomer's Telegram, 7646

Nelemans G., 2005, in Astronomical Society of the Pacific Conference Series, Vol. 330, The Astrophysics of Cataclysmic Variables and Related Objects, Hameury J.-M., Lasota J.-P., eds., p. 27

Nelemans G., Siess L., Repetto S., Toonen S., Phinney E. S., 2016, ApJ, 817, 69

Nicholl M. et al., 2016, ArXiv e-prints, 1603.04748

O'Donoghue et al D., 2006, MNRAS, 372, 151

Osaki Y., 1996, PASP, 108, 39

Osborn J., Wilson R. W., Dhillon V. S., Avila R., Love G. D., 2011, MNRAS, 411, 1223

Palaversa L. et al., 2013, AJ, 146, 101

Parsons S. G. et al., 2015, MNRAS, 449, 2194

Parsons S. G., Marsh T. R., Copperwheat C. M., Dhillon V. S., Littlefair S. P., Gänsicke B. T., Hickman R., 2010a, MNRAS, 402, 2591

Parsons S. G. et al., 2010b, MNRAS, 407, 2362

Patterson J. et al., 2005, PASP, 117, 1204

Patterson J. et al., 2013, MNRAS, 434, 1902

Pearson K. J., 2007, MNRAS, 379, 183

Pejcha O., Metzger B. D., Tomida K., 2016, MNRAS, 455, 4351

Petroff E. et al., 2015a, MNRAS, 447, 246

Petroff E. et al., 2016, ArXiv e-prints, 1601.03547

Petroff E. et al., 2015b, MNRAS, 454, 457

Petroff E. et al., 2015c, MNRAS, 451, 3933

Phillips M. M., 1993, ApJ, 413, L105

Pickles A. J., 1998, PASP, 110, 863

Popov S. B., Postnov K. A., 2010, in Evolution of Cosmic Objects through their Physical Activity, Harutyunian H. A., Mickaelian A. M., Terzian Y., eds., p. 129

Popov S. B., Pshirkov M. S., 2016, ArXiv e-prints, 1605.01992

Poznanski D. et al., 2010, Science, 327, 58

Prialnik D., Kovetz A., 1995, ApJ, 445, 789

Pshirkov M. S., Postnov K. A., 2010, Ap&SS, 330, 13

Pyrzas S. et al., 2012, PASP, 124, 204

Qian S. B., Han Z. T., Fernández Lajús E., Zhu L. Y., Li L. J., Liao W. P., Zhao E. G., 2015, ApJS, 221, 17

Quimby R. M., 2012, in IAU Symposium, Vol. 279, Death of Massive Stars: Supernovae and Gamma-Ray Bursts, Roming P., Kawai N., Pian E., eds., p. 22

Ramanjooloo Y., Jones G. H., Coates A. J., Owens M. J., 2015, in AAS/Division for Planetary Sciences Meeting Abstracts, Vol. 47, AAS/Division for Planetary Sciences Meeting Abstracts, p. 506.01

Ramsay G., Doyle J. G., Hakala P., Garcia-Alvarez D., Brooks A., Barclay T., Still M., 2013, MNRAS, 434, 2451

Rane A., Lorimer D. R., Bates S. D., McMann N., McLaughlin M. A., Rajwade K., 2016, MNRAS, 455, 2207

Rau A. et al., 2009, PASP, 121, 1334

Ravi V., Lasky P. D., 2014, MNRAS, 441, 2433

Ravi V., Shannon R. M., Jameson A., 2015, ApJL, 799, L5

Rees M. J., 1977, Nat, 266, 333

Rees M. J., 1988, Nat, 333, 523

Rest A., Mndermann L., Widenhorn R., Bodegom E., McGlinn T. C., 2002, Review of Scientific Instruments, 73, 2028

Retter A., Marom A., 2003, MNRAS, 345, L25

Rich R. M., Mould J., Picard A., Frogel J. A., Davies R., 1989, ApJL, 341, L51

Richichi A., Irawati P., Soonthornthum B., Dhillon V. S., Marsh T. R., 2014, AJ, 148, 100

Richichi A., Tasuya O., Irawati P., Soonthornthum B., Dhillon V. S., Marsh T. R., 2016, AJ, 151, 10

Ritter H., Kolb U., 2003, A&A, 404, 301, (Update RKcat 7.23, 2015)

Rixon G. et al., 2014, The Astronomer's Telegram, 6593, 1

Rodrigues C. V. et al., 2012, in American Institute of Physics Conference Series, Vol. 1429, American Institute of Physics Conference Series, Hoffman J. L., Bjorkman J., Whitney B., eds., p. 252

Sackett P. D., 1999, in Series C, Vol. 532, NATO Advanced Science Institutes (ASI), Mariotti J.-M., Alloin D., eds., p. 189

Saito R. K., Baptista R., 2006, AJ, 131, 2185

Sanduleak N., 1972, Information Bulletin on Variable Stars, 663

Savoury C. D. J. et al., 2011, MNRAS, 415, 2025

Savoury C. D. J., Littlefair S. P., Marsh T. R., Dhillon V. S., Parsons S. G., Copperwheat C. M., Steeghs D., 2012, MNRAS, 422, 469

Scargle J. D., 1982, ApJ, 263, 835

Schaefer B. E., King J. R., Deliyannis C. P., 2000, ApJ, 529, 1026

Schenker K., King A. R., 2002, in Astronomical Society of the Pacific Conference Series, Vol. 261, The Physics of Cataclysmic Variables and Related Objects, Gänsicke B. T., Beuermann K., Reinsch K., eds., p. 242

Schmidtobreick L., Rodríguez-Gil P., Gänsicke B. T., 2012, Mem. Soc. Astron. Ital., 83, 610

Schmitt J. H. M. M. et al., 2014, Astronomische Nachrichten, 335, 787

Schneider J., Dedieu C., Le Sidaner P., Savalle R., Zolotukhin I., 2011, A&A, 532, A79

Scholz P. et al., 2016, ArXiv e-prints, 1603.08880

Schreiber M. R., Zorotovic M., Wijnen T. P. G., 2016, MNRAS, 455, L16

Scott E. R. D., 2006, in Lunar and Planetary Science Conference, Vol. 37, 37th Annual Lunar and Planetary Science Conference, Mackwell S., Stansbery E., eds., p. 2367

Seaman R., Allan A., Williams R., 2008, in Astronomical Society of the Pacific Conference Series, Vol. 394, Astronomical Data Analysis Software and Systems XVII, Argyle R. W., Bunclark P. S., Lewis J. R., eds., p. 213

Shahbaz T. et al., 2015, MNRAS, 453, 3461

Shappee B. J. et al., 2014, ApJ, 788, 48

Shara M. M., Yaron O., Prialnik D., Kovetz A., Zurek D., 2010a, ApJ, 725, 831

Shara M. M., Zurek D., Prialnik D., Yaron O., Kovetz A., 2010b, ApJ, 725, 824

Shears J., Hambsch F.-J., Littlefield C., Miller I., Morelle E., Pickard R., Pietz J., Sabo R., 2012, ArXiv e-prints, 1209.4062

Shen K. J., Bildsten L., 2014, ApJ, 785, 61

Shumkov V. et al., 2014a, The Astronomer's Telegram, 6117, 1

Shumkov V. et al., 2014b, The Astronomer's Telegram, 6130

Shurpakov S. et al., 2014a, The Astronomer's Telegram, 5774, 1

Shurpakov S. et al., 2014b, The Astronomer's Telegram, 6257, 1

Sivakoff G. R., Bahramian A., Altamirano D., Beardmore A. P., Kuulkers E., Motta S., 2015, The Astronomer's Telegram, 7959

Smak J., 1984, PASP, 96, 5

Smith A. M. S., Anderson D. R., Skillen I., Collier Cameron A., Smalley B., 2011, MNRAS, 416, 2096

Smith D. A., Dhillon V. S., 1998, MNRAS, 301, 767

Smith N. et al., 2016, MNRAS, 458, 950

Southworth J., Copperwheat C. M., 2011, The Observatory, 131, 66

Southworth J., Copperwheat C. M., Gänsicke B. T., Pyrzas S., 2010, A&A, 510, 100

Southworth J., Hickman R. D. G., Marsh T. R., Rebassa-Mansergas A., Gänsicke B. T., Copperwheat C. M., Rodríguez-Gil P., 2009, A&A, 507, 929

Southworth J., Marsh T. R., Gänsicke B. T., Aungwerojwit A., Hakala P., de Martino D., Lehto H., 2007, MNRAS, 382, 1145

Southworth J., Tappert C., Gänsicke B. T., Copperwheat C. M., 2015, A&A, 573, 61

Spitler L. G. et al., 2014, ApJ, 790, 101

Spitler L. G. et al., 2016, Nature, Online, doi:10.1038/nature17168

Spitzer L., 1998, Physical Processes in the Interstellar Medium. Wiley Classics Library, New York

Staley T. D., Fender R., 2016, ArXiv e-prints, 1606.03735

Staley T. D. et al., 2013, MNRAS, 428, 3114

Stanek K. Z. et al., 2014a, The Astronomer's Telegram, 6233, 1

Stanek K. Z. et al., 2014b, The Astronomer's Telegram, 6479, 1

Starrfield S., Iliadis C., Hix W. R., 2016, PASP, 128, 051001

Starrfield S., Sparks W. M., Truran J. W., 1985, ApJ, 291, 136

Steeghs D., Perryman M. A. C., Reynolds A., de Bruijne J. H. J., Marsh T., Dhillon V. S., Peacock A., 2003, MNRAS, 339, 810

Steele I. A. et al., 2004, in SPIE Conference Series, Vol. 5489, Ground-based Telescopes, Oschmann Jr. J. M., ed., pp. 679–692

Stellingwerf R. F., 1978, ApJ, 224, 953

Street R. A., Tsapras Y., Hundertmark M. P. G., Horne K. D., Dominik M., Snodgrass C., Bramich D. M., Figuera Jaimes R., 2014, in The Third Hot-wiring the Transient Universe Workshop, Wozniak P. R., Graham M. J., Mahabal A. A., Seaman R., eds., p. 189

Strolger L. G. et al., 2014, AJ, 147, 49

Szkody P. et al., 2002, AJ, 123, 430

Szkody P. et al., 2011, AJ, 142, 181

Szkody P. et al., 2009, AJ, 137, 4011

Szkody P. et al., 2003, AJ, 126, 1499

Szkody P. et al., 2006, AJ, 131, 973

Szkody P. et al., 2004, AJ, 128, 1882

Szkody P. et al., 2005, AJ, 129, 2386

Szkody P. et al., 2007, AJ, 134, 185

Tendulkar S. P., Kaspi V. M., Patel C., 2016, ArXiv e-prints, 1602.02188

Tetarenko A., Sivakoff G., Miller-Jones J., Coulson I., 2016, The Astronomer's Telegram, 8499, 1

Thompson T. A., Prieto J. L., Stanek K. Z., Kistler M. D., Beacom J. F., Kochanek C. S., 2009, ApJ, 705, 1364

Thornton D. et al., 2013, Science, 341, 53

Thorstensen J. R., 2000, PASP, 112, 1269

Thorstensen J. R., 2016, The Astronomer's Telegram, 9141

Thorstensen J. R., Ringwald F. A., Wade R. A., Schmidt G. D., Norsworthy J. E., 1991, AJ, 102, 272

Thorstensen J. R., Skinner J. N., 2012, AJ, 144, 81

Tingay S. J., Kaplan D. L., 2016, ArXiv e-prints, 1602.07643

Tiurina N. et al., 2013, The Astronomer's Telegram, 4871, 1

Topinka M., Hanlon L., Meehan S., Tisdall P., Jelínek M., Kubánek P., van Heerden H., Meintjes P., 2014, in Revista Mexicana de Astronomia y Astrofisica Conference Series, Vol. 45, Revista Mexicana de Astronomia y Astrofisica Conference Series, p. 65

Totani T., 2013, PASJ, 65, L12

Tovmassian G., Stephania Hernandez M., González-Buitrago D., Zharikov S., García-Díaz M. T., 2014, AJ, 147, 68

Townsley D. M., Bildsten L., 2004, ApJ, 600, 390

Tsapras Y. et al., 2009, Astronomische Nachrichten, 330, 4

Tubbs R. N., 2003, PhD thesis, University of Cambridge

Tulloch S. M., 2004a, in SPIE Conference Series, Vol. 5490, Advancements in Adaptive Optics, Bonaccini Calia D., Ellerbroek B. L., Ragazzoni R., eds., p. 1167

Tulloch S. M., 2004b, in SPIE Conference Series, Vol. 5492, Ground-based Instrumentation for Astronomy, Moorwood A. F. M., Iye M., eds., p. 604

Tulloch S. M., 2010, PhD thesis, University of Sheffield

Tulloch S. M., Dhillon V. S., 2011, MNRAS, 411, 211

Tulloch S. M., Rodríguez-Gil P., Dhillon V. S., 2009, MNRAS, 397, L82

Tuntsov A. V., 2014, MNRAS, 441, L26

Tylenda R., Górny S. K., Kamiński T., Schmidt M., 2015, A&A, 578, A75

Tylenda R. et al., 2011, A&A, 528, A114

Tylenda R., Soker N., 2006, A&A, 451, 223

Udalski A., Kubiak M., Szymanski M., 1997, Acta Astron., 47, 319

Udalski A. et al., 2002, Acta Astron., 52, 1

Uemura M. et al., 2000, PASJ, 52, L9

Vachaspati T., 2008, Physical Review Letters, 101, 141301

van der Horst A. J. et al., 2015, MNRAS, 446, 4116

van Paradijs J. et al., 1997, Nat, 386, 686

van Velzen S., Körding E., Falcke H., 2011, MNRAS, 417, L51

Verhoeve P., Martin D. D. E., Hijmering R. A., Verveer J., van Dordrecht A., Sirbi G., Oosterbroek T., Peacock A., 2006, Nuclear Instruments and Methods in Physics Research A, 559, 598

Voges W. et al., 1999, A&A, 349, 389

Voges W. et al., 2000, IAU Circular, 7432, 1

von Essen C. et al., 2014, A&A, 561, A48

Vrielmann S., Stiening R. F., Offutt W., 2002, MNRAS, 334, 608

Wang B., Han Z., 2012, New Astron. Rev., 56, 122

Wang J.-S., Yang Y.-P., Wu X.-F., Dai Z.-G., Wang F.-Y., 2016, ArXiv e-prints, 1603.02014

Wang S.-Y. et al., 2014, in SPIE Conference Series, Vol. 9154, High Energy, Optical, and Infrared Detectors for Astronomy VI, p. 91542I

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

Warner B., Cropper M., 1983, MNRAS, 203, 909

Whitehurst R., 1988, MNRAS, 232, 35

Wijnen T. P. G., Zorotovic M., Schreiber M. R., 2015, A&A, 577, A143

Williams P. K. G., Berger E., 2016, ArXiv e-prints, 1602.08434

Williams P. K. G., Berger E., Chornock R., 2016, The Astronomer's Telegram, 8752, 1

Wils P., Gänsicke B. T., Drake A. J., Southworth J., 2010, MNRAS, 402, 436

Wils P., Krajci T., Hambsch F.-J., Muyllaert E., 2011, Information Bulletin on Variable Stars, 5982

Wilson R. W., 2002, MNRAS, 337, 103

Wilson R. W., O'Mahony N., Packham C., Azzaro M., 1999, MNRAS, 309, 379

Woosley S., Janka T., 2005, Nature Physics, 1, 147

Woosley S. E., Bloom J. S., 2006, ARA&A, 44, 507

Woudt P. A., Warner B., 2010, MNRAS, 403, 398

Woudt P. A., Warner B., 2011, The Astronomer's Telegram, 3705

Woudt P. A., Warner B., de Budé D., Macfarlane S., Schurch M. P. E., Zietsman E., 2012, MNRAS, 421, 2414

Wright D. E. et al., 2015, MNRAS, 449, 451

Wu X.-F. et al., 2016, ArXiv e-prints, 1602.07835

Wu Z.-Y., Zhou X., Ma J., Du C.-H., 2009, MNRAS, 399, 2146

Wyrzykowski Ł., 2016, ArXiv e-prints, 1601.02827

Wyrzykowski L., Hodgkin S., Blogorodnova N., Koposov S., Burgon R., 2012, ArXiv e-prints, 1210.5007

Yakin D. G., Suleimanov V. F., Shimansky V. V., Borisov N. V., Bikmaev I. F., Sakhibullin N. A., 2010, in American Institute of Physics Conference Series, Vol. 1273, American Institute of Physics Conference Series, Werner K., Rauch T., eds., p. 346

Yaron O., Prialnik D., Shara M. M., Kovetz A., 2005, ApJ, 623, 398

Yecheistov V. et al., 2014, The Astronomer's Telegram, 6300, 1

Yi S.-X., Gao H., Zhang B., 2014, ApJL, 792, L21

Zacharias N., Monet D. G., Levine S. E., Urban S. E., Gaume R., Wycoff G. L., 2004, in Bulletin of the American Astronomical Society, Vol. 36, American Astronomical Society Meeting Abstracts, p. 1418

Zampieri L. et al., 2015a, in SPIE Conference Series, Vol. 9504, Photon Counting Applications 2015, p. 95040C

Zampieri L., Verroi E., Barbieri M., Naletto G., Barbieri C., 2015b, The Astronomer's Telegram, 6875

Zampieri L., Verroi E., Naletto G., Barbieri C., Barbieri M., 2015c, The Astronomer's Telegram, 7351

Zejda M., Pejcha O., 2016, The Astronomer's Telegram, 9132

Zengin Çamurdan D., İbanoğlu C., Çamurdan C. M., 2010, New Astronomy, 15, 476

Zerbi R. M. et al., 2001, Astronomische Nachrichten, 322, 275

Zhang B., 2014, ApJL, 780, L21

Zhang B., 2016a, ArXiv e-prints, 1602.04542

Zhang B., 2016b, ArXiv e-prints, 1602.08086

Zhou B., Li X., Wang T., Fan Y.-Z., Wei D.-M., 2014, Phys. Rev. D, 89, 107303

Zorotovic M., Schreiber M. R., Gänsicke B. T., 2011, A&A, 536, A42

Życki P. T., Done C., Smith D. A., 1999, MNRAS, 309, 561

Appendix A

Classes of Astronomical Transients

Figure 1.1 shows the majority of known or predicted transient events in a plot of peak magnitude as a function of decay time. There are many types of transient event, and possibly even more which we are yet to discover or conceive. Here I provide a brief overview of the most common and most relevant classes of transient shown in Figure 1.1, with regards to the development of our follow-up facilities.

Cataclysmic Variables - Dwarf Novae

One of the best studied classes of transients are cataclysmic variables (CVs). CVs are compact binary systems in which a white dwarf (WD) accretes matter from a (usually cool) main sequence companion, via Roche lobe overflow. Often the transfer of gas is via an accretion disk around the WD, conserving angular momentum. Sometimes, particularly in highly magnetic systems, the gas streams directly on to the WD at the poles, following the magnetic field lines. These systems are known as polars. Some CVs undergo 'dwarf nova outbursts', when the viscosity in the accretion disk changes and it heats up. The system can brighten by several magnitudes, and the outburst might usually last a few days. Outbursts can re-occur semi-periodically, every few weeks, months or years. These systems are known as dwarf novae (DNe), and the dwarf nova outburst is the most common way to discover new CVs, via transient surveys. DNe are discussed in more detail in Chapter 4.

Cataclysmic Variables - Classical and Recurrent Novae

Classical or recurrent novae also originate in CVs. Nova eruptions occur when the WD accretes sufficient gas to trigger a thermonuclear runaway explosion on the surface of the WD (Hellier, 2001; Starrfield, Iliadis & Hix, 2016). The result is a brightening of

up to 15 magnitudes, which can live for up to a few hundred days, driven mostly by radioactivity as the explosion creates a shell of expanding material. Novae are thought to be responsible for part of the chemical enrichment of the galaxy, particularly of ^{13}C, ^{15}N, and ^{17}O (Starrfield, Iliadis & Hix, 2016; Li et al., 2016a) and of other peculiar isotopes (Gehrz et al., 1998). Classical novae are, by definition, not seen to re-occur, however some 'recurrent novae' are known, with amplitudes and repetition timescales which bridge the gap between novae and dwarf novae (i.e. a few hundred years, Hellier, 2001). Nova explosions may cause a net loss of mass, suggesting that recurrent novae see their WDs lose mass over time, despite the ongoing accretion (Prialnik & Kovetz, 1995; Townsley & Bildsten, 2004; Yaron et al., 2005). This would rule out the single degenerate scenario for Type Ia supernovae (SNIa, Wang & Han, 2012), which requires a white dwarf to grow in mass and exceed the Chandrasekhar limit at 1.4 $M_{\odot}$. Despite this, some argue that WDs can in fact steadily grow in mass, even through many nova cycles (Hillman et al., 2015, 2016). Again, this is discussed further in Chapter 4.

Type Ia Supernovae

Type Ia SNe are produced by the total thermonuclear deflagration of a white dwarf, and form a key component in the cosmic distance ladder, being a 'standardisable candle' after scaling the shape of the light curve and the colours (della Valle & Panagia, 1992; Phillips, 1993). Type Ia SNe have since been famously used to detect and characterise the acceleration in the expansion of the Universe, and of course the presence of dark energy (see Frieman, Turner & Huterer, 2008 for a review). New SNe Ia are being detected on a daily basis, but dedicated and detailed photometric and spectroscopic follow-up observations are needed for each event in order for them to be used as distance indicators (see e.g. Betoule et al., 2014). Still, there are several systematic uncertainties remaining in these calibrations, particularly regarding accurate flux calibration for low-redshift SNe (Howell et al., 2009). Studying Type Ia SNe at a wider range of redshifts should help in understanding their progenitors, which in turn would help greatly with reducing the systematic uncertainties (Howell et al., 2009). This thesis does not study Type Ia SNe in detail, though we have observed a few long-term, multi-colour light curves - see Chapter 3.

Core-Collapse Supernovae

More common and often more energetic than Type Ia SNe, core-collapse supernovae represent the violent deaths of the most massive stars, releasing 10^{46}J of energy, although

usually most of this is in the form of neutrinos (Woosley & Janka, 2005). They also produce most of the heavy element content in the Universe. There is much still to be learnt regarding the end stages of stellar evolution, and many peculiar core-collapse SNe events continue to be detected - for example, superluminous supernovae (SLSNe), which are discussed in more detail in Chapter 2.

Peculiar Supernovae: Type Iax and Type .Ia

A recently discovered class of transient are the peculiar supernovae which are observationally linked to Type Ia supernovae. The Type Iax class are thought to be about one third as common as normal SNe Ia, although only a few tens have so far been discovered (see Foley et al., 2013 and references therein). Type Iax SNe appear spectroscopically similar to Type Ia SNe, with no sign of hydrogen, and strong SiII features, but have significantly lower maximum-light velocities ($v < 8000$km s^{-1}) and lower peak magnitudes ($M_V > -18$). Signs of CO burning, low ejecta masses and strong Fe lines, as well as no evidence of massive stars existing nearby to any progenitor locations, all suggest the progenitor is a WD, and leading theories for the formation of Type Iax SNe involve a primary WD accreting gas from a Helium star (either another WD or a non-degenerate He-star). Deflagration of the accreted helium layers is thought to drive the supernova explosion, which may also involve full detonation of the WD core (Foley et al., 2013).

Observationally, and perhaps physically similar to Type Iax SNe, are the Type .Ia supernovae. First hypothesised by Bildsten et al. (2007), Type .Ia SNe are predicted to occur in some AM CVn-type binaries, in which a CO WD accretes helium from its WD companion. Once the orbital period drops below 25 minutes, the rate of mass accretion is insufficient to drive stable helium burning on the surface of the WD, and it burns instead via unstable thermonuclear flashes. These flashes may widen the binary orbit, which reduces the rate of accretion, and induces more violent flashes. The final flash may be sufficiently energetic to burn helium into radioactive elements which drive a faint supernova. Such an event would shine approximately one-tenth as bright ($M_V \ -16$) and for one-tenth the duration as a normal SN Ia, hence the name '.Ia' (Bildsten et al., 2007; Brooks et al., 2015). Type .Ia supernovae may occur at a rate of approximately 5% that of Type Ia SN, though only a few possible candidates have so far been observed (Poznanski et al., 2010; Kasliwal et al., 2010; Inserra et al., 2015). It is thought that some AM CVn systems undergoing these events may also trigger a secondary detonation of the CO core itself, which could drive a full blown Type Ia supernova (Shen & Bildsten, 2014).

Luminous Red Novae

Luminous red novae (LRNe), the first of which was discovered serendipitously in a study of red supergiants in M31 (Rich et al., 1989), remain an understudied and mysterious class of transients. These red transients bridge the luminosity gap between classical novae and supernovae, as seen in Figure 1.1. LRNe tend to undergo slow outbursts, with rise times of tens of days, and long optical plateaus, particularly in the red, of up to 100 days. The spectra of these red transients show H and CaII emission lines, as well as NaI absorption lines. Some events have shown multiple peaks in their light curves (e.g V838 Mon, Munari et al., 2002) and some appear to be deeply embedded in large dust clouds (Thompson et al., 2009). There are many competing theories which might explain the progenitors of LRNe. Top contenders include an extension of the classical nova scenario (Shara et al., 2010b,a) and a 'mergeburst' scenario in which a main sequence star of considerable mass (e.g. $8M_{\odot}$) merges with a low mass star, pulled into an eccentric orbit by a wide companion (Tylenda & Soker, 2006; Tylenda et al., 2015; Smith et al., 2016). Other possibilities that have been discussed include low-energy supernovae depleted by neutrinos (Lovegrove & Woosley, 2013), mass-losing binary outflows (Pejcha, Metzger & Tomida, 2016), and the swallowing of close-in giant planets by an inflated RGB star (Retter & Marom, 2003).

AGN/Blazars

Active Galactic Nuclei (AGN) are highly luminous sources, powered by the accretion of matter onto a supermassive black hole (BH) in the centres of galaxies. The most powerful of AGN (distant, high redshift quasars) are the most persistently luminous objects in the Universe ($4 \times 10^{12} \mathrm{L}_{\odot}$, Greenstein & Schmidt, 1964). AGN display irregular and non-periodic optical variability on timescales of minutes to years, with higher amplitude variations occuring on longer timescales (Fitch, Pacholczyk & Weymann, 1967). As well as optical variability, AGN are violently variable in the X-ray regime. In fact, because the observed rapid X-ray variability requires a relatively small emitting region (due to the short light travel time), these observations became a key driver for the supermassive blackhole accretion theory for the engines of AGN (see Mushotzky, Done & Pounds, 1993 and references therein). In addition, variability can be a powerful tool for avoiding selection effects in discovering/selecting AGN for futher study, since it is much less biased against low-luminosity sources (Heinis et al., 2016), with lower-luminosity AGN providing relatively higher amplitude variability above the host galaxy background.

Blazars are a particular class of AGN, in which we observe the powerful relativistic

jets from the host galaxy directly along the line of the jet. Instabilities in these outflows are the main source of variability in blazars. Blazars are major contributors to the cosmic gamma-ray background (Fornasa & Sánchez-Conde, 2015), as well as being a source of ultra-high energy cosmic rays (Dermer, 2007). Selecting AGN by their blazar variability is also useful in removing some selection effects (Rau et al., 2009) The stochastic nature of their variability means that AGN are often detected in transient surveys.

Tidal Disruption Events

Tidal disruption events (TDEs, also called tidal disruption flares), occur when a star passes too close to a supermassive black hole and is torn apart by the extreme gravity gradient. Approximately half of the mass of the star remains bound to the BH, and powers the long-lived (few year) flare as it falls into the BH (Rees, 1988; van Velzen, Körding & Falcke, 2011). TDEs also display extremely long-lived (several months) gamma-ray counterparts (Burrows et al., 2011), and this is often the source of their original detection. TDEs provide a vital tool for detecting and probing quiescent black holes in other galaxies, especially in dwarf galaxies where black holes might be expected to be less massive (Levan et al., 2016). However, it is very difficult to correctly detect, classify and study TDEs in detail, because they are difficult to distinguish from nuclear supernovae and variable AGN. They are also around 100 times rarer than SNe (Rau et al., 2009; Levan et al., 2016).

Gravitational Microlensing

When an object of substantial mass passes in front of a background star, a magnification of the light source occurs due to the gravitational lensing of the light beam. This can be observed as a brightening of the background star, lasting several weeks, with a light curve shape that can be used to infer information about the lensing body. This technique can be used to discover lone brown dwarfs, as well as binary systems and extra-solar planets (Mao & Paczynski, 1991). The ideal target fields for such surveys are in the Galactic plane and bulge, since there are lots of potential background sources which a foreground body could gravitationally lens. This technique has been used in several surveys, such as the Optical Gravitational Lensing Experiment (OGLE, see e.g. Udalski, Kubiak & Szymanski, 1997; Udalski et al., 2002).

Stellar Flares

Flaring activity on cool stars has been studied for many decades (e.g. Moffett, 1974; Gershberg & Shakhovskaia, 1983), but recent transient surveys have provided an enormous increase in the number of known flare stars (Davenport et al., 2012). Stellar flares are believed to be due to magnetic reconnection events like those exhibited by the Sun (Haisch, Strong & Rodono, 1991). Flaring events are likely to have profound effects on any forms of life on exoplanets around cool stars (Ramsay et al., 2013), and flares are also known to produce signals at other frequencies, such as in the radio, which can impart information about particle and plasma physics (Bastian, Benz & Gary, 1998). Since flaring activity is so common (3 events per second per sky, Kulkarni & Rau, 2006), and high amplitude flares can occur on stars with a range of spectral types (Schaefer, King & Deliyannis, 2000), it is vital to understand stellar flares in order to separate them from other, perhaps more interesting transients. Most flares are very blue in colour, sometimes even peaking in the UV (Kulkarni & Rau, 2006), and thus multi-colour follow-up of transient candidates is a useful way to identify flares.

Near-Earth Objects and other Solar System objects

Several transient surveys, such as the Catalina Sky Survey (CSS) and Pan-STARSS, included the detection and characterisation of comets, asteroids and other Solar System objects as part of their primary objectives (Larson et al., 1998, 2003; Kaiser et al., 2002). Near-Earth objects (NEOs) are of particular interest to global defence considerations and space exploration safety, due to the risk of collisions. Aside from the safety aspect of studying NEOs, they can also help us better understand the formation of the Solar System (Scott, 2006) and the behaviour of the Solar wind (Ramanjooloo et al., 2015). Solar System objects also contribute to a substantial transient detection foreground, and automatic transient-detection pipelines must be able to distinguish between NEOs and genuine astrophysical transients. One way to do this, when the image resolution is sufficient, is to check for source elongation/ellipticity, as NEOs usually have some transverse motion relative to the stellar background (Bloom et al., 2012). Even better, surveys can take several images of each field, a few minutes apart, rather than just a single image. This allows NEOs to be detected and removed (Mahabal et al., 2011), since any moving objects found in these succesive images are likely to be NEOs.

Gamma-Ray Bursts and afterglows

Gamma-Ray Bursts (GRBs) are some of the brightest electromagnetic events in the known Universe, releasing up to 10^{50} ergs (Frail et al., 2001). They were first discovered by primitive satellites designed to discover the gamma-ray radiation that might be associated with the testing of nuclear weaponry in space during the Cold War (Klebesadel, Strong & Olson, 1973). Their origins were unknown for a long time, but Galactic sources began to be discounted when the BATSE instrument on board the Compton Gamma Ray Observatory collected a large sample of GRBs with an isotropic distribution, likely associated with cosmological distances (Meegan et al., 1992).

Later, the BeppoSAX satellite observed the first fading X-ray counterpart to a GRB, and optical follow-up enabled a precise location to be pinpointed, which associated the GRB with a faint galaxy (van Paradijs et al., 1997). Soon after this, the first redshift of a GRB afterglow was measured, conclusively placing it at a cosmological distance (Metzger et al., 1997). In 1998 an unusual GRB was directly associated with an unusual supernova explosion just days after the gamma-ray detection (Kulkarni et al., 1998). These observations linked GRBs with known optical transients for the first time. Now the X-ray/γ-ray satellite *Swift* is the major source of new GRBs, detecting several bursts per week and often pinpointing their locations using rapid X-ray and UV follow-up (Gehrels et al., 2004).

In general there are two main classes of GRB - the short, hard GRBs lasting less than two seconds and making up around 30% of the total population, and the long GRBs with durations longer than two seconds, being the other 70% of the population. Short GRBs, often found in older elliptical galaxies, are suspected to be produced by the merging of binary neutron stars, or a neutron star and black hole (Nakar, 2007), with the possible production of 'kilonovae' as a result (kilonovae are expected to be ~ 1000 times brighter than novae, Metzger et al., 2010). Long GRBs on the other hand tend to be found in star-forming regions of galaxies, and are thought to be produced by energetic jets escaping along the rotation axes of core-collapse supernovae (Woosley & Bloom, 2006).

In order to probe the emission mechanisms and energetics of these events, we need comprehensive broadband follow-up observations at all wavelengths. Modelling of the light curves across the electromagnetic spectrum allows some contraints to be made on the properties of the event, including the host galaxy extinction, extent of dust scattering, presence of reverse shocks, jet opening angle, jet kinetic energy, ambient medium density, and magnetic field energy density (van der Horst et al., 2015).

Fast Radio Bursts

Fast Radio Bursts (FRBs) are a recently discovered class of transient with only a handful of events detected, and so far seen only in the radio (Lorimer et al., 2007; Petroff et al., 2016). Lasting only a few milliseconds, but at suspected extragalactic distances, they have impressively large fluxes. The progenitors of FRBs are still unknown, but if localised and a redshift measured, they could become powerful tools for probing the inter-galactic medium, and perhaps even the physics of dark energy (Keane et al., 2016). In Chapter 5 I detail our attempts to observe optical afterglows of FRBs, and present a full discussion of their history, properties, possible progenitors, and recent developments in the field.

Gravitational Wave Transients

The exciting recent discovery of transient sources of gravitational waves by Abbott et al. (2016b) opens up a new window of astronomy in a totally different regime. This is equally exciting for electromagnetic transient astronomy, since many of the sources which are predicted to produce gravitational waves, may also produce optical, gamma-ray, or radio counterparts. This field is particularly young, but much effort is already going in to preparing wide-field optical follow-up of gravitational wave transients (see GOTO for example[1]).

Unknown and Unclassified Transients

As alluded to by Figure 1.1, some transient classes have been theoretically predicted but not yet observed. Of course, there could be new classes of transients that we may one day observe which have no theoretical predictions. Many historical transient detections remain unclassified, and although most of these will belong to known classes, some may be related to unknown phenomena. FRBs were only discovered within the last decade, and LRNe are also a fairly recent discovery. As we probe the transient universe to deeper limits in terms of brightness and cadence, we will surely find new and exciting types of transients.

[1] `http://www.goto-observatory.org/`

Appendix B

Extended Tables and Figures of Eclipsing CVs

Table B.1: Journal of Observations. The systems are ordered alphabetically, and some system names have been shortened. Start(MJD) is the start time of each observing run, given in MJD(UTC). Mid-eclipse times are given in BMJD(TDB) and the number in parentheses is the uncertainty in the last digit. T_{exp} is the exposure time in seconds and is supplied as $T_{blue}/T_{green}/T_{red}$ for the three beams of ULTRACAM. ΔT is the duration of the observing run in minutes. Mag. is the estimated out-of-eclipse magnitude.

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T$_{exp}$ (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
1RXS J180834.7+101041	55316.33464	55316.3613(3)	6/2/2	141	NTT/UCAM	$u'/g'/r'$	16.9/16.9/16.6
		55316.4310(3)					
	55334.37156	55334.4299(3)	10/2/2	101	NTT/UCAM	$u'/g'/r'$	17.0/17.0/16.8
ASASSN-13cx	56943.91536	56943.9710(3)	150	137	*pt5m*	V	16.1
	56968.81315	56968.8218(3)	150	206	*pt5m*	V	18.0
		56968.9015(3)					
	56973.81141	56973.8401(7)	150	406	*pt5m*	V	18.3
		56973.9186(7)					
	56989.65725	56989.68988(3)	3.9	60	TNT/USPEC	g'	18.9
	56993.59267	-	3.9	53	TNT/USPEC	*KG5*	17.9
	57014.83182	57014.8596(2)	120	69	*pt5m*	V	16.2
	57015.86791	57015.89471(7)	120	69	*pt5m*	V	16.1
	57023.51317	57023.54121(3)	3.9	57	TNT/USPEC	*KG5*	18.5
	57024.56187	57024.57665(3)	4.9	46	TNT/USPEC	g'	18.3
	57027.56687	57027.60339(3)	5.9	69	TNT/USPEC	*KG5*	18.9
	57198.12288	57198.13413(15)	12/4/4	31	WHT/UCAM	$u'/g'/r'$	18.4/18.4/18.3

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
ASASSN-14cl	56827.07177	-	10	188	*pt5m*	*V*	11.3
	56829.08098	-	10	180	*pt5m*	*V*	11.6
	56852.08564	-	10	184	*pt5m*	*V*	16.7
ASASSN-14ds	56853.00227	-	60	304	*pt5m*	*V*	16.7
	56854.99663	-	60	314	*pt5m*	*V*	16.7
	56855.99399	-	60	318	*pt5m*	*V*	16.9
	56856.99123	-	60	183	*pt5m*	*V*	16.8
	56858.98604	-	60	310	*pt5m*	*V*	17.1
	56859.98344	-	60	170	*pt5m*	*V*	17.0
ASASSN-14gl	56995.80900	-	60	124	*pt5m*	*V*	18.2
	56998.80903	-	60	147	*pt5m*	*V*	18.6
	57016.88881	-	120	129	*pt5m*	*V*	19.1
	57018.81524	-	120	348	*pt5m*	*V*	19.3
	57016.88881	-	120	129	*pt5m*	*V*	18.2
ASASSN-14gu	57057.83442	-	120	141	*pt5m*	*V*	18.0
	57058.83810	-	120	105	*pt5m*	*V*	17.3
	57084.89539	-	120	253	*pt5m*	*V*	16.4
ASASSN-14hk	56922.12637	-	60	172	*pt5m*	*V*	15.8

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	56932.99821	-	60	128	*pt5m*	V	17.1
ASASSN-14mv	57026.59217	-	1.0	122	TNT/USPEC	g'	11.8
ASASSN-15au	57059.83530	57059.8614(1)	120	226	*pt5m*	V	18.1
		57059.9302(3)					
	57067.00211	57067.0322(2)	120	148	*pt5m*	V	18.0
		57067.1011(2)					
	57067.89695	57067.9283(2)	120	69	*pt5m*	V	18.4
	57067.97227	57067.9974(3)	120	69	*pt5m*	V	18.1
	57068.03489	57068.0665(2)	120	69	*pt5m*	V	18.1
	57076.52507	57076.54688(5)	7.0	40	TNT/USPEC	g'	18.3
	57077.62875	57077.65003(5)	7.0	39	TNT/USPEC	r'	18.5
	57106.52037	57106.54011(5)	4.0	37	TNT/USPEC	g'	18.2
	57367.70969	57367.72165(5)	10	36	TNT/USPEC	g'	17.8
ASASSN-15bu	57079.51212	57079.52541(5)	3.9	38	TNT/USPEC	g'	17.6
	57080.51411	57080.52418(5)	3.8	51	TNT/USPEC	g'	17.8
	57081.51652	57081.52293(5)	3.9	48	TNT/USPEC	g'	17.9
	57082.51514	57082.52163(5)	4.9	35	TNT/USPEC	g'	18.0
ASASSN-15ni	57231.92628	-	10	114	*pt5m*	V	13.2
	57232.90219	-	20	311	*pt5m*	V	13.6

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	57235.90511	-	30	287	*pt5m*	V	14.1
	57236.89403	-	30	138	*pt5m*	V	14.2
CSS080227:112634-100210	55333.97234	55333.98693(3)	9/3/3	154	NTT/UCAM	$u'/g'/r'$	18.6/18.7/18.4
		55334.06433(3)					
	55711.01379	55711.03000(3)	12/4/4	41	NTT/UCAM	$u'/g'/r'$	18.3/18.5/18.1
	56046.90666	-	8.5/4.2/4.2	25	WHT/UCAM	$u'/g'/r'$	18.4/18.5/18.3
	56685.91260	56685.92246(3)	4.9	31	TNT/USPEC	*KG5*	18.3
	57076.88173	57076.90134(3)	9.9	41	TNT/USPEC	*KG5*	19.1
	57077.81348	57077.83043(3)	9.9	39	TNT/USPEC	*KG5*	19.0
	57078.80800	57078.83694(3)	9.9	56	TNT/USPEC	*KG5*	19.0
CSS080306:082655-000733	56045.86872	56045.87387(3)	6/2/2	14	WHT/UCAM	$u'/g'/r'$	20.3/19.9/20.0
	56045.90974	56045.93365(3)	12/4/4	42	WHT/UCAM	$u'/g'/r'$	20.2/20.0/19.9
	56046.86831	56046.88989(3)	12/4/4	36	WHT/UCAM	$u'/g'/r'$	19.8/19.8/19.6
	56689.94364	56689.95507(3)	3.0	26	SALT/SALTICAM	r'	19.7
	56693.87936	56693.89955(3)	3.0	39	SALT/SALTICAM	r'	20.1
	56694.95652	56694.97525(3)	3.0	33	SALT/SALTICAM	g'	20.1
	56695.96725	56695.99128(3)	3.0	44	SALT/SALTICAM	g'	20.3
	56988.76551	56988.77724(5)	5.9	31	TNT/USPEC	*KG5*	19.7
CSS080623:140454-102702	55329.21114	55329.23537(3)	13/3.3/3.3	41	NTT/UCAM	$u'/g'/r'$	19.8/19.8/19.6
	55334.09572	55334.12090(3)	15/5/5	59	NTT/UCAM	$u'/g'/r'$	20.1/19.8/19.8

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	55334.15978	55334.18049(3)	16/4/4	60	NTT/UCAM	$u'/g'/r'$	20.1/19.9/19.8
	55334.20345	55334.24003(3)	12/4/4	67	NTT/UCAM	$u'/g'/r'$	20.1/19.9/19.9
	55355.01296	55355.03309(3)	7.7/3.8/3.8	41	NTT/UCAM	$u'/g'/r'$	19.7/19.7/19.5
	55355.12826	55355.15230(3)	7.7/3.8/3.8	49	NTT/UCAM	$u'/g'/r'$	20.0/19.8/19.7
	55709.03970	55709.05135(3)	6/3/3	21	NTT/UCAM	$u'/g'/r'$	19.9/19.8/19.9
	55711.95380	55711.97074(3)	12/4/4	29	NTT/UCAM	$u'/g'/r'$	19.8/19.7/19.9
	55711.99633	55712.03031(3)	12/4/4	53	NTT/UCAM	$u'/g'/r'$	20.1/19.8/19.9
	55712.11378	55712.14946(3)	12/4/4	67	NTT/UCAM	$u'/g'/r'$	20.0/19.7/19.9
CSS081220:011614+092216	56489.16741	56489.17702(3)	9.8/3.2/3.2	28	WHT/UCAM	$u'/g'/i'$	19.5/19.6/19.3
	56500.15832	56500.17282(3)	9.8/3.2/3.2	35	WHT/UCAM	$u'/g'/r'$	19.2/19.4/19.3
	56504.17318	56504.18927(3)	9.8/3.2/3.2	33	WHT/UCAM	$u'/g'/i'$	19.1/19.2/18.9
	56510.10544	56510.11505(3)	6.5/3.2/3.2	30	WHT/UCAM	$u'/g'/r'$	17.7/17.8/17.8
	56657.84860	56657.86688(3)	9.8/3.2/3.2	48	WHT/UCAM	$u'/g'/r'$	18.8/18.9/18.7
	56657.92579	56657.93272(3)	9.8/3.2/3.2	24	WHT/UCAM	$u'/g'/r'$	19.1/19.3/19.1
	56658.89473	56658.92033(3)	12/4/4	50	WHT/UCAM	$u'/g'/r'$	19.2/19.4/19.2
	56683.54083	56683.54564(3)	5.9	21	TNT/USPEC	*KG5*	19.5
	56685.57737	56685.58670(8)	9.4	33	TNT/USPEC	*KG5*	19.6
	56689.53020	56689.53730(4)	9.4	22	TNT/USPEC	*KG5*	19.2
	56873.08525	56873.10772(3)	9/3/3	57	WHT/UCAM	$u'/g'/r'$	19.0/19.0/18.8
	56874.07296	56874.09532(3)	9/3/3	41	WHT/UCAM	$u'/g'/r'$	18.8/18.9/18.8
	56988.71680	56988.72799(3)	4.1	24	TNT/USPEC	*KG5*	19.0
	56989.50680	56989.51817(3)	4.1	20	TNT/USPEC	*KG5*	19.2

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	57366.64759	57366.66722(3)	7.9	48	TNT/USPEC	*KG5*	16.7
CSS090102:132536+210037	55711.97933	55711.98534(2)	15/5/5	19	NTT/UCAM	$u'/g'/r'$	19.9/19.9/19.7
	55712.03626	55712.04774(2)	15/5/5	26	NTT/UCAM	$u'/g'/r'$	19.8/19.9/19.6
	55714.01658	55714.04404(2)	15/5/5	100	NTT/UCAM	$u'/g'/r'$	19.8/19.9/19.6
		55714.10641(2)					
	55943.10286	55943.12141(1)	18/6/6	50	WHT/UCAM	$u'/g'/r'$	20.4/20.2/20.1
	55943.21976	55943.24618(1)	18/6/6	55	WHT/UCAM	$u'/g'/r'$	19.9/19.8/19.8
	56873.87607	56873.90434(1)	9/3/3	57	WHT/UCAM	$u'/g'/r'$	20.2/20.1/20.2
	56874.89075	56874.90250(1)	9/3/3	31	WHT/UCAM	$u'/g'/i'$	20.1/20.1/20.2
CSS090219:044027+023301	56960.08801	-	150	50	*pt5m*	V	18.1
	56961.02405	-	150	340	*pt5m*	V	18.5
	57003.94554	-	150	54	*pt5m*	V	18.6
	57015.91831	-	150	295	*pt5m*	V	18.6
CSS090419:162620-125557	56486.92629	56486.93325(5)	12/4/4	18	WHT/UCAM	$u'/g'/i'$	21.3/21.1/20.6
	56498.90349	56498.92860(5)	12/4/4	58	WHT/UCAM	$u'/g'/i'$	20.7/20.6/20.0
	56499.87896	56499.90938(5)	12/4/4	55	WHT/UCAM	$u'/g'/r'$	20.7/20.8/20.4
	56501.92501	56501.94635(5)	15/3.9/3.9	50	WHT/UCAM	$u'/g'/i'$	20.9/20.7/20.1
	56508.87512	56508.88707(5)	12/4/4	30	WHT/UCAM	$u'/g'/r'$	20.5/20.4/20.2
	56508.95486	56508.96249(5)	12/4/4	26	WHT/UCAM	$u'/g'/r'$	21.1/21.0/20.5
	56872.87452	56872.89818(5)	4.2/4.2	51	WHT/UCAM	g'/r'	21.8/20.4

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
CSS090622:215636+193242	56873.91771	56873.9890(5)	120	243	*pt5m*	V	19.2
		56874.0621(5)					
	56874.10512	56874.13184(5)	9/3/3	49	WHT/UCAM	$u'/g'/r'$	19.6/19.5/19.1
	56874.19057	56874.20268(5)	9/3/3	24	WHT/UCAM	$u'/g'/r'$	19.5/19.6/19.2
	56874.96361	56874.98303(5)	9/3/3	36	WHT/UCAM	$u'/g'/i'$	19.2/19.2/18.0
	56877.93877	56877.96200(5)	9/3/3	39	WHT/UCAM	$u'/g'/i'$	20.0/19.9/18.7
	56878.15537	56878.17474(5)	9/3/3	39	WHT/UCAM	$u'/g'/i'$	19.7/19.6/18.3
	56880.20188	56880.23174(5)	9/3/3	47	WHT/UCAM	$u'/g'/i'$	19.7/19.8/18.3
	56992.49766	56992.5128(2)	5.9	37	TNT/USPEC	*KG5*	19.5
CSS091116:232551-014024	56924.94657	-	100	289	*pt5m*	V	18.9
	56939.84366	-	150	369	*pt5m*	V	18.4
CSS100218:043829+004016	56657.88694	56657.89355(2)	12/4/4	11	WHT/UCAM	$u'/g'/r'$	19.3/19.5/19.3
	56657.94634	56657.95904(1)	12/4/4	26	WHT/UCAM	$u'/g'/r'$	19.3/19.5/19.3
	56658.99048	56659.00696(1)	12/4/4	20	WHT/UCAM	$u'/g'/r'$	20.2/19.9/19.7
	56991.77750	-	6.0	53	TNT/USPEC	*KG5*	19.9
	57081.57619	-	30	50	TNT/USPEC	*KG5*	20.2
CSS100508:085604+322109	57091.85566	-	150	141	*pt5m*	V	19.5
	57097.91622	-	150	288	*pt5m*	V	19.8

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
CSS100520:214426+222024	56938.83073	-	120	353	*pt5m*	V	18.0
CSS110113:043112-031452	55580.09688	55580.12266(5)	7.2/2.4/2.4	47	NTT/UCAM	$u'/g'/i'$	16.1/15.8/15.9
	55940.94016	55940.95858(3)	12/4/4	37	WHT/UCAM	$u'/g'/r'$	19.9/19.9/19.5
	55942.00021	55942.01540(3)	20/5/5	32	WHT/UCAM	$u'/g'/r'$	20.1/20.0/19.7
	55942.85431	55942.87405(3)	12/4/4	34	WHT/UCAM	$u'/g'/r'$	20.0/20.0/19.8
	56180.11393	56180.12877(3)	12/4/4	30	WHT/UCAM	$u'/g'/r'$	19.7/19.8/19.6
	56214.05905	56214.07895(3)	9/3/3	35	WHT/UCAM	$u'/g'/r'$	19.8/19.8/19.6
	56659.90460	56659.92234(3)	12/4/	27	WHT/UCAM	$u'/g'/r'$	19.1/19.0/18.7
	56685.54492	56685.55006(3)	9.4	19	TNT/USPEC	*KG5*	20.0
	56685.60074	56685.61612(3)	9.4	35	TNT/USPEC	*KG5*	20.1
	56686.50253	56686.54082(3)	9.4	81	TNT/USPEC	g'	19.8
	56689.56937	56689.57917(3)	9.0	44	TNT/USPEC	g'	19.8
	56730.84840	56730.86095(3)	8.0/2.6/2.6	42	WHT/UCAM	$u'/g'/r'$	19.7/19.7/19.4
CSS110114:091246-034916	57046.98180	-	120	105	*pt5m*	V	18.0
	57051.94906	-	120	105	*pt5m*	V	16.4
	57053.99243	-	120	295	*pt5m*	V	17.8
	57054.99178	-	120	291	*pt5m*	V	18.4
CSS110226:112510+231036	57049.06558	-	150	318	*pt5m*	V	19.5
	57068.08501	-	120	121	*pt5m*	V	19.5
	57070.05512	-	120	304	*pt5m*	V	19.7

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
CSS110513:210846-035031	56176.01207	56176.0486(4)	120	126	*pt5m*	V	17.7
	56176.91862	56176.9894(3)	120	239	*pt5m*	V	17.9
	56177.85543	56177.9306(3)	120	157	*pt5m*	V	17.9
	56486.97477	56487.0769(3)	60	209	*pt5m*	V	18.3
	56489.07325	56489.11534(8)	11/3.8/3.8	80	WHT/UCAM	$u'/g'/i'$	18.6/18.3/17.4
	56508.97803	56509.04487(8)	11/3.8/3.8	115	WHT/UCAM	$u'/g'/r'$	18.5/18.3/17.7
	56871.98925	56872.01611(5)	8/2/2	64	WHT/UCAM	$u'/g'/r'$	18.0/17.9/17.5
	56872.94317	56872.95764(5)	9.1/2.2/2.2	40	WHT/UCAM	$u'/g'/r'$	18.0/17.9/17.5
CSS111003:054558+022106	56976.09353	56976.1086(6)	150	61	*pt5m*	V	18.6
	56987.07470	-	150	135	*pt5m*	V	18.2
	57017.93566	57017.9642(5)	150	76	*pt5m*	V	18.8
	57020.91573	57020.9908(8)	150	342	*pt5m*	V	19.1
		57021.1107(7)					
	57024.71894	57024.7369(2)	3.9	78	TNT/USPEC	*KG5*	18.7
	57027.61787	57027.6402(2)	9.7	53	TNT/USPEC	*KG5*	18.7
	57028.93920	57028.9727(5)	120	69	*pt5m*	V	19.2
	57077.58668	57077.6014(2)	8.9	57	TNT/USPEC	g'	19.1
	57078.53462	57078.5690(2)	8.9	87	TNT/USPEC	r'	18.5
	57365.74728	-	15	38	TNT/USPEC	g'	19.1
	57367.68063	57367.6917(2)	16	36	TNT/USPEC	g'	19.0

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
CSS111019:233313-155744	55867.80455	55867.8116(6)	20/10/10	53	WHT/UCAM	$u'/g'/r'$	20.1/19.7/19.4
	56874.14164	56874.1486(1)	9/3/3	130	WHT/UCAM	$u'/g'/r'$	18.4/18.2/18.2
		56874.1914(1)					
		56874.2342(1)					
	56878.18496	56878.2193(2)	7/2.3/2.3	54	WHT/UCAM	$u'/g'/i'$	20.0/19.5/19.3
	56879.18534	56879.2049(5)	9/3/3	80	WHT/UCAM	$u'/g'/i'$	20.1/19.7/19.4
CSS111101:233003+303301	56874.09009	56874.155(1)	120	191	*pt5m*	V	18.1
	56874.95930	56875.092(1)	120	256	*pt5m*	V	18.2
	56876.95942	56877.118(1)	120	259	*pt5m*	V	18.1
	57016.81470	57016.8744(1)	150	104	*pt5m*	V	17.6
	57017.83588	-	120	140	*pt5m*	V	18.4
	57020.81609	-	120	140	*pt5m*	V	18.4
	57026.53984	57026.5459(3)	7.9	71	TNT/USPEC	*KG5*	18.5
	57026.83120	57026.8580(5)	120	69	*pt5m*	V	18.0
	57028.53937	57028.5731(3)	7.9	94	TNT/USPEC	*KG5*	18.4
	57028.85964	57028.8857(6)	120	69	*pt5m*	V	18.5
CSS130906:064726+491542	56542.20638	-	2	49	*pt5m*	V	13.8
	57043.96248	-	120	321	*pt5m*	V	17.5
CSS131106:052412+004148	56604.20805	-	35	99	*pt5m*	V	15.6
	56940.10939	56940.181(9)	120	204	*pt5m*	V	17.8

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	56943.14634	56943.150(9)	150	153	*pt5m*	V	18.2
	56946.10810	56946.117(1)	150	210	*pt5m*	V	15.8
	56956.18535	56956.248(9)	150	105	*pt5m*	V	18.1
	56970.03219	56970.0459(2)	150	142	*pt5m*	V	16.2
	56976.15126	56976.160(3)	150	152	*pt5m*	V	18.1
	56987.00005	-	150	104	*pt5m*	V	18.2
	56990.78805	56990.83142(5)	4.9	96	TNT/USPEC	g'	18.6
	56992.71556	56992.75265(5)	4.9	86	TNT/USPEC	*KG5*	17.9
	57023.63566	57023.66863(5)	7.9	68	TNT/USPEC	*KG5*	18.1
	57027.65700	57027.68606(5)	8.0	73	TNT/USPEC	*KG5*	18.3
	57080.59617	57080.60994(5)	15	65	TNT/USPEC	g'	18.0
	57361.78732	57361.82302(5)	5.0	95	TNT/USPEC	g'	18.0
	57365.64517	57365.666(2)	15	55	TNT/USPEC	u'	18.3
	57417.49956	57417.54151(5)	8.0	100	TNT/USPEC	g'	18.2
CSS140402:173048+554518	56754.06043	-	15	78	*pt5m*	V	14.6
CSS140901:013309+133234	56904.17086	-	170	97	*pt5m*	V	16.7
	56908.07756	-	150	233	*pt5m*	V	17.1
	56913.97688	-	150	177	*pt5m*	V	17.8
CSS141005:023428-045431	56969.93998	-	60	120	*pt5m*	V	14.6
	56995.95685	-	60	174	*pt5m*	V	14.7

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	56997.88016	-	60	276	*pt5m*	*V*	14.9
	56998.91431	-	60	107	*pt5m*	*V*	14.7
	57027.88719	-	60	148	*pt5m*	*V*	14.8
CSS141117:030930+263804	57038.82444	-	120	51	*pt5m*	*V*	18.2
	57041.82943	-	120	158	*pt5m*	*V*	18.5
	57052.85159	-	120	48	*pt5m*	*V*	18.3
	57053.83222	-	120	160	*pt5m*	*V*	18.4
	57054.83256	-	120	225	*pt5m*	*V*	18.3
CzeV404 Her	56740.83024	56740.86289(8)	4.4	77	TNT/USPEC	g'	17.1
	56742.88593	56742.9211(2)	3.4	82	TNT/USPEC	g'	15.1
	56854.02607	56854.0773(2)	60	98	*pt5m*	*V*	14.8
	56856.89675	56856.9201(2)	60	69	*pt5m*	*V*	14.1
	56871.87078	56871.91729(5)	3.6/1.8/1.8	80	WHT/UCAM	$u'/g'/r'$	17.1/17.0/16.6
	57079.90232	57079.91840(5)	3.4	58	TNT/USPEC	*KG5*	16.1
	57083.92911	57083.9373(2)	3.4	33	TNT/USPEC	*KG5*	14.9
Gaia15aan	57054.25900	-	30	40	*pt5m*	*V*	18.3
	57055.25668	-	30	42	*pt5m*	*V*	18.6
	57057.21084	-	60	107	*pt5m*	*V*	18.6
	57059.17340	-	60	40	*pt5m*	*V*	18.8

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
GALEX J003535.7+462353	55940.81525	55940.8250(1)	5.6/1.8/1.8	50	WHT/UCAM	$u'/g'/r'$	17.1/17.0/16.6
	55941.83001	55941.8592(5)	5.6/1.8/1.8	65	WHT/UCAM	$u'/g'/r'$	16.5/15.9/15.7
	56179.91441	56179.9424(1)	4.6/1.5/1.5	65	WHT/UCAM	$u'/g'/r'$	16.7/16.4/16.1
	56209.90088	56209.9185(1)	6/2/2	55	WHT/UCAM	$u'/g'/r'$	16.4/16.3/15.9
	56210.93470	56210.9523(1)	5/1/1	39	WHT/UCAM	$u'/g'/r'$	16.6/16.3/16.1
	56500.18602	56500.2015(1)	6/2/2	78	WHT/UCAM	$u'/g'/r'$	16.4/16.3/16.0
	56505.17727	56505.1977(1)	6/2/2	53	WHT/UCAM	$u'/g'/i'$	17.0/16.9/16.3
	56509.99699	56510.0215(1)	6/2/2	75	WHT/UCAM	$u'/g'/r'$	16.9/16.8/16.4
	56510.16914	56510.1936(1)	6/2/2	65	WHT/UCAM	$u'/g'/r'$	16.7/16.5/16.2
	56657.79859	56657.8333(1)	6/2/2	68	WHT/UCAM	$u'/g'/r'$	16.6/16.3/16.0
	56658.83889	56658.8669(1)	6/2/2	75	WHT/UCAM	$u'/g'/r'$	16.4/16.4/15.9
	56872.09341	56872.1433(1)	6/2/2	100	WHT/UCAM	$u'/g'/r'$	16.6/16.4/16.1
	56873.12930	56873.1770(1)	6/2/2	117	WHT/UCAM	$u'/g'/r'$	16.6/16.4/16.0
	56988.57231	56988.6012(1)	1.9	87	TNT/USPEC	*KG5*	16.4
	57023.55486	57023.5727(1)	1.9	45	TNT/USPEC	*KG5*	16.4
	57365.69822	57365.7111(1)	10	56	TNT/USPEC	g'	17.2
GY Cnc	55938.24332	55938.26366(4)	12/4/4	56	WHT/UCAM	$u'/g'/r'$	16.7/16.6/16.2
	55941.22656	55941.24626(4)	6.2/3.1/3.1	52	WHT/UCAM	$u'/g'/r'$	16.7/16.8/16.2
	55942.98285	55943.00068(4)	7.5/2.5/2.5	46	WHT/UCAM	$u'/g'/r'$	17.0/16.9/16.3
	55943.14044	55943.17605(4)	7.5/2.5/2.5	71	WHT/UCAM	$u'/g'/r'$	16.7/16.5/16.1
	55947.18590	55947.21130(4)	12/4/4	51	WHT/UCAM	$u'/g'/r'$	16.6/16.6/16.1
	56657.18967	56657.22682(5)	12/4/4	74	WHT/UCAM	$u'/g'/r'$	16.3/16.4/16.0

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	56683.69249	56683.71865(5)	1.3	53	TNT/USPEC	*KG5*	16.5
	57080.55294	57080.56902(5)	2.5	50	TNT/USPEC	*KG5*	16.7
	57367.73864	57367.76812(5)	4.0	91	TNT/USPEC	*KG5*	16.6
HS 2325+8205	55941.87952	55941.89303(5)	5.6/1.8/1.8	48	WHT/UCAM	$u'/g'/r'$	14.7/14.4/14.4
	56487.16448	56487.19622(5)	5.6/1.8/1.8	94	WHT/UCAM	$u'/g'/i'$	17.1/16.7/16.1
	56509.13073	56509.15629(5)	5.6/1.8/1.8	76	WHT/UCAM	$u'/g'/r'$	15.2/14.9/14.8
HT Cas	54400.23091	54400.24274(5)	9.4/2.3/2.3	23	WHT/UCAM	$u'/g'/r'$	16.4/16.6/16.5
	56872.19923	56872.21011(5)	6.6/2.2/2.2	46	WHT/UCAM	$u'/g'/r'$	16.5/16.6/16.4
	56874.03047	56874.05133(5)	6.3/2.1/2.1	56	WHT/UCAM	$u'/g'/r'$	16.2/16.5/16.3
	56988.48215	56988.49911(5)	1.0	36	TNT/USPEC	g'	16.8
	56993.48629	56993.50711(5)	2.0	32	TNT/USPEC	g'	16.5
	57023.61271	57023.62882(5)	2.0	29	TNT/USPEC	*KG5*	16.8
	57024.65162	57024.65988(5)	2.0	22	TNT/USPEC	g'	16.7
	57285.06842	57285.07646(5)	6.6/2.2/2.2	15	WHT/UCAM	$u'/g'/r'$	16.0/16.3/16.1
IY UMa	56746.62944	56746.64028(2)	2.2	46	TNT/USPEC	*KG5*	17.6
	56746.70652	56746.71419(2)	2.2	28	TNT/USPEC	*KG5*	17.7
	56746.78005	56746.78810(2)	2.2	33	TNT/USPEC	*KG5*	17.6
	56991.93503	56991.94411(5)	3.4	28	TNT/USPEC	*KG5*	14.5
	57025.92282	57025.94209(3)	3.4	40	TNT/USPEC	g'	17.4
	57028.88243	57028.89840(5)	4.0	30	TNT/USPEC	r'	17.3

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	57076.85098	57076.86533(3)	3.9	41	TNT/USPEC	r'	17.2
MAST003059.39+301634.3	56904.95602	56904.9976(3)	90	407	*pt5m*	V	19.4
		56905.0685(3)					
		56905.1396(5)					
		56905.2084(5)					
	56907.94797	-	120	117	*pt5m*	V	19.7
	56921.95125	56922.002(1)	120	148	*pt5m*	V	19.7
	56922.95876	56922.984(1)	120	343	*pt5m*	V	19.8
		56923.0559(9)					
		56923.1263(3)					
		56923.1965(8)					
	56992.64865	56992.6851(1)	5.9	61	TNT/USPEC	*KG5*	19.9
	57023.58867	57023.6008(1)	6.9	33	TNT/USPEC	*KG5*	19.2
	57199.17958	57199.18681(4)	10.5/3.5/3.5	29	WHT/UCAM	$u'/g'/r'$	19.6/19.7/19.5
	57284.89172	57284.90695(8)	20/6.7/6.7	28	WHT/UCAM	$u'/g'/r'$	19.5/19.5/19.3
MAST034045.31+471632.2	56936.97202	-	120	402	*pt5m*	V	18.7
MAST041923.57+653004.3	56944.02661	-	120	326	*pt5m*	V	17.2
	56960.12942	-	120	132	*pt5m*	V	17.3
	56967.11017	-	120	129	*pt5m*	V	17.4

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
MAST171921.40+640309.8	56732.06257	-	30	100	*pt5m*	V	16.2
	56735.02752	-	40	155	*pt5m*	V	16.6
MAST192328.22+612413.5	56774.04124	56774.148(5)	120	263	*pt5m*	V	19.3
	56779.11794	56779.1753(2)	35/7/7	164	WHT/UCAM	$u'/g'/r'$	18.2/17.9/17.7
	56799.97068	-	180	176	*pt5m*	V	17.4
	56825.92616	56825.9488(3)	180	80	*pt5m*	V	18.8
	56826.92700	56826.9551(3)	180	119	*pt5m*	V	18.9
	56834.92790	56835.0026(3)	180	398	*pt5m*	V	17.2
		56835.1704(3)					
	56873.03328	56873.05815(5)	13/3.3/3.3	69	WHT/UCAM	$u'/g'/r'$	19.3/19.1/18.8
MAST194955.17+455349.6	56874.88619	-	90	106	*pt5m*	V	18.6
	56878.95944	-	90	292	*pt5m*	V	18.7
MAST201121.95+565531.1	56836.00727	-	60	203	*pt5m*	V	17.2
MAST202157.69+212919.4	56894.91143	-	120	304	*pt5m*	V	20.3
MAST203421.90+120656.9	56849.02716	-	100	264	*pt5m*	V	17.9
	56849.96328	-	100	262	*pt5m*	V	17.8
MAST210316.39+314913.6	56858.03866	-	80	255	*pt5m*	V	18.0

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	56879.88002	-	80	425	*pt5m*	*V*	18.3
MAST232100.42+494614.0	56799.13236	56799.152(9)	60	111	*pt5m*	*V*	18.9
	56854.16419	-	90	61	*pt5m*	*V*	18.7
	56878.02050	56878.1546(4)	100	295	*pt5m*	*V*	16.9
	56880.96494	56881.128(9)	100	377	*pt5m*	*V*	18.5
	56881.99945	56882.191(6)	100	329	*pt5m*	*V*	18.7
	56884.95665	56885.160(9)	150	392	*pt5m*	*V*	18.8
	56895.12463	56895.143(4)	150	157	*pt5m*	*V*	19.0
	56921.88054	56921.902(9)	150	99	*pt5m*	*V*	19.0
	56922.05619	56922.124(9)	150	98	*pt5m*	*V*	18.7
	56943.07065	56943.136(3)	150	105	*pt5m*	*V*	18.9
	56956.03000	56956.0949(3)	150	105	*pt5m*	*V*	16.8
	56988.63813	-	4.0	87	TNT/USPEC	*KG5*	17.9
	56992.57897	56992.6242(1)	5.0	97	TNT/USPEC	g'	19.3
	57025.50124	57025.5411(1)	6.0	96	TNT/USPEC	*KG5*	19.0
	57028.49920	57028.5144(1)	8.0	55	TNT/USPEC	*KG5*	19.1
	57366.60377	-	14	56	TNT/USPEC	g'	18.5
MLS101226:072033+172437	56770.86964	56770.9453(8)	120	132	*pt5m*	*V*	18.0
	56772.87068	56772.9016(8)	120	156	*pt5m*	*V*	18.0
	57045.01865	57045.1377(6)	120	232	*pt5m*	*V*	18.3
	57048.84401	57048.897(5)	120	69	*pt5m*	*V*	18.3

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	57049.01683	57049.0479(6)	120	67	*pt5m*	V	18.4
	57051.87424	57051.9049(4)	120	69	*pt5m*	V	18.6
	57056.83720	57056.8694(7)	120	69	*pt5m*	V	18.4
	57056.98798	57057.0204(4)	120	69	*pt5m*	V	18.4
	57079.71591	57079.7303(2)	6.9	52	TNT/USPEC	g'	18.7
MLS120517:152507-032655	56772.03312	-	180	92	*pt5m*	V	16.3
	56772.98238	56773.0340(5)	120	343	*pt5m*	V	16.5
		56773.0981(5)					
		56773.1630(5)					
	56775.95890	56775.9956(5)	120	308	*pt5m*	V	16.5
		56776.0600(5)					
		56776.1244(4)					
	56864.90455	56864.9069(5)	120	104	*pt5m*	V	17.0
		56864.9710(3)					
	56874.87061	56874.8857(3)	3.6/1.7/1.7	25	WHT/UCAM	$u'/g'/i'$	18.0/17.4/16.3
SDSS J040714.78-064425.1	54392.19660	-	10/2.5/2.5	109	WHT/UCAM	$u'/g'/r'$	17.8/17.7/17.3
	54396.19804	54396.23455(8)	20/5/5	98	WHT/UCAM	$u'/g'/r'$	17.3/17.0/16.7
	55939.82002	55939.9837(4)	10	281	*pt5m*	R	15.7
	56215.00769	56215.03344(8)	2.8/1.4/1.4	78	WHT/UCAM	$u'/g'/r'$	16.0/15.8/15.7
	57026.70536	57026.7366(2)	6.9	75	TNT/USPEC	*KG5*	15.6
	57082.54167	57082.5629(1)	6.9	74	TNT/USPEC	*KG5*	18.4

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	57083.56625	57083.5841(1)	6.9	64	TNT/USPEC	*KG5*	17.8
	57367.63138	57367.6551(2)	5.1	64	TNT/USPEC	g'	15.6
SDSS J075059.97+141150.1	55940.89611	55940.92964(4)	12/4/4	58	WHT/UCAM	$u'/g'/r'$	18.9/19.1/18.9
	55940.99589	55941.02271(4)	12/4/4	66	WHT/UCAM	$u'/g'/r'$	19.2/19.3/19.1
	55942.12618	55942.14077(4)	12/4/4	41	WHT/UCAM	$u'/g'/r'$	19.3/19.6/16.4
	55943.05921	55943.07241(4)	12/4/4	30	WHT/UCAM	$u'/g'/r'$	19.3/19.7/19.5
	56211.19435	56211.20263(5)	9/3/3	29	WHT/UCAM	$u'/g'/r'$	19.1/19.4/19.3
	56657.98984	56658.02444(4)	12/4/4	69	WHT/UCAM	$u'/g'/r'$	18.9/19.2/19.0
	56658.93310	56658.95609(4)	12/4/4	39	WHT/UCAM	$u'/g'/r'$	18.8/19.0/18.8
	56659.02043	56659.04924(4)	12/4/4	47	WHT/UCAM	$u'/g'/r'$	18.8/19.0/18.8
	56683.62169	56683.64491(4)	10	45	TNT/USPEC	*KG5*	18.7
	56684.64142	56684.66967(5)	7.0	53	TNT/USPEC	*KG5*	18.4
SDSS J090103.93+480911.1	53803.84353	53803.90711(2)	5/5/5	114	WHT/UCAM	$u'/g'/r'$	19.5/19.7/19.5
	53804.98276	53804.99742(2)	5/5/5	25	WHT/UCAM	$u'/g'/r'$	19.2/19.4/19.2
	53805.12890	53805.15321(2)	5/5/5	49	WHT/UCAM	$u'/g'/r'$	19.1/19.3/19.1
	55203.93743	55203.96531(5)	6.8/1.7/1.7	48	WHT/UCAM	$u'/g'/r'$	19.4/19.6/19.5
	55942.08419	55942.11712(2)	13/4.5/4.5	56	WHT/UCAM	$u'/g'/r'$	20.0/20.0/19.8
	55942.15686	55942.19500(2)	13/4.5/4.5	74	WHT/UCAM	$u'/g'/r'$	20.2/20.2/20.0
	55942.95787	55942.97383(2)	13/4.5/4.5	31	WHT/UCAM	$u'/g'/r'$	19.5/19.6/19.5
	55943.01684	55943.05169(2)	13/4.5/4.5	56	WHT/UCAM	$u'/g'/r'$	19.7/19.8/19.6
	55943.19203	55943.20746(2)	13/4.5/4.5	34	WHT/UCAM	$u'/g'/r'$	19.5/19.6/19.5

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	55943.26137	55943.28535(2)	13/4.5/4.5	51	WHT/UCAM	$u'/g'/r'$	19.7/19.7/19.5
SDSS J090403.49+035501.2	56045.88320	56045.89912(2)	12/12/12	26	WHT/UCAM	$u'/g'/r'$	19.6/19.3/19.3
SDSS J092009.54+004245.0	55700.02192	55700.10597(5)	8.7/2.9/2.9	154	NTT/UCAM	$u'/g'/r'$	17.7/18.0/17.7
	55708.95938	55708.97858(5)	5.6/2.8/2.8	57	NTT/UCAM	$u'/g'/r'$	18.2/18.0/17.7
	57026.77655	57026.84657(5)	5.9	118	TNT/USPEC	*KG5*	17.9
SDSS J092444.48+080150.9	55709.00465	55709.02440(7)	15/5/5	42	NTT/UCAM	$u'/g'/r'$	20.0/20.1/19.2
SDSS J093537.46+161950.8	55710.94728	55710.98340(6)	15/5/5	88	NTT/UCAM	$u'/g'/r'$	19.5/19.2/18.7
SDSS J100658.40+233724.4	55942.02575	55942.0530(1)	12/4/4	79	WHT/UCAM	$u'/g'/r'$	18.4/18.5/17.9
	56682.70289	56682.7302(2)	3.4	65	TNT/USPEC	*KG5*	18.2
	56682.87478	56682.9163(2)	3.4	85	TNT/USPEC	*KG5*	18.2
	56683.81609	56683.8457(2)	3.4	75	TNT/USPEC	*KG5*	18.2
	56685.85274	56685.8907(2)	3.4	81	TNT/USPEC	*KG5*	18.1
	56690.66858	56690.7244(2)	5.9	118	TNT/USPEC	g'	18.5
	56690.87758	56690.9103(2)	3.4	70	TNT/USPEC	r'	18.6
	57362.77084	57362.8006(2)	4.9	75	TNT/USPEC	g'	18.6
SDSS J115207.00+404947.8	56730.88610	56730.90933(1)	12/4/4	45	WHT/UCAM	$u'/g'/r'$	19.6/19.5/19.5
	56746.66410	56746.69503(1)	5.4	57	TNT/USPEC	*KG5*	19.7

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	56746.73621	56746.76277(1)	5.4	57	TNT/USPEC	*KG5*	19.4
	56746.80606	56746.83053(1)	5.4	56	TNT/USPEC	*KG5*	19.5
SDSS J125023.84+665525.4	53798.98823	53799.00334(4)	4/4/4	32	WHT/UCAM	$u'/g'/r'$	18.2/18.3/18.1
	53799.22336	53799.23838(4)	4/4/4	23	WHT/UCAM	$u'/g'/r'$	18.3/18.4/18.2
	53799.93218	53799.94313(4)	4/4/4	23	WHT/UCAM	$u'/g'/r'$	18.6/18.6/18.5
	53802.92144	53802.93866(4)	4/4/4	37	WHT/UCAM	$u'/g'/r'$	18.6/18.7/18.6
	53802.97429	53802.99743(4)	4/4/4	40	WHT/UCAM	$u'/g'/r'$	18.5/18.7/18.5
	53803.10330	53803.11493(4)	4/4/4	23	WHT/UCAM	$u'/g'/r'$	18.6/18.8/18.7
	53803.15397	53803.17368(4)	3/3/3	41	WHT/UCAM	$u'/g'/r'$	18.6/18.7/18.6
	53803.21193	53803.23237(4)	4/4/4	37	WHT/UCAM	$u'/g'/r'$	18.6/18.9/18.6
	55202.24092	55202.25763(4)	9/2.2/2.2	61	WHT/UCAM	$u'/g'/z'$	18.6/19.0/18.7
	55204.10992	55204.13721(4)	11/3.8/3.8	52	WHT/UCAM	$u'/g'/r'$	18.4/18.7/18.5
	55204.17561	55204.19589(4)	11/3.8/3.8	33	WHT/UCAM	$u'/g'/r'$	18.5/18.8/18.6
	55204.24567	55204.25463(4)	11/3.8/3.8	22	WHT/UCAM	$u'/g'/r'$	18.6/18.9/18.8
	57076.91313	57076.92429(4)	9.9	32	TNT/USPEC	*KG5*	17.9
	57077.78863	57077.80532(4)	9.9	32	TNT/USPEC	*KG5*	18.0
SDSS J152419.33+220920.1	55709.06004	55709.09910(3)	20/5/5	80	NTT/UCAM	$u'/g'/r'$	19.3/19.3/18.9
	55709.14436	55709.16439(3)	20/5/5	35	NTT/UCAM	$u'/g'/r'$	19.3/19.2/18.9
	55712.08247	55712.10372(3)	12/3/3	39	NTT/UCAM	$u'/g'/r'$	19.2/19.2/18.9
	55714.11519	55714.12865(3)	12/3/3	26	NTT/UCAM	$u'/g'/r'$	19.0/19.1/18.8
	56046.05227	56046.07845(1)	12/4/4	51	WHT/UCAM	$u'/g'/r'$	18.8/19.1/19.0

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	56046.12728	56046.14374(1)	12/4/4	28	WHT/UCAM	$u'/g'/r'$	18.6/18.9/18.7
	56046.95909	56046.99289(1)	12/4/4	58	WHT/UCAM	$u'/g'/r'$	18.5/18.9/18.6
	56486.89370	56486.91454(1)	10.6/3.5/3.5	39	WHT/UCAM	$u'/g'/i'$	19.1/19.3/18.8
	56494.86606	56494.88342(1)	10.6/3.5/3.5	40	WHT/UCAM	$u'/g'/i'$	19.3/19.3/19.0
	56503.87498	56503.89742(1)	13.6/4.5/4.5	49	WHT/UCAM	$u'/g'/i'$	19.3/19.4/19.1
	56509.95159	56509.97206(1)	10.6/3.5/3.5	60	WHT/UCAM	$u'/g'/r'$	19.2/19.5/19.1
	56719.23748	56719.25331(4)	12/4/4	37	WHT/UCAM	$u'/g'/r'$	19.5/19.7/19.6
	56746.87519	56746.88311(3)	3.4	23	TNT/USPEC	*KG5*	19.4
	56871.95572	56871.96849(4)	8/2/2	43	WHT/UCAM	$u'/g'/r'$	19.0/19.2/18.9
	56873.91985	56873.92807(4)	9/2.2/2.2	24	WHT/UCAM	$u'/g'/r'$	19.0/19.2/18.8
SDSS J155531.99-001055.0	54260.20835	54260.22695(1)	6/2/2	25	VLT/UCAM	$u'/g'/r'$	19.1/19.4/19.2
	54260.27914	54260.30580(1)	6/2/2	54	VLT/UCAM	$u'/g'/r'$	19.0/19.4/19.2
	54272.98205	54272.99993(1)	6/2/2	27	VLT/UCAM	$u'/g'/r'$	19.0/19.2/19.0
	55334.26865	55334.33987(2)	15/5/5	123	NTT/UCAM	$u'/g'/r'$	19.0/19.3/19.0
	55712.20386	55712.24662(2)	11/3.5/3.5	71	NTT/UCAM	$u'/g'/r'$	19.0/19.4/19.2
	56046.21549	56046.23642(2)	12/4/4	27	WHT/UCAM	$u'/g'/r'$	19.0/19.3/19.0
	56176.86930	56176.88349(2)	14/3.5/3.5	40	WHT/UCAM	$u'/g'/r'$	19.0/19.4/19.2
	56489.01055	56489.03300(2)	9/3/3	45	WHT/UCAM	$u'/g'/i'$	19.0/19.3/19.0
	56498.88028	56498.88873(2)	9.8/3.2/3.2	23	WHT/UCAM	$u'/g'/i'$	19.1/19.4/19.1
	56501.87310	56501.88484(2)	9/3/3	24	WHT/UCAM	$u'/g'/i'$	18.9/19.3/18.9
	56872.91245	56872.93200(2)	9.8/3.2/3.2	35	WHT/UCAM	$u'/g'/r'$	18.8/19.2/18.9
	56873.93949	56873.95698(2)	9.8/3.2/3.2	48	WHT/UCAM	$u'/g'/r'$	19.0/19.2/18.9

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	**T_{exp} (s)**	**ΔT (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	57081.86284	57081.87272(2)	15	32	TNT/USPEC	*KG5*	19.6
	57198.07630	57198.09102(2)	9/3/3	24	WHT/UCAM	$u'/g'/r'$	18.8/19.1/18.8
SDSS J155656.92+352336.6	54683.88334	54683.95644(4)	3.4/3.4/3.4	121	WHT/UCAM	$u'/g'/r'$	17.5/17.2/17.2
	55942.22648	55942.25538(4)	12/4/4	63	WHT/UCAM	$u'/g'/r'$	17.7/17.4/17.4
	56046.01331	56046.02708(8)	12/4/4	34	WHT/UCAM	$u'/g'/r'$	18.9/18.8/18.7
	56046.09093	56046.11505(8)	12/4/4	49	WHT/UCAM	$u'/g'/r'$	18.9/18.7/18.6
	56046.18489	56046.20324(8)	12/4/4	41	WHT/UCAM	$u'/g'/r'$	18.9/18.8/18.8
	56499.92019	-	12/4/4	51	WHT/UCAM	$u'/g'/r'$	19.1/19.1/19.0
	56501.89201	56501.9004(3)	120	156	*pt5m*	V	18.9
		56501.9885(3)					
	56873.97532	56873.9990(1)	6.5/2.1/2.1	46	WHT/UCAM	$u'/g'/r'$	19.0/18.8/18.8
	57079.85619	57079.8685(1)	8.9	39	TNT/USPEC	*KG5*	18.8
SSS100615:200331-284941	56872.97323	56873.024417(5)	12/3/3	78	WHT/UCAM	$u'/g'/r'$	19.9/19.7/19.6
	56874.01104	56874.022391(5)	9/3/3	21	WHT/UCAM	$u'/g'/r'$	19.7/19.6/19.4
	56874.94861	56874.961664(5)	15/5/5	18	WHT/UCAM	$u'/g'/i'$	19.4/19.5/19.8
SSS120402:134015350512	57027.91986	57027.95944(5)	6.9	71	TNT/USPEC	*KG5*	18.6
SSS130413:094551-194402	56683.65823	56683.67469(5)	5.8	35	TNT/USPEC	*KG5*	17.0
	56683.72993	56683.74054(5)	5.8	21	TNT/USPEC	*KG5*	17.1
	56683.78617	56683.80628(5)	5.8	40	TNT/USPEC	*KG5*	17.0

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	56684.68236	56684.72712(5)	5.8	90	TNT/USPEC	*KG5*	17.1
	56685.69475	56685.71357(5)	2.9	37	TNT/USPEC	*KG5*	17.1
	56689.63539	56689.65982(5)	2.9	46	TNT/USPEC	*KG5*	17.1
	56689.84202	56689.85708(5)	2.9	31	TNT/USPEC	*KG5*	17.1
	56690.75257	56690.77779(5)	2.9	49	TNT/USPEC	g'	17.1
	56690.82365	56690.84355(5)	2.9	33	TNT/USPEC	r'	17.0
	56699.05164	56699.06472(3)	1.0	25	SALT/SALTICAM	g'	17.0
	56739.62758	56739.6449(1)	2.9	31	TNT/USPEC	*KG5*	17.2
	56777.84234	56777.85636(3)	1.0	35	SALT/SALTICAM	g'	17.0
	57023.75635	57023.76773(3)	2.9	29	TNT/USPEC	i'	16.8
	57416.70864	57416.73923(5)	9.3	52	TNT/USPEC	u'	16.9
	57418.69218	57418.71238(5)	8.1	29	TNT/USPEC	u'	16.6
V2051 Oph	54261.28607	54261.31011(3)	1/1/1	45	VLT/UCAM	$u'/g'/r'$	15.5/15.6/15.4
	54264.16307	54264.18181(3)	1/1/1	30	VLT/UCAM	$u'/g'/r'$	15.6/15.6/15.5
	54268.27875	54268.30206(3)	3/1/1	41	VLT/UCAM	$u'/g'/r'$	15.8/15.9/15.7
	54268.34368	54268.36449(3)	3/1/1	46	VLT/UCAM	$u'/g'/r'$	16.0/16.0/15.9
	55314.14160	55314.15612(3)	6/2/2	128	NTT/UCAM	$u'/g'/r'$	15.5/15.4/15.1
		55314.21854(3)					
	55314.30287	55314.34340(3)	6/2/2	86	NTT/UCAM	$u'/g'/r'$	15.7/15.5/15.4
	55315.38811	55315.40460(3)	6/2/2	59	NTT/UCAM	$u'/g'/r'$	15.7/15.6/15.4
	55329.30344	55329.32629(3)	3/1.5/1.5	55	NTT/UCAM	$u'/g'/r'$	15.2/15.0/14.8
	55353.32613	55353.36077(3)	4/1.3/1.3	69	NTT/UCAM	$u'/g'/r'$	15.6/15.5/15.3

Object	Start (MJD)	Mid-eclipse time (BMJD)	T_{exp} (s)	ΔT (min)	Tel./Inst.	Filter	Mag.
	55354.01710	55354.04752(3)	4/1.3/1.3	57	NTT/UCAM	$u'/g'/r'$	15.8/15.7/15.5
	55700.37290	55700.39738(3)	9/3/3	44	NTT/UCAM	$u'/g'/r'$	16.0/15.7/15.5
	55701.32116	55701.33373(3)	9/3/3	26	NTT/UCAM	$u'/g'/r'$	15.8/15.7/15.5
	55701.38384	55701.39615(3)	9/3/3	57	NTT/UCAM	$u'/g'/r'$	15.5/15.5/15.3
	55705.28872	55705.32906(3)	9/3/3	77	NTT/UCAM	$u'/g'/r'$	15.7/15.6/15.4
	55706.29035	55706.32798(3)	10/3.3/3.3	91	NTT/UCAM	$u'/g'/r'$	16.0/15.9/15.6
	55709.31110	55709.32452(3)	6/2/2	95	NTT/UCAM	$u'/g'/r'$	16.0/15.8/15.6
		55709.38697(3)					
	55710.23956	55710.26091(3)	10/2/2	34	NTT/UCAM	$u'/g'/r'$	15.7/15.5/15.4
	55713.23903	55713.25747(3)	9/3/3	34	NTT/UCAM	$u'/g'/r'$	15.8/15.8/15.6
	55713.30486	55713.31989(3)	6/2/2	31	NTT/UCAM	$u'/g'/r'$	15.8/15.7/15.5
	55713.36174	55713.38235(3)	7.4/2.4/2.4	40	NTT/UCAM	$u'/g'/r'$	15.7/15.7/15.5
	55714.24128	55714.25632(3)	8.4/2.8/2.8	64	NTT/UCAM	$u'/g'/r'$	15.8/15.6/15.5
	57080.91855	57080.92712(3)	1.9	29	TNT/USPEC	*KG5*	15.6
	57082.91711	57082.92489(3)	2.0	37	TNT/USPEC	*KG5*	15.6
	57084.91585	57084.92251(3)	2.0	22	TNT/USPEC	*KG5*	15.6
V713 Cep	55801.14595	55801.18114(2)	4.4/2.2/2.2	64	WHT/UCAM	$u'/g'/i'$	18.9/18.8/18.7
	56046.15044	-	6.3/2.1/2.1	46	WHT/UCAM	$u'/g'/r'$	18.7/18.6/18.5
	56176.90266	56176.93715(2)	10/3.4/3.4	61	WHT/UCAM	$u'/g'/r'$	18.4/18.3/18.2
	56177.84463	56177.87676(2)	10/3.4/3.4	65	WHT/UCAM	$u'/g'/r'$	18.3/18.2/18.1
	56179.99153	56180.01222(2)	10/3.4/3.4	47	WHT/UCAM	$u'/g'/r'$	18.3/18.3/18.2
	56488.17430	56488.20223(2)	10/3.4/3.4	74	WHT/UCAM	$u'/g'/z'$	18.6/18.5/17.6

Object	**Start (MJD)**	**Mid-eclipse time (BMJD)**	$\mathbf{T}_{exp}$ **(s)**	Δ**T (min)**	**Tel./Inst.**	**Filter**	**Mag.**
	56489.04428	56489.05642(2)	10/3.4/3.4	35	WHT/UCAM	$u'/g'/i'$	18.6/18.4/18.2
	56489.13138	56489.14184(2)	10/3.4/3.4	26	WHT/UCAM	$u'/g'/i'$	18.6/18.5/18.3
	56489.18909	56489.22726(2)	10/3.4/3.4	66	WHT/UCAM	$u'/g'/i'$	18.7/18.5/18.3
	56499.19598	56499.22121(2)	10/3.4/3.4	59	WHT/UCAM	$u'/g'/z'$	18.4/18.4/17.6
	56509.18559	56509.21517(2)	10/3.4/3.4	77	WHT/UCAM	$u'/g'/r'$	18.5/18.5/18.4
	56510.12816	56510.15480(2)	10/3.4/3.4	56	WHT/UCAM	$u'/g'/r'$	18.6/18.5/18.4
	56872.03795	56872.07302(3)	6/2/2	74	WHT/UCAM	$u'/g'/r'$	18.2/18.2/18.0
	56880.16994	56880.18780(3)	10/3.4/3.4	41	WHT/UCAM	$u'/g'/i'$	18.5/18.4/18.2
	57198.09643	57198.11544(2)	10/3.4/3.4	35	WHT/UCAM	$u'/g'/r'$	19.2/19.0/19.0
	57282.84397	57282.85057(2)	15/5/5	20	WHT/UCAM	$u'/g'/r'$	18.5/18.6/18.4

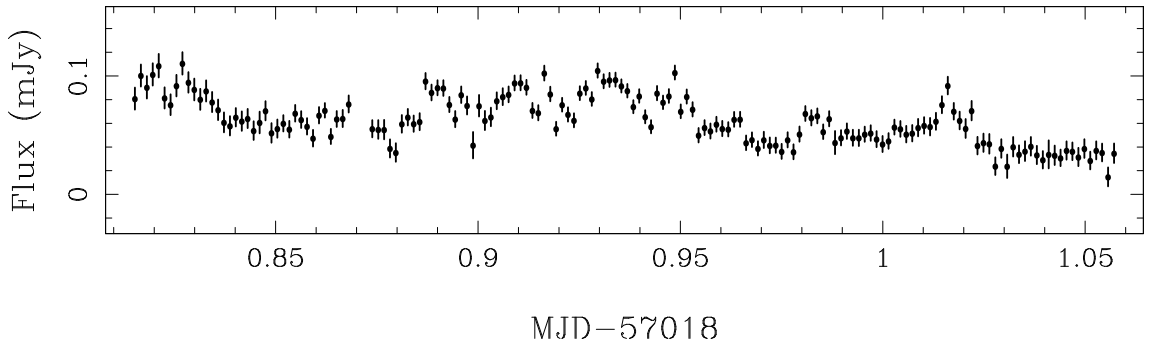

Figure B.1: ASASSN-14gl light curve observed during quiescence on 2014-12-27 with *pt5m*.

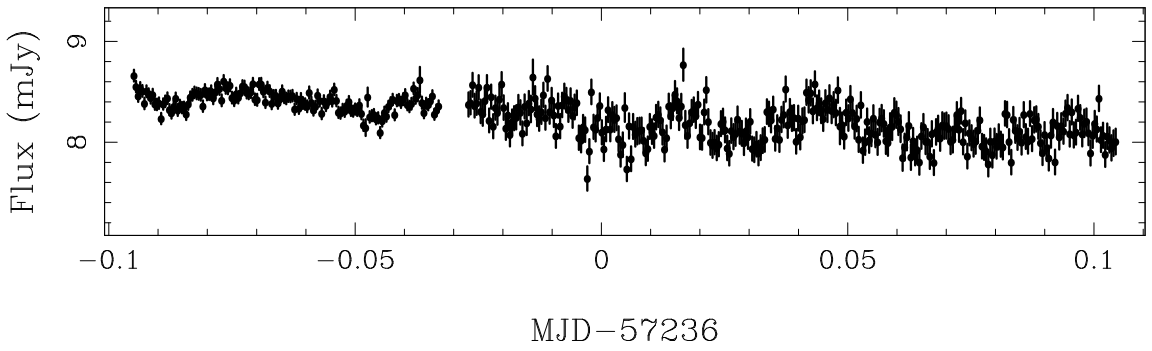

Figure B.2: ASASSN-15ni light curve observed on 2015-08-01 during the decline from outburst with *pt5m*. The gap in the light curve is due to the telescope conducting a pier-flip as the target transited the observers meridian. After the pier-flip, a different (fainter) comparison star had to be used. This explains the difference in the size of the error bars before and after the gap.

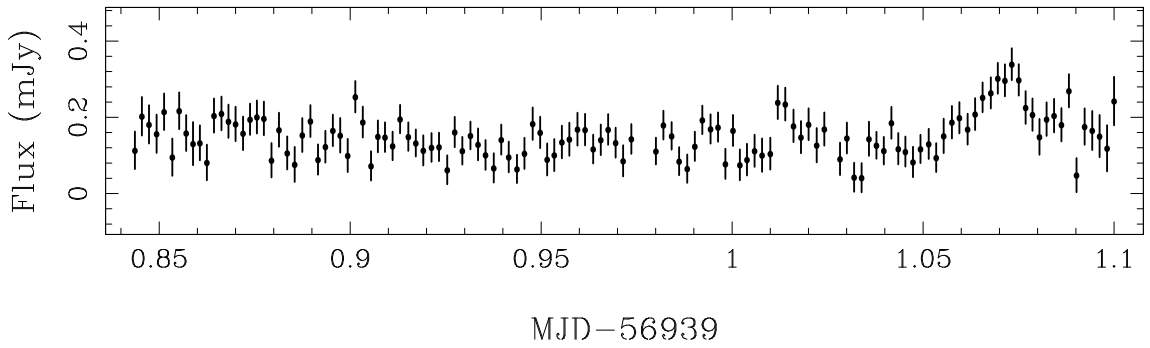

Figure B.3: CSS091116:232551-014024 light curve observed during quiescence on 2014-10-09 with *pt5m*.

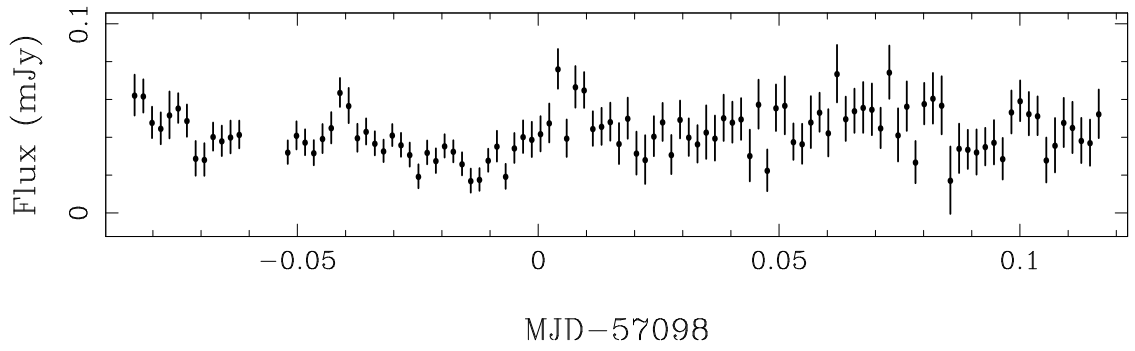

Figure B.4: CSS100508:085604+322109 light curve observed in quiescence on 2015-03-16 with *pt5m* in windy conditions. The sporadic variability is attributed to flickering.

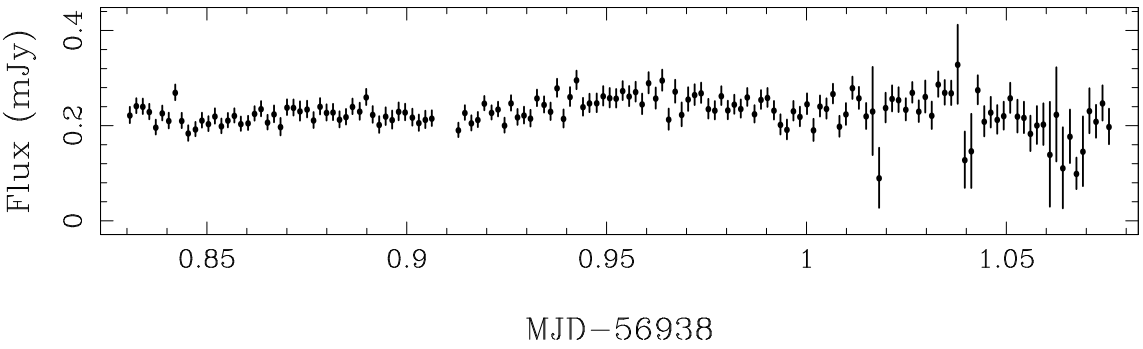

Figure B.5: CSS100520:214426+222024 light curve observed during quiescence on 2014-10-08 with *pt5m*. The disruption towards the end is due to patchy cloud cover.

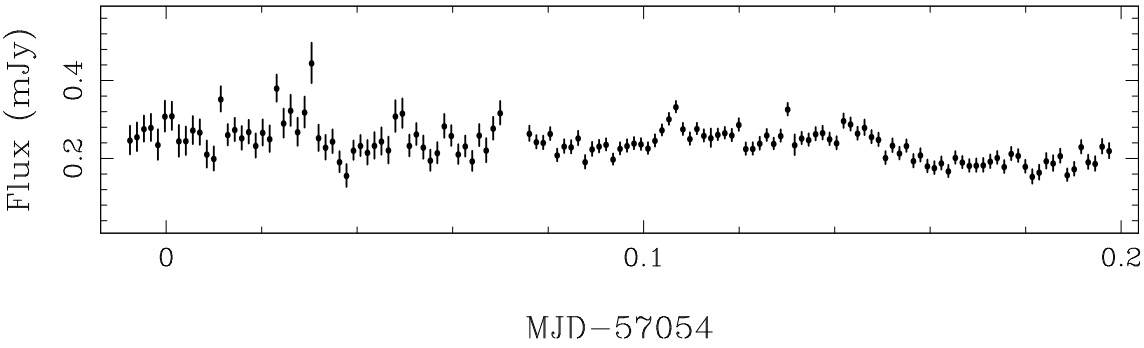

Figure B.6: CSS110114:091246-034916 light curve observed during quiescence on 2015-01-31 with *pt5m*.

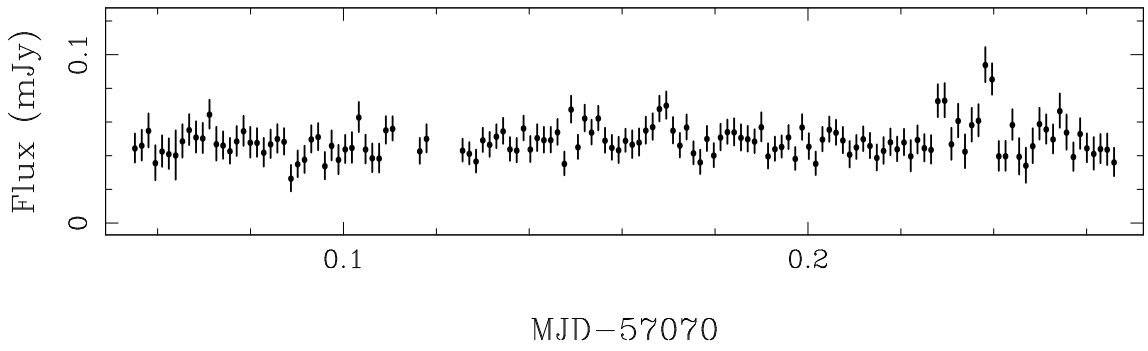

Figure B.7: CSS110226:112510+231036 light curve observed during quiescence on 2015-02-16 with *pt5m*.

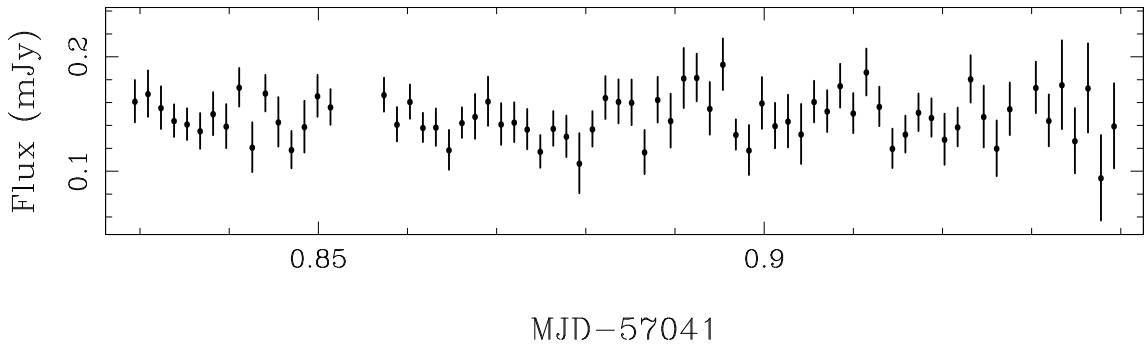

Figure B.8: CSS141117:030930+263804 light curve observed in quiescence on 2015-01-19 with *pt5m*. Some thin overhead cirrus affects the quality of the light curve.

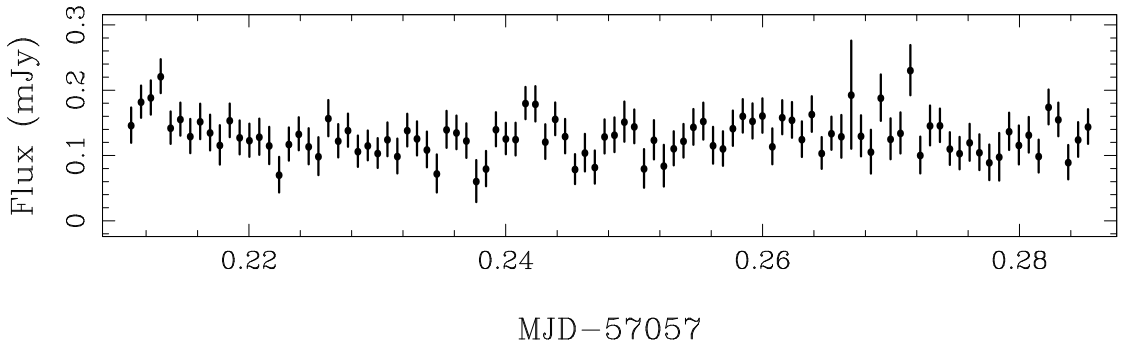

Figure B.9: Gaia15aan/ASASSN-14mo light curve observed during the decline from outburst, on 2015-02-03 with *pt5m*.

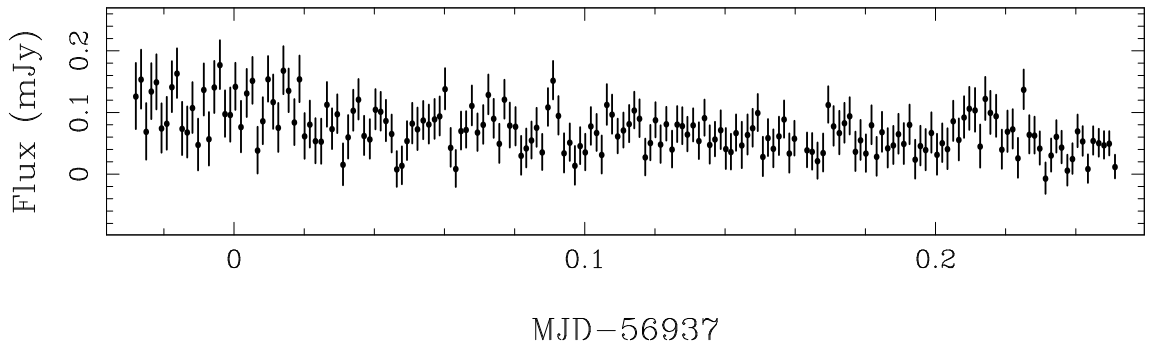

Figure B.10: MASTER OT J034045.31+471632.2 light curve observed during quiescence on 2014-10-06 with *pt5m*.

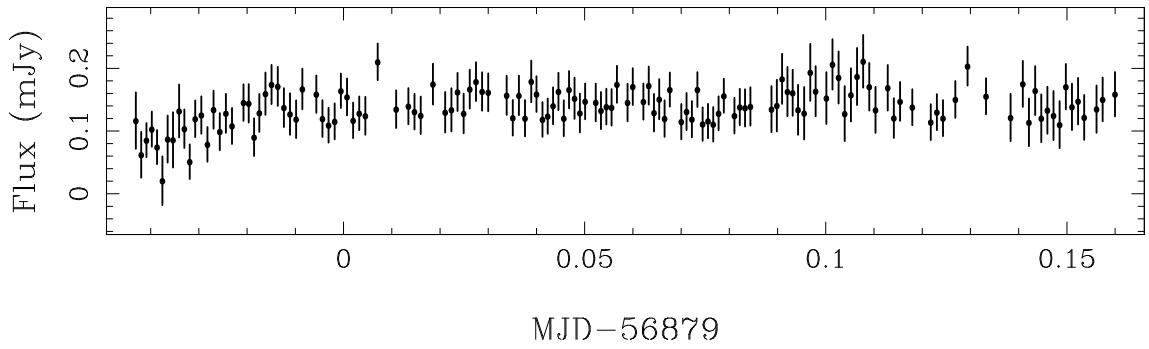

Figure B.11: MASTER OT J194955.17+455349.6 light curve observed in quiescence on 2014-08-09 with *pt5m*.

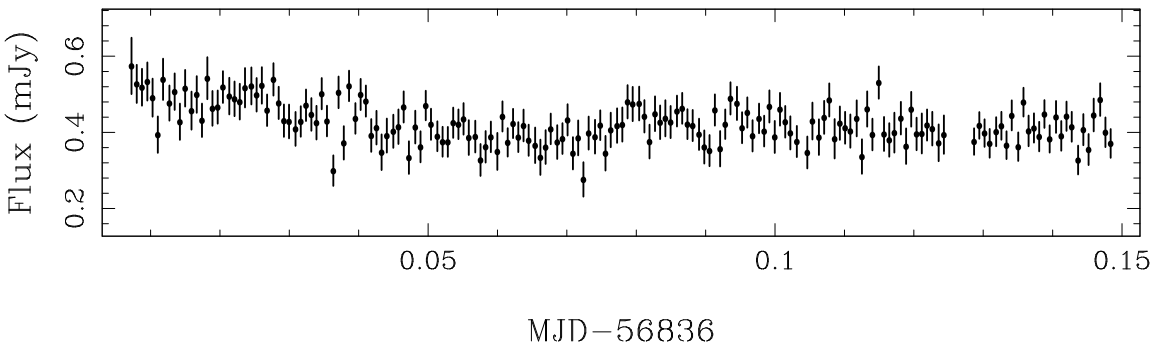

Figure B.12: MASTER OT J201121.95+565531.1 light curve observed during the decline from outburst on 2014-06-27 with *pt5m*.

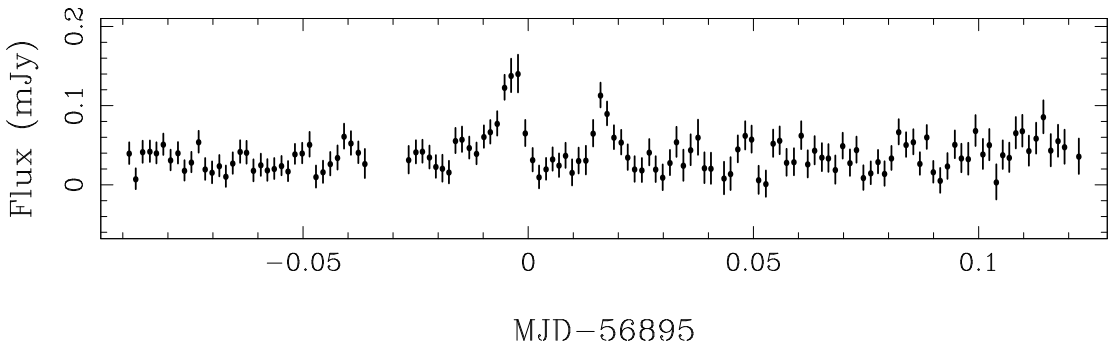

Figure B.13: MASTER OT J202157.69+212919.4 light curve observed on 2014-08-25 with *pt5m*. The two peaks in the centre are caused by the apertures tracing across hot pixels on the CCD.

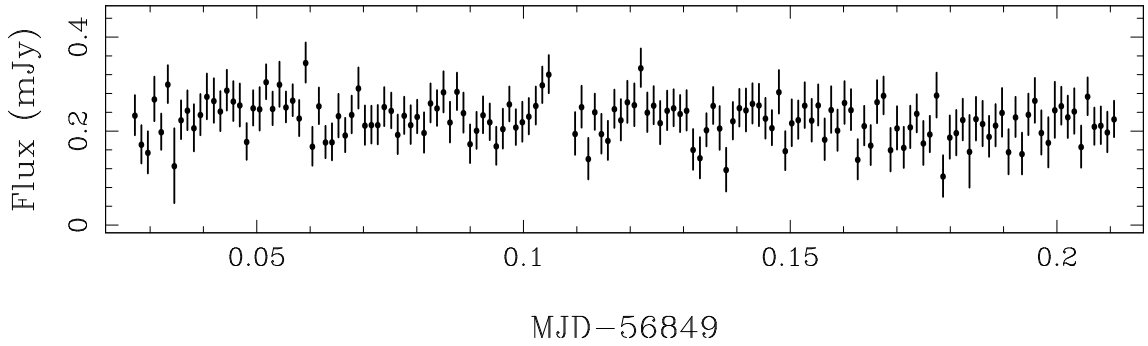

Figure B.14: MASTER OT J203421.90+120656.9 light curve observed during outburst on 2014-07-10 with *pt5m*.

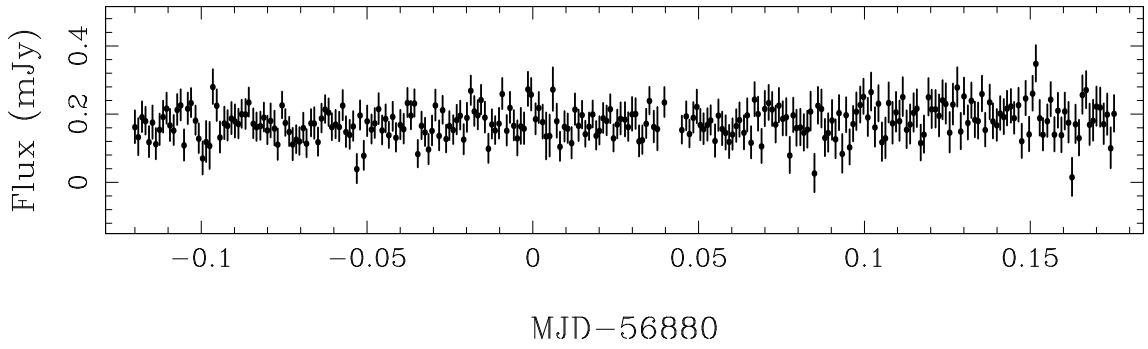

Figure B.15: MASTER OT J210316.39+314913.6 light curve observed during quiescence on 2014-08-10 with *pt5m*.