ULTRACAM – an ultra-fast, triple-beam CCD camera

Mark John Stevenson

Department of Physics and Astronomy, University of Sheffield

September 2004

A thesis submitted to the University of Sheffield
in partial fulfilment of the requirements of the
degree of Doctor of Philosophy

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 in this University or elsewhere.

This thesis is the result of my own work unless otherwise stated.

Summary

Charge-Coupled Devices (CCD) have revolutionised observational astronomy since they were introduced in the 1970s. Their only limitation is their inability to operate at high frame rates. This has meant that faint objects that vary on timescales of less than of order seconds cannot be temporally resolved, ruling out, for example, the study of variability in compact objects, such as white dwarfs, neutron stars and black holes.

ULTRACAM is a triple-beam camera designed to overcome this limitation and enable the study of astrophysics on fast timescales using CCDs. The project was awarded £292,034 in July 1999 by the Particle Physics and Astronomy Research Council and the instrument saw first light on 16 May 2002 on the William Herschel Telescope, on budget and three months ahead of schedule.

This thesis describes my contribution to the ULTRACAM project. It begins with a description of the motivation for building ULTRACAM and lists its functional and performance requirements. An instrument design which meets these requirements is then presented, followed by a description of the manufacture and integration phase of the project. A whole chapter is then devoted to an in-depth analysis of the commissioning data obtained on the WHT, which verifies that ULTRACAM performs to specification. The thesis concludes with some suggestions for enhancements and future work.

Acknowledgements

I would like to convey my immense thanks and gratitude to Vikram Dhillon for his expert supervision throughout my Ph.D. His enthusiasm and knowledge have not only provided me with a wealth of observing experience but have provided me with skills that I shall carry with me for the rest of my life. It is also purely thanks to his supervision and keen eye for detail that I have finished this thesis.

Thanks should also go to the rest of the ULTRACAM project team, who have made working on ULTRACAM interesting and more importantly fun. My thanks go all those at the United Kingdom Astronomy Technology Centre, especially to Tully Peacocke for his wisdom in optics and their mounts, Dave Atkinson for his hard work on the CCD detectors and Andy Vick for his work on the software. I would also like to thank Paul Kerry, our Starlink manager, computer expert and thoroughly good bloke. He has without a doubt added to the success of ULTRACAM on more occasions than I can remember. I would also like to thank him for all of his help and support throughout the rest of my Ph.D.

If it was not for the hard work and effort of pretty much the entire staff in the Central Mechanical Workshop at the University of Sheffield, ULTRACAM would never have been completed on time. My personal thanks go to John Kelly, Richard 'Dicky' Nicholson, Richard 'Frenchie' French, Richard Pashley, Paul Russell, Simon Dixon, Martin Bently, Tim Haycock, and many others who I have failed to mention. It is only with effort above and beyond the call of duty that we were successful, so thanks guys.

My thanks now turn to those who made my Ph.D. fun by brightening up the office. I would like to thank Mike Bellamy for his sarcastic wit, Lucy Hadfield for her soothing Yorkshire tones, Jo Holt for keeping a modicum of sanity in the office and Will Feline for being Will. They made life bearable (minus the odd hangover) when times were low. I would also like to thank Chris Watson for his help in FORTRAN and being a good friend and succumbing to a good thrashing at tennis and Mark Taylor for also being a good friend and entertaining us with his amazing repertoire of gaffs and blunders, not withstanding his prowess on the ski slopes – not. My final thanks here go to Tim Thoroughgood for being a good mate and drinking partner. Cheers Woody.

Of course I wouldn't be where I am without my mum and dad, Sylvia and John. There must have been times when they wondered whether I was doing the right thing and where all this work was taking me. Some of the choices that I made must have appeared strange, but my mum and dad never questioned them and allowed me the freedom to make my own decisions. The unflinching support and love they have given me are the only reason I have finished this work – I'm only here because of you.

My final thanks go to Kirsty, who has suffered my many moans and panics, especially during thesis submission. She has never doubted that I could do this, even when I felt it was all too much. Her undeniable confidence in me gave me the will to carry on. My thanks and love go to you – I couldn't have done it without you.

To my Mum and Dad

For their unwavering support, love and most of all, belief in me

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Charge-Coupled Devices (CCDs) have revolutionised astronomy in recent years. They were introduced in 1970, and were originally intended as memory storage devices where a packet of charges acted as a physical quantity to determine the presence or absence of a bit of information. These charges are stored in a depletion region in the CCD architecture and can be moved from one site to the next by manipulating voltages, hence the name 'charge-coupled device'. The structure and operation of CCDs is described in section 1.5.1.

It soon became obvious to those working in the semiconductor field that CCDs could have other applications. Photons incident on the CCD structure liberate electrons due to the photoelectric effect. From this it was realised that, coupled with their small size, low weight, low-power consumption, robustness, linearity and sensitivity over a large wavelength, CCDs would make ideal scientific detectors. The only disadvantage with CCDs is their low frame rates, which results from the finite time required to shift the photo-generated electrons from the detection site and digitise them. In a typical CCD this process can take from a few seconds to minutes, meaning that any variability in an observed object below this timescale will not be detected.

This thesis will describe ULTRACAM, an ULTRA-fast CAMera designed to utilise the sensitivity and versatility of CCDs for high-speed imaging. ULTRACAM overcomes the low frames rates normally associated with CCDs by using frame transfer devices, which are described in section 2.3. This chapter will begin by reviewing the science that ULTRACAM will undertake, followed by an analysis of strengths and weaknesses of other instruments designed for high-time resolution astrophysics. Using the science requirements and the strengths of other instruments, functional and performance requirements will be defined for an ideal high-speed imaging instrument. An overview of the construction and operation

of CCDs will then be presented, followed by some examples of CCDs that have been employed in high-speed imaging.

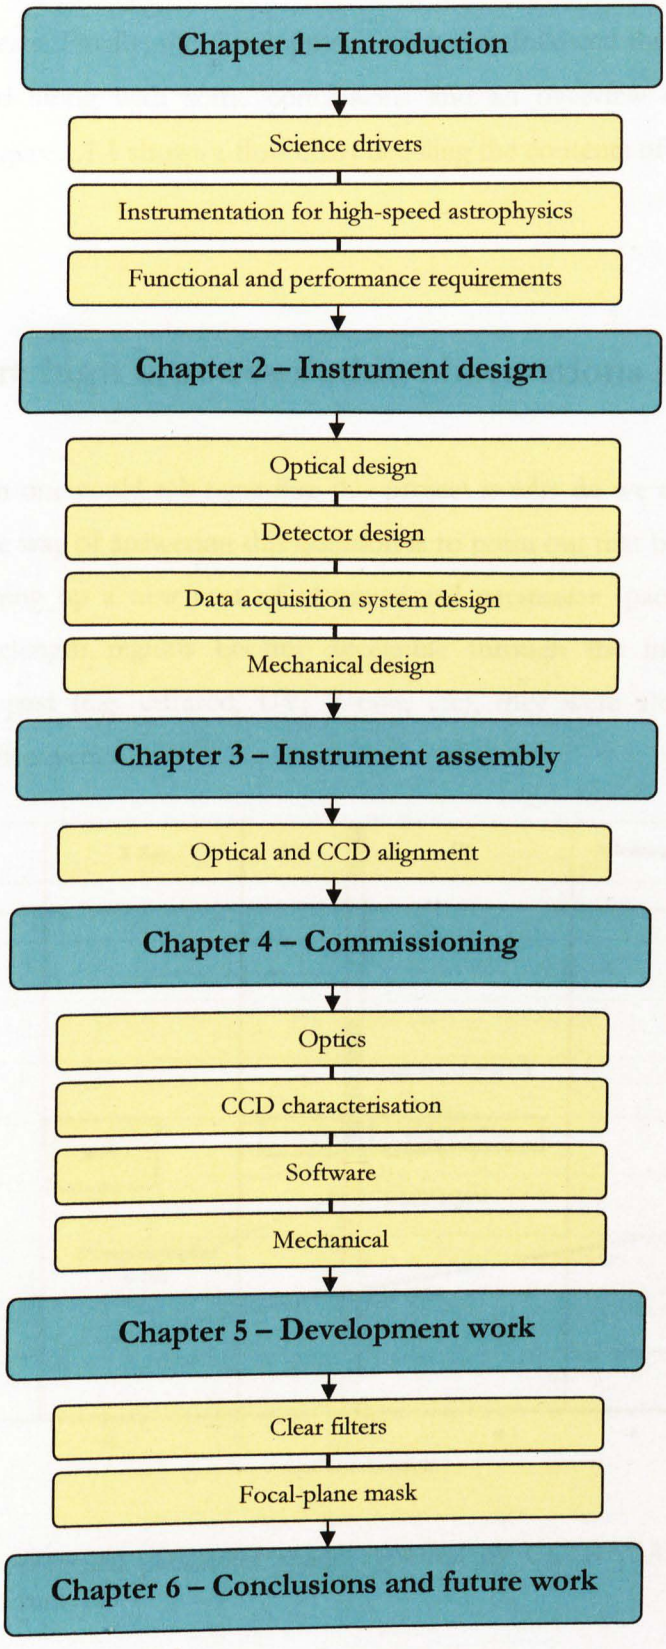

Figure 1.1.1: Flow chart detailing the contents of this thesis.

The remaining chapters of this thesis will then describe the design of ULTRACAM, covering the optics, detectors, data acquisition software and the mechanical chassis. This will be followed by a description of the assembling of ULTRACAM and the analysis of the commissioning tests. Finally, the development work that followed the commissioning run will be described along with some conclusions and an overview of future plans for ULTRACAM. Figure 1.1.1 shows a flow chart detailing the contents of this thesis.

1.1 Why are high time-resolution observations important?

The first question one could ask regarding this project is why do we need to make CCDs image faster? One way of answering this question is to point out that by imaging at a faster rate we are opening up a new area of observational parameter space (see figure 1.1.2). When new wavelength regions became accessible through the introduction of new detectors in the past (e.g. infrared, UV, X-rays, etc), they were always followed by a plethora of new discoveries.

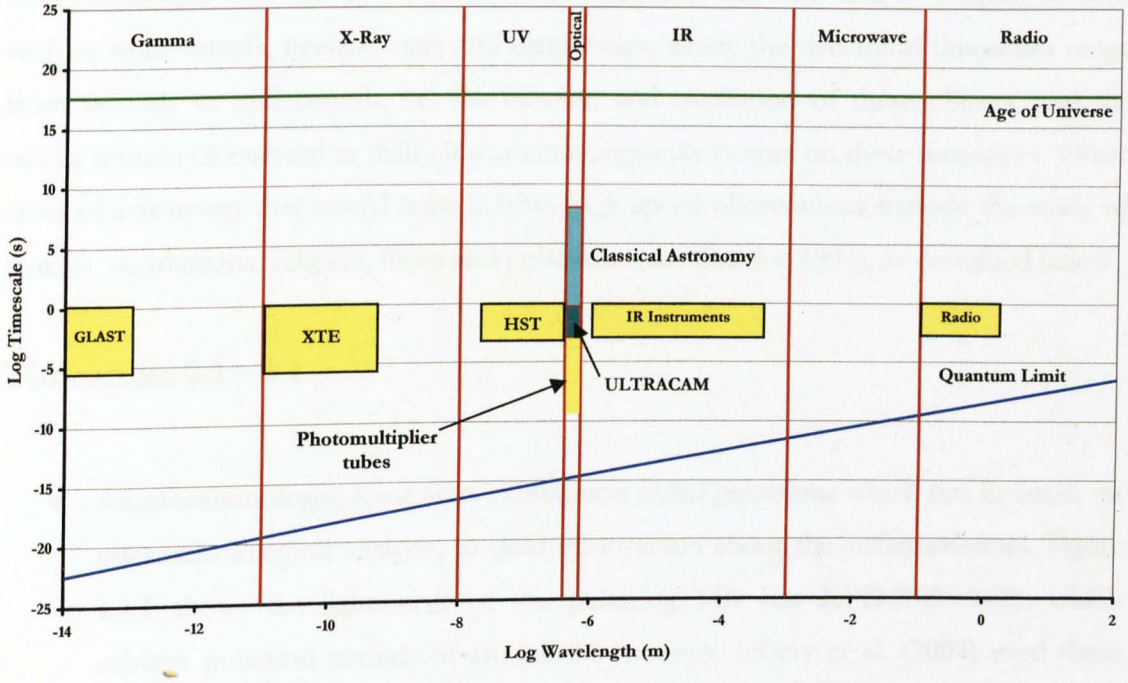

Figure 1.1.2: Observational parameter space explored by ULTRACAM and other high time-resolution instruments.

Figure 1.1.2 shows two-dimensions of observational parameter space - the timescale of observations as a function of wavelength. The upper limit to the timescale of any observation is the age of the universe, whilst the lower limit is defined by the quantum limit $\Delta t = 1/\nu$, where Δt is the observed timescale in seconds and ν is the frequency of the observed light in hertz. In the visible part of the spectrum, classical astronomy has typically observed on timescales of seconds to years. The gamma ray detector (GLAST), the X-ray Timing Explorer (XTE), the Hubble Space Telescope (HST) and radio telescopes have all contributed to high-time resolution astrophysics at gamma to radio wavelengths, as shown in figure 1.1.2. With the exception of Photomultiplier Tubes (PMT) which will be discussed later, there have, however, been few equivalent optical instruments - until ULTRACAM.

1.2 Science Drivers

The study of a new region of observational parameter space, like the optical high-time resolution domain, offers not only the chance of serendipitous discoveries but also allows closer investigation of known fast time-scale variations. One such area is compact objects, such as white dwarfs, neutron stars and black holes, where the dynamical timescales range from seconds to milliseconds, i.e. the rotation and oscillation of these objects, and the orbital motion of material in their close vicinity, typically occurs on these timescales. Other areas of astronomy that would benefit from high-speed observations include the study of transits, occultations, eclipses, flares and pulsations (see Dravins 1994), as described below.

Timescale: 0.1 – 1 s

- **Asteroseismology**: some stars exhibit non-radial pulsations which can be used, via asteroseismological analysis, to yield information about the stellar interiors. Figure 1.2.1 shows the light-curve of the pulsating sdB star KPD2109+4401, which exhibits pulsation periods of around 3.3 minutes. Jeffery et al. (2004) used these data to identify the pulsation modes and, by comparing with theoretical models of the star's interior, deduced that the envelope mass is ~$0.0002M_{solar}$ and the star has a post-horizontal branch age of approximately 47 Myr.

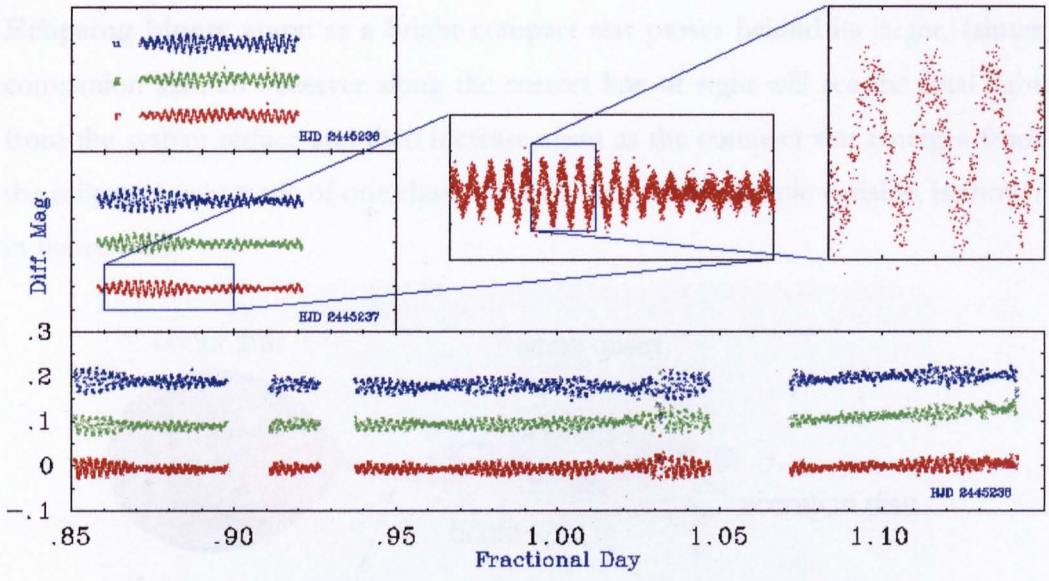

Figure 1.2.1: ULTRACAM light-curve of the sdB star KPD2109+4401, illustrating the 3.3 minute pulsations in the *u'*, *g'* and *r'* bands (Jeffery et al. 2004).

- **Planetary transits:** a number of extra-solar planets are known to transit their host stars (e.g. Brown et al. 2001). This provides a unique opportunity to determine the mass and radius of the planet via a fit to the observed light-curve (see figure 1.2.2). High time-resolution observations of these transits may yield evidence of moons by timing the ingress/egress over many cycles and searching for periodicities (e.g. Charbonneau 2003).

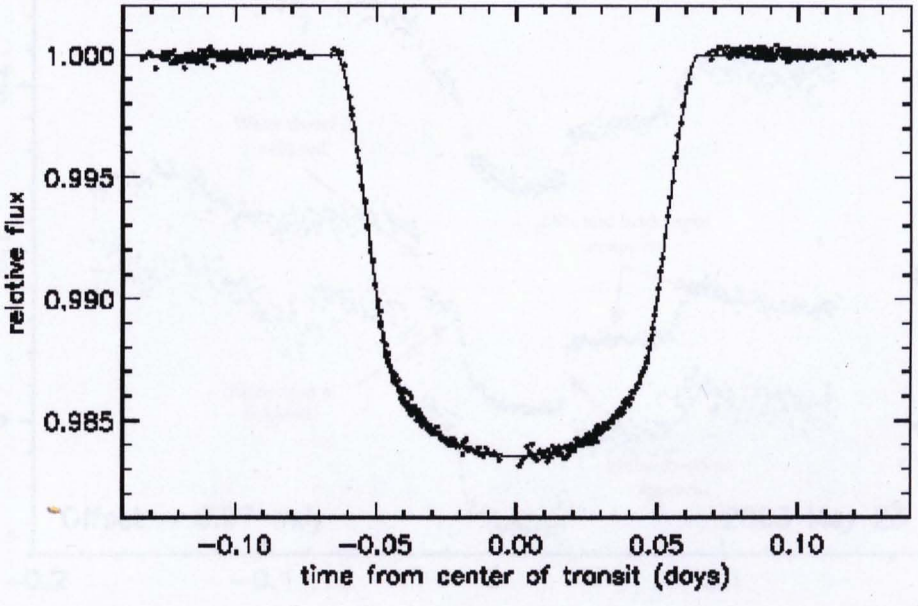

Figure 1.2.2: HST/STIS light-curve of the extra-solar planet HD 209458b (Brown et al. 2001). The solid line shows the fit to the data used to determine the mass and radius of the planet.

- **Eclipsing binary stars**: as a bright compact star passes behind its larger, fainter companion star, an observer along the correct line of sight will see the total light from the system reduce and then increase again as the compact star emerges from the eclipse. A schematic of one class of binary star - a cataclysmic variable, is shown in figure 1.2.3.

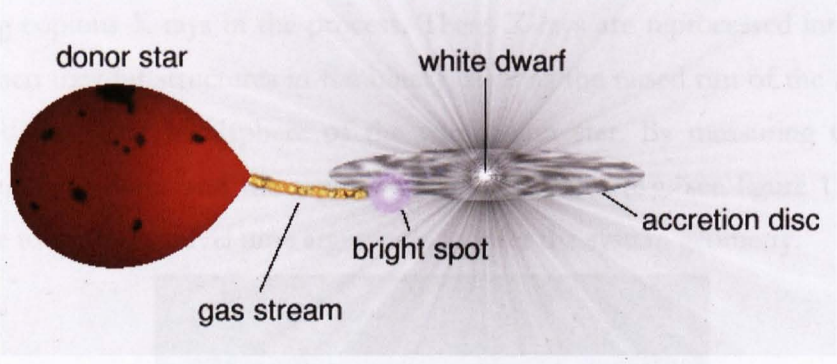

Figure 1.2.3: Schematic of a cataclysmic variable star (Christopher Watson – private communication).

- **In close binary systems,** the entire eclipse can take of the order of minutes. High-speed instruments can measure the sharp ingress and egress of the eclipse, as shown in figure 1.2.4. This information can be used to determine the stellar masses and radii (Feline et al. 2004a, b).

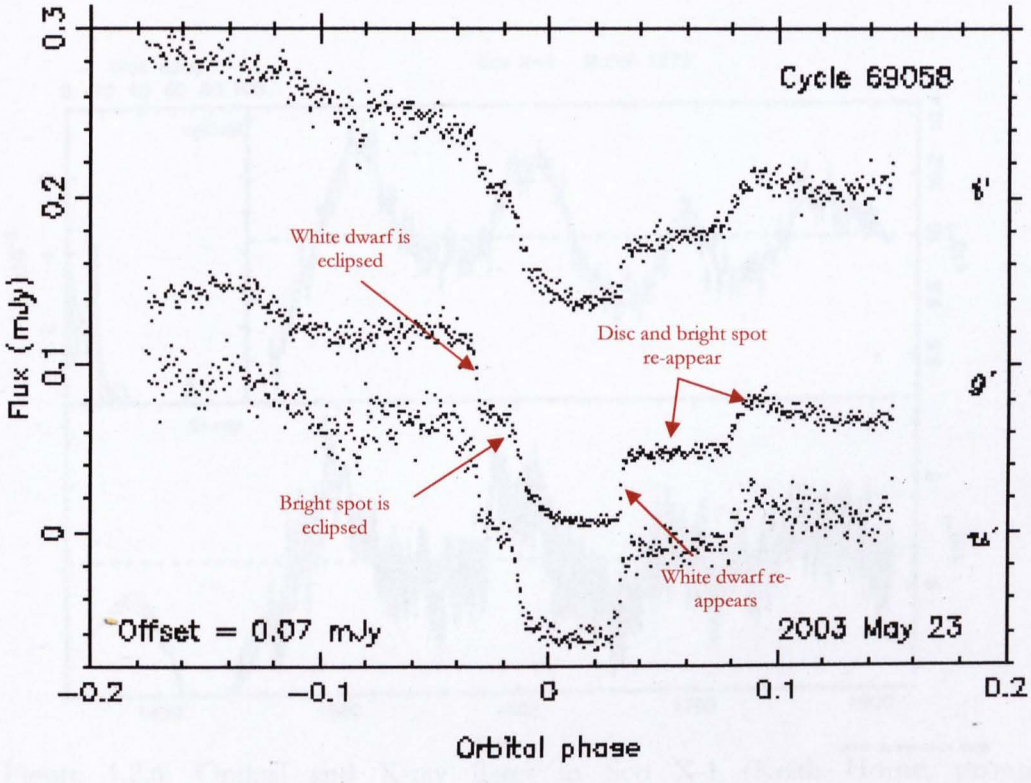

Figure 1.2.4: Light-curve of the cataclysmic variable DV UMa (Feline et al. 2004b).

Timescale: 0.01 - 0.1 s

- **Echo mapping**: X-ray binary stars consist of a main sequence or evolved companion star in orbit with a neutron star or black hole (see figure 1.2.5). The compact object accretes material from the companion via an accretion disc, emitting copious X-rays in the process. These X-rays are reprocessed into optical light when they hit structures in the binary, such as the raised rim of the accretion disc and the inner hemisphere of the companion star. By measuring the delay between X-ray flares and the reprocessed optical flares (e.g. see figure 1.2.6) it is possible to use light-travel time arguments to infer the system geometry.

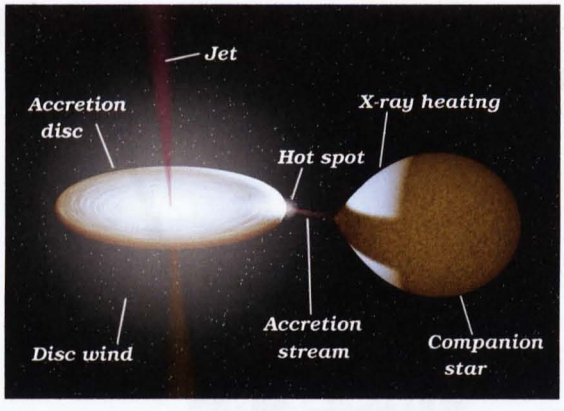

Figure 1.2.5: Schematic of an X-ray binary system (Hynes 2001).

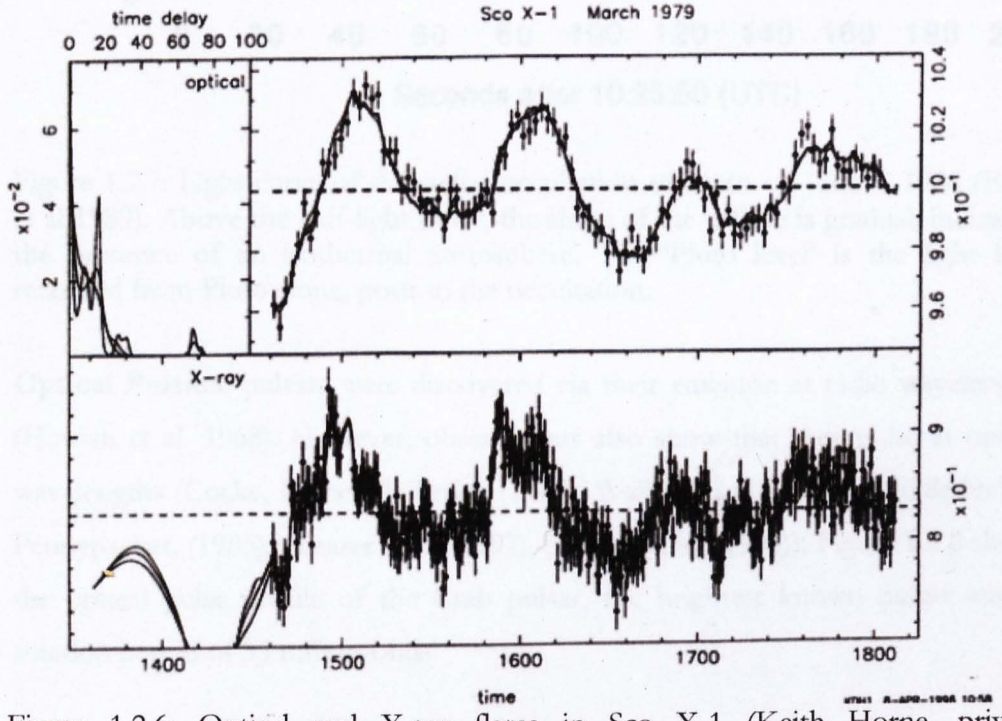

Figure 1.2.6: Optical and X-ray flares in Sco X-1 (Keith Horne, private communication).

Timescale: 0.001 - 0.01 s

- **Occultations**: observations of a star as it is occulted by a planet/satellite can reveal the existence and physical properties of the planet/satellite atmosphere (e.g. Elliot et al. 1989). Figure 1.2.7 shows a stellar occultation by Pluto. This light-curve (obtained at sub-second time resolution) was used to discover the existence of Pluto's atmosphere and measure its temperature (67 ± 6 K) and pressure (0.78 microbar).

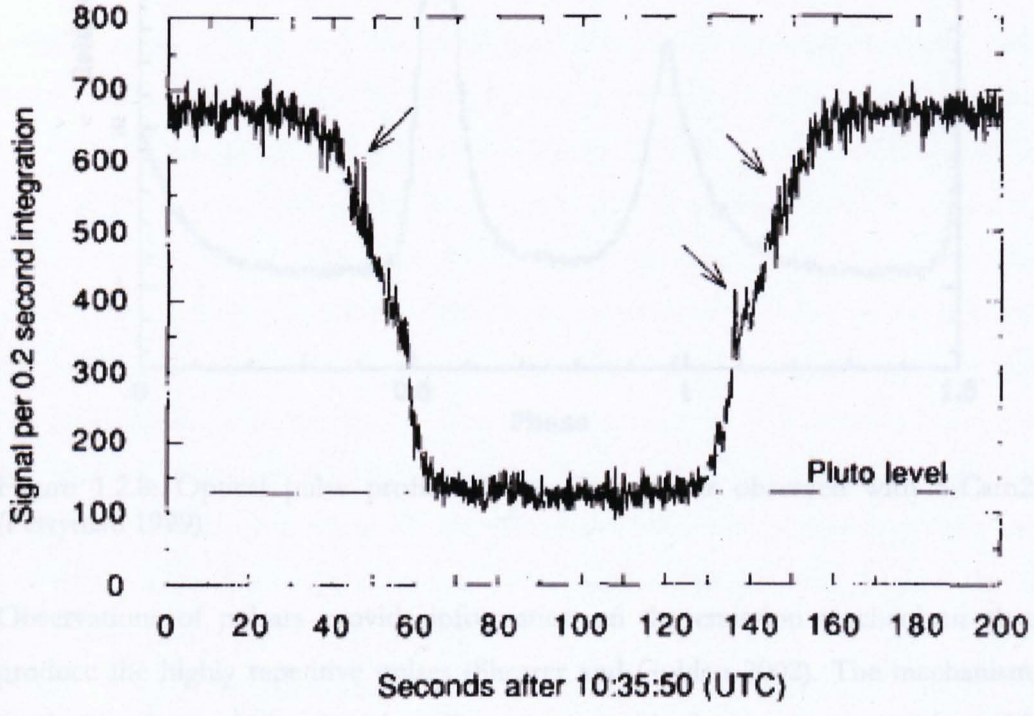

Figure 1.2.7: Light-curve of the stellar occultation of Pluto on June 9, 1988 (Elliot et al 1989). Above the half-light point, the slope of the eclipse is gradual, indicating the presence of an isothermal atmosphere. The 'Pluto level' is the light level recorded from Pluto alone, prior to the occultation.

- **Optical Pulsars**: pulsars were discovered via their emission at radio wavelengths (Hewish et al. 1968). However, observations also show that they pulse at optical wavelengths (Cocke, Disney & Taylor (1969), Wallace et al. (1977), Middleditch & Pennypacker, (1985), Shearer et al. (1997), Shearer et al. (1998)). Figure 1.2.8 shows the optical pulse profile of the Crab pulsar, the brightest known pulsar with a rotation period of 33 milliseconds.

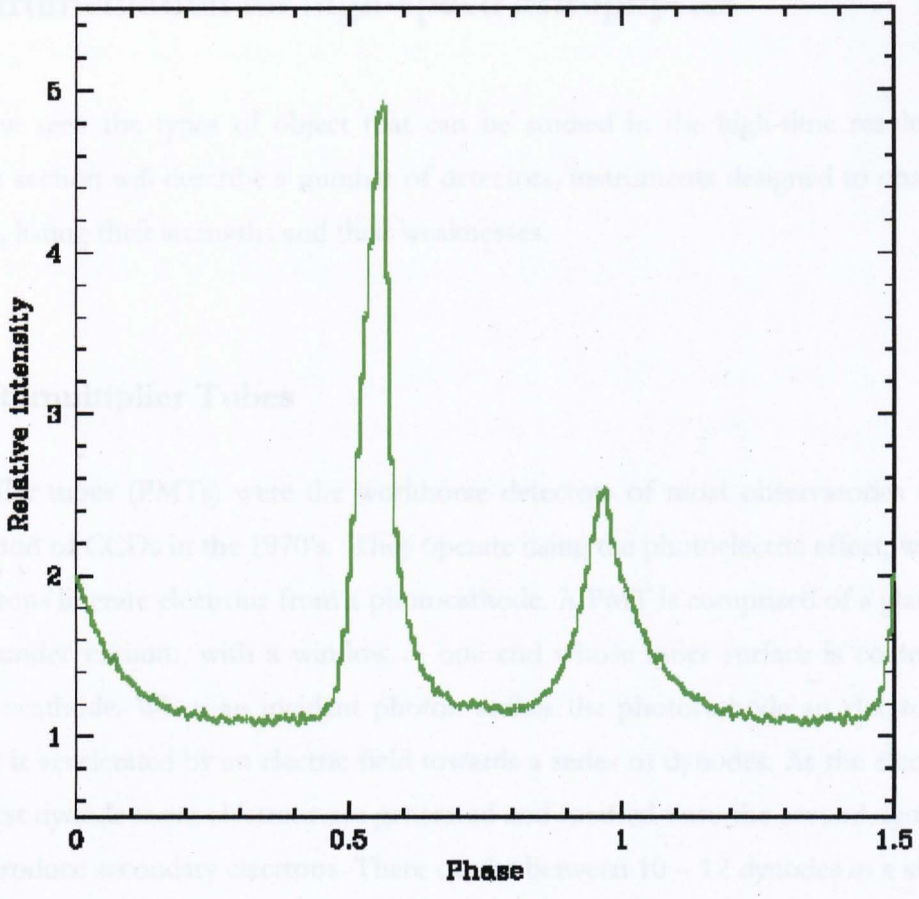

Figure 1.2.8: Optical pulse profile of the Crab pulsar observed with S-Cam2 (Perryman 1999)

- Observations of pulsars provide information on the emission mechanisms that produce the highly repetitive pulses (Shearer and Golden 2002). The mechanism producing the optical pulses is still uncertain, with the most successful model proposed by Pacini (1971) and revised later by Pacini and Salvati (1983; PS83, 1987; PS87). The model proposes that the optical emission is due to synchrotron radiation, but there is uncertainty as to where this originates from. Two models based on this propose the emission originates either close to the neutron star in the Polar gap model (Daugherty and Harding 1996) or further out in the magnetosphere in the Outer gap model (Cheng et al 2000). Further observation of pulsars will provide evidence so these models can determine the emission regions more precisely. However, most pulsars are extremely faint at optical wavelengths (e.g. Geminga, 26th magnitude) and have very short rotation periods (e.g. Crab pulsar, 33 ms), making them challenging targets for high-speed instrumentation, even on the world's largest telescopes.

1.3 Instrumentation for high-speed astrophysics

We have now seen the types of object that can be studied in the high-time resolution domain. This section will describe a number of detectors/instruments designed to observe these objects, listing their strengths and their weaknesses.

1.3.1 Photomultiplier Tubes

Photomultiplier tubes (PMTs) were the workhorse detectors of most observatories until the introduction of CCDs in the 1970's. They operate using the photoelectric effect, where incident photons liberate electrons from a photocathode. A PMT is comprised of a glass or quartz tube under vacuum, with a window at one end whose inner surface is coated to make a photocathode. When an incident photon strikes the photocathode an electron is liberated and is accelerated by an electric field towards a series of dynodes. As the electron strikes the first dynode more electrons are generated and focused onto the second dynode, where they produce secondary electrons. There can be between 10 – 12 dynodes in a single PMT, each separated in electrical potential by about 150 V. One electron can produce 3 – 8 electrons at each dynode, depending on the potential between each dynode. When the electron shower reaches the anode there can be as many as 10^7 electrons (Rapkin 1997). Figure 1.3.1 shows a schematic of a typical PMT.

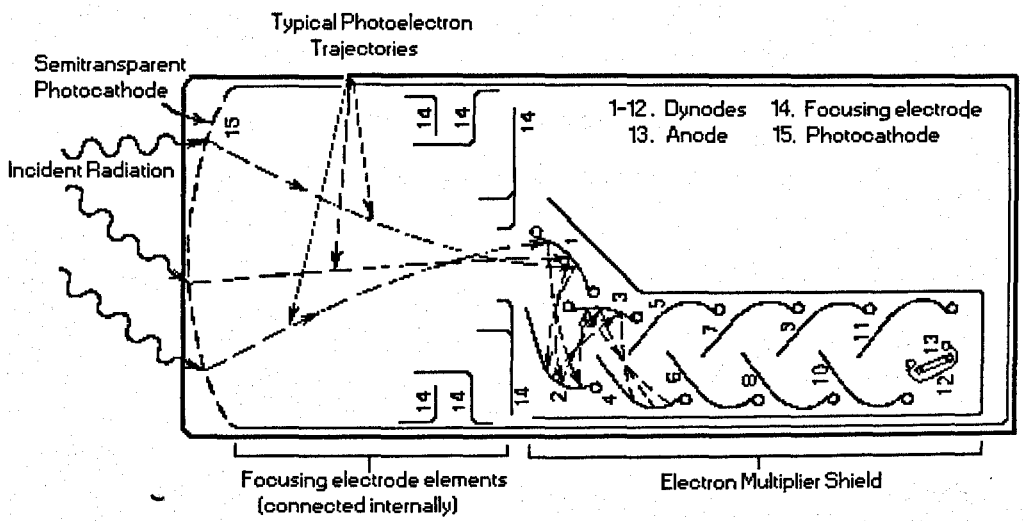

Figure 1.3.1: Schematic of a photomultiplier tube. Incident photons enter the tube from the left and strike the first dynode, liberating electrons. The electrons are then accelerated down the tube, striking more dynodes and liberating more electrons until the sizeable electron shower reaches the anode and is measured.

The dead time of a PMT is approximately 20 – 80 ns, making it ideal for observing rapidly changing phenomena. High voltages are critical at the photocathode (typically 1500 V) to ensure that there are enough electrons in the shower when it reaches the anode. The signal at the anode is typically a few mV, but if it is much lower then the signal will require electronic amplification, which introduces noise. However, if the voltage is too high then dark current inherent to PMTs increases exponentially (Rapkin 1997).

Dark current in a PMT arises from several sources. At room temperature electrons can be released from the photocathode and dynodes, which will be detected as a signal. Only electrons released by the photocathode and first few dynodes are a problem because they travel through the whole PMT and are amplified up to a detectable signal. A second dark current source is due to electrons colliding with residual gas in the tube causing ionisation. The ions can drift to the photocathode and dynodes, liberating electrons, which are seen as a pulse that closely follows the primary pulse. It is possible to reduce this latter type of noise by imposing an artificial dead time (20 – 80 ns) so that the secondary pulse is not recorded.

The quantum efficiency of a PMT is defined by the photocathode. The photocathode (and dynodes) is designed to have a low work function, which is a measure of how easy it is to liberate electrons when struck by a single photon (or electron). The lower the work function, the easier it is to liberate electrons. The photocathode coating can be modified to make it sensitive to different wavelengths. However, the cathode and coatings are inefficient at liberating photo-electrons and quantum efficiency is typically 25 – 30% at peak, as shown in figure 1.3.2. It can be seen that the photocathode is most sensitive at blue wavelengths, becoming significantly less sensitive towards red wavelengths.

PMTs offer negligible readout time and low noise if the voltage step between dynodes is optimised to ensure the resulting electron shower is large enough to be detected. One problem faced by PMTs is the observation of multiple targets. Because a single PMT is not an area detector (i.e. it is effectively one large pixel) it just counts photons regardless of the spatial characteristics of the source. The observer has to place an aperture (i.e. a mask with a hole in it) in front of the PMT to select the target. This means that sky background and comparison stars have to be recorded separately or with another PMT. There are other problems associated with observing with PMTs, such as coincidence losses. Coincidence

losses occur because each pulse has a finite length, so if two pulses arrive close together they will be counted as one. This is a particular problem when observing bright objects that have a high incident photon flux. Another weakness of PMTs is that they are fragile, both physically (glass tubes) and to their environment (dome lights will damage them). PMTs also require high voltages which present a danger to people working in their vicinity.

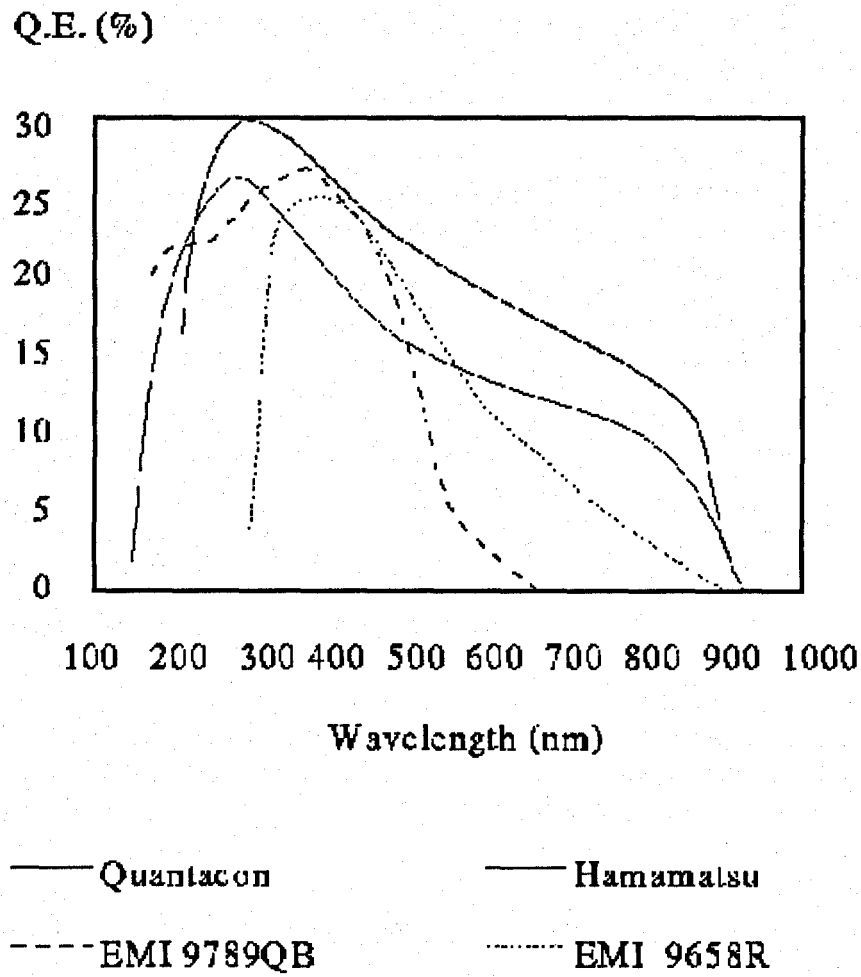

Figure 1.3.2: Quantum efficiencies of four different photomultiplier tubes (ESO 0.5 m Telescope).

1.3.2 TRIFFID

The Transputer For Fast Image Deconvolution (TRIFFID) was specifically designed to perform high time-resolution photometry. The detector used is the ESO Multi-Anode Micro-channel Array (MAMA), which is based on a Micro-Channel Plate (MCP). An MCP is effectively an array of small photomultiplier tubes, each 5 – 6 μm in diameter, forming a

plate that is approximately 0.5 mm thick (Siegmund 2000). The upper surface of the plate is coated with a photocathode material that releases an electron when it is struck by a photon. The liberated photoelectron is accelerated down the nearest tube by a large potential. As it travels through the tube it generates secondary electrons due to a coating on the tube wall. The probability of collisions is increased by curving the tube wall. Figure 1.3.3 shows a cross-section of an MCP. These devices are usually stacked to increase amplification of the detected signal, producing an emerging electron cloud of up to 10^7 electrons.

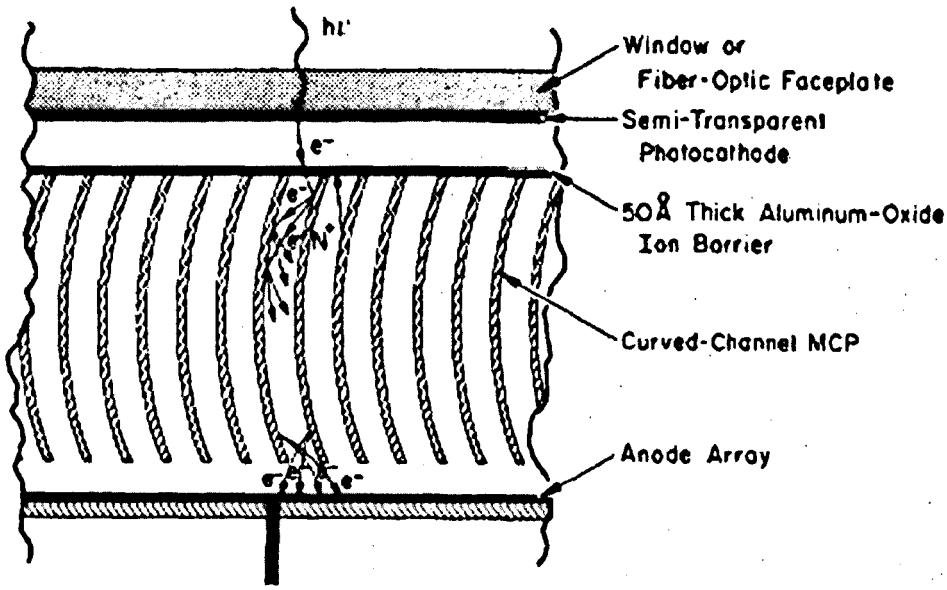

Figure 1.3.3: Cross-section of a Micro-Channel Plate (MCP). An incident photon strikes a photocathode liberating electrons which are accelerated down curved tubes by a strong electric field. The tube walls are coated so that electron collisions produce secondary electrons (Clampin 1997).

The emerging electron cloud is detected using a multi-anode array, which consists of a series of crossed electrode strips forming a grid that measures the spatial co-ordinate of the electron cloud (see figure 1.3.4). The incident electron cloud strikes both the upper and lower electrodes, which enables the spatial location of the electron cloud to be determined.

As expected, the quantum efficiency and time resolution of MAMA detectors are very similar to PMTs, given that an MCP is effectively an array of PMTs. The advantage of the MAMA detectors is that spatial imaging can be performed, potentially allowing simultaneous observation of a target, sky and comparison star. However, when TRIFFID is mounted at GHRIL (the Nasmyth port) on the William Herschel Telescope (WHT) the field of view is only 14 arcseconds (Redfern et al. 1992), meaning that it is almost

impossible to find a comparison star in the same field as the target. Besides the limitation on simultaneously imaging more than one object, MAMA detectors also suffer from the same disadvantages as PMTs, i.e. low quantum efficiency, high voltages and coincidence losses. An added problem for TRIFFID is that the throughput at GHRIL is poor compared to the Cassegrain focus due to the tertiary mirror and de-rotation optics required to feed the light to the Nasmyth focus.

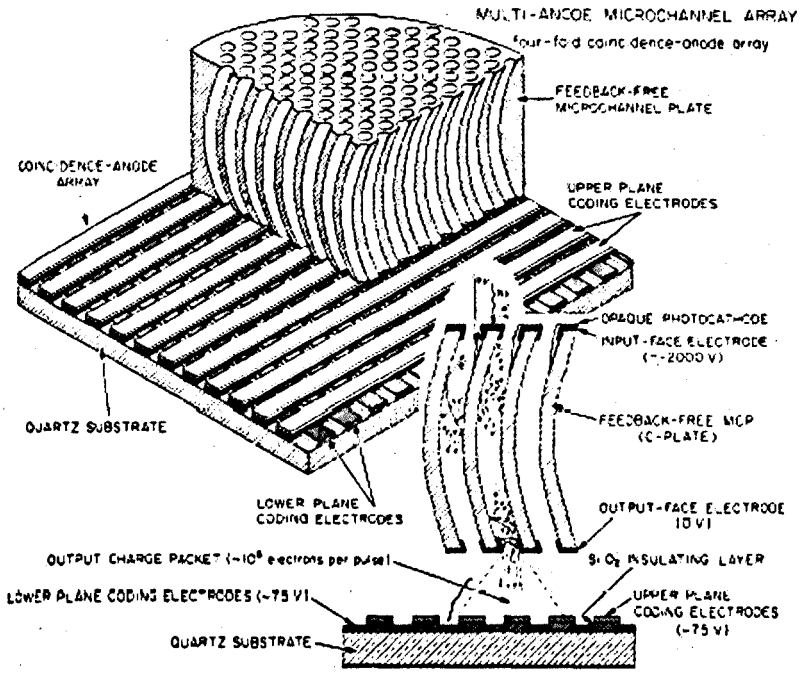

Figure 1.3.4: Diagram of a Multi-Anode Micro-channel Array (MAMA). The MCP stack sits on top of the multi-anode array. When an electron cloud emerges, it is detected by one upper and one lower electrode, allowing the position of the event to be determined (Clampin 1997).

1.3.3 OPTIMA

Optical Pulsar Timing Analyser (OPTIMA) is a high-speed photometer designed to observe rapidly rotating pulsars using Avalanche Photo-Diodes (APD). APDs are a solid-state alternative to PMTs made from semiconductor materials such as silicon or gallium arsenide. They operate using pn-junctions, which comprise of an *n*-type and *p*-type doped piece of semiconductor joined together, as shown in figure 1.3.5. The free electrons and holes in each type of the semiconductor cancel each other at the junction interface, forming a region where there are no moving charges, called the *depletion region*. The donor

and acceptor atoms are fixed in the solid and, due to their respective charges, keep the depletion region clear of any free charges. A free charge requires some extra energy to overcome the forces from the donor and acceptor atoms to be able to cross the junction, which acts like a barrier.

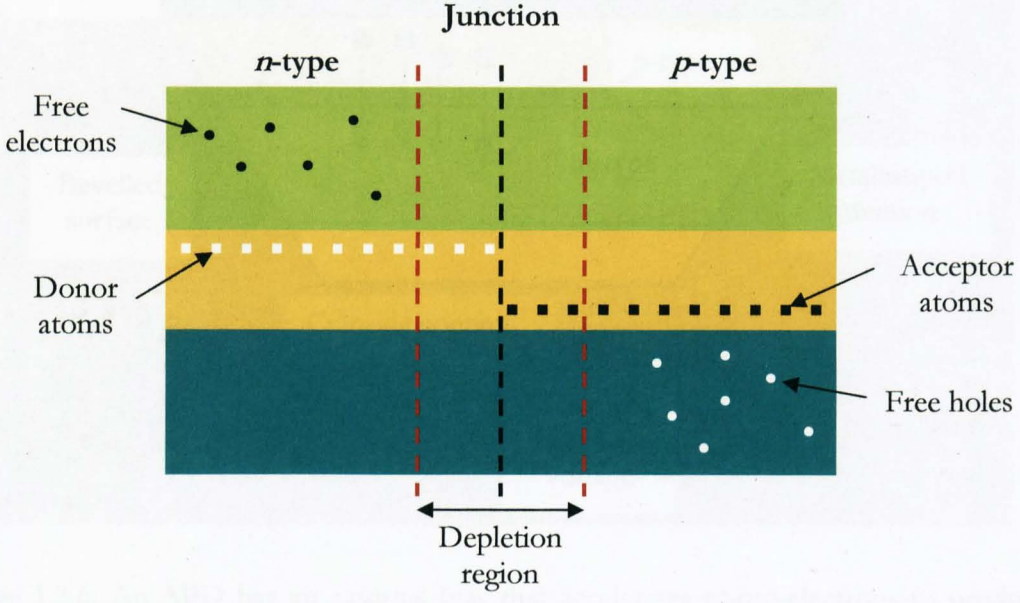

Figure 1.3.5: Schematic of a pn-junction.

The internal structure of an APD is a pn-junction. When a photon hits the diode an electron is excited and accelerated through the junction by a large external bias of approximately 2000 V (Koren 2001), as shown in figure 1.3.6. During this time the electrons collide with the lattice and release many more secondary electrons forming an avalanche of charge carriers, which are detected as a weak photo-current. The detected current is electronically amplified and is directly proportional to the light falling onto the detector. Once the avalanche has been detected, the APD is quenched to make it ready for the next detection. There are two types of quenching: passive and active. Passive quenching involves a resistor in series with the diode, which slows the recharge of a capacitor that enables the APD to operate. In this regime dead times of up to 1 microseconds are attainable (Dravins et al. 2000). Active quenching detects the rise of the avalanche and reduces the bias to avoid complete ionisation of the APD and in this regime nanosecond dead times are achievable.

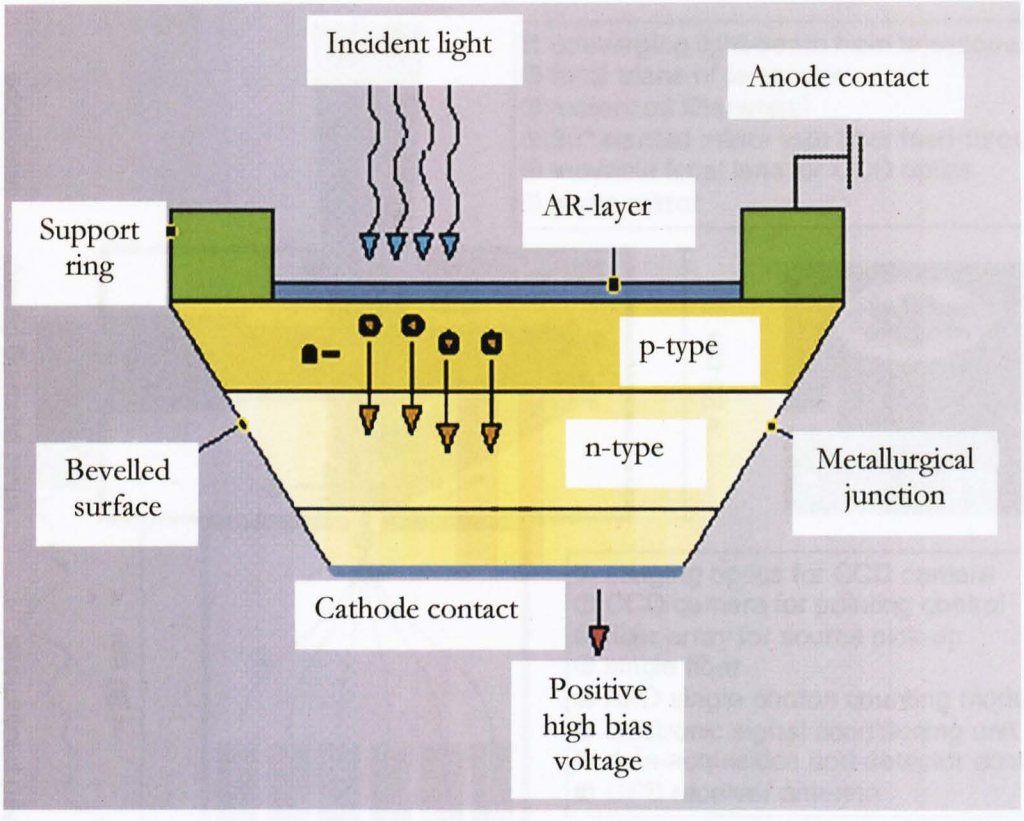

Figure 1.3.6: An APD has an external bias that accelerates photo-electrons to produce a shower of electrons at the electrode, where they can be detected (Koren 2001).

OPTIMA comprises of eight APDs to perform photometry at time resolutions of milliseconds and below (Straubmeier et al. 2001). Figure 1.3.7 shows a drawing of the instrumental layout. Incident light from the telescope falls onto an angled mirror, which directs it to a targeting CCD. In a hole through the mirror is a bundle of seven optical fibres, which feed the light to each APD (see inset of figure 1.3.7). The target star is positioned on the central fibre of the group of seven. The fibres are fixed at a size that corresponds to two seeing discs at the plate scale of the telescope, meaning that 98% of the target flux falls onto the central fibre. The surrounding six fibres collect the rest of the target flux and, along with an eighth fibre positioned further away, are used to construct a time-dependent sky background model. Each APD is cooled to a temperature of –30°C to reduce dark current. The arrival time of each photon is time-stamped using the highly accurate clock signal from GPS satellites.

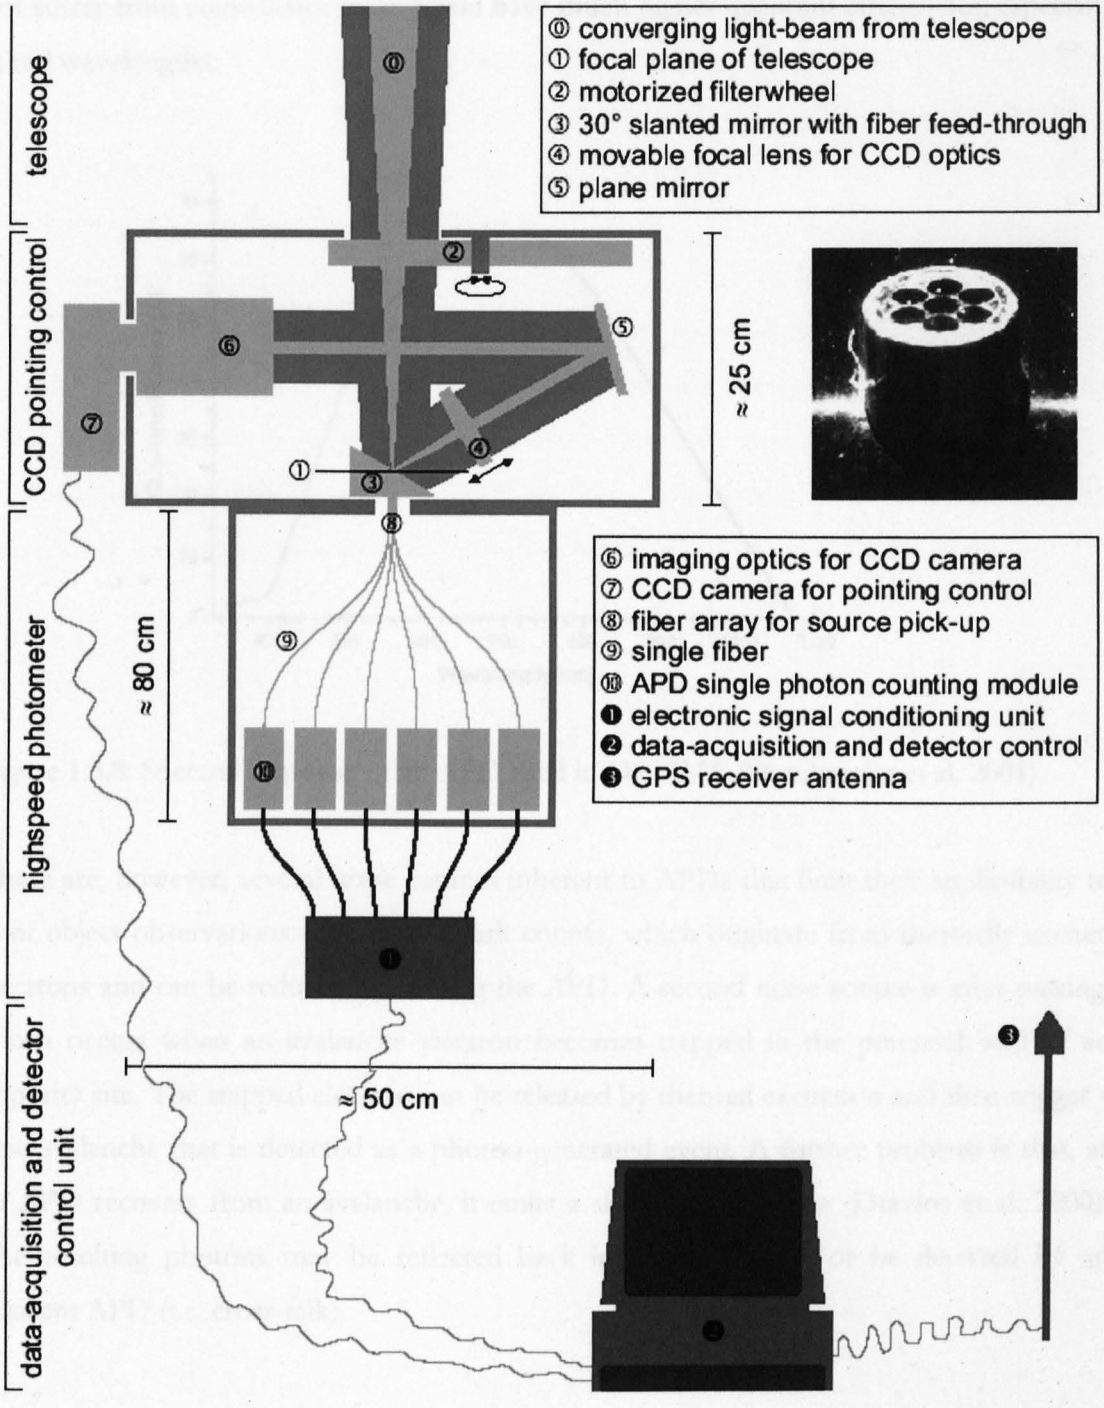

Figure 1.3.7: Sketch of the components of OPTIMA (Straubmeier et al. 2001).

APDs have similar spectral ranges and quantum efficiencies to CCDs (given they are manufactured from similar material), as shown in figure 1.3.8. The diameter of the first generation of APD detectors was approximately 0.2 mm, whereas more recently diameters of up to 16 mm have become available. As a replacement for photomultiplier tubes, APDs offer the advantage that they are much smaller, have a better linear response (i.e. they do

not suffer from coincidence losses), and have much higher quantum efficiencies, especially at red wavelengths.

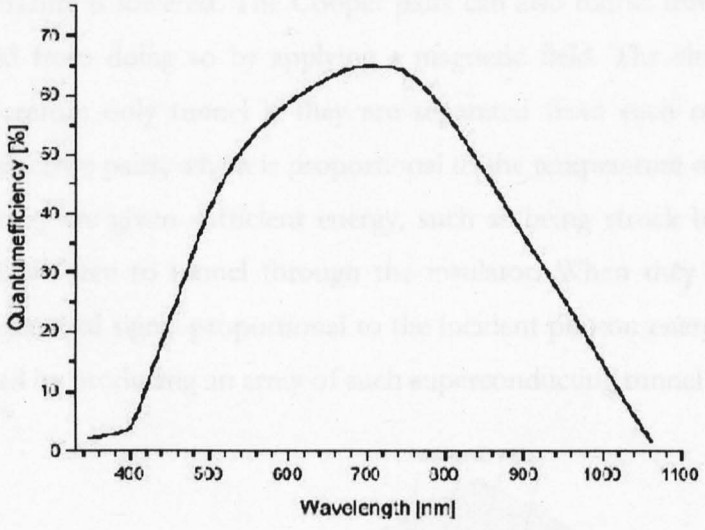

Figure 1.3.8: Spectral response of an APD used in OPTIMA (Straubmeier et al. 2001).

There are, however, several noise sources inherent to APDs that limit their applicability to faint object observations. The first is dark counts, which originate from thermally excited electrons and can be reduced by cooling the APD. A second noise source is after-pulsing, which occurs when an avalanche electron becomes trapped in the potential well of an impurity site. The trapped electron can be released by thermal excitation and then trigger a new avalanche that is detected as a photon-generated event. A further problem is that, as an APD recovers from an avalanche, it emits a shower of photons (Dravins et al. 2000). The resulting photons may be reflected back into the detector or be detected by an adjacent APD (i.e. cross-talk).

1.3.4 Superconducting Tunnel Junctions

When a superconducting material is cooled below its critical temperature, electrons form into Cooper pairs. These pairs remain together by overcoming Coulomb repulsion through interactions with the crystal lattice. When in this state, the pairs have a much lower energy which inhibits collisions and consequently reduces the electrical resistivity, i.e. the material becomes superconducting.

The Cooper pairs can be used to form a junction called a Superconducting Tunnel Junction (STJ, or Josephson junction). If an insulating layer separates two pieces of superconducting material, the electrons will quantum tunnel from one piece of superconductor to the other as the temperature is lowered. The Cooper pairs can also tunnel through the insulator but are prevented from doing so by applying a magnetic field. The electrons in the Cooper pairs can therefore only tunnel if they are separated from each other. There is energy binding the electron pairs, which is proportional to the temperature of the superconducting material. If they are given sufficient energy, such as being struck by a photon, the pairs separate and are free to tunnel through the insulator. When they do this they create a measurable electrical signal proportional to the incident photon energy. A photon detector can be formed by producing an array of such superconducting tunnel junctions.

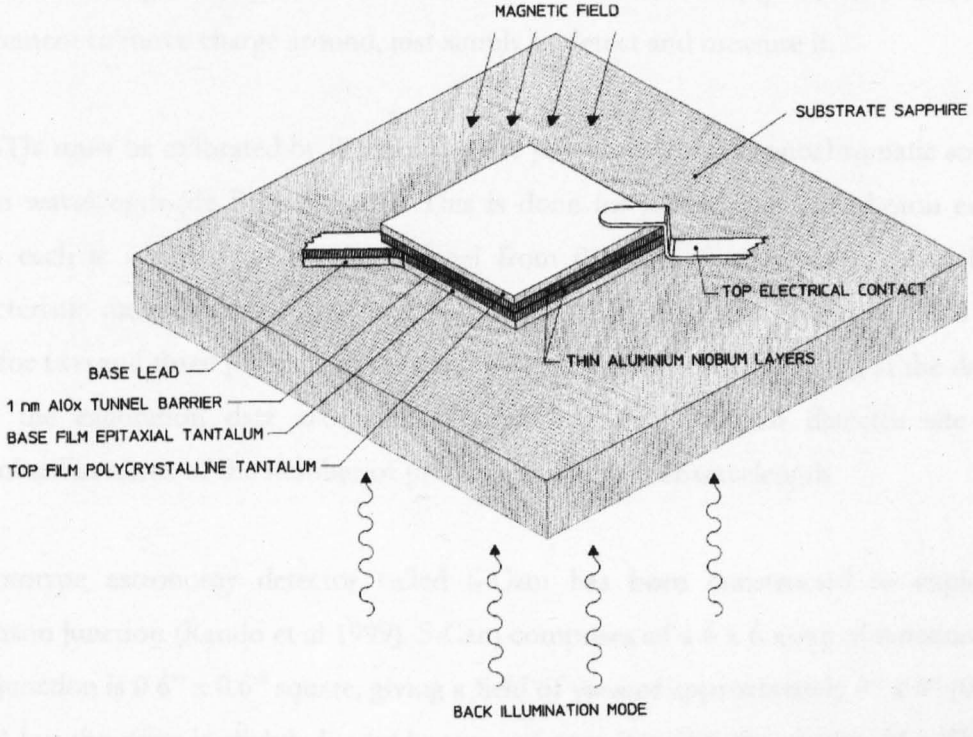

Figure 1.3.9: Schematic of one Superconducting Tunnel Junction (Peacock et al, 1998).

Figure 1.3.9 shows a schematic of one STJ. Each junction is made from the superconducting material, tantalum (Peacock et al, 1998). Two layers of tantalum are deposited onto a sapphire substrate with a layer of insulation between them (see figure 1.3.9). Two leads are then attached to form electrical connectors. A niobium layer is introduced to prevent the broken Cooper pairs from travelling down the leads and being lost and a magnetic field is applied across the STJ to suppress the Josephson supercurrent,

which is due to Cooper pairs tunnelling through the insulation layer. The light is incident onto the back of the junction so that none of the photons are blocked by the top lead.

The energy gap of the Cooper pairs (i.e. the energy required to split up the pair) is approximately three orders of magnitude lower than in semiconductor CCD detectors. A single detected optical photon can release the equivalent of 1000 electrons in an STJ (by breaking 500 Cooper pairs), depending on its initial energy. Measuring the number of electrons then determines the energy of the incident photon. This means that there is no requirement for filters or multi-channel devices because the STJ can determine the wavelength of all incident photons. The spectral range is dependent on the superconducting material of the STJ and ranges from X-ray to near-UV and near-UV to near-IR wavelengths. High-time resolution is achievable simply because there is no requirement to move charge around, just simply to detect and measure it.

The STJs must be calibrated by illuminating the junctions with a monochromatic source of known wavelength (de Bruijne 2001). This is done for several incident photon energies, where each is assigned an energy channel from 0 – 255. Each energy channel has a characteristic mean charge output and Full Width Half Maximum (FWHM). This is also done for two and three photons of the same energy simultaneously arriving at the detector. Using the calibration data allows the energy deposited in each detector site to be determined in terms of the number of photons at a particular wavelength.

A prototype astronomy detector called S-Cam has been constructed to exploit the Josephson junction (Rando et al 1999). S-Cam comprises of a 6 x 6 array of tantalum STJs. Each junction is 0.6" x 0.6" square, giving a field of view of approximately 4" x 4" (0.6" x 6 = 3.6" but the array is slightly bigger because of gaps between the pixels). The STJ array operates in a vacuum at a temperature of 0.3 K, cooled by liquid helium. The quantum efficiency is 60 – 70% over the 310 – 720 nm wavelength range and the time resolution is approximately 5 microseconds. S-Cam2 is an enhanced version of S-Cam with the same array size, but improved wavelength resolution (de Bruijne 2001). The wavelength resolution of S-Cam2 is $\lambda/\Delta\lambda = 20$ at $\lambda = 300$ nm.

The development of S-Cam is clearly a major advance in observational astronomy, and with it new data reduction techniques are required, such as energy calibration. Currently

STJ technology is still in its primary stages and array sizes are very small making observations difficult. The detector itself is also quite bulky due to the refrigeration required and needs a large amount of support (both physically, in terms of the mechanical structure required, and operationally, in terms of the staff effort required). A further disadvantage is that STJs suffer coincidence losses, similar to PMTs. If the incident photon flux is high, too many photons incident on each pixel will not be resolved, i.e. there will be so much energy deposited in each pixel that it will be impossible to determine the number and wavelength of the incident photons. Larger area arrays with improved energy resolution are being developed, as shown in figure 1.3.10, but they have yet to be tested on a telescope.

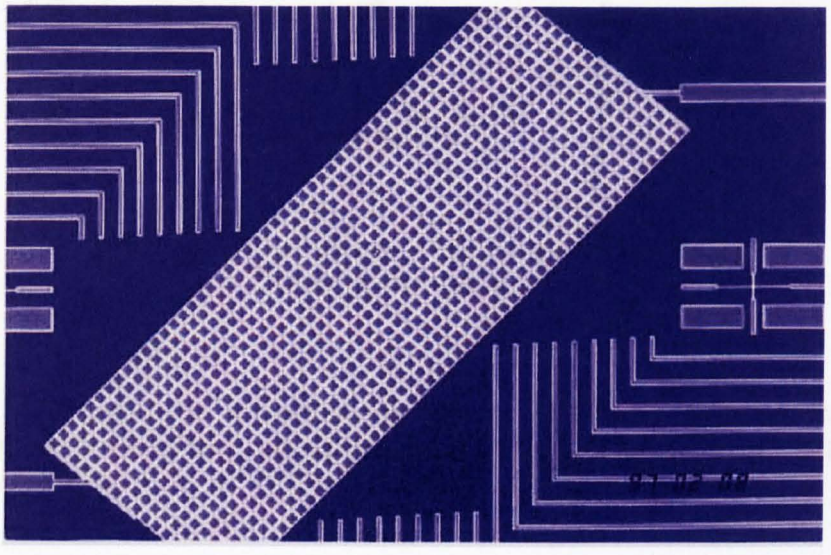

Figure 1.3.10: Prototype 18 x 50 element STJ array.

1.3.5 Transition-Edge Sensors

Transition-Edge Sensors (TES) are another type of superconducting device capable of accurately measuring the energy and arrival time of a single photon (Cabrera 1999). Thin films of tungsten deposited on silicon or germanium substrates exhibit a sharp superconducting to non-superconducting transition in the region between 70 mK to 130 mK, as shown in figure 1.3.11. This transition is exploited to form a TES, by manufacturing the tungsten films and substrates into small pixels. Applying a voltage bias to the pixel sets up a stable equilibrium between energy dissipated in the TES (i.e. current

passing through it) and heat lost to the substrate. If the temperature is raised in the pixel by the impact of an incident photon, then this causes the resistance to increase (see figure 1.3.11) and is measured as a reduction in the current flow. This is used to detect a photon, where the amount of resistance is proportional to the incident photon energy. The heated TES pixel is then cooled after the detection event by heat sinks.

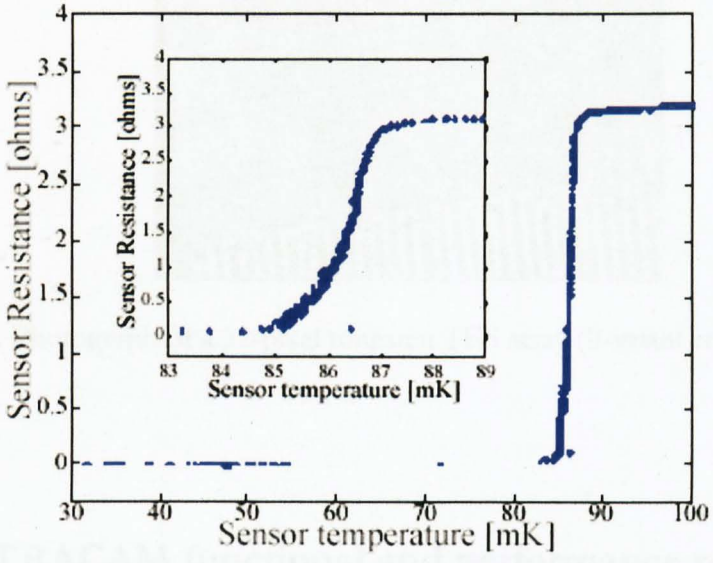

Figure 1.3.11: Resistance as a function temperature for a TES. There is a sharp transition between the superconducting and non-superconducting state at approximately 86 mK (see inset) (Cabrera 1999).

An optical imaging TES detector has been developed by Romani et al (1999). The detector is a 32-pixel array of 18 μm square, 40 nm thick tungsten pixels cooled to the superconducting transition temperature. The TES array is shown in figure 1.3.12. The wavelength range of the detector is near-UV to near-IR, with a time resolution of 0.1 μs and an energy resolution of 0.15 eV. Similar problems to the STJ exist in determining whether multiple low energy photons or one high energy photon has been detected, so a similar energy calibration technique to that used in STJs is employed to discriminate between these events.

As with STJs the pixel arrays in TES detectors are still relatively small, although, large format arrays (1024 pixels) are being developed (SAFIRE II (Benford et al. 2002) and MIRCam (Mather et al. 1999) to name but a few). A disadvantage with TES detectors is that they suffer from coincidence losses. If the detector is exposed to excessive light, then

the heat will drive the superconductor into its non-superconducting state, and no further incident photons will be detected in the time it takes to cool the detector to below the transition temperature.

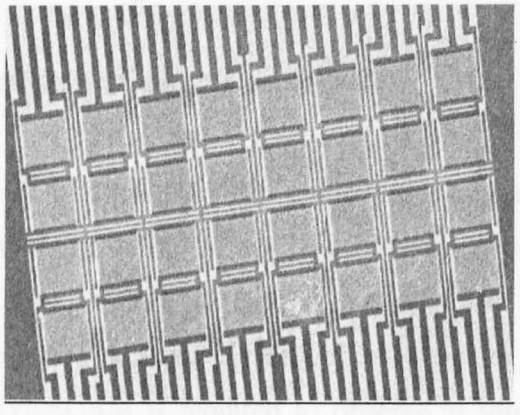

Figure 1.3.12: Photograph of a 32-pixel tungsten TES array (Romani et al. 1999).

1.4 ULTRACAM functional and performance requirements

The instruments described in section 1.3 have advantages and disadvantages. The ideal high-speed detector would only have their advantages, so we have specified ULTRACAM by adopting these as the basis of the functional and performance requirements listed below.

1.4.1 Short exposure times and negligible dead-times

The fastest variability likely to be observed in an astrophysical environment corresponds to the dynamical timescales of compact objects, i.e. white dwarfs (t_D ~ seconds), neutron stars/black holes (t_D ~ milliseconds). This sets the minimum required exposure times for ULTRACAM (~ 1 millisecond), with negligible dead-time (i.e. << 1 millisecond) between them.

Dead time is the amount of time that a detector cannot gather data due to a process in the data collection sequence, such as data archiving or signal digitisation. ULTRACAM is required to be a 'detector limited' system to keep dead-time to a minimum. In a detector

limited system, the dead time only exists because of processes intrinsic to the detector, such as clocking and digitisation (see section 1.5.1), i.e. transfer and archiving of the data occurs quickly enough to not affect the frame rate.

1.4.2 Multi-colour observations (3 or more)

Normally spectroscopic data is used to determine the temperature of an object The variability in the objects that ULTRACAM will observe are faint and spectroscopy will not be possible. A suitable alternative is to observe an object in at least three different wavelength bands covering as wide a part of the optical spectrum as possible to determine the stellar spectrum. ULTRACAM is required to observe in three colour bands ranging from 300 – 1100 nm. One channel will be optimised for the far blue (300 – 400 nm), since flickering and oscillations in many binaries exhibit features at these wavelengths (e.g. Marsh & Horne 1998).

1.4.3 Simultaneous measurement of each wavelength band

In order to obtain multi-colour photometry, single channel detectors must perform regular filter changes. This means that different colour-band observations are taken at different times. This makes comparison of the different colour-band data difficult because the characteristics of the observed variability may change between filter changes. This is also an inefficient method of observing compared to simultaneous multi-colour imaging. ULTRACAM is required to observe simultaneously with three colour bands.

1.4.4 Imaging capability

An imaging capability is required to improve observing efficiency. Observing variable stars requires not only data from the variable, but also sky background and comparison stars. Without an area detector this involves moving the detector away from the target to a blank part of the sky and a nearby star. Also, for non-imaging detectors (PMTs for example) an aperture has to be placed in front of the detector. The aperture is over-sized to ensure that all the star-light enters the detector, but this also allows excess sky to enter. An imaging

detector allows simultaneous observation of the sky and the reduction process allows a more accurately sized aperture to be fitted around the target. The final advantage of area detectors is that they can be used in poor observing conditions, when simultaneous observing a comparison star allows transparency variations to be corrected for.

To optimally resolve an object, the pixel scale must be half the seeing value. The median seeing value at the WHT has been measured as 0.69 arcseconds (Wilson et al. 1999), so the ULTRACAM pixel scale has been set to 0.3 arcseconds per pixel. A detector of 1024 x 1024 pixels will then give a field of view of 5 arcminutes, which is sufficient to give a 96% chance of finding a comparison star of R=13 magnitude (which is at the bright end of the magnitude range of typical ULTRACAM targets) at a galactic latitude of 10 to 90 degrees (Simons 1995).

1.4.5 High efficiency and portability

Observing faint objects at high temporal resolutions requires an efficient instrument. The detectors need to have low noise and high quantum efficiency, the throughput of the instrument needs to be as high as possible and large telescopes need to be used. More quantitatively, we might wish to observe a ~10 s quasi-periodic oscillation (QPO) of amplitude ~30% in a cataclysmic variable. In this case we should aim for a signal-to-noise of at least 10. If we are observing with a CCD camera at the Cassegrain focus of the WHT, the limiting magnitude in a 1 s exposure would be V = 20, assuming a new moon, a detector quantum efficiency of 90%, readout noise (see section 1.5.2) of 4 electrons and an instrumental throughput of 50%[1].

This limiting magnitude corresponds to the observed magnitudes of some of the faintest known cataclysmic variables (e.g. OU Vir; Feline et al. 2004), and therefore sets our minimum throughput, quantum efficiency and noise requirements. Note that we have assumed here that dark current is negligible, which will be the case if it is at least a factor of ~3 less than the sky. Taking the worst case of a *u'* filter observation on a 1 m telescope in dark time, the sky contributes 0.3 electrons/pixel/second. This sets our dark current

[1] Also assuming a pixel size of 0.3 arcseconds, seeing of 1 arcsecond and an airmass of 1.

requirement at 0.1 electrons/pixel/second, i.e. as long as our dark current is below this level, it will be an insignificant source of noise compared to the sky.

The objects that will be observed with ULTRACAM range in brightness from magnitude 8 – 25, making it necessary to select a telescope of appropriate aperture to suit the scientific programme. ULTRACAM is therefore specified to be portable so it can be fitted to any suitable telescope, which requires the instrument to be compact, low-mass and stand-alone (in terms of its software, computing and electronic interfaces with the telescope). Having a number of telescopes that ULTRACAM can visit has the added advantage that the instrument spends a greater fraction of its time on a telescope rather than sitting in an instrument store.

1.5 Charge Coupled Devices (CCDs)

The functional and performance requirements in section 1.4 prompted the ULTRACAM project team to adopt CCDs as the detector for ULTRACAM. This section will describe what CCDs are and their noise sources, followed by a description of how they can be used to perform high-speed observations.

1.5.1 CCD construction

Before reviewing the techniques employed to use CCDs for high-speed imaging this section will briefly outline the basics of CCD operation and construction. The basis of a CCD is a semiconductor substrate, usually silicon but other materials such as germanium can be used depending on the specific requirements of the CCD. An incident photon will, according to the photoelectric effect, be absorbed into the silicon crystalline structure and transfer its energy, elevating an electron from its normal location in the valence band into the conduction band where it is free to move throughout the crystal lattice. The energy gap between the valence and conduction band is approximately 1.14 eV (Janesick and Elliot 1992). Movement of the photo-generated electron can now controlled by an electric field. Columns called *channel stops* are defined by p-n junctions (see section 1.3.3) which confines

the electron to move only in the *y*-direction (see figure 1.5.1). A silicon-dioxide insulating layer is formed on the silicon substrate onto which metal electrodes called *gates* are added. The gates are arranged in a three-phase configuration and confine the electron in the *y*-direction, defining a location on the silicon called a *pixel*. On one of the electrodes in each gate a bias level is applied to attract all the electrons produced in that detector site (see figure 1.5.2). Figure 1.5.1 shows the architecture of a CCD.

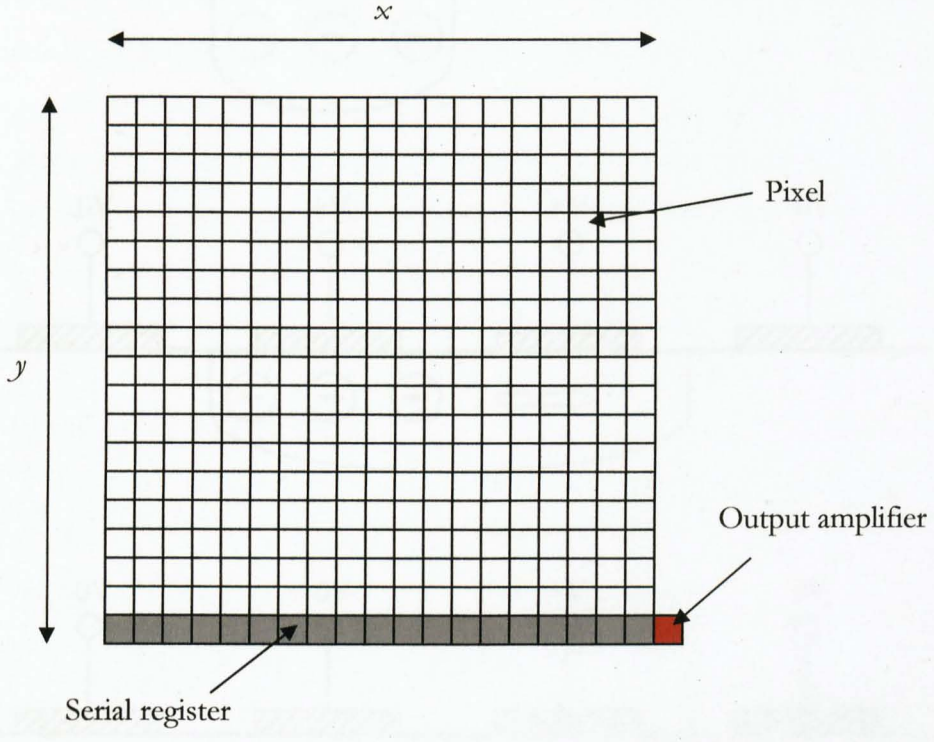

Figure 1.5.1: Schematic showing the architecture of a Charge-Coupled Device. Charge in the array of pixels is confined to move only in the *y*-direction (V-clocking), except when in the serial register. Here the charge can shift in the *x*-direction (H-clocking) to the output amplifier where it is digitised.

A CCD has an array of many pixels each capable of accumulating photo-generated electrons proportional to the number of incident photons. The total amount of charge collected in each pixel can be translated into an image of what photons the array has been exposed to. But to view this image the charge at each pixel must be measured. Movement of the charge in each pixel is done by three gates, as shown in figure 1.5.2. Sequentially raising and lowering voltages on adjacent gates shifts the charge in the *y*-direction and is called *V-clocking*. Charge is V-clocked until it reaches the bottom of the array and a row called the *serial register* (see figure 1.5.1). It is then shifted in the *x*-direction, called *H-clocking*, to the *output amplifier*, where it is amplified and digitised by an analogue-to-digital-converter

(A/D converter) using *correlated double sampling*. This is done for all the pixels in the CCD array by V-clocking each row in turn onto the serial register and H-clocking each pixel onto the output register.

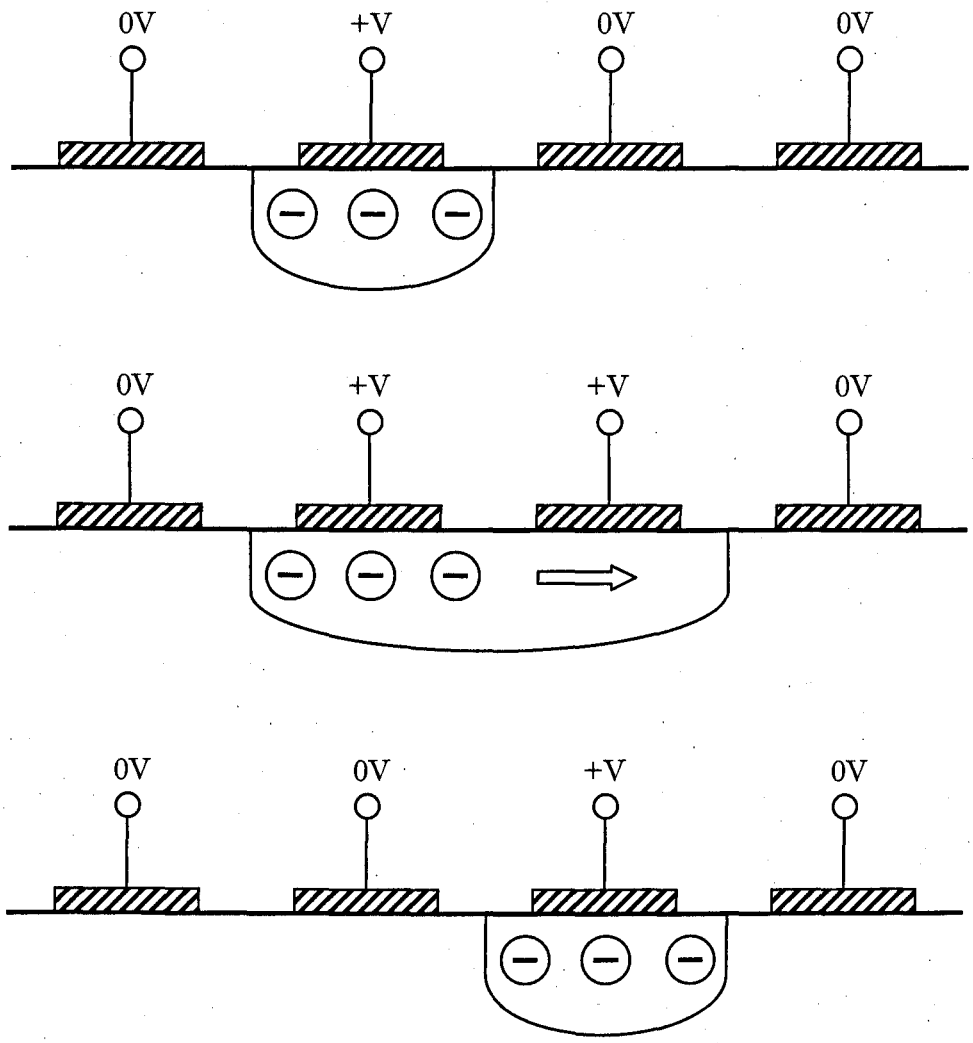

Figure 1.5.2: Schematic of the gate structure used to confine and shift photo-generated electrons through sequential raising and lowering of an electric field.

Correlated Double Sampling (CDS) is the method by which the charge in each pixel is measured. In early CCDs, the signal electrons in a pixel were transferred into the output register (a capacitor) giving the total charge as a combination of the reference level of the capacitor and the pixel charge. This is amplified and the pixel charge recorded. Before the next pixel can be measured the output register is reset, but this operation introduces noise called *reset noise*. This noise results in a random variation of the reset level from pixel to pixel. CDS reduces this noise by measuring the reference charge value (C_{ref}) and the pixel charge plus reference charge value ($C_{ref} + C_{pixel}$) and taking the difference between the two ($C_{ref} - (C_{ref} + C_{pixel})$). This is done several times to obtain the average reference and pixel

charge values and hence calculate the charge in the pixel charge. This reduces the reset noise because the reference value is not assumed to be the same level as in the earlier digitisation method and it is averaged several times to obtain good statistics. This digitisation process for each pixel takes a finite time and is dependent on the desired readout noise performance. The more time taken, i.e. the more times the two levels are compared, the better the noise performance (McClean 1997).

The frame rate of a CCD is limited by the time it takes to shift all pixels to the output amplifier (clocking) and the time it takes to digitise the photo-generated charge in each pixel. To increase the frame rate of a CCD the digitisation speed can be increased, but this results in increased readout noise. There is also a limit to increasing the clocking speed defined by the *charge transfer efficiency* (CTE) of the CCD. The CTE is a measure of how well charge in each pixel is transferred as it is clocked. In a typical CCD, 99% of the original charge in a pixel will reach the output amplifier where the remaining 1% is spread over the pixels along the charge transfer path.

1.5.2 Noise sources in CCDs

In the description of CCDs in section 1.5.1, we have touched on the subject of noise sources, particularly in the digitisation process. There are other significant noise sources in CCDs which will be described in this section. The main contributors to noise in CCDs are (Gilliland, 1992):

- **Poisson noise.** Light from a non-varying source will arrive in a non-regular way. That is to say, a non-varying star with a mean number of detected photons, N, will have a noise level, $N^{1/2}$, associated due to the Poisson nature of light. This is a fundamental limit and is irreducible.
- **Read noise.** This results from converting the analogue signal in a pixel to a digitised signal (see section 1.5.1 for one contribution to this noise – the reset noise). Read noise is inherent to each CCD and can range from typically 2 – 3 electrons to tens of electron depending on the quality of the CCD and the speed of digitisation.

- **Dark noise.** This arises due to thermally generated electrons in the silicon of the CCD and is proportional to the exposure time. Without cooling, a CCD can saturate purely from thermally generated electrons. Advanced Inverted Mode Operation (AIMO, also referred to as Multi-Phase Pinned - MPP) has been developed to dramatically reduce dark noise. There are three main sources of dark noise: thermal generation and diffusion in the neutral bulk, thermal generation in the depletion region (see figure 1.3.5) and thermal generation from surface states in the silicon-silicon dioxide interface (Janesick and Elliot 1992). It is the latter of these sources, the *surface dark current*, that contributes the most to dark current. At the interface, dark current generation is dependent on the density of interface states and free carriers (holes and electrons) that populate it. Electrons will hop (thermally) from the valence band to an interface state and then to the conduction band, producing a dark electron-hole pair. Normal CCD operation depletes the interface of free carriers and therefore maximise surface dark current generation. AIMO reduces this by filling the silicon-silicon dioxide interface with holes, suppressing hopping and therefore reducing surface dark current.

1.5.3 High Time Resolution Imaging with CCD detectors

Now that the basic construction and operation of a CCD has been explained the rest of this section will discuss high-speed imaging with CCDs. A typical CCD can be modified to image rapidly varying optical sources with low dead-time. The first example uses the CCD array as a storage area whilst a small section is dedicated to imaging. This example is called **Drift mode** and was implemented on the ISIS spectrograph of the WHT (Rutten et al. 1997; note that this mode is no longer available). Data is accumulated on a small section of the chip and is V-clocked towards the serial register whilst the shutter remains open. A stack of data windows is formed on the chip whilst an imaging window is exposed further up the array and a window on the serial register is read out. This mode is used for both imaging and spectroscopy and can achieve frames rates of 4 Hz (Rutten et al 1997). The total number of exposures that drift mode can gather is limited to the detector memory (16Mb), so after a period of time observing has to stop so that the data can be archived. A very similar mode of operation is **Time series mode** used at the Anglo-Australian Observatory (Stathakis & Johnston 2002), where the CCD is continuously exposed and

clocked after a set amount of time. This mode is primarily designed for spectroscopy and can achieve frame rates of 100 Hz.

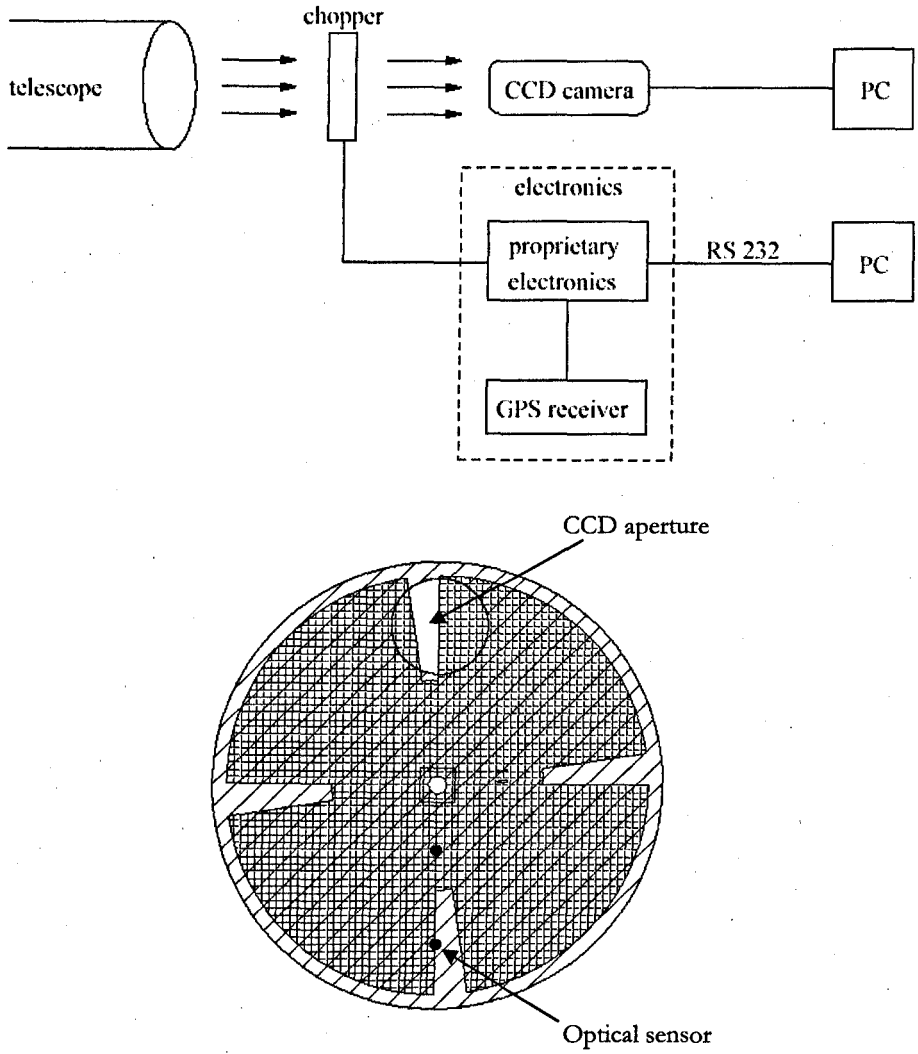

Figure 1.5.3: Schematic of the stroboscopic system with a drawing of the chopper (Kotar et al 2003). The CCD aperture allows the light from the pulsar onto the CCD when the slit is positioned over it. The optical sensor monitors the revolution rate of the chopper.

In **stroboscopic mode** (Čadež et al. 2001, Kotar et al. 2003), used at several different telescopes, a CCD is operated conventionally to observe pulsars. Temporal resolution is achieved using a specially built rotating shutter with slits to allow light to fall onto the CCD. The slits are precisely separated and the shutter is rotated at a controlled rate so each slit passes over the pulsar at the same pulsar phase. Figure 1.5.3 shows a schematic of the instrumental layout and the rotating shutter (known as the 'chopper'). There are four slits in the chopper and the widths correspond to the width of the pulsar main pulse. The CCD is exposed as the shutter rotates and the faint signal from the pulsar peak accumulates. One exposure can contain several pulsar periods stacked on top of each other to give improved

signal to noise (because very little sky accumulates). The disadvantage with this method is that it relies on knowing the precise pulsar period so that the shutter is rotated at the correct rate for accurate phase binning.

High time resolution (HIT) mode has been developed for FIERO, the European Southern Observatory (ESO) new generation CCD controller (Cumani & Mantel 2001). This mode is very similar to drift and time series mode. There are two methods of operating HIT mode, both of which can undertake imaging and spectroscopy. The first method called *period shift*, exposes a small slit or window on the CCD located very close to the serial register. The CCD is continuously read out during the exposure so each window is separated from the next by the time it takes to shift and read out a window. The second method, called *one shift* mode, is used when the periodicity of the variable source is known. The slit or window is positioned at the middle of the chip and the exposure begins. During the exposure the CCD is clocked so that a stack of windows forms up the chip from the centre, covering the duration of the pulse period. The shutter closes when the first window reaches the top of the chip. The whole CCD array is clocked down the chip so that the very first window is in its original position in the middle of the chip. The shutter opens and the cycle repeats until there is sufficient signal from the variable source to give good signal-to-noise. Masks are used to define the window locations and prevent charge from accumulating elsewhere on the chip due to other stars in the field. The mode can operate at a frame rate of 833 Hz, but the disadvantage is the exposure needs to stop frequently to archive the data. It also relies on knowing the period of the observed object.

Identical to one-shift mode, **Phase binning mode** employed at the Palomar Observatory Hale 5 m telescope (Kern 2002), defines a slit in the middle of the CCD. During the exposure the CCD is clocked to illuminate each window. The image data is not read out after one cycle of exposing all of the windows. Instead the CCD returns to the window position settings at the start of the run and repeats the exposure sequence until a good signal-to-noise level is attained. This mode can achieve frame rates of 1000 Hz, but has to stop to archive the data and needs to know the period of the observed object.

Freerun mode at the 4 m Mayall telescope at Kitt Peak National Observatory uses an intensifier to boost signal from a faint varying source and then channels that enhanced signal onto a frame transfer CCD (see later in this section), which is continuously read out

(Fordham et al. 2000, Fordham et al. 2002). Figure 1.5.4 shows a schematic of the instrumental set-up. This detector is designed for spectroscopy. The CCD reads out a single row of pixels every 41.4 μs. During operation, the CCD is continuously V-clocking at a rate which is limited by the readout speed of a single row. This is called *continuous readout mode.* The image projected onto the CCD is intensified by an MCP because the exposure time of the CCD is effectively 41.4 μs per row and the signal will be faint without amplification.

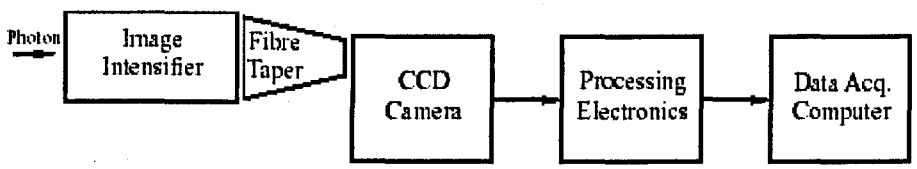

Figure 1.5.4: Schematic of the instrumental set-up of Freerun mode (Fordham el al 2000). The incident photons are first intensified by an MCP and then projected onto a frame transfer CCD.

Figure 1.5.5 shows the CCD illuminated by a pin-hole mask during normal operation and during continuous readout. During continuous readout the signal is smeared vertically down the chip. The smear of the target star changes in intensity due to variability of the source. This type of readout can achieve frame rates of 24000 Hz, but this only exposes one row of pixels at a time.

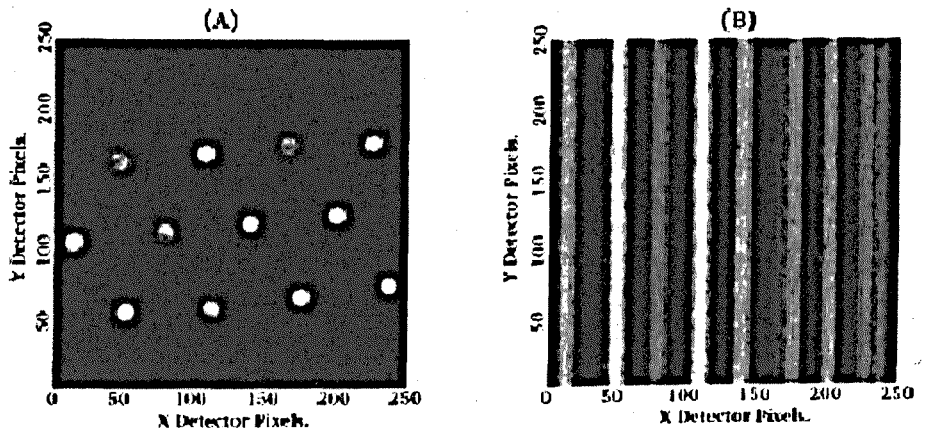

Figure 1.5.5: Illumination of the CCD using a pin hole during (A) normal CCD readout and (B) continuous CCD readout (Fordham et al. 2000).

Most of the CCDs mentioned previously have been standard arrays (with the exception of free-run mode) optimised for high speed imaging by special modifications. **Frame transfer**

CCDs are specifically designed to improve frame rates. The CCD is divided into two sections, an imaging array and a storage section. The storage section is identical to the imaging array but is masked to prevent light from falling on it. An image is accumulated on the imaging array and is then shifted or 'frame transferred' to the storage section. Once in the storage section, the image can be read out whilst a second image is being exposed. The dead-time (or time when no data can be accumulated) is the time it takes to frame transfer the image to the storage section. Figure 1.5.6 shows a schematic of a frame transfer CCD.

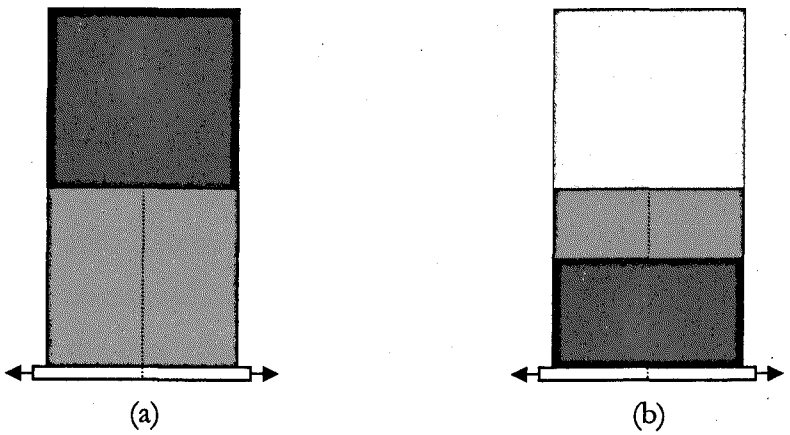

Figure 1.5.6: Schematic of a frame transfer CCD. An image is accumulated (a) and then frame transferred to a storage section that is masked to prevent illumination. A second exposure (b) can begin whilst the previous image is read out in the store section.

An example of a frame transfer CCD is the **Gemini acquisition camera**, which can achieve a frame rate of 6 Hz (Hippel & Winge 2002). The **JOSE camera** is also a frame transfer CCD 512 x 512 array, operating on the Nordic Optical Telescope (NOT) on La Palma. The CCD can achieve frame rates of 150 Hz when windowed and with binning (Baldwin et al 2001). The disadvantage with this detector is that it has to stop after typically 160 seconds to archive data.

1.6 Summary

This chapter has described the importance of high time resolution observations and the range of science that can be undertaken in this area of astrophysics. I have also described different methods of imaging at high speeds covering many different detectors, from

photomultipliers to superconducting tunnel junctions, culminating in a list of requirements for an ideal high-speed instrument. These requirements form the basis of the functional and performance requirements of ULTRACAM.

The functional and performance requirements support the use of CCDs as the detector of choice for ULTRACAM. A description of CCD manufacture and operation was then presented, followed by examples of CCDs used in high-speed observing. Nearly all of these examples tend to have one underlying limitation – they are not detector limited systems. This is one area in which ULTRACAM differs from most other CCD-based high speed imagers.

The ULTRACAM project began in July 1999 thanks to the award of £292,034 by the Particle Physics and Astronomy Research Council (PPARC). Shortly after, tender documents were distributed to procure the optical design, CCD camera design and manufacture, and the data acquisition system design and delivery. On August 2000, the Edinburgh-based United Kingdom Astronomy Technology Centre (UKATC) were finally awarded the contract.

Chapter 2

Instrument Design

2.1 Overview

The aim of this chapter is to give a detailed account of the whole ULTRACAM design. The design has been completed by several groups and I will identify my contributions throughout this section. My main role has been the design of the opto-mechanical structure and integrating the optics, detectors and electronics. However, during the course of the project I have input work into several other areas. The rest of this chapter will be divided according to the main sections in the ULTRACAM design:

i. Optics
ii. Detectors
iii. Data acquisition and camera control software
iv. Mechanical

2.1.1 Design philosophy

The first purpose of ULTRACAM is high-speed imaging. The philosophy has therefore been that the rate at which ULTRACAM can record data is limited only by the performance of the detectors. This of course assumes that the detectors alone are capable of high frame rates, i.e. it is a "detector-limited" system. Processes such as archiving, data transfer and data reduction should in no way impede imaging speed.

To utilise its primary purpose fully, ULTRACAM should be a travelling instrument so that the telescope aperture can be matched to the scientific programme being undertaken. By being able to operate on several telescopes ULTRACAM will spend less time in an instrument store, as many private instruments do, and more time on a telescope gathering data. Additional to this requirement ULTRACAM has been designed as a stand alone instrument where the only interface with an observatory is the mechanical interface between the telescope and ULTRACAM. The list of telescopes that ULTRACAM will visit includes the Aristarchos telescope (2.3 m) in Greece, the William Herschel Telescope (WHT, 4.2 m) on La Palma, and the Very Large Telescope (VLT, 8.2 m) in Chile.

2.2 Optics

2.2.1 Optical design

The optic layout has been designed by Tully Peacocke at the United Kingdom Astronomy Technology Centre (UKATC) and is divided into four main groups:

i. Collimating fore-optics
ii. Dichroic beam-splitters
iii. Re-imaging cameras
iv. Filters and CCD windows

The collimating fore-optics are positioned 317.4 mm below the focal plane of the telescope (WHT). They receive the diverging light and make it parallel. It then falls onto two dichroic beam-splitters. The dichroic beam-splitters divide the light into three broad wavelength bands, hereafter known as the "blue", "green" and "red" channels. After the dichroic beam-splitters are the re-imaging cameras which focus the light onto each CCD. A set of Sloan Digital Sky Survey filters (SDSS), *u'g'r'i'z'*, have been procured to fit between the re-imaging cameras and CCDs and there is a window immediately before each CCD. Figure 2.2.1 shows the optical layout of ULTRACAM.

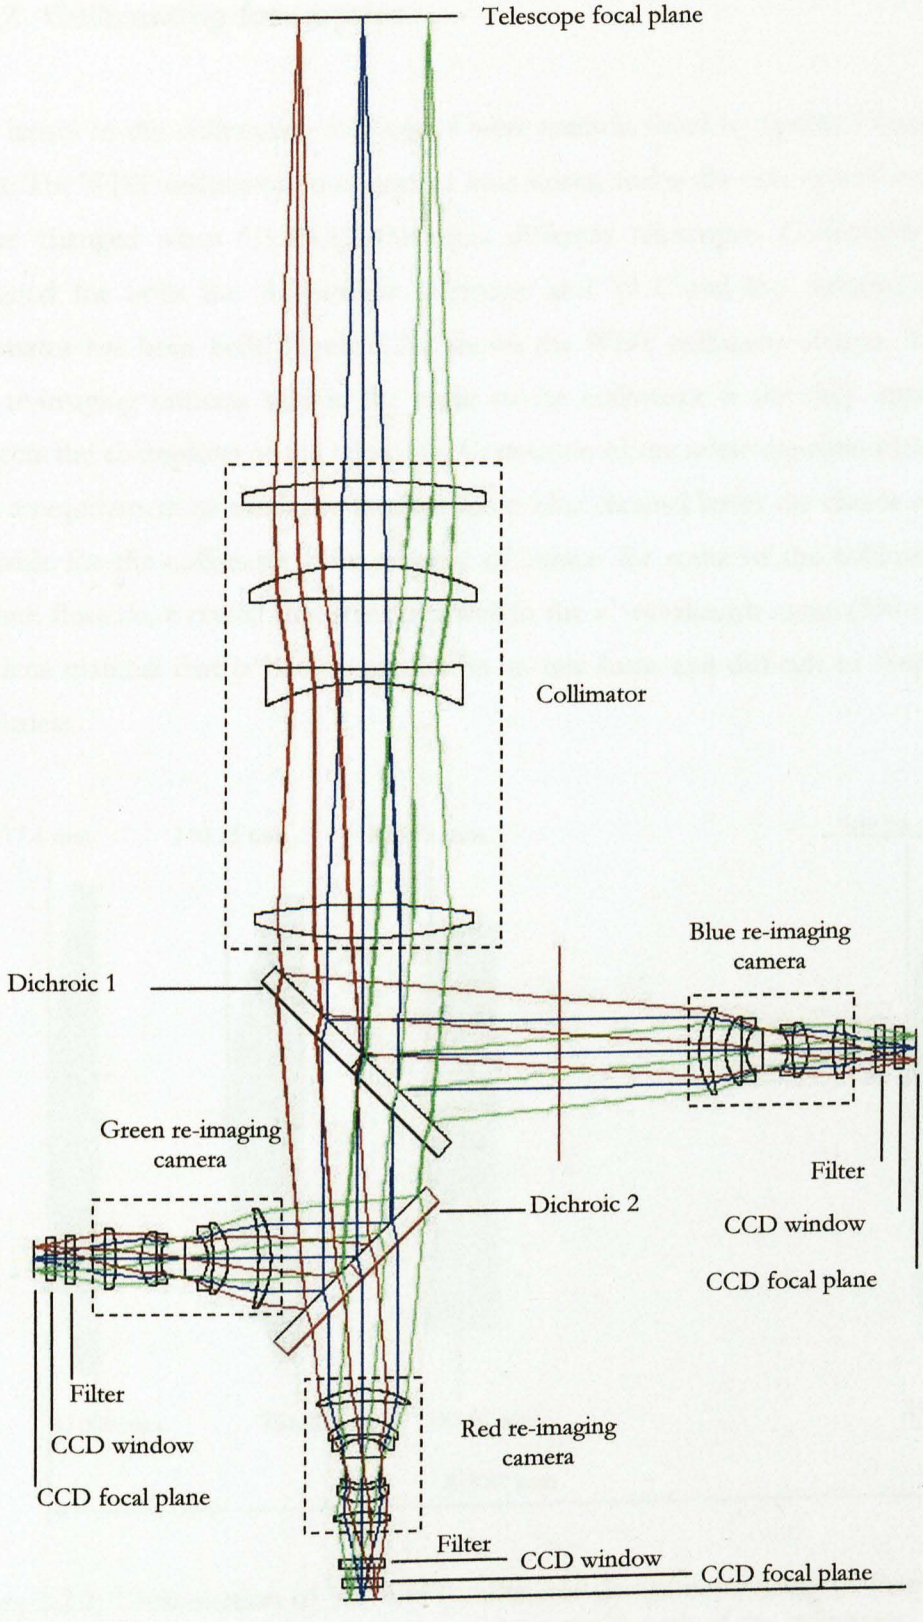

Figure 2.2.1: Optical layout of ULTRACAM. Diverging light from the telescope is collimated by the fore optics then divided into three wavelength bands by two dichroic beam-splitters. Each band of light is then re-imaged onto its respective CCD through a filter and CCD window.

2.2.2 Collimating fore-optics

The lenses in the collimating fore-optics were manufactured by Specac Limited, based in Kent. The WHT collimator comprises of four lenses, and is the only optical unit that needs to be changed when ULTRACAM visits different telescopes. Collimators have been designed for both the Aristarchos telescope and VLT and the Aristarchos telescope collimator has been built. Figure 2.2.2 shows the WHT collimator design. The dichroics and re-imaging cameras remain the same so the collimator is the only optical unit that corrects the aberrations of the telescope. Correction of the telescope aberrations combined with a requirement to maximise throughput in blue channel limits the choice of glass types available for the collimator. The material of choice for some of the collimator lenses is calcium fluoride, a crystal that transmits well in the *u'* wavelength range (300 – 400 nm). It is a lens material that is both expensive in its raw form and difficult to shape due to its brittleness.

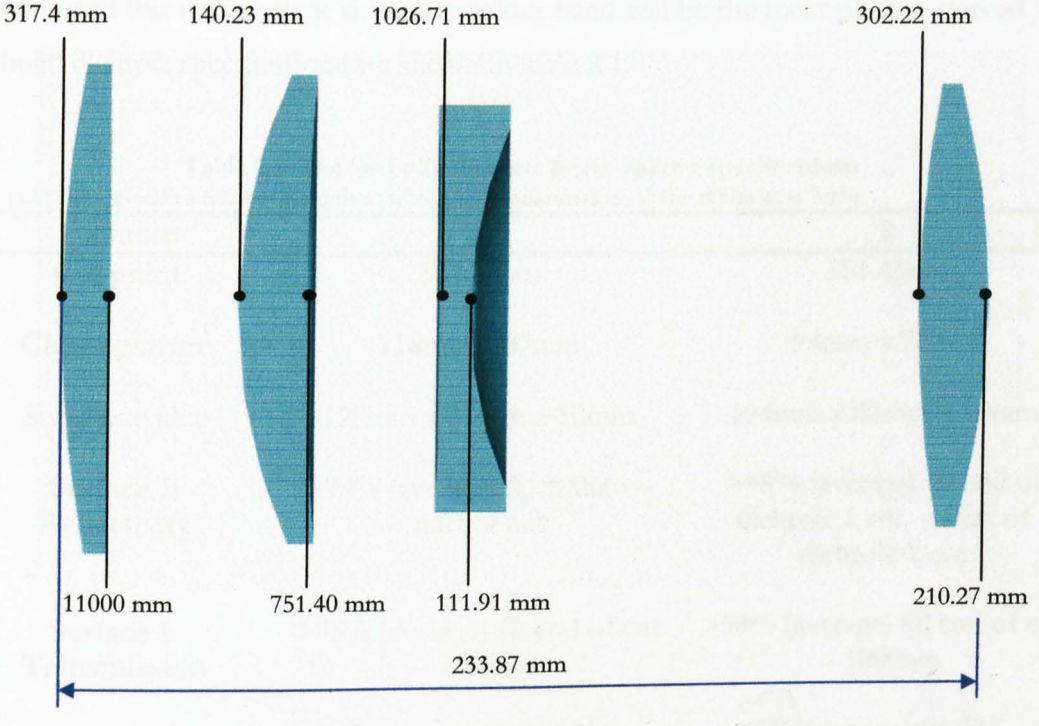

Figure 2.2.2: Cross-section of the WHT collimator design comprising of four lenses. The radius of curvature of each lens is shown and the total length of the collimator.

The field of view is determined by the collimator design and is required to be approximately 5 arcminutes based on the probability of finding a comparison star

(Simmons, 1995). To optimally resolve an object the pixel scale must be half the seeing value. The pixel scale was chosen to be half the median seeing value at the WHT in La Palma, which is typically 0.6 arcseconds (Wilson et al, 1999) so a pixel scale of 0.3 arcseconds per pixel was selected.

2.2.3 Dichroic beam-splitters

CVI Technical Optics based on the Isle of Man were selected to supply the dichroic beam-splitters. Each dichroic beam-splitter is a flat piece of UV-grade fused silica with a dielectric coating that reflects light in a particular wavelength band whilst transmitting the rest. The dichroics in ULTRACAM are long wave pass (LWP) meaning that long wavelengths are transmitted whilst short wavelengths are reflected. This is desirable for ULTRACAM because dichroic beam-splitters reflect a higher percentage of light than they transmit (CVI Laser). By reflecting short wavelengths, throughput in blue channel is kept as high as possible and this is important since this colour band will be the most photon starved. The 'as built' dichroic specifications are shown in table 2.1.

Table 2.1: The 'as built' dichroic beam-splitter specifications.
The cut point specifies the wavelength at which the transmission of the dichroic is 50%.

Dichroic	1	2
Cut point	387.79nm	551.43nm
Clear aperture	118mm x 89mm	94mm x 70mm
Substrate size	128mm x 99mm x 10mm	104mm x 80mm x 10mm
Surface 1: Reflectivity	>99% (average) @ 300nm – start of cut	>98% (average) @ end of dichroic 1 cut – start of dichroic 2 cut
Surface 1: Transmission	>95% (average) @ end of cut – 1000nm	>94% (average) @ end of cut – 1000nm
Surface 2: Reflectivity	<2% (average) @ 335 – 1000nm	<2% (average) @ 335 – 1000nm

The dichroic cut points were originally specified to fall between the *u'-g'* (dichroic 1) and *g'-r'* (dichroic 2) SDSS filter profiles. I determined the position of the cut points of dichroic 1 and 2 by calculating the wavelength at which the transmission of the *u'-g'* and *g'-r'* filters

cross-over. There is an added complication to this because the central wavelengths for the filters were specified for operation in collimated light, but in ULTRACAM, the filters are placed after the re-imaging cameras where light is converging. The central wavelength shift is described in section 2.2.5.

The cut points in the original specification were required to be as sharp as possible, lying at 388 nm and 551 nm for dichroics 1 and 2 respectively and are shown in figure 2.2.3. The dichroic coatings were required to operate in a temperature range of -10 to +30 degrees C and in a humidity range from 0 to 100%. Each dichroic substrate was specified to be very flat and smooth (form error of $\lambda/4$ @ 633 nm over the whole surface and scratch-dig 10 – 5)[3]

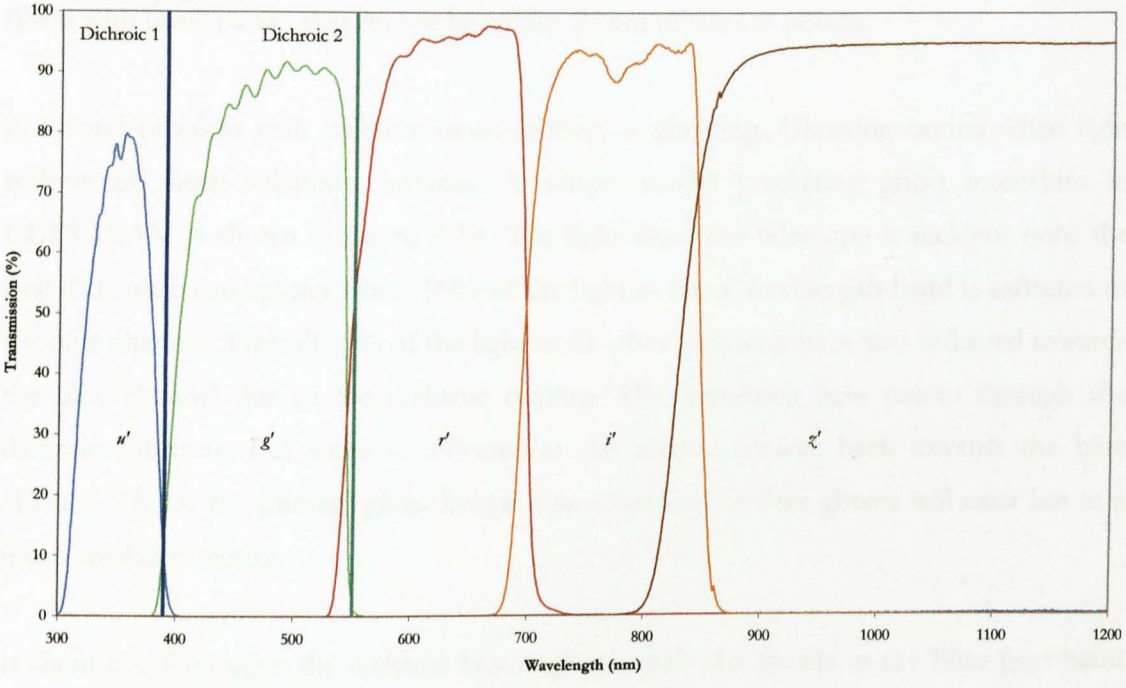

Figure 2.2.3: Theoretical dichroic cut points and SDSS filter profiles.

The overall substrate sizes were calculated by myself using a Computer Aided Design package (CAD). The optical designer provided the dimensions of the size of the optical field on the last lens of the collimator and on the first lens of each re-imaging camera. Using CAD, a scale drawing was produced that accurately represented the size of the light

[3] Scratch refers to a mark on the optical surface, where the number refers to the width of the scratch in 0.1 micron intervals. For example, a scratch of 10 is equivalent to a 1 micron scratch. Dig refers to a pit or crater on the optical surface and is a measure of the pit diameter in one-hundredth of a millimetre. For example, a dig of 5 is equivalent to a pit diameter of 0.05 mm.

beam between the collimator and the re-imaging cameras. From this the clear apertures of the dichroic beam-splitters were measured.

One problem with the dichroics is that they are designed to operate in randomly polarised light. Linearly polarised light with polarisation angles perpendicular to the fast axis of the dichroic beam splitter (i.e. parallel to the long axis) will cause the cut points to shift by up to 20 nm (Helmut Kessler – private communication). Observing variably polarised objects will cause photon flux to increase in one colour band and decrease in the other, which may be misinterpreted as variability in the total flux of the observed object. Options to eliminate this effect when observing polarised objects are to depolarise the light from the object before it enters ULTRACAM (e.g. using a rotating half wave plate) or use narrow band filters with band-passes that do not lie within 20 nm of the cut points.

A second problem with dichroic beam-splitters is ghosting. Ghosting occurs when light reflects off both substrate surfaces. A simple model predicting ghost intensities in ULTRACAM is shown in figure 2.2.4. The light from the telescope is incident onto the first dichroic beam-splitter where 99% of the light in the *u'* wavelength band is reflected to the blue channel. Typically 5% of the light at all other wavelengths is also reflected towards the blue channel due to the dichroic coating. The remaining light passes through the dichroic substrate, but some is reflected at the second surface back towards the blue channel. This is the primary ghost image. Secondary and tertiary ghosts will exist but at a much weaker intensity.

If there is a filter after the dichroic beam-splitters only the ghosts in the filter pass-band will be present. The ghost model predicts that the primary ghost intensity in the blue, green and red channels will be 0.0002%, 0.0008% and 0.098% of the total light respectively. However, the ghosts are formed in a collimated beam, meaning they fall on top of the main image and are in focus. Figure 2.2.5 shows the spot diagram of ghost images in the red CCD, the channel with the highest intensity ghost images. Ghost images at the centre of the field will not be noticed, but because of chromatic aberration ghosts at the extreme corners will be smeared into a disc of approximately 6 pixels (or 1.8 arcseconds on the WHT). However, they will be of extremely low intensity and will be undetectable.

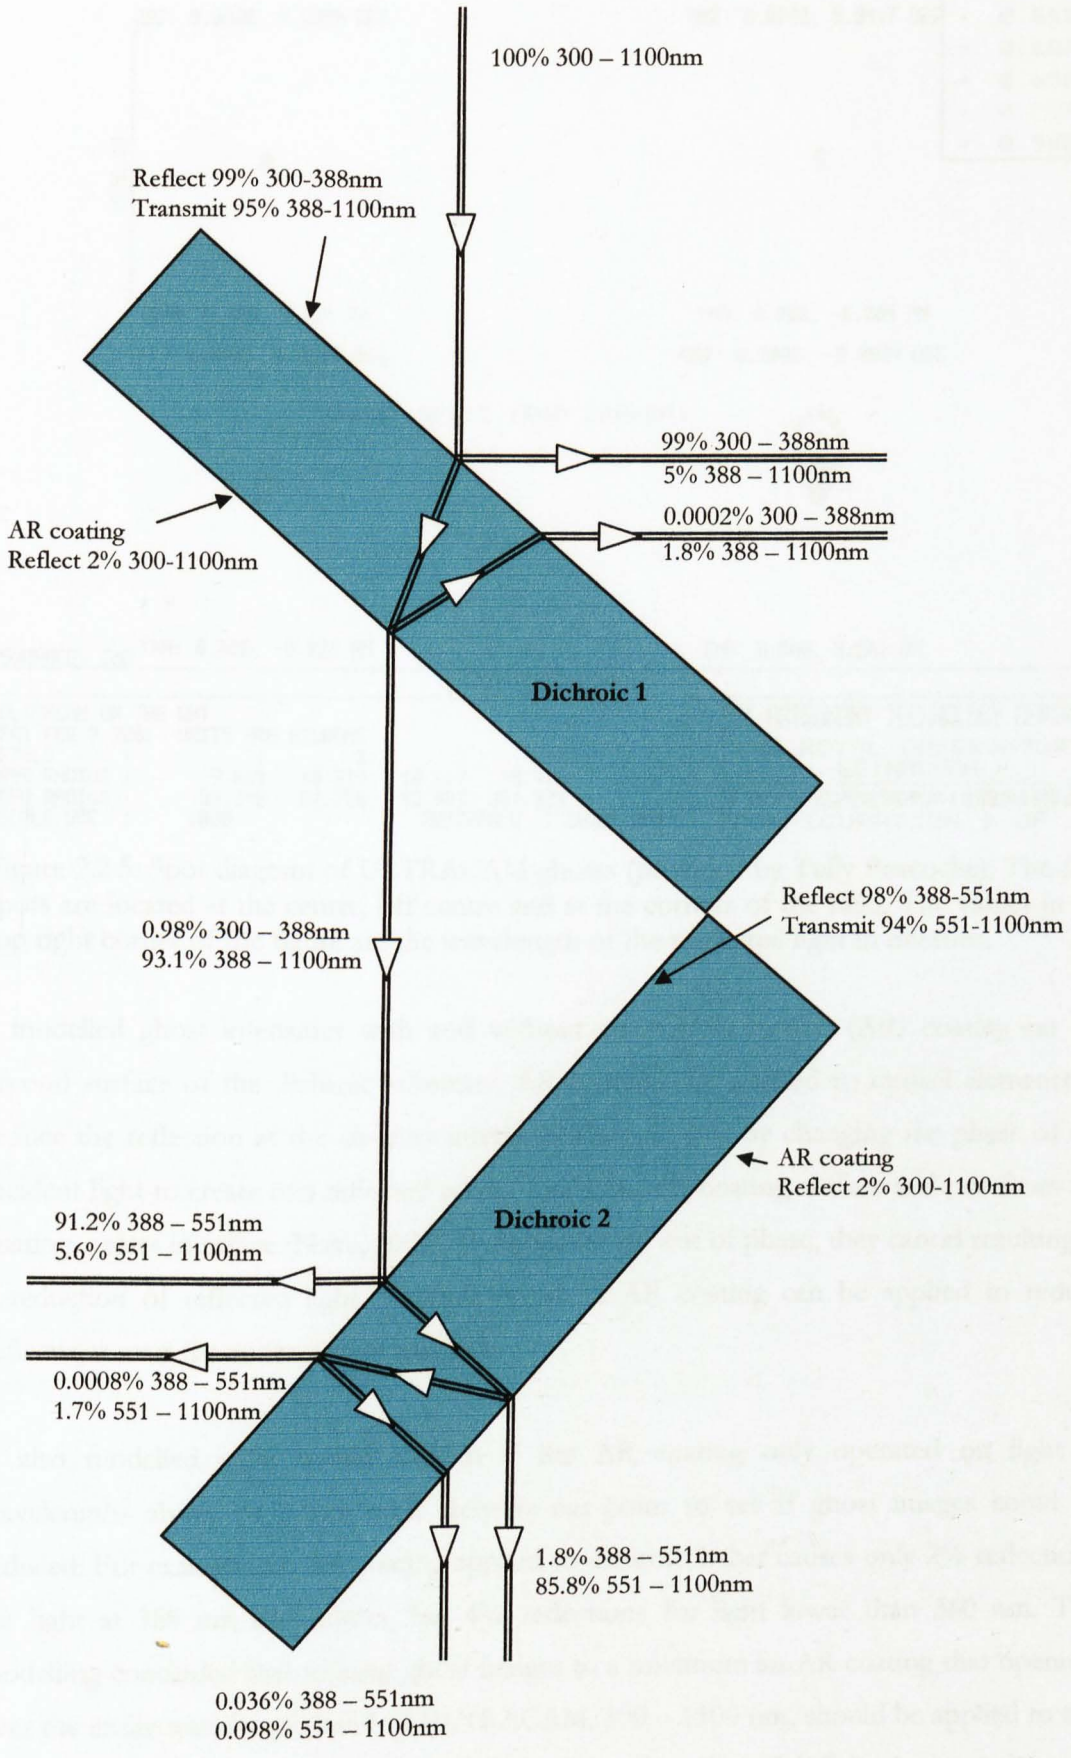

Figure 2.2.4: Simple model of ghosting caused by dichroic beam splitters in ULTRACAM.

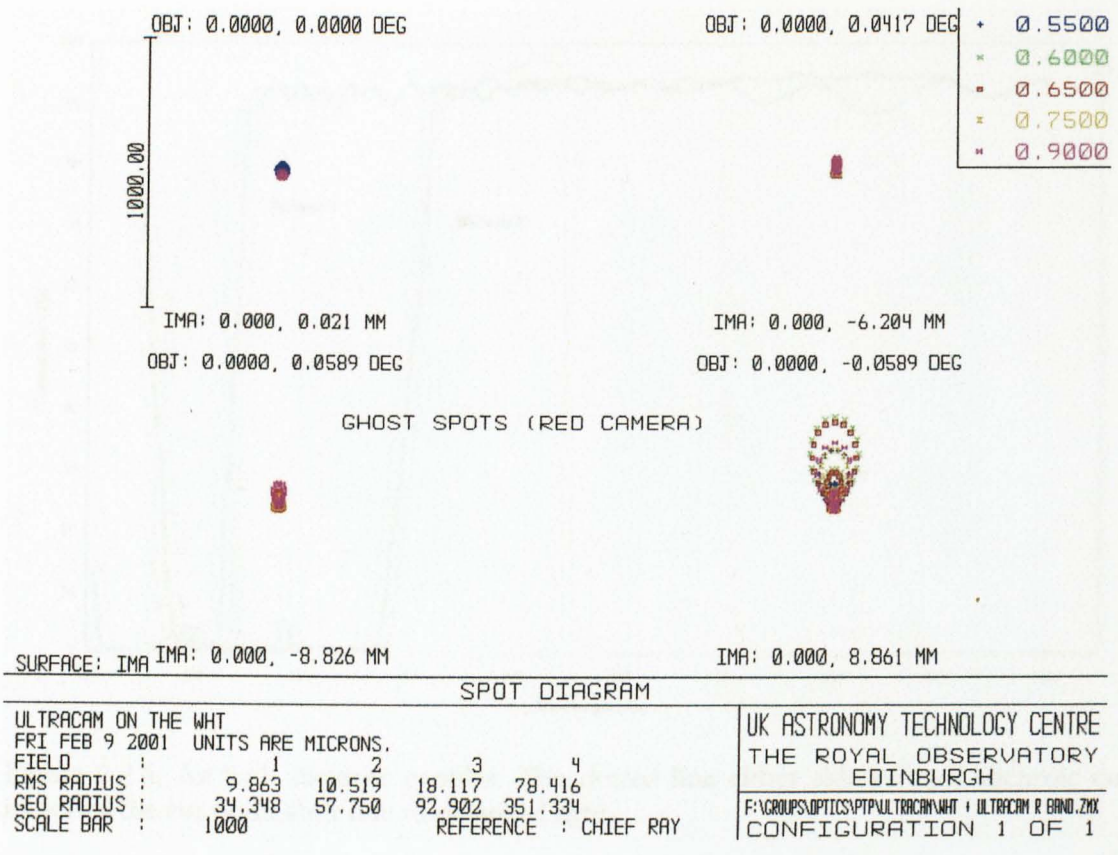

Figure 2.2.5: Spot diagram of ULTRACAM ghosts (provided by Tully Peacocke). The four spots are located at the centre, off centre and at the corners of the field. The values in the top right corner of the figure are the wavelength of the projected light in microns.

I modelled ghost intensities with and without an Anti-Reflection (AR) coating on the second surface of the dichroic substrate. AR coatings are applied to optical elements to reduce the reflection at the air-glass interface. They do this by changing the phase of the incident light to create two reflected waves, one from the coating surface and one from the coating – glass interface (Nave, 2004). If the waves are out of phase, they cancel resulting in a reduction of reflected light. Multiple layers of AR coating can be applied to reduce reflections over the entire visible spectrum.

I also modelled what would happen if the AR coating only operated on light at wavelengths above or below each dichroic cut point to see if ghost images could be reduced. For example, an AR coating applied to dichroic 1 that causes only 2% reflections for light at 388 nm and higher, but 4% reflections for light lower than 360 nm. The modelling concluded that to keep ghost images to a minimum an AR coating that operates over the entire wavelength band of ULTRACAM, 300 – 1100 nm, should be applied to the second surface of each dichroic beam-splitter. The final 'as built' dichroic cut points are shown in figure 2.2.6.

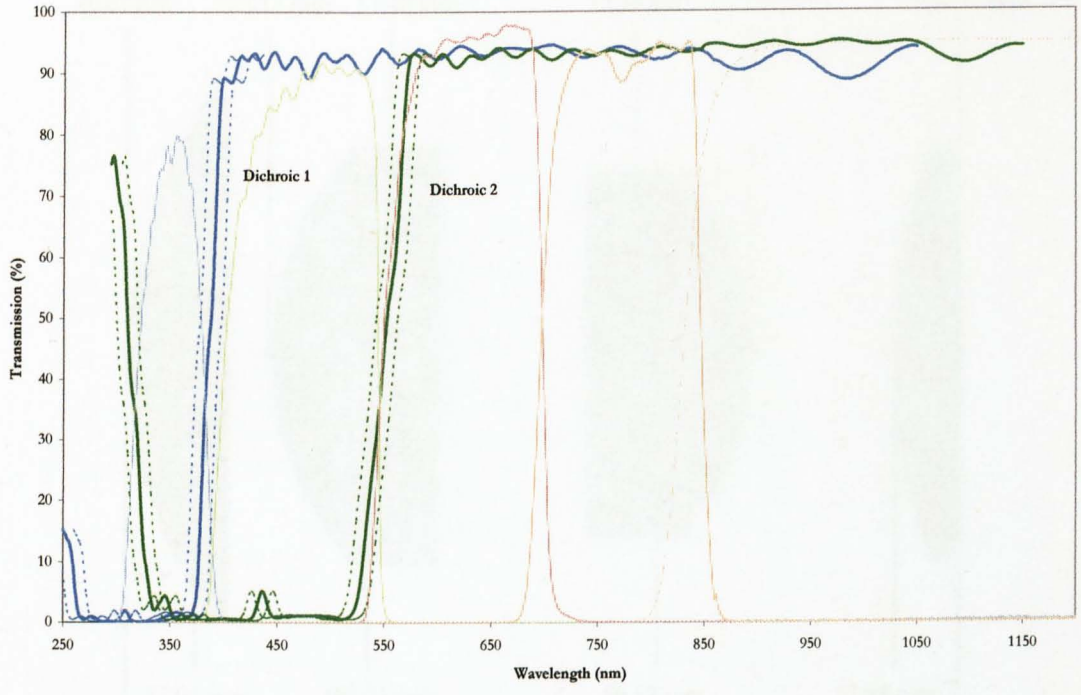

Figure 2.2.6: As built dichroic profiles. The dotted line either side of each dichroic cut indicates the cut point shift due to polarised light.

2.2.4 Re-imaging cameras

The lenses in the re-imaging cameras were manufactured by Specac Limited. Each re-imaging camera is comprised of six lenses that re-focus light from the dichroic beam-splitters onto a CCD chip. The lens layout is singlet – doublet – doublet – singlet (shown in figures 2.2.7a and 2.2.7b) where each doublet is glued using cement that is transparent to the wavelength of light that the re-imaging camera operates in and matched in refractive index. Each re-imaging camera is optimised for the wavelength band that it transmits, thereby maximising throughput and reducing chromatic aberrations.

Tolerances on the lens drawings were as small as the optics budget would allow, so that image quality was kept as high as possible (see section 4.2.1 for image quality). The glass types in the re-imaging cameras were standard types so that extra money could be spent adhering to the low manufacturing tolerances on the lenses.

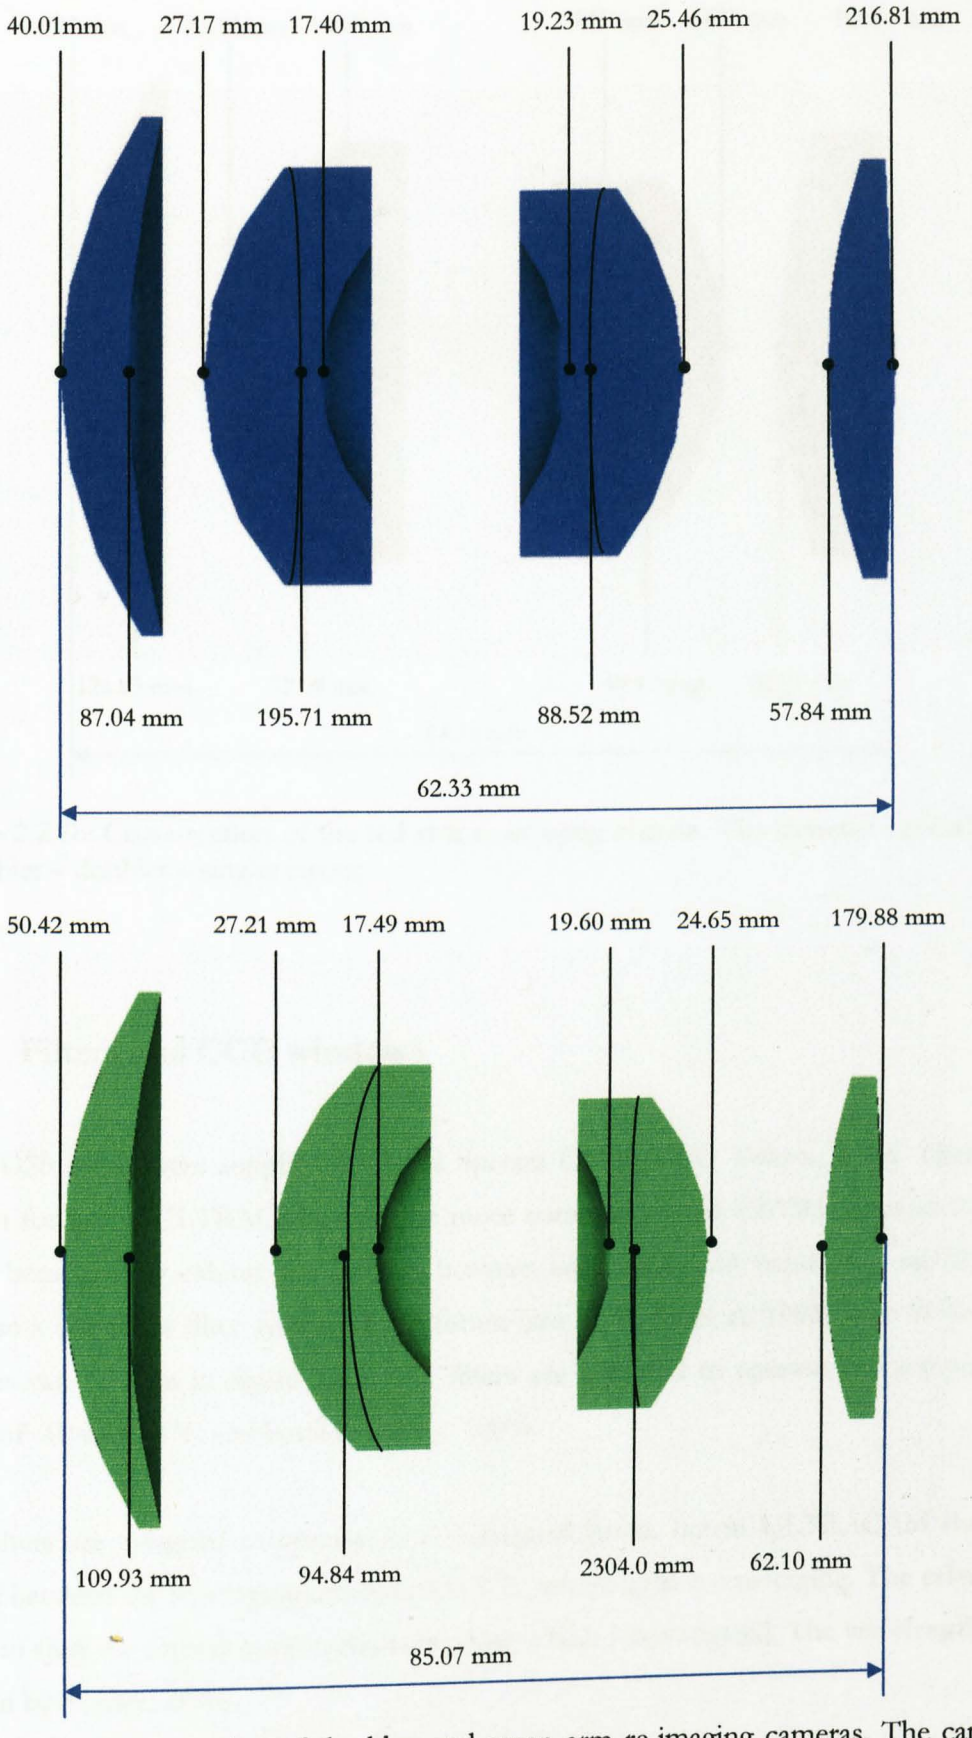

Figure 2.2.7a: Cross-section of the blue and green arm re-imaging cameras. The cameras have a singlet – doublet – doublet – singlet layout. Note that the lens in the first doublet of the blue camera is tapering at the edges, forming a knife edge (see section 2.2.8).

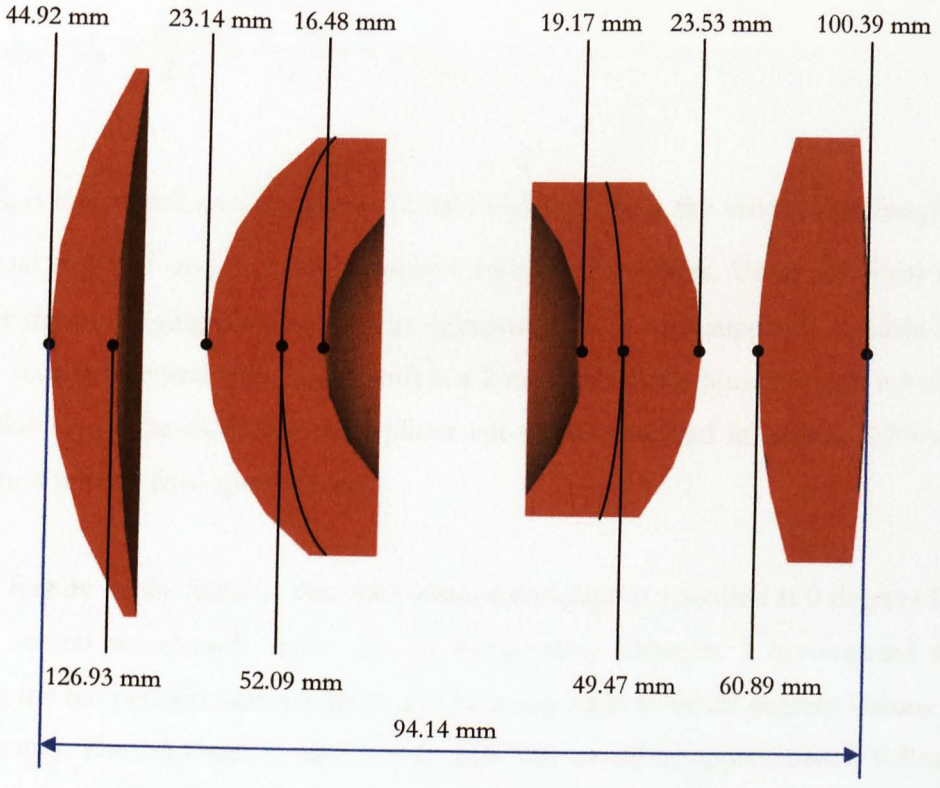

Figure 2.2.7b: Cross-section of the red arm re-imaging camera. The cameras have a singlet – doublet – doublet – singlet layout.

2.2.5 Filters and CCD windows

The SDSS filters were supplied by Asahi Spectra Co. based in Tokyo, Japan. They were chosen for use in ULTRACAM over the more commonly-used UBVRI Johnson-Cousins filters because they exhibit less overlap between adjacent colour bands and are likely to become a dominant filter system of the future (see Fukugita et al, 1996). The SDSS filter profiles can be seen in figure 2.2.3. The filters are specified to operate in a temperature range of -10 to +30 °C and humidity of 0 – 100%.

The filters are designed to operate in a collimated beam, but in ULTRACAM they are placed between the re-imaging camera and CCD, where light is converging. The effect this has is to shift the central wavelength, something which I investigated. The wavelength shift is given by (Carter, 1995):

$$\lambda_\theta - \lambda_0 = \frac{\lambda_0}{2}\left(\frac{\sqrt{n^2 - \sin^2\theta}}{n} - 1\right), \quad (2.1)$$

where λ_0 is the central wavelength at normal incidence, λ_θ is the central wavelength at the off-normal angle, θ and n is the refractive index of the filters. Using the focal ratio of f/2.3 for the re-imaging cameras, θ was calculated and n was supplied by Asahi Spectra Co. The resulting central wavelength shift is a 2 nm shift to the blue in *u'* and a 4 nm shift to the blue in *g'*. The dichroic beam-splitter cut points specified in section 2.2.3 use this information in their final specification.

Another feature of the filters is that the central wavelength is specified at 0 degrees Celsius, but the central wavelength shifts due to temperature changes. I investigated this by assuming the temperature differential in an observing night to be 20 degrees Celsius (at the very extreme). The calculated central wavelength shift would be approximately 0.4 nm at λ_0 = 750 nm (Parker & Bland-Hawthorn 1998). This effect is insignificant compared to the shift due to the converging beam and has therefore been ignored.

2.2.6 CCD windows

The CCD windows are made from UV-grade fused silica and fit into the CCD head-lid to allow light from the re-imaging cameras to focus onto the CCD chip. They are plane-parallel discs manufactured by Specac Limited. Each window is 50 mm in diameter and 5 mm thick with flatness of λ/4 @ 633 nm, parallelism < 1 arcminute[4] and scratch – dig of 20 – 10.

One concern which I investigated was flexing of this window due to the vacuum in the CCD head. If this occurred then the window would act as a meniscus-type lens, causing the focal plane to not be at the surface of the CCD. To analyse this flexure, the window was treated as a uniformly loaded circular plate supported at the edges (Roark & Young 1975). The maximum deflection, *y* is given by:

[4] Parallelism or wedge refers to how parallel the two surfaces of a flat piece of glass are, and is specified as an angle measured in arcminutes.

$$y = \frac{(5+\nu)\rho r^4}{64(1+\nu)D}, \tag{2.2}$$

where ρ is the uniform loading pressure, ν is Poisson's ratio, r is the radius of the hole that the plate covers and D is flexural rigidity given by:

$$D = \frac{Et^3}{12(1-\nu^2)}, \tag{2.3}$$

where E is Young's modulus and t is the thickness of the circular plate. For fused silica E = 7.3 x 10^{10} Pa and ν = 0.164. The resulting deflection is y = (1.43 ± 0.24) x 10^{-6}m, corresponding to a lens with a radius of curvature of 140 ± 25 m. The focal length of such a lens, calculated using the lens makers formula (Nolan, 1995), is f = 27 ± 4.5 km. It is concluded that a window flexing into a lens with such a large focal length will have a negligible effect on the focus and can be treated as a flat piece of glass.

2.2.7 Anti-reflection coatings

Reflection losses at each optical surface can severely reduce instrumental throughput. The exact loss at an air/glass interface, R is given by:

$$R = \left[\frac{(n-1)}{(n+1)}\right]^2, \tag{2.4}$$

where n is the refractive index of the glass. Typical loss values at each lens surface are 4%, making losses in systems with many optical elements significant. Light passing through the blue arm of ULTRACAM is incident upon 21 optical surfaces, reducing the instrumental throughput to just 42%. Note that this figure does not take the additional losses dues to absorption in the glass into account. The anti-reflection coating used in ULTRACAM reduces reflection losses to approximately 2% per surface increasing the instrumental throughput in the blue arm to 65% (not taking absorption into account). Table 2.2 shows the throughput for all three ULTRACAM channels. There is therefore a clear requirement

for anti-reflection coatings to be applied to the ULTRACAM optics. The coatings have been supplied by CVI Technical Optics.

Table 2.2: Throughput of the ULTRACAM optics with and without AR coating

Channel	Number of optical surfaces	Throughput (%)	
		Uncoated	Coated
Blue	21	42	65
Green	23	39	63
Red	24	38	62

Four different AR coatings were required to accommodate the eight types of glass used in the ULTRACAM optics (see figure 2.2.8). Unfortunately, these coatings are expensive and there was insufficient money in the optics budget to provide a broad-band AR coating on the collimator and AR coatings on each re-imaging camera optimised for its wavelength of operation. As a compromise a high-quality multi-layer broad-band coating was applied to all of the optics. The second surface of each dichroic, CCD windows and the SDSS filters were also AR coated in this way (the former two by CVI Technical Optics and the latter by Asahi Spectra Co.).

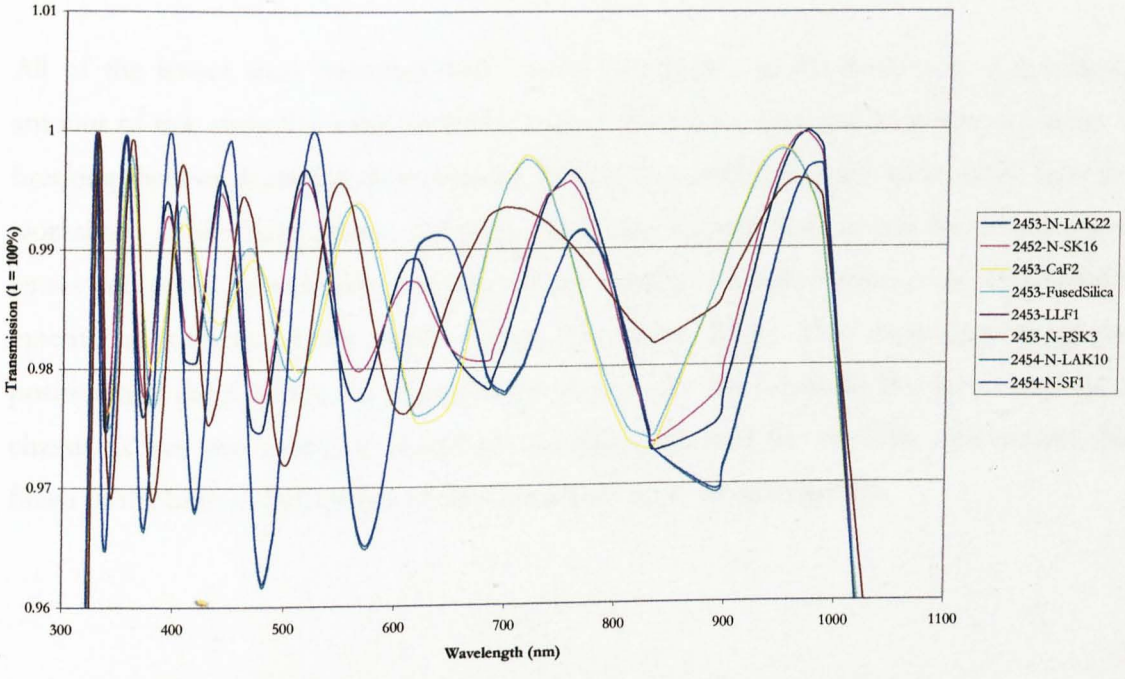

Figure 2.2.8: Theoretical AR coating performance from CVI Technical Optics.

2.2.8 Optics procurement and mounting

Procurement and mounting of the optical elements was the responsibility of the Sheffield project team. Tender documents were sent to several lens manufacturers and the contract was awarded to Specac of Kent, to manufacture lenses for the WHT and Aristarchos collimators, three re-imaging cameras and windows for the CCDs. My task was to produce engineering drawings of each lens according to the optical prescription, a document specifying each lens dimension and material. The drawings incorporated an addition to the clear aperture to allow the lenses to be mounted and dimensioning tolerances. Drawings were also compiled for the CCD windows and dichroic beam splitter substrates.

To mount each lens, 2 mm was added to the diameter of the clear aperture. In several cases however, this caused the lens to form a knife edge (see figure 2.2.7a), where the lens edge thins because the radius of curvature on each surface is too severe for the overall lens diameter and thickness. Knife edges cause problems during manufacture, as the lens is susceptible to chipping and can be difficult to mount in a barrel. To overcome this problem, the lens thickness is usually increased, but this means re-optimising the optical layout. Fortunately, the worst knife edges occurred in the re-imaging camera doublets, which meant that these lenses could be mounted by gluing them to their doublet partners.

All of the lenses were mounted into barrels by Specac, which transferred a significant amount of risk away from the Sheffield project team. The main risk in mounting lenses is breaking them or damaging their coatings. Giving Specac this responsibility, rather than the workshops at Sheffield, meant that the project team did not have to pay for any damaged lenses and saved considerable time and effort. I had to provide Specac with details of the external appearance of the optics barrels (see figure 2.5.6). This meant specifying the position and detail of the mounting flanges used to fix the barrels to the opto-mechanical chassis. It was also necessary to specify the flange required for the filter mounts and the finish of the barrels. Full details of the optics barrels are in section 2.5.8.

2.2.9 As built spot diagrams

Spot diagrams theoretically model the path of a monochromatic light beam travelling through the ULTRACAM optics. The spots are a measure of the distribution of the light falling on the CCD and hence give an indication of the image quality. Three spots have been modelled at different locations in the field, as shown in figure 2.2.9(a). There is a spot at the chip centre, half way from the centre to the chip corner and at the chip corner. The ULTRACAM optics have been designed so that the best image quality is in an annulus around the chip centre, as shown in figure 2.2.9b. This is because in drift mode windows will be placed nearer to the edges and corners of the CCD (see section 2.3.3).

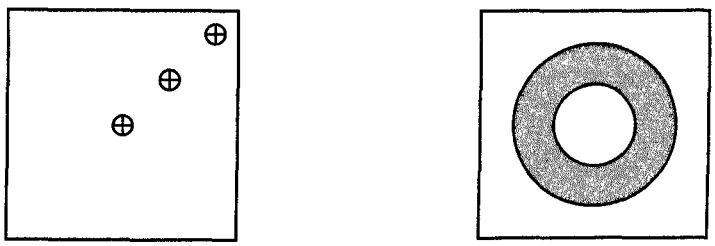

Figure 2.2.9: (a) The location of the spots on the CCD chip. (b) The region of optimum image quality on the CCD chip.

The spot diagrams for the ULTRACAM 'as built' optics are shown in figures 2.2.10, 2.2.11 and 2.2.12. The optical design went through a number of iterations to arrive at the final design. The initial optical design was handed to SPECAC who compared the optical elements with a tool plate list to identify which lens they did not have a tool plate for. The tool plate is used to shape the lens to the correct radius of curvature. Manufacturing new tool plates is expensive, so it is desirable to match the optical design to the existing tool plates. This was done with no loss of image quality, but it was then discovered that the expensive CaF2 lenses in the collimator required so much glass that the cost was beyond the optics budget. Another iteration of the optical design alleviated this. An error was noticed by the optical designer in the collimator design so a third design iteration was required to overcome this. A final iteration was required because Schott AG of Mainz, Germany, were unable to supply the N-SK16, N-LAK10 and N-LAK22 glasses, which are lead free. These were originally offered by Schott but were subsequently unavailable so the old versions of the glass, SK16, LAK10 and LAKN22 containing lead were used and the final design was complete.

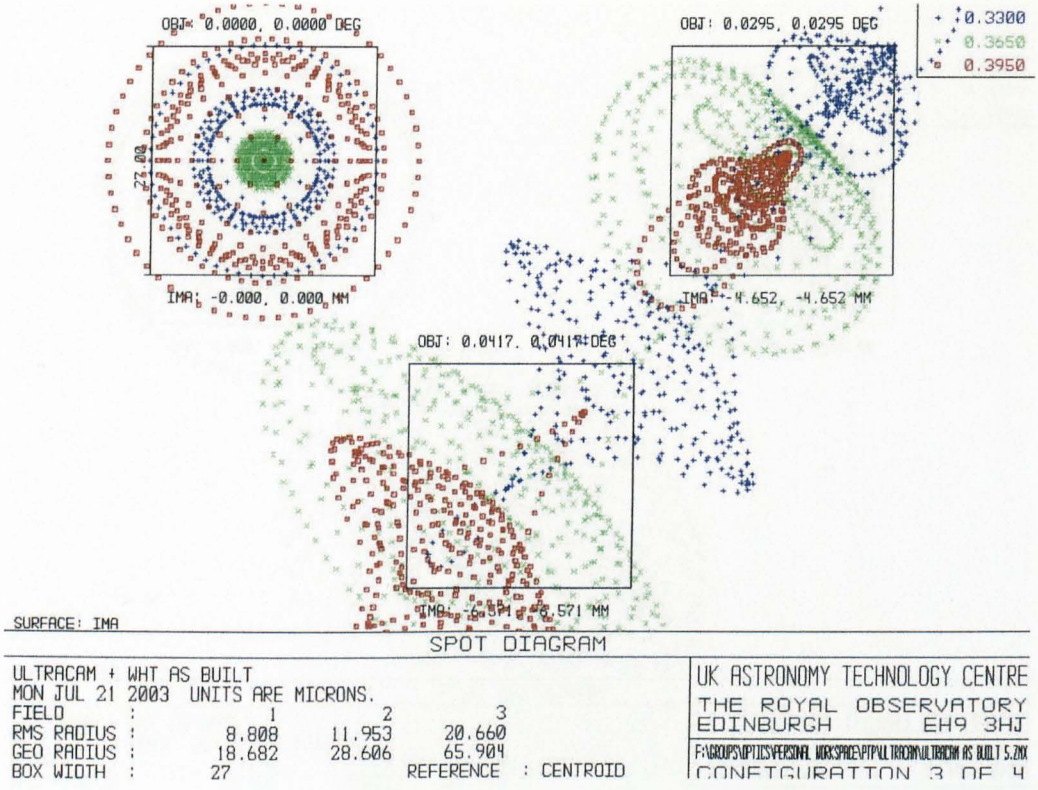

Figure 2.2.10: Blue CCD spot diagram, 'as built'. The top left spot is at the centre of the field, the top right is mid way to the corner of the field and the bottom spot is at the corner of the field. The wavelength of light used is shown in the top right corner.

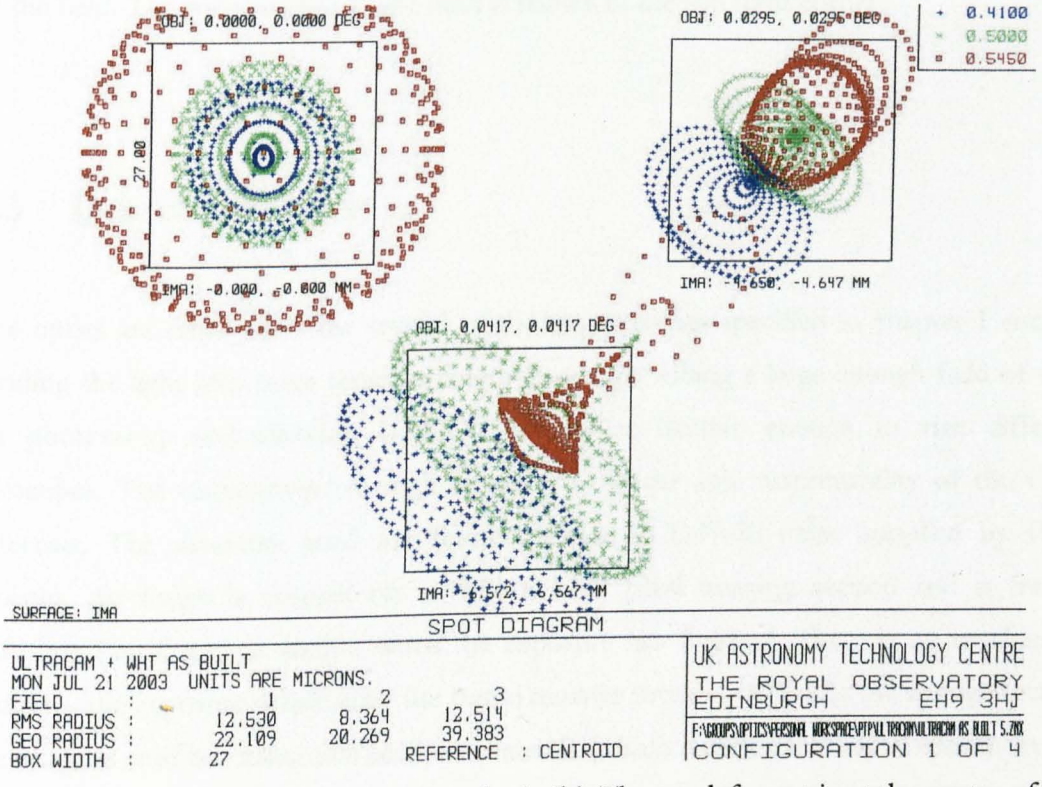

Figure 2.2.11: Green CCD spot diagram, 'as built'. The top left spot is at the centre of the field, the top right is mid way to the corner of the field and the bottom spot is at the corner of the field. The wavelength of light used is shown in the top right corner.

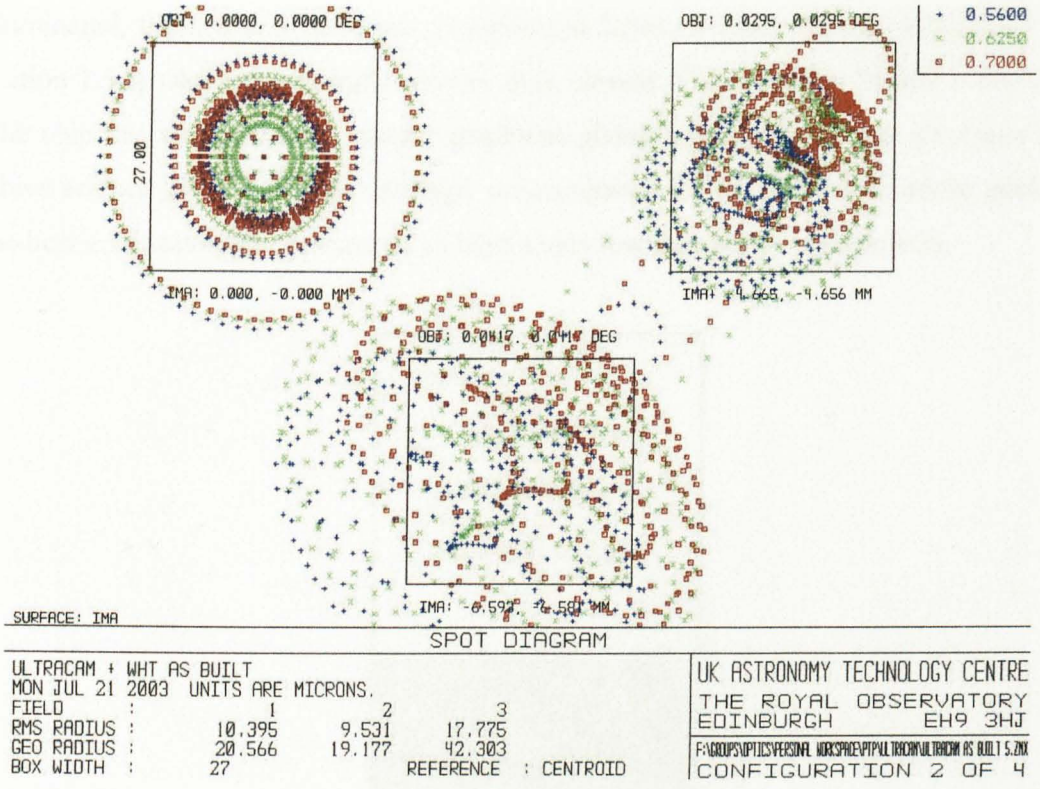

Figure 2.2.12: Red CCD spot diagram, 'as built'. The top left spot is at the centre of the field, the top right is mid way to the corner of the field and the bottom spot is at the corner of the field. The wavelength of light used is shown in the top right corner.

2.3 Detector design

The optics are responsible for several of the requirements specified in chapter 1 such as dividing the light into three separate colour bands, providing a large enough field of view for photometry and allowing ULTRACAM to be flexible enough to visit different telescopes. The requirement of high frame rates is the sole responsibility of the CCD detectors. The detectors used are frame transfer CCD47-20 chips supplied by E2V, Lincoln. An image is formed on a 1024 x 1024 pixel imaging section and is frame-transferred to a storage section when the exposure has finished. There is no mechanical shutter – the exposure is halted by the frame transfer process. Whilst in the storage section the image is read out using two readout channels (which halves the readout speed). In the meantime, the imaging section is taking the next exposure. The minimum dead time is the time it takes to frame transfer the image to the storage section.

Figure 2.3.1 shows a schematic of a 47-20 frame transfer CCD. The CCD47-20 is back-illuminated, thinned and functions in Advanced Inverted Mode Operation (AIMO – see section 1.5.2) which dramatically reduces dark current (E2V CCD47-20 data sheet, 2000). The chip has very low-noise output amplifiers giving typically 2 – 4 rms electrons/pixel. Three science grade 0 CCD47-20 chips were acquired for ULTRACAM, where grade 0 is the best grade available, guaranteed to have a very low number of pixel defects.

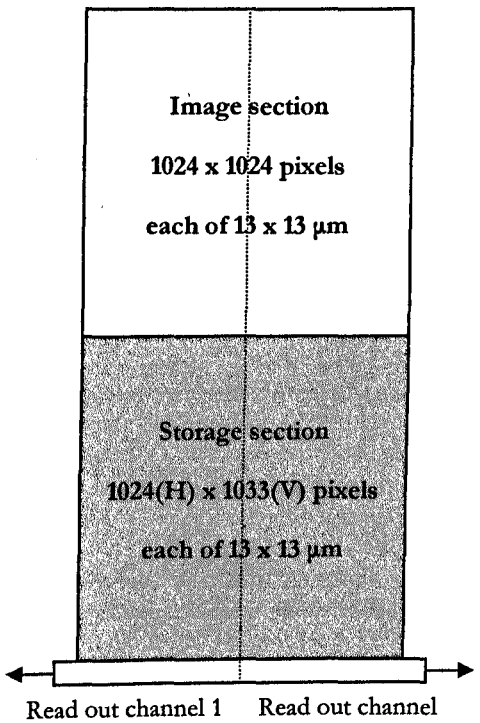

Figure 2.3.1: E2V CCD 47-20 frame-transfer chip used in ULTRACAM. There is a 1024 x 1024 pixel array used for imaging and a similarly sized array to store the image prior to being read out. Each pixel is 13 microns square.

There are three CCD readout modes for ULTRACAM: full frame, window mode and drift mode, each of which is described in more detail in sections 2.3.1, 2.3.2 and 2.3.3. The exposure time with ULTRACAM is not as simple a concept as with non-frame transfer chips. The CCDs in ULTRACAM try to frame as fast as they can. This sets the minimum exposure time which can be added to by setting a parameter called the 'exposure delay'. For each readout mode, the user can specify the number of exposures to be taken, the exposure delay in milliseconds and the binning factor in x and y. There is also a gain-speed parameter which allows the user to choose between slow and fast digitisation times, a V-clock parameter to adjust vertical clocking speeds and several gain settings. There is increased noise when operating with fast readout since the sampling time is 3.4μs/pixel compared to

11.2μs/pixel for slow readout. Each frame is time-stamped using the time signal from a GPS receiver.

Each CCD47-20 has an over-scan region, as shown in figure 2.3.2. This is a region on the CCD that is shielded from incident light and can be used for measuring the bias levels and dark current.

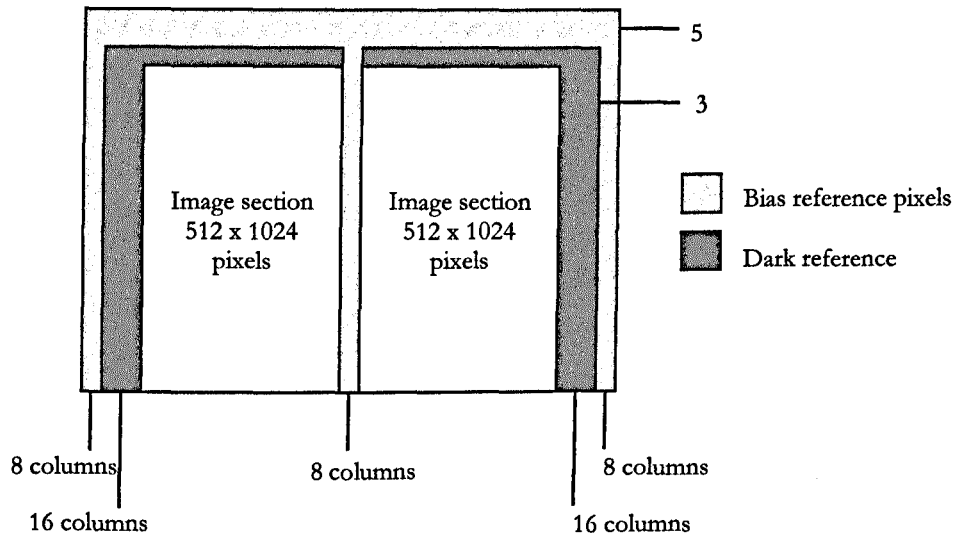

Figure 2.3.2: Layout of the E2V CCD 47-20 array, with over-scan regions yielding a total array size of 1080 x 1033 (E2V CCD47-20 data sheet, 2000).

2.3.1 Full frame read out

Full frame readout is the simplest and slowest readout mode, where the entire 1024 x 1024 imaging array is exposed and read out. The readout sequence is shown in figure 2.3.3. Time-stamping of each ULTRACAM frame occurs at the beginning of each exposure. There are three readout modes in full frame mode:

1. Full frame clear (FF_CL)
2. Full frame no clear (FF_NOCL)
3. full frame overscan (FF_OVER)

In full frame clear mode, the CCD is cleared before each exposure, of any charge accumulated whilst the previous frame was reading out. The dead-time of the CCD is therefore the clearing time plus the readout time plus the time to V-clock the image into

the storage section. The advantage of this mode is that the exposure time can be anything and is not limited by the readout time. Full frame no clear mode does not involve clearing the CCD before each exposure. The advantage here is that the CCD is exposed whilst the previous image is being readout and the dead-time is the time it takes to frame transfer the image to the storage section (1024 V-clocks). The disadvantage is that the minimum exposure time is limited to the readout time. Full frame over-scan operates in the same way as full-frame no clear but accumulates data from the over-scan regions shown in figure 2.3.2.

The frame rate in full-frame mode is relatively simple to calculate. For example, the total time taken for the sequence in figure 2.3.3 is the addition of the following components:

Exposure delay (1 ms)

Time to V-clock array in to storage section
1024 x 23.3 μs = 0.02386 s

Time to V-clock rows to serial register
1024 x 23.3 μs = 0.02386 s

Time to H-clock pixels along serial register

(1024 x 512) x 0.48 μs = 0.25166 s

Time to digitise pixels
$(1024^2$ x 3.4 μs)/2 = 1.78258 s (fast)
$(1024^2$ x 11.2 μs)/2 = 5.87202 s (slow)

Dead time = 0.02386 s

Total read out time = 2.083 s (fast)
Total read out time = 6.172 s (slow)

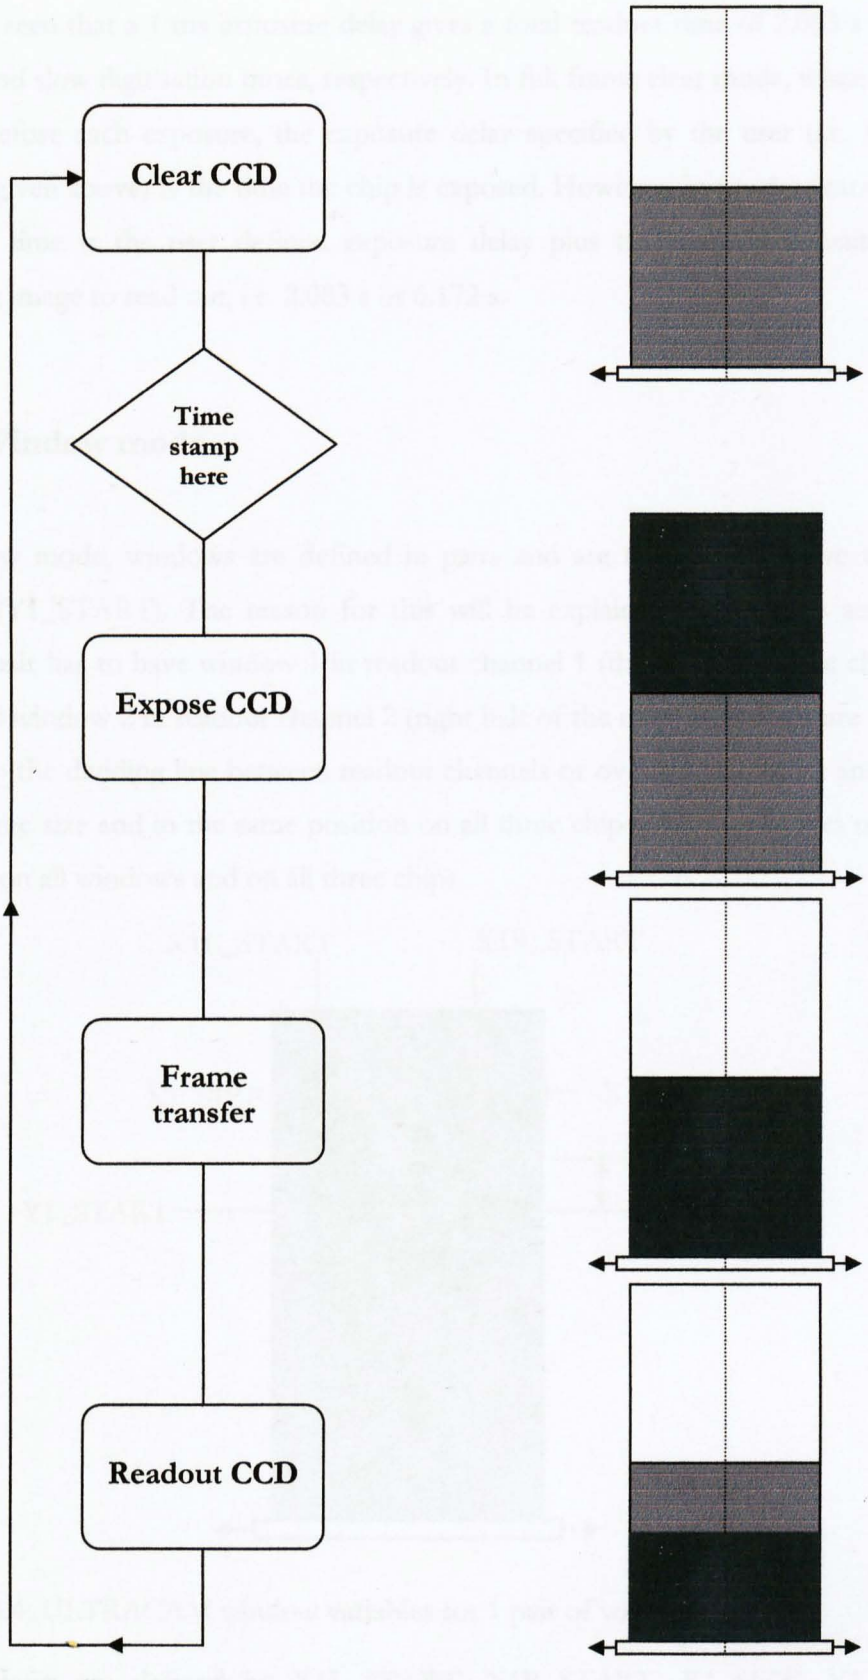

Figure 2.3.3: One full-frame sequence. Min-dead mode clears the CCD on the first frame but then skips clearing for the rest of the run.

It can be seen that a 1 ms exposure delay gives a total readout time of 2.083 s and 6.172 s for fast and slow digitisation times, respectively. In full frame clear mode, when the CCD is cleared before each exposure, the exposure delay specified by the user (i.e. 1 ms in the example given above) is the time the chip is exposed. However, in no-clear mode, the total exposure time is the user defined exposure delay plus the extra time waiting for the preceding image to read out, i.e. 2.083 s or 6.172 s.

2.3.2 Window mode

In window mode, windows are defined in pairs and are restricted to have the same y-position (Y1_START). The reason for this will be explained later in this section. Each window pair has to have window 1 in readout channel 1 (the left half of the chip in figure 2.3.4) and window 2 in readout channel 2 (right half of the chip). Windows are not allowed to overlap the dividing line between readout channels or overlap each other and they must be the same size and in the same position on all three chips. Binning factors must also be the same on all windows and on all three chips.

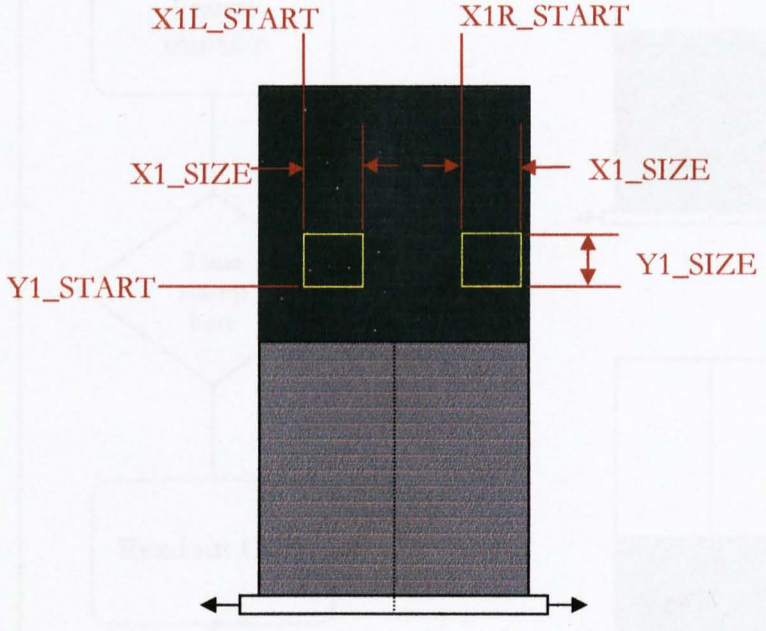

Figure 2.3.4: ULTRACAM window variables for 1 pair of windows.

The windows are defined by X1L_START, X1R_START, X1_SIZE, Y1_SIZE and Y1_START, which are shown on figure 2.3.4. A maximum of three window pairs are allowed and each pair is separately defined.

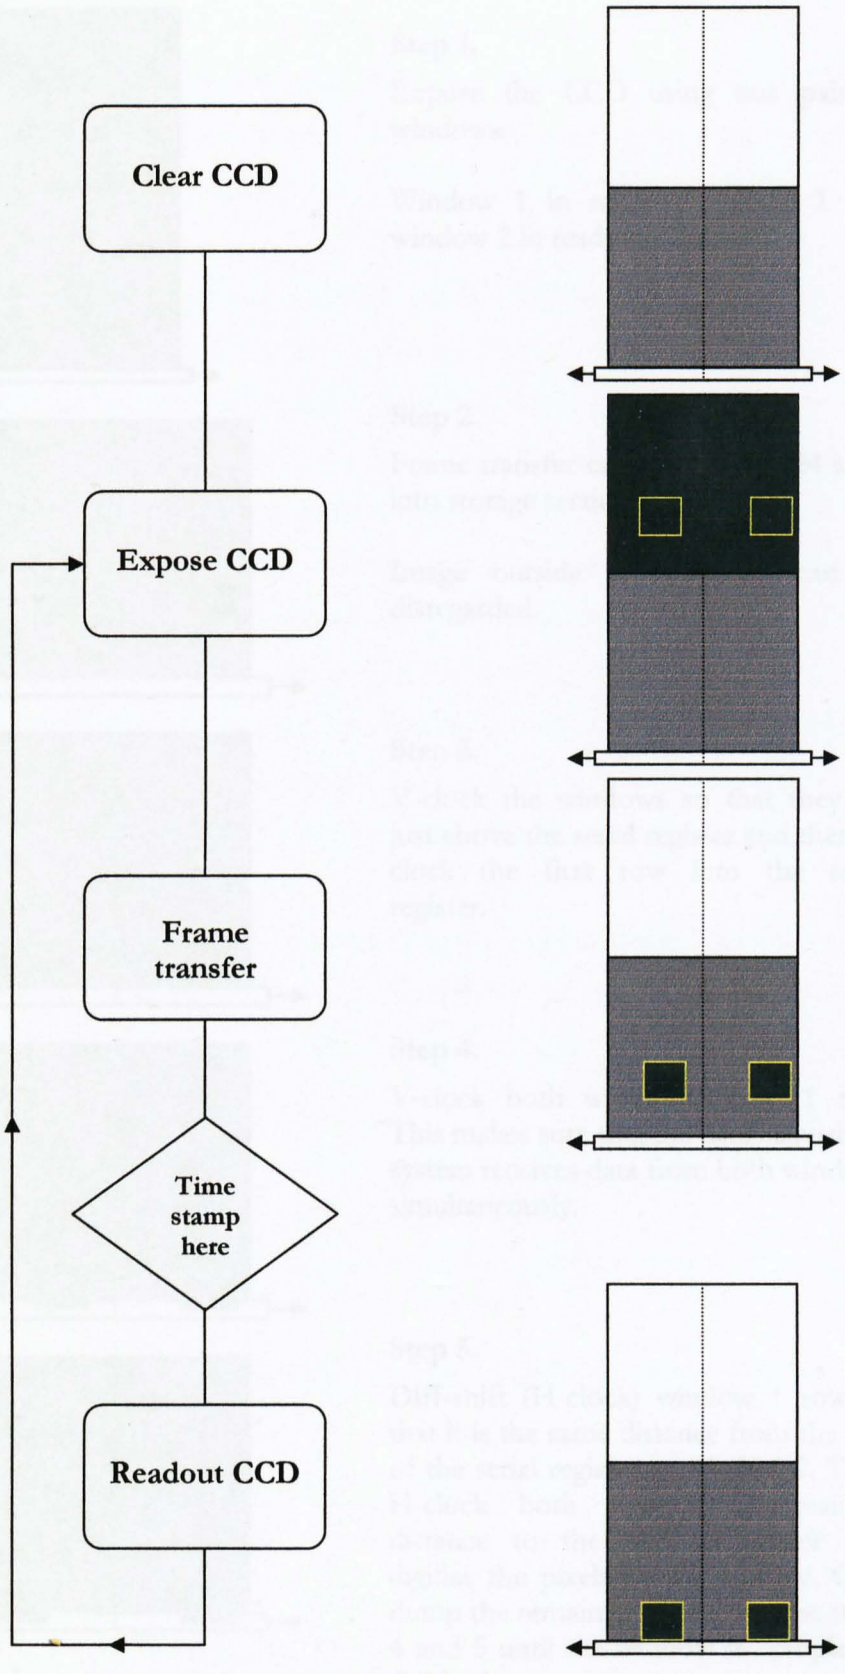

Figure 2.3.5: One sequence for window mode readout, where CCD clearing occurs just once at the beginning of the run. The entire imaging array is frame transferred, but data outside the windows is disregarded. Note that time-stamping occurs after frame transfer, as it denotes the start of a new exposure.

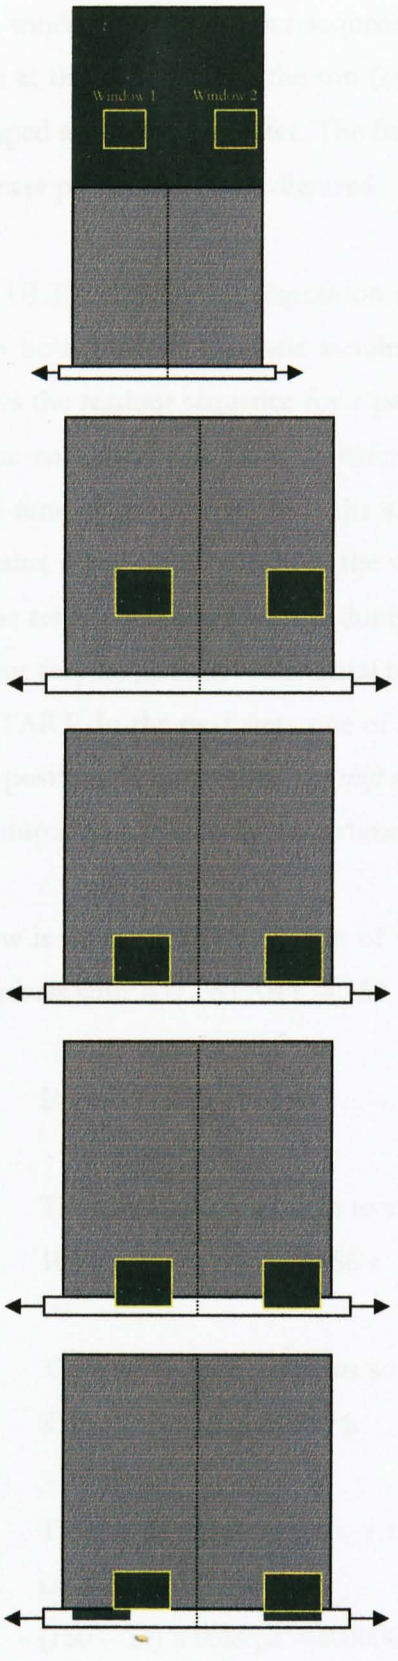

Step 1.

Expose the CCD using one pair of windows.

Window 1 in readout channel 1 and window 2 in readout channel 2.

Step 2.

Frame transfer entire 1024 x 1024 array into storage section.

Image outside of windows can be disregarded.

Step 3.

V-clock the windows so that they are just above the serial register and then V-clock the first row into the serial register.

Step 4.

V-clock both windows down 1 row. This makes sure that the data acquisition system receives data from both windows simultaneously.

Step 5.

Diff-shift (H-clock) window 1 row so that it is the same distance from the end of the serial register as window 2. Then H-clock both rows the remaining distance to the output register and digitise the pixels in the window. Gate dump the remaining pixels. Repeat steps 4 and 5 until the window is completely digitised.

Figure 2.3.6: Readout sequence in window mode.

The window-mode readout sequence is shown in figure 2.3.5. The CCD is cleared only once at the beginning of the run (as in full frame no-clear mode) and each frame is time-stamped after frame transfer. The frame rate is considerably higher than in full-frame mode as fewer pixels need to be digitised

The ULTRACAM data acquisition system (section 2.4) has been designed to expect data from both readout channels simultaneously (i.e. diff shift, see figure 2.3.6). Figure 2.3.6 shows the readout sequence for 1 pair of windows. When the CCD exposure finishes (step 1) the entire array is frame transferred (step 2) to the storage section. Consequently the dead-time in window mode is the same as for full frame readout (min-dead). Although it contains signal, data outside of the windows is disregarded, but it still has to be V-clocked to the serial register where it is dumped. Step 3 in the sequence is to V-clock the windows so that they lie just above the serial register. This is why each window in a pair has the same Y_START. In the next step, one of the windows is H-clocked so that it mirrors its partner in x-position. This is called *diff shift* and ensures that the pixels of each window both reach the output amplifier at the same time. The windows are then read out in parallel.

Below is an example calculation of window mode readout speed. The windows are 250 x 150 pixels with X1L_START = 150, X1R_START = 750 and Y1_START = 275 pixels.

Exposure delay (1 ms)

Time to V-clock array in to storage section (step 2)
1024 x 23.3 μs = 0.02386 s

Time to V-clock windows so they are on the serial register (step 3)
275 x 23.3 μs = 0.00641 s

Time to H-clock window 1 to mirror window 2 in the along the centre of the chip, i.e. diff shift (step 4)
(150 – 24) x 0.48 μs = 0.00006s

Time to H-clock the remaining distance and gate dump the pixels outside the window (step 5)

(150 x 512) x 0.48 μs = 0.03686 s

Time to V-clock the entire window onto serial register (step 5)
150 x 23.3 μs = 0.00350 s

Time to digitise pixels (step 5)
(250 x 150) x 3.4 μs = 0.1275 s (fast)
(250 x 150) x 11.2 μs = 0.4200 s (slow)

Dead time = 0.02386 s

Total read out time = 0.199 s (fast)
Total read out time = 0.492 s (slow)

The time taken to complete steps 3 and 4 can be minimised to increase the frame rate by optimising the placement of windows. A window pair should be placed at the lowest y-value possible, close to the storage section, to minimise step 3. Both windows should be positioned so that they mirror each other down the centre of the chip, hence eliminating the diff shift (step 4 of figure 2.3.6). Multiple window pairs will decrease the frame rate: steps 3 and 4 have to be performed for each window pair. The dead-time for window mode is always the same, the time to V-clock 1024 rows into the storage section (0.02386s).

2.3.3 Drift mode

Drift mode is the fastest readout mode, using only one window pair positioned optimally on the CCD to eliminate steps 3 and 4 of figure 2.3.6. The main difference in drift mode is a reduction in the dead-time due to the frame transfer process. The window pair is simply V-clocked until it is just covered by the mask defining the storage section. Dead-time is therefore significantly reduced from 0.02386 seconds in full frame and window modes, especially when the window pair is positioned at Y1_START = 1 and the window size is small. For example, the dead time incurred by frame transferring a 24 x 24 pixel window positioned at Y1_START = 1 is 0.0006 seconds (see example calculation).

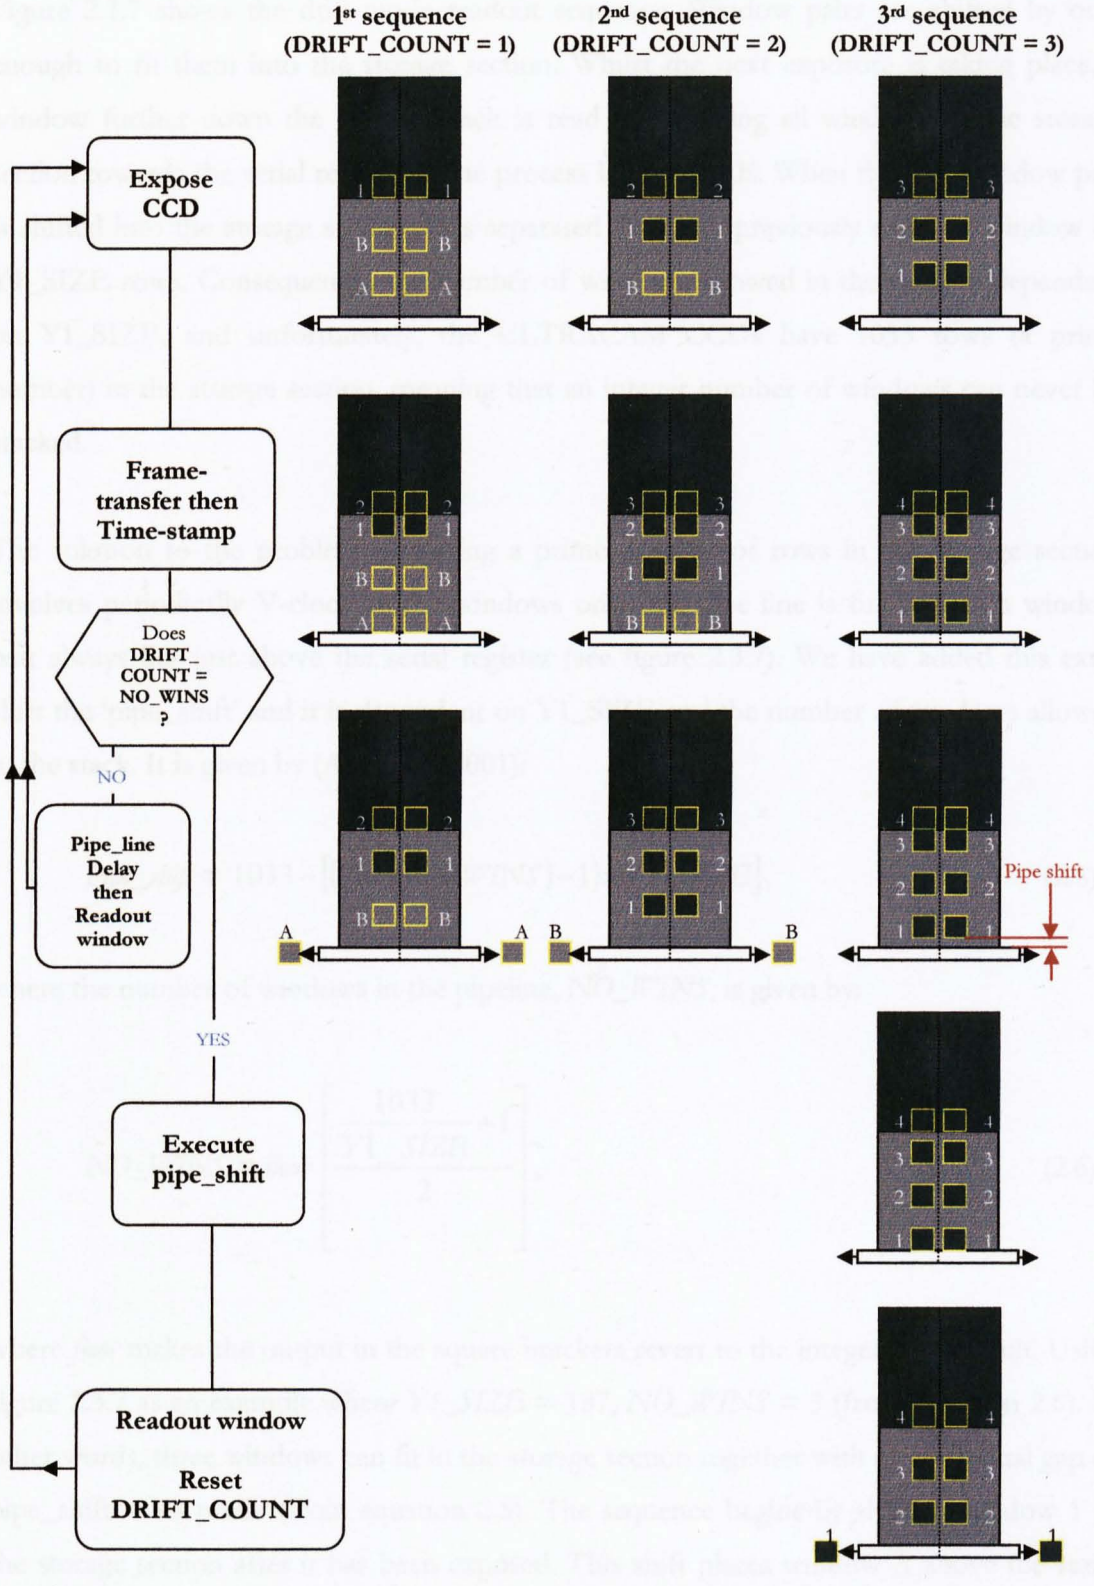

Figure 2.3.7: Readout sequence in drift mode. Exposed windows form a vertical stack in the storage section, but because the storage section has 1033 rows, the vertical gaps between the windows can never all be the same. To overcome this problem, a pipe shift is introduced (3rd sequence) when the stack is full. After the pipe shift the sequence resets (returns to the 1st sequence) and cycles until the required number of frames (windows) have been accumulated.

Figure 2.3.7 shows the drift-mode readout sequence. Window pairs are shifted by only enough to fit them into the storage section. Whilst the next exposure is taking place, a window further down the vertical stack is read out, shifting all windows in the storage section towards the serial register in the process by Y1_SIZE. When the next window pair is shifted into the storage section, it is separated from the previously exposed window by Y1_SIZE rows. Consequently, the number of windows allowed in the stack is dependent on Y1_SIZE, and unfortunately, the ULTRACAM CCDs have 1033 rows (a prime number) in the storage section, meaning that an integer number of windows can never be stacked.

The solution to the problem of having a prime number of rows in the storage section involves periodically V-clocking the windows once the pipe line is full so that a window pair always lies just above the serial register (see figure 2.3.7). We have added this extra shift the 'pipe_shift' and it is dependent on Y1_SIZE and the number of windows allowed in the stack. It is given by (Atkinson, 2001):

$$Pipe\_shift = 1033 - [((2 \times NO\_WINS) - 1) \times Y1\_SIZE], \tag{2.5}$$

where the number of windows in the pipeline, *NO_WINS*, is given by:

$$NO\_WINS = floor\left[\frac{\frac{1033}{Y1\_SIZE} + 1}{2}\right], \tag{2.6}$$

where *floor* makes the output in the square brackets revert to the integer value result. Using figure 2.3.7 as an example where *Y1_SIZE* = 187, *NO_WINS* = 3 (from equation 2.6). In other words, three windows can fit in the storage section together with an additional gap of pipe_shift = 98 pixels (from equation 2.5). The sequence begins by shifting window 1 to the storage section after it has been exposed. This shift places window A above the serial register, where it is read out whilst the exposure of window 2 takes place. Windows A and B are junk generated at the very start of each run in order to produce the first drift cycle. After window A has been read out, windows 1 and B are shifted down the chip by Y1_SIZE pixels. The same sequence repeats for window 2 until the exposure of window 3

ends. When window 3 is frame transferred to the store section, window 1 should ideally be above the serial register waiting to be read out. It is not, because a pipe_shift of 98 pixels is required to shift it down to the serial register.

Below is an example calculation of drift mode readout speed. The windows are 24 x 24 pixels with X1L_START = 100, X1R_START = 901 and Y1_START = 1 pixels.

Exposure delay (1 ms)

Time to V-clock array into storage section
25 x 23.3 μs = 0.0006 s

Time for pipe_shift (pipe_shift calculated using equation 2.5)
1 x 23.3 μs = 0.00002 s

Time to H-clock window 1 to mirror window 2 along the centre of the chip, i.e. diff shift
(100 – 100) x 0.48 μs = 0 s

Time to H-clock the remaining distance and gate dump the pixels outside the window
(24 x 124) x 0.48 μs = 0.0014 s

Time to V-clock the entire window onto serial register
24 x 23.3μs = 0.0006 s

Time to digitise pixels
(24 x 24) x 3.4 μs = 0.0020 s (fast)
(24 x 24) x 11.2 μs = 0.0064 s (slow)

Dead time = 0.0006 s

Total read out time = 0.0056 s (fast)
Total read out time = 0.0100 s (slow)

In practice, the windows defined in drift mode are much smaller than in the example in figure 2.3.7. The pipe_shift not only causes a reduction in frame rate, it also causes a slightly longer exposure time for the window integrating on the chip whilst the pipe_shift is being executed. Such variable exposure times are undesirable, so to make all windows have the same integration time, a delay equal to the time required for the pipe_shift is added (called pipe-shift delay – see figure 2.3.7) to every window pair, except when a pipe_shift is actually occurring. To maintain high frame rates, optimisation of Y1_SIZE is necessary, since this ultimately defines the size of the pipe_shift and pipe_shift delay. For example, a 24 pixel square window pair will require a pipe_shift of only 1 pixel whereas a 25 pixel square window requires a pipe_shift of 8 pixels. Figure 2.3.8 shows a plot of pipe_shift as a function window size, and is useful for choosing optimum window sizes to keep the pipe_shift to a minimum.

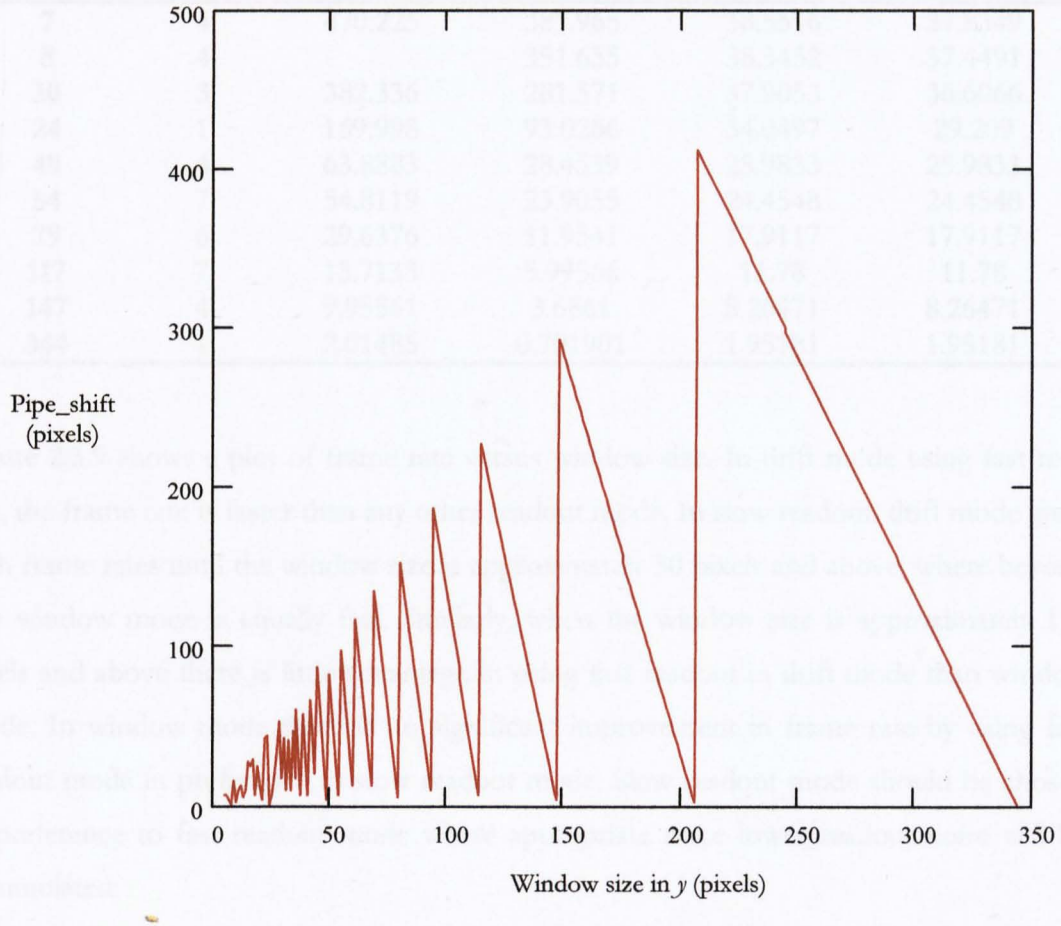

Figure 2.3.8: Plot of pipe_shift versus window size in Y (Atkinson 2001).

2.3.4 Frame rates

The frame rate has been compared for window and drift mode as a function of window size. Optimum window sizes that reduce the pipe shift have been used for the comparison, and these are 7, 8, 10, 24, 49, 54, 79, 117, 147 and 344 pixels. The frame rate has been calculated by an exposure time calculator compiled by Vikram Dhillon. The output from the exposure time calculator is shown in table 2.3.

Table 2.3: Frame rate comparison
Comparison of ULTRACAM frame rate as a function of window size. The frame rate has been calculated using the exposure time calculator on the ULTRACAM web page.

Window size (pixels)	Pipe shift (pixels)	Drift mode frame rate (Hz)		Window mode frame rate (Hz)	
		Fast	Slow	Fast	Slow
7	4	470.225	381.965	38.5516	37.8349
8	4	-	351.655	38.3452	37.4491
10	3	382.336	281.571	37.9053	36.6066
24	1	169.998	93.0266	34.0497	29.209
49	4	63.8883	28.4539	25.9833	25.9833
54	7	54.8119	23.9055	24.4548	24.4548
79	6	29.6376	11.9541	17.9117	17.9117
117	7	15.7133	5.99566	11.78	11.78
147	4	9.95561	3.6861	8.26471	8.26471
344	1	2.01485	0.701901	1.95181	1.95181

Figure 2.3.9 shows a plot of frame rate versus window size. In drift mode using fast read out, the frame rate is faster than any other readout mode. In slow readout, drift mode gives high frame rates until the window size is approximately 50 pixels and above, where beyond this window mode is equally fast. Similarly, when the window size is approximately 117 pixels and above there is little advantage in using fast readout in drift mode than window mode. In window mode there is no significant improvement in frame rate by using fast readout mode in preference to slow readout mode. Slow readout mode should be chosen in preference to fast readout mode where appropriate since lower readout noise will be accumulated.

A maximum frame rate of 1.9 KHz has been achieved in the lab, by placing one pixel in each bottom corner of the CCDs. The highest frame rate achieved during observing is 360 Hz. This was achieved with a 24 x 40 pixel window with 8 x 8 binning. The windows were

located at X1L_START = 159, X1R_START = 825 and Y1_START = 1, with a 1 millisecond exposure delay.

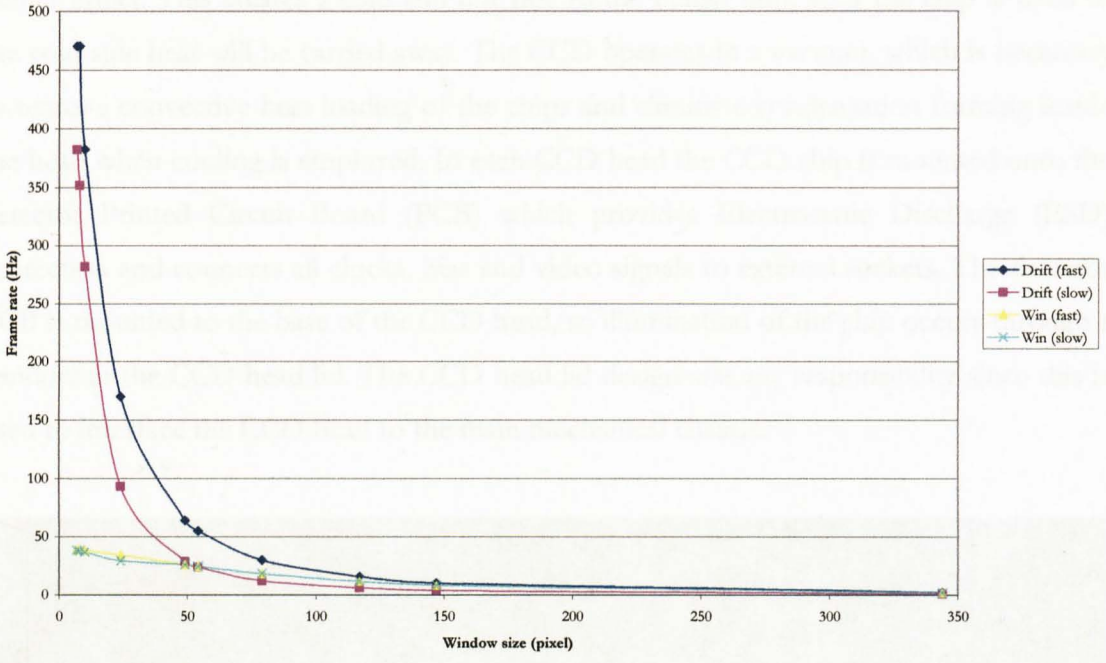

Figure 2.3.9: Frame rate as a function of window size for window and drift mode.

2.3.5 SDSU controller

Readout of the CCDs is controlled by a San Diego State University (SDSU) controller. The CCDs are read out in parallel by the SDSU controller, which has to be positioned close to all three CCDs to minimise noise. In the controller, there is a timing board which sequences the readout procedure. There are also clock boards which control the clock voltages for each CCD, and video boards which digitise data from each CCD. The power supply for the SDSU controller is located in the ULTRACAM electronics enclosure. The whole SDSU unit is cooled by water flowing through a heat exchanger fitted to the SDSU base (see figure 2.5.19).

2.3.6 CCD head

The design and manufacture of the CCD heads was undertaken by the UKATC. Each CCD chip (and its associated electronics) is contained in one of three identical CCD heads,

which control the environment that the chips operate in. The CCDs are cooled by contact with thermoelectric coolers (TECs), which are solid state heat pumps that operate using the Peltier effect. This creates a cold and hot side to the Peltier unit, so if the chip is fixed to the cold side heat will be carried away. The CCD operates in a vacuum, which is necessary to remove convective heat loading of the chips and eliminate condensation forming inside the head when cooling is employed. In each CCD head the CCD chip is mounted onto the detector Printed Circuit Board (PCB) which provides Electrostatic Discharge (ESD) protection and connects all clocks, bias and video signals to external sockets. The detector PCB is mounted to the base of the CCD head, so illumination of the chip occurs through a window in the CCD head lid. The CCD head lid design was my responsibility since this is used to interface the CCD head to the main mechanical chassis.

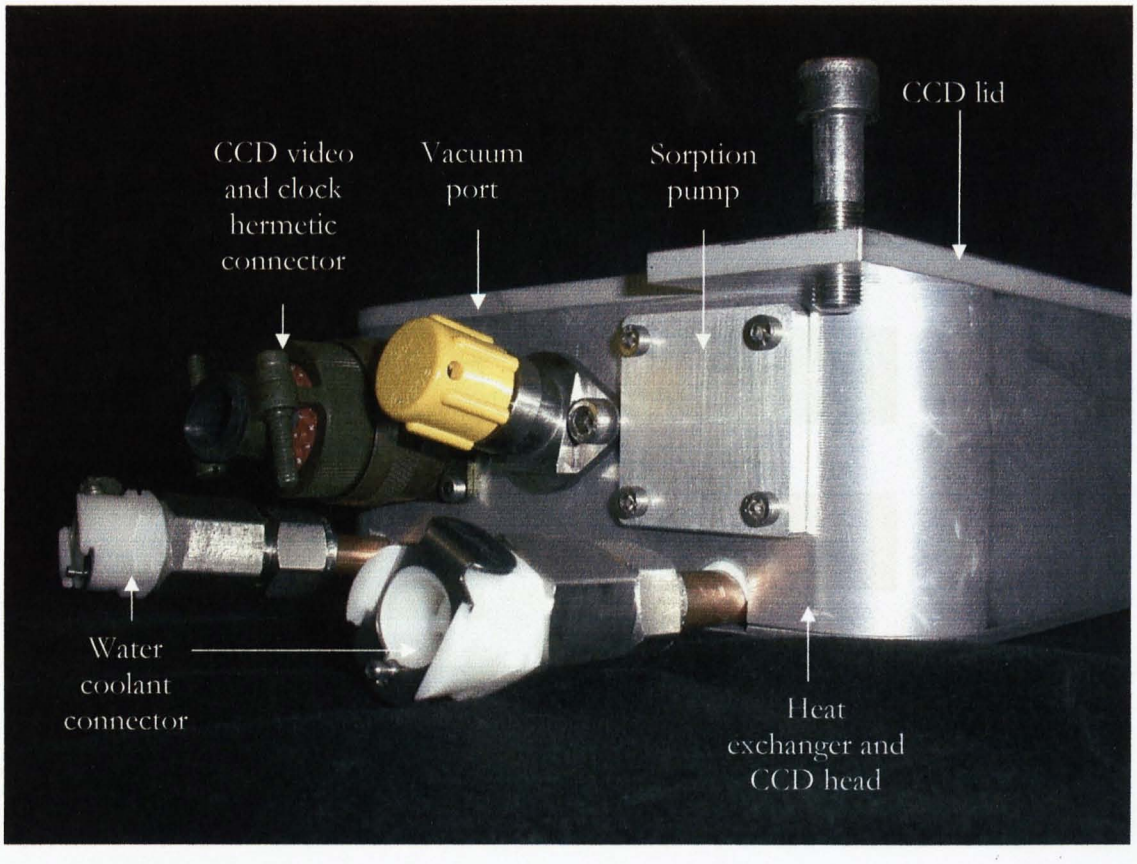

Figure 2.3.10: ULTRACAM CCD head showing one hermetic connector, vacuum port, sorption pump and water coolant connectors.

The detector PCB provides an interface between the CCD chip and its controller via two hermetic connectors (one is the green connector in figure 2.3.10). The smaller hermetic connector (not shown in figure 2.3.10) provides power to the TECs and returns the value

of the temperature of the cold side of the TEC (i.e. the CCD). Each of the three TECs has its own power supply and temperature control unit located in the ULTRACAM electronics enclosure. The second, larger, hermetic connector transfers video and clock signals between the CCD and SDSU controller. Each CCD head also has a vacuum port, which is used to pump the air out of the head before each observing run (typical hold times are approximately 1 month). A sorption pump is also installed to each head to further remove condensation when the vacuum is poor.

The CCDs are cooled to -40°C using the TECs, reducing dark current to 0.1 electrons/pixel/second. The TEC fits into the detector PCB and the cold side forms a contact with the underside of the CCD chip. Both the TEC and the detector PCB mount to the base of the CCD head which forms a heat exchanger. The heat exchanger is simply a block of aluminium embedded with copper tubing through which water can flow to carry away heat from the hot side of the TEC. The water temperature is maintained at a constant level of typically 10°C (but this depends on the dew point in the telescope dome) using a re-circulation chiller. All three CCD heads and the SDSU controller are cooled by the same re-circulation chiller – figure 2.3.11 shows the flow diagram of the water cooling system.

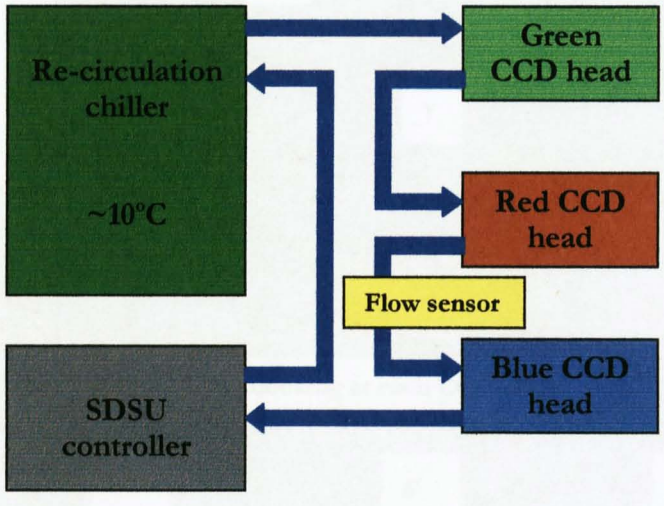

Figure 2.3.11: Water cooling flow diagram for ULTRACAM.

When cooling the CCDs before a run, they need to be cooled at a rate of no greater than 5°C / minute, otherwise the chips may suffer thermal damage. The re-circulation chiller has to maintain water temperature above the dew point temperature otherwise water will condense on the CCDs and SDSU controller, making short circuits likely.

The water pipes are arranged in series with the SDSU last in the chain since it produces the most heat and is the least important thing to keep cool. Within the system there is a flow sensor to ensure that the TECs are not operated without a water supply, otherwise they could be damaged. If the water flow is cut off for any reason this fail safe mechanism automatically shuts down the TECs.

2.3.7 Orientation of the CCD heads

The blue and green CCDs in ULTRACAM each see light reflected off a dichroic, where the red channel does not. This means the orientation of the field on the red CCD is different to the blue and green CCDs, which is a problem given that the CCD windows must be in the same location on each chip. I investigated the consequences of this and how this would affect the integration of the CCD heads with the mechanical chassis.

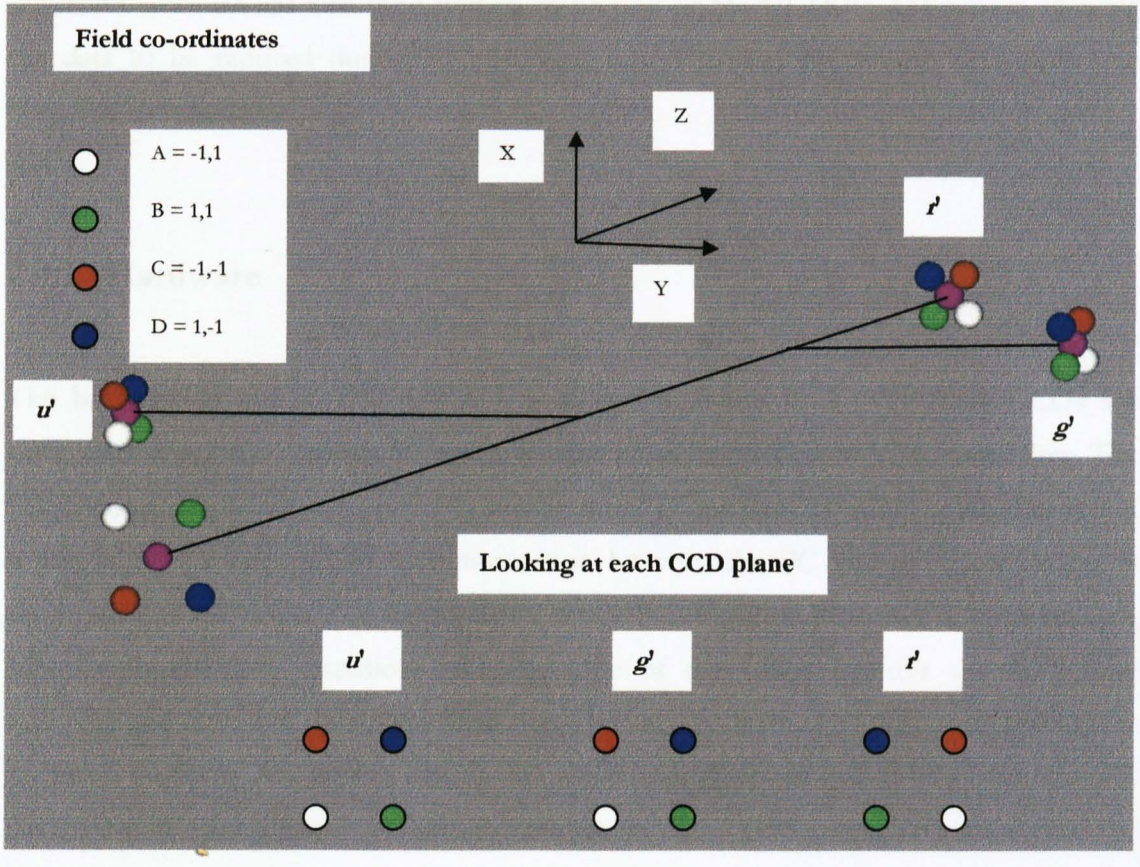

Figure 2.3.12: Image orientation on the 3 ULTRACAM CCDs. Coloured spots are propagated from the corners (A, B, C and D) and centre of the field and are shown how they are observed on each CCD, blue (*u'*), green (*g'*) and red (*r'*).

The coloured spots in figure 2.3.12 are propagated from the collimator through the system, reflecting off the dichroic beam-splitters and falling onto each CCD. When looking at the orientation of the spots on each CCD the blue and green cameras are the same, but the red image is a mirror of the other two. Consequently, the red readout electronics were rewired to remove this effect with the result that it now forms images of the same orientation as the blue and green CCDs.

2.4 Data acquisition and camera control software

The UKATC procured all the hardware necessary for controlling ULTRACAM and produced the software to perform the data acquisition process and camera control. This was performed under contract and is described in more detail by Beard et al. (2002). Data reduction software has also been written to be part of the ULTRACAM system allowing the data to be reduced during an observing run. This was written by Tom Marsh at Warwick University

2.4.1 Hardware

The hardware layout of ULTRACAM is shown in figure 2.4.1. The three CCDs are controlled by an SDSU controller and are temperature controlled by TEC controllers. The SDSU controller is linked to the instrument control PC by two 250 MHz optical fibres and a parallel port interrupt line. Communication between the PC and SDSU controller is through an SDSU PCI. The instrument control PC is a dual processor system running Linux with real-time extensions (RTAT). One of the CPUs controls the GPS time stamping and the other does everything else. One task performed by the latter processor is archiving ULTRACAM data to a SCSI disk array for high-speed data storage and retrieval. Accurate time signals from GPS satellites are received by a GPS antenna placed outside the dome and are relayed to the instrument control PC for time-stamping ULTRACAM data frames. A data reduction PC situated in the telescope control room remotely connects to the instrument control PC, providing a user interface to control ULTRACAM CCDs and the pipeline data reduction system (see section 2.4.4).

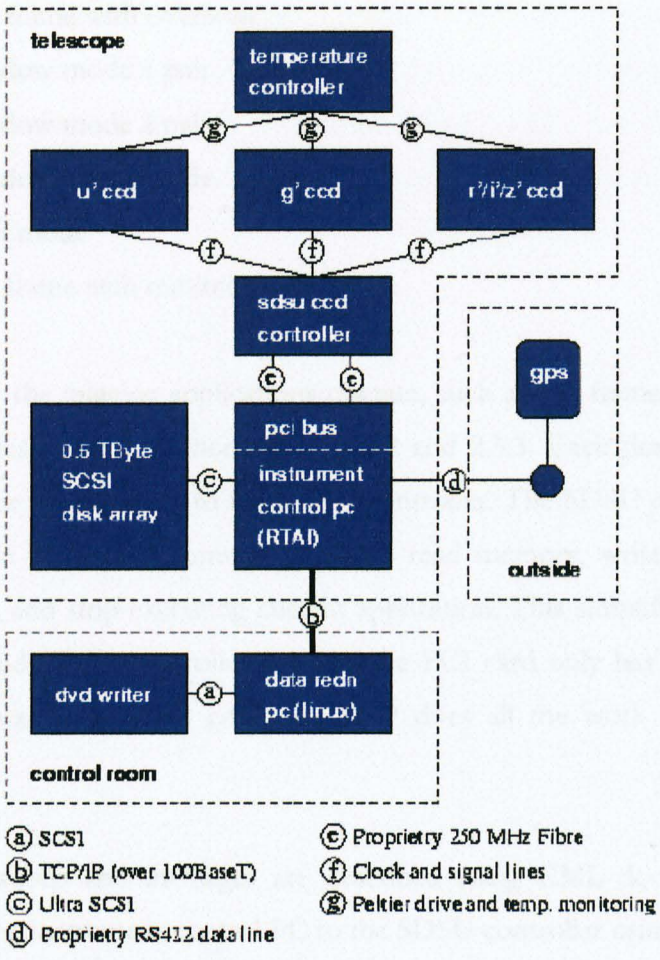

Figure 2.4.1: ULTRACAM hardware layout.

2.4.2 Camera control software

The philosophy of the software design has been to avoid the more standard SDSU camera control approach of having one large multi-purpose application downloaded into the SDSU memory, but instead to have several simple downloadable applications. Applications are basically programmes that give a specific function to the CCD, such as power on or operate in full frame mode, etc. Having a number of smaller, simpler applications makes the digital signal processing (DSP) software in which they are written more simple and easy to develop. The ULTRACAM applications are:

1. power on
2. power off
3. full frame

4. full frame with overscan
5. window mode 1 pair
6. window mode 2 pair
7. window mode 3 pair
8. drift mode
9. full frame with minimum dead-time

The details of how the imaging applications operate, such as full frame, window and drift mode have been described in sections 2.3.1, 2.3.2 and 2.3.3. Each time an application is required it has to be downloaded to the SDSU controller. The SDSU controller DSP and PCI DSP have the same five commands: reset, read memory, write memory, execute current application and stop executing current application. This simplifies communication between the PC and SDSU controller because the PCI card only has to relay these five commands and data. The SDSU controller DSP does all the work in configuring the CCDs.

The SDSU applications and messages are described using XML documents which are transmitted from the instrument control PC to the SDSU controller using an internet based interface. XML (Extensible Markup Language) is a standardised language, very similar to HTML, specially designed for transmitting data to web applications. For the applications, the XML document contains a link to the DSP code and the values of all the application parameters, such as exposure time, window positions, etc, so the camera controller can find them. The parameters are tested to see if they comply with allowed values before being executed. For example, if the parameters cause windows to overlap, an error message will be displayed.

The software architecture for ULTRACAM is shown in figure 2.4.2. An engineer creates several applications which are stored in the XML repository. A user selects the appropriate application from the XML repository through an interface called the sequencer. The application is downloaded through the camera control server where the XML document is decoded and the application code is passed to the SDSU controller. A second XML document is downloaded from the XML repository to the data handling server which prepares the PCI card for the data blocks that it is about to receive. The operator can then send a "GO" command (see figure 2.4.3) and the application will be executed. A

synchronisation pulse is sent by the SDSU controller when the exposure begins and the CCD frame is time-stamped (see section 2.4.3). At the end of the exposure, the image data is transferred via a shared data buffer (see figure 2.4.2) where it is combined with the information used to configure the cameras (exposure time, windows positions, etc) and is then archived.

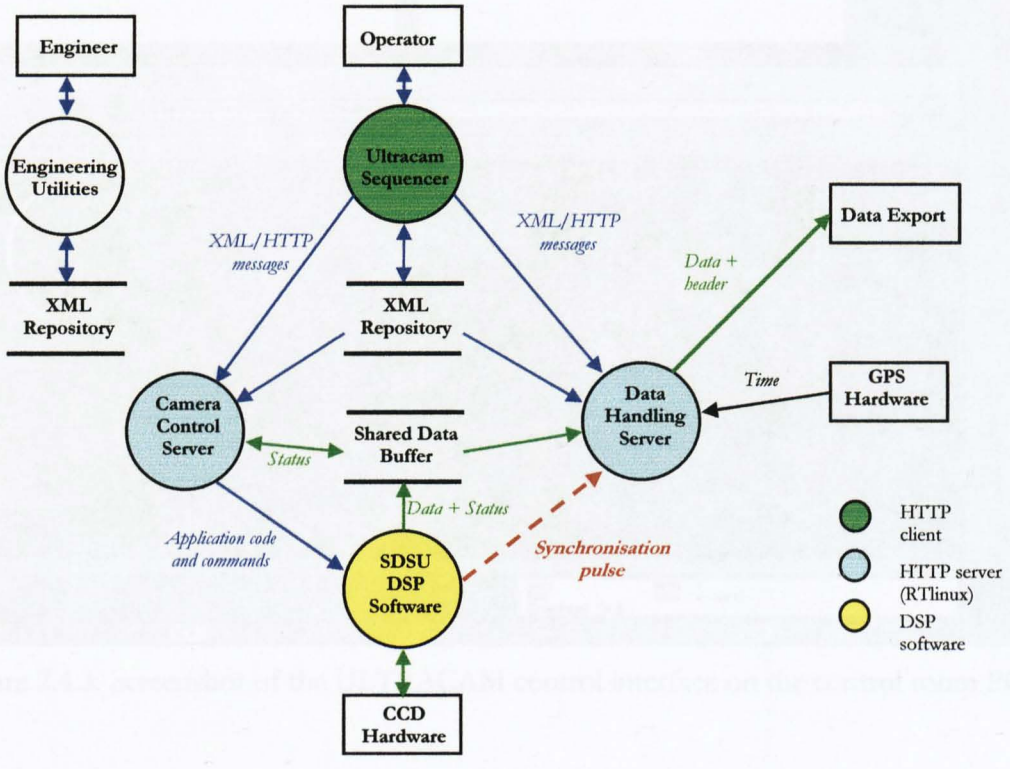

Figure 2.4.2: Architecture of ULTRACAM software (Beard 2002).

Figure 2.4.3 shows a screenshot of the ULTRACAM control interface that is displayed on the PC located in the telescope control room. The power on applications are shown under the heading 'engineering', whilst the CCD mode applications are shown under the 'generic modes' heading. A third column called 'user-defined modes' is used for observations where the window parameters are specifically set-up and stored. The applications are executed and stopped (if necessary) using the 'GO' and 'STOP' command.

ULTRACAM data is written to two files, an 'XML' file describing the CCD parameters and a 'DAT' file containing the raw data and timestamps, e.g. run022.xml and run022.dat for example. The DAT file can be very large (maximum of 2GB) if a target is observed for a long period of time. A file server on the instrument control PC allows data to be transferred across an ULTRACAM network provided for sole use by ULTRACAM and viewed whilst being recorded.

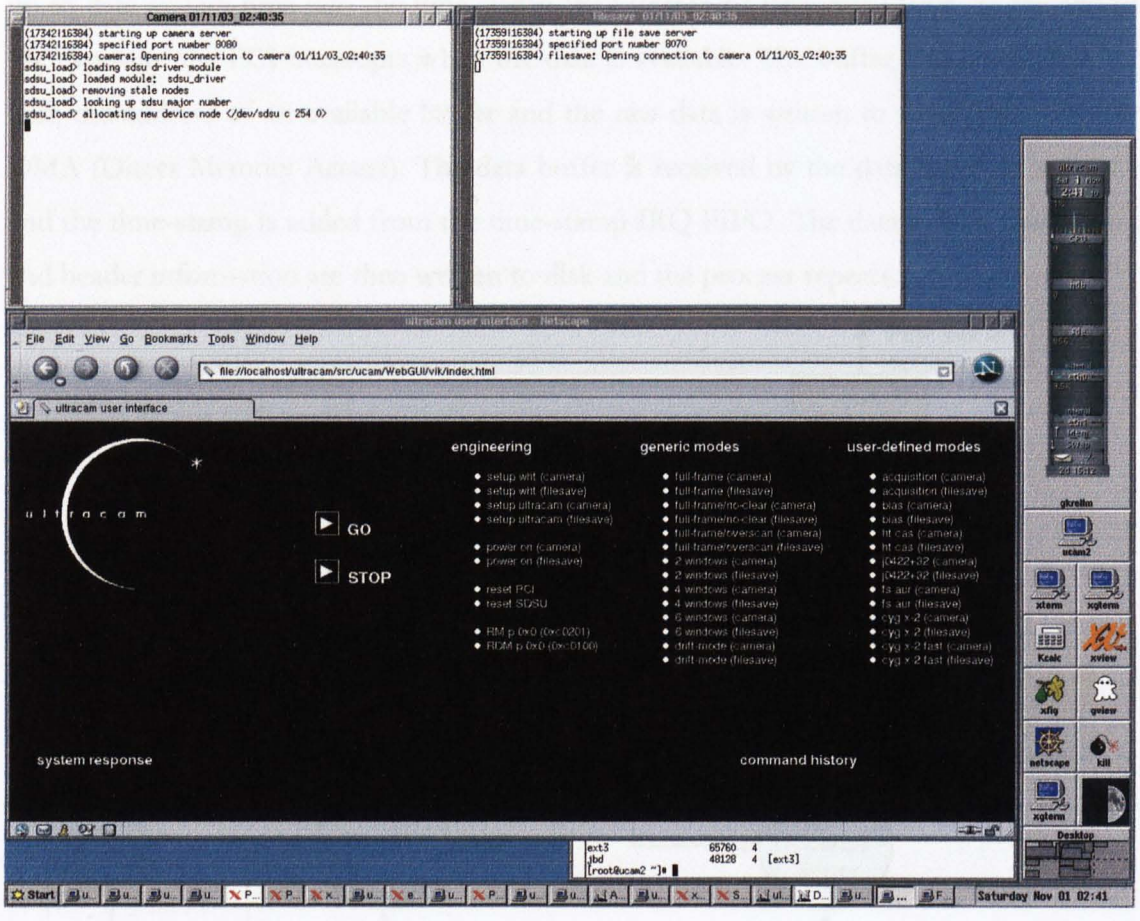

Figure 2.4.3: Screenshot of the ULTRACAM control interface on the control room PC.

2.4.3 Time stamping and real-time interfaces

Accurate time-stamping of each frame is a key requirement for ULTRACAM. The time signal from the GPS satellite system is received through an antenna, converted to UTC and then used to time-stamp each frame. The GPS thread reads the GPS time every 10 seconds, and uses it to update the instrument control PC motherboard clock. The GPS time stamps are then taken from the PC clock. Figure 2.4.4 shows the real time interfaces in the ULTRACAM control software.

The data handling and camera control software communicate with the buffer handler and camera interface through FIFOs (First In First Out). Commands and responses are passed between the SDSU PCI card and the camera control by the camera interface. The buffer handler FIFO manages the shared memory area by addressing each buffer item. When an exposure starts, a synchronisation pulse is sent by the SDSU controller which writes the

up-to-date system time into the time-stamp IRQ (Interrupt Request) FIFO. At the end of the exposure the PCI interrupts when the data is available. The buffer handler provides it with the address of an available buffer and the raw data is written to the PC memory by DMA (Direct Memory Access). The data buffer is received by the data handling software and the time-stamp is added from the time-stamp IRQ FIFO. The data buffer, time-stamp and header information are then written to disk and the process repeats.

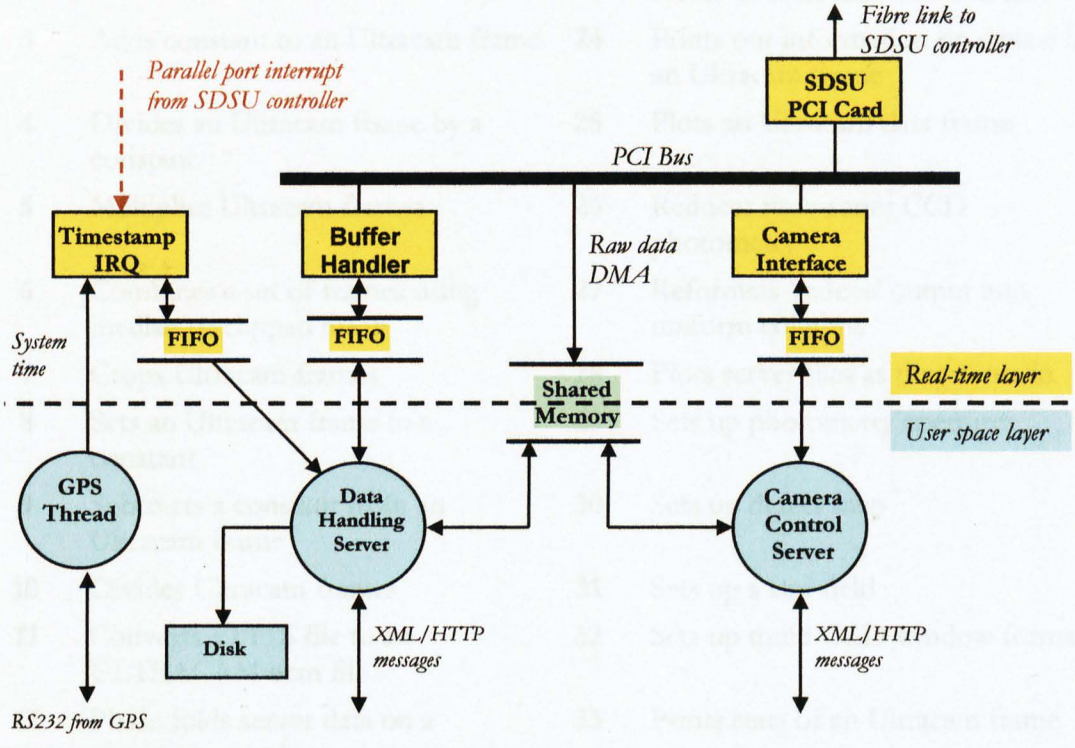

Figure 2.4.4: Real time interfaces with GPS time-stamping (Beard 2002).

2.4.4 Data reduction software

The data reduction software has been written by Tom Marsh at Warwick University (Marsh 2004) and is specially designed to handle high-speed, multi-CCD data such as that output by ULTRACAM. The software is installed on the data reduction PC in the telescope control room. The reduction software has been designed to fully reduce ULTRACAM data whilst it is being taken, providing the observer with a real-time view of the data so that the observing strategy can be revised if necessary. Pipeline data reduction is an important part of the ULTRACAM system since typically 10GB of data can be produced in a single night. The raw and reduced data are archived to DVDs after each night. The reduction software

enables the user to flat field, de-bias and dark-frame correct their data. Table 2.4 shows a list of all the routines the data reduction software can perform.

Table 2.4: Reduction software routines

	Routine		Routine
1	Adds Ultracam frames	22	Adds noise to an Ultracam frame
2	Adds a star field to a data frame	23	Prints one line of useful info from server or local .dat and .xml files
3	Adds constant to an Ultracam frame	24	Prints out information on a pixel in an Ultracam frame
4	Divides an Ultracam frame by a constant	25	Plots an Ultracam data frame
5	Multiplies Ultracam frames	26	Reduces time-series CCD photometry
6	Combines a set of frames using median or clipped mean	27	Reformats 'reduce' output into uniform columns
7	Crops Ultracam frames	28	Plots server files as they come in
8	Sets an Ultracam frame to a constant	29	Sets up photometry apertures
9	Subtracts a constant from an Ultracam frame	30	Sets up defect map
10	Divides Ultracam frames	31	Sets up a star field
11	Converts a FITS file to an ULTRACAM ucm file	32	Sets up multi-CCD window format
12	Phase folds server data on a specified period	33	Prints stats of an Ultracam frame
13	Grabs server or local .dat files, writing them as .ucm files	34	Subtracts Ultracam frames
14	Grabs server files, spits out FITS	35	Subtracts one frame from many
15	Edits the headers of Ultracam frames	36	Equivalent of 'cat' for server .xml files
16	Lists headers of Ultracam frames	37	Converts an ULTRACAM ucm file to FITS
17	Plots histograms of Ultracam frame	38	Prints information on an Ultracam frame
18	Generates a log of a set of runs using oneline	39	Creates a blank Ultracam frame
19	Plots an ultracam image plus light curve	40	Equivalent of 'ls' for server files
20	Multiplies Ultracam frames	41	Sets/splices Ultracam frames
21	Analyses noise in Ultracam frames	42	Joins all windows into 1 large one

```
1 52415.22888267527 1 8 0.1 1 1 400.481 549.525 585129 817.889 41.8486 33 0 2 741.002 328.261 6919.92 160.396 36.7627 10 0
3 915.657 326.149 8737.88 165.643 36.4692 11 0
1 52415.22888267527 1 8 0.1 2 1 392.991 547.831 550083 857.978 112.711 23 0 2 736.957 325.126 3902.71 197.085 106.089 12 0
3 913.303 323.569 5286.88 200.65 105.967 16 0
1 52415.22888267527 1 8 0.1 3 1 393.04 549.044 18133.6 194.349 8.02995 8 0
2 52415.22973040854 1 8 0.1 1 1 400.14 550.888 583339 816.616 41.7207 32 0 2 741.405 329.467 7295.71 162.236 37.2167 9 0 3
915.774 326.278 8845.91 166.722 37.0962 13 0
2 52415.22973040854 1 8 0.1 2 1 392.585 548.431 548366 857.586 112.972 31 0 2 736.984 325.794 3961.63 198.037 107.322 7 0
3 913.263 323.777 5421.64 201.762 107.348 10 0
2 52415.22973040854 1 8 0.1 3 1 392.745 550.292 18178.1 195.539 8.54596 12 0
```

Figure 2.4.5: Example of a reduction log file output by the ULTRACAM pipeline data reduction system. The first line has been annotated, indicating the meaning of each parameter in the log file. The CCD numbers 1, 2, 3 correspond to red, green and blue, respectively.

The software allows variable and fixed-sized apertures for target and comparison stars to be defined with surrounding annuli for measuring the sky level. Both normal and optimal extraction methods can be selected. The resulting light-curve can be displayed in real-time, along with the raw images with the aperture overlaid. The reduced data is written to a log file of the format shown in figure 2.4.5.

In the reduction log file there are three lines per frame, one for each colour CCD. The CCDs are numbered 1, 2 and 3, corresponding to the red, green and blue CCDs respectively. The basic data for each frame gives the time in Mean Julian Date (MJD), the number of GPS satellites visible to the GPS antenna, the CCD number and the time flag. The time flag indicates whether the quoted time is reliable or not (1 = reliable, 0 = unreliable). The exposure delay (in seconds) is the time input into the CCD configuration file and is not the actual exposure time (see section 2.3). Data for each aperture is then listed. The aperture data includes target counts, x and y position of the star, sky counts and target count uncertainty.

2.5 Mechanical design

In this section the mechanical design and integration of optics, CCD heads and electronics will be described. This constitutes the majority of my effort during the design stage of the ULTRACAM project. The vast majority of the mechanical components of ULTRACAM were manufactured by staff in the Central Mechanical Workshop (CMW) at the University of Sheffield.

All of the mechanical design work was performed using AutoDesk Mechanical Desktop release 5. It is 3-dimensional design software that is optimised for such work because it allows parts to be built component by component. Each designed component is editable, so minor changes during the design process are relatively simple to perform. The software also allows 3-dimensional viewing of the design from any angle, making module integration straightforward.

2.5.1 Mechanical design requirements

ULTRACAM has to fit on several different-sized telescopes. The smallest aperture telescope is the Aristarchos telescope (2.3 m). The size and mass limits of instruments on this telescope have therefore been used to define the size and mass limits of ULTRACAM, as larger telescopes generally have larger size and mass envelopes. Using this as a basis, a list of key mechanical requirements was defined. The mechanical chassis should:

1. Provide a platform on which to mount the optics, CCD heads and SDSU controller
2. Allow easy access to the optics and CCD heads for alignment, filter changes and vacuum pumping
3. Provide alignment mechanisms for the optics and CCD heads
4. Weigh less than 100 kg
5. Have an overall length of less than 850 mm
6. Exhibit low thermal expansion
7. Exhibit low flexure at any orientation
8. Be electrically and thermally isolated from the telescope

The requirement for low thermal expansion derives from the optical design. The re-imaging cameras have a fast focal-ratio of 2.3, meaning that any slight movement in the relative positions of the cameras and the CCD heads will be immediately detectable as a loss of image quality. Thermal expansion in the opto-mechanical chassis onto which the optics and CCDs are mounted may cause such a relative movement and therefore, materials which have low thermal expansion coefficients have been used in the most critical areas (see section 2.5.3).

The requirement for low flexure arises from the fact that each CCD has to be aligned such that a star falls on the same pixel on each CCD. This is necessary because when a set of small windows are defined in drift mode, the target needs to remain within the same windows on all three CCD as the telescope tracks the target over the sky.

2.5.2 Opto-mechanical chassis

The mechanical design was split into two major parts: the optics hull, which integrates the optics and the opto-mechanical chassis which forms a rigid structure to which all other modules interface. This section will describe the opto-mechanical chassis. The type of structure chosen for ULTRACAM is known as a double octopod and is based on Serrurier trusses (see figure 2.5.1). It provides a strong open framework meeting all of the mechanical requirements defined in section 2.5.1. The Serrurier truss is a tried-and-tested rigid, mechanical arrangement of beams which has been employed extensively in telescope design, such as the structures used to mount secondary mirrors in Cassegrain telescopes.

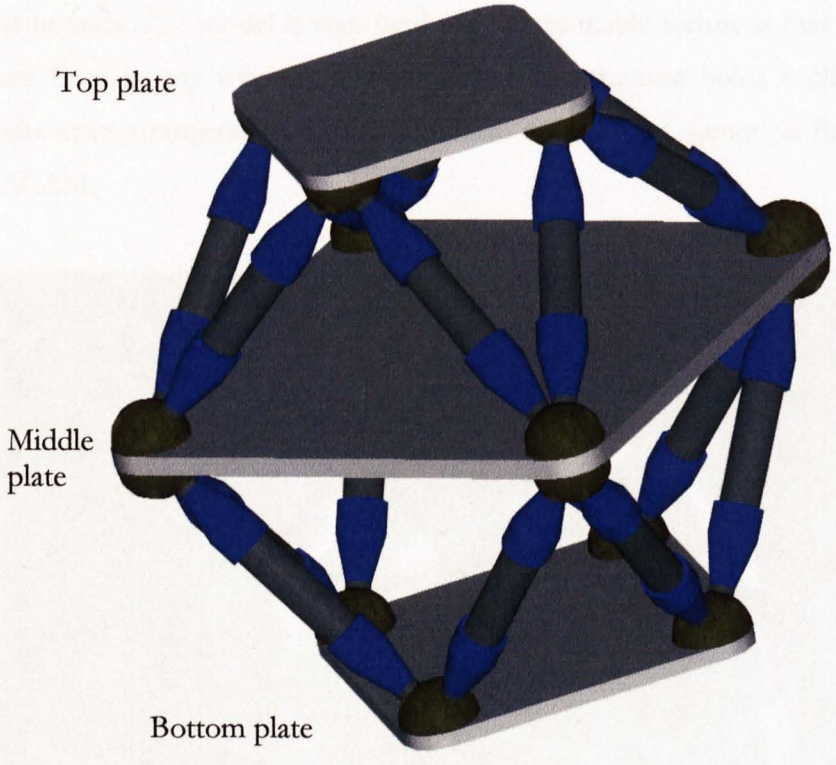

Figure 2.5.1: The double octopod structure used by ULTRACAM for its opto-mechanical chassis. The octopod comprises of three parallel plates onto which modules can be attached, rather like an optical bench. The top plate forms the telescope interface.

Each parallel plate in figure 2.5.1 is made from aluminium tool plate onto which hemispherical nodes are bolted. Each hemispherical node is machined aluminium into which the struts are fixed by a single bolt per interface. Between each strut and node there is a special washer with one hemispherical face to match the radius of the nodes. The three plates are parallel which makes aligning the optics to the double octopod and interfacing

the instrument to telescopes straightforward. The thicknesses of the special washers were varied to increase and decrease the effective length of individual struts, allowing enough correction to make all three plates parallel. The struts are composite parts of two aluminium strut ends glued and pinned to a strut tube. The strut tube material is carbon fibre and will be discussed in section 2.5.3. All of the aluminium components are anodised to protect them from the environment.

A quarter-size model of the ULTRACAM double octopod was constructed in the CMW to provide some idea of the structural rigidity of the design. Figure 2.5.2 shows a picture of the model. The model was also very useful to the engineers who had to build the full-size octopod. They could test and practice manufacturing techniques for producing parts such as the strut ends. The model is very rigid and a remarkable feature is that none of the strut tubes are fixed in any way to the strut ends. The structure holds itself together by the composite strut arrangement. The next section outlines the numerical flexure analysis for ULTRACAM.

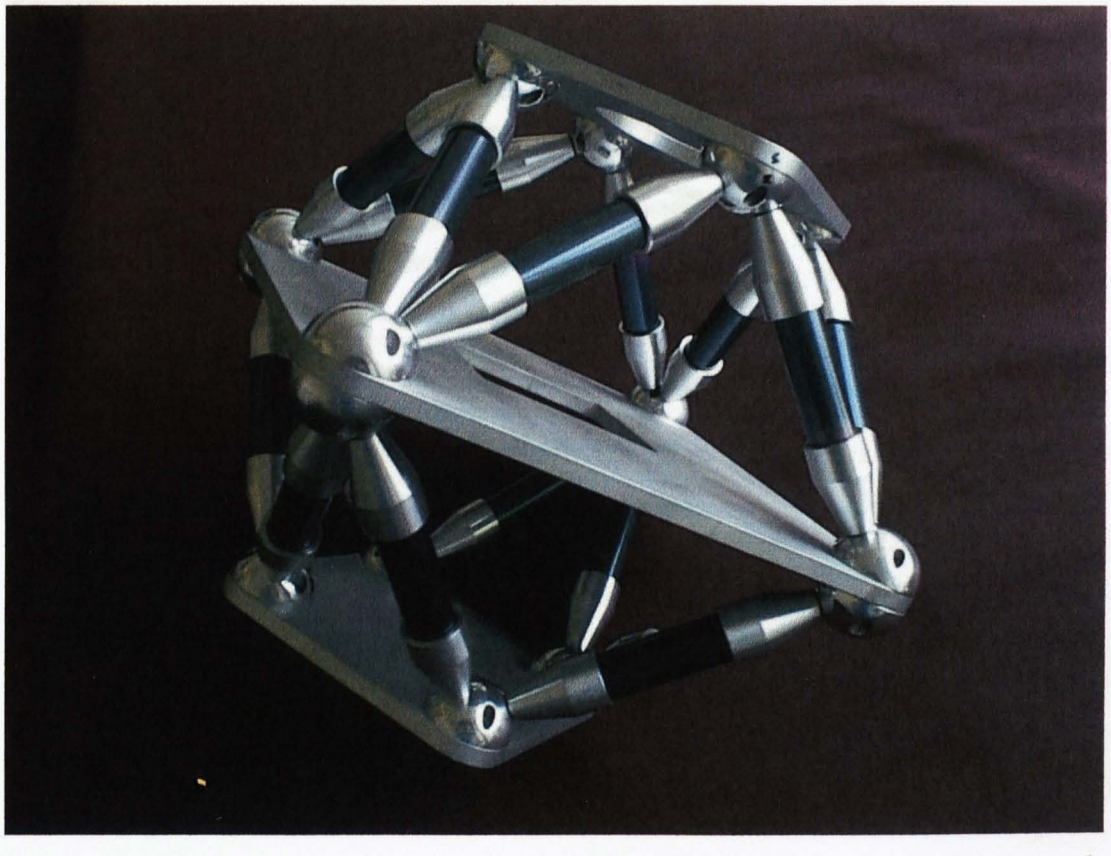

Figure 2.5.2: Photograph of the quarter-size model of the ULTRACAM double octopod.

2.5.3 Flexure

The impetus for the flexure analysis was to see how rigid the double octopod would be, determine what load would cause the double octopod to fail and determine an appropriate material for the strut tubes.

ULTRACAM will be tilted through various orientations as the telescope points to different parts of the sky. As ULTRACAM moves all of the struts in the double octopod share the instrument load. The effect on each triangular strut section is to put one strut under compressive forces and the other under tensile forces. The forces are effectively trying to change the length of each strut, not bend them and this requires great force. This is what makes the serrurier truss arrangement so strong and rigid. However, overloading in compression will eventually result in the struts buckling. Figure 2.5.3 shows the simplest case for calculating flexure of a Serrurier truss.

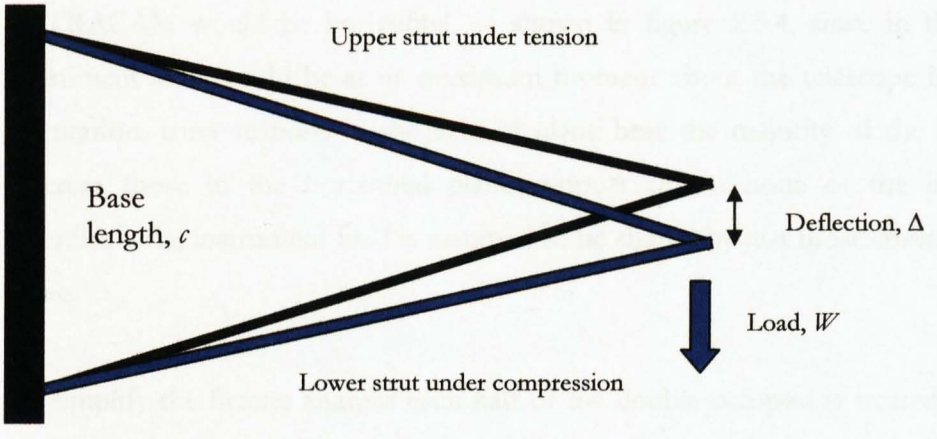

Figure 2.5.3: Loading of a Serrurier-truss section illustrating tensile and compressive forces.

Flexure of the double octopod has been calculated using engineering formulae (Sinnott 1994). The following equations and analysis assume that all joints in the structure are rigid. Flexure measured as deflection Δ, at the apex of the truss and is calculated by:

$$\Delta = \frac{WL^3}{c^2 AE}, \qquad (2.7)$$

where W is the total loading force, L is the truss tube length, c is the base length of the triangle formed by the truss, A is the cross-sectional area of the tube and E is the modulus of elasticity. A solution for the buckling, P, of long columns is:

$$P = \frac{\pi^2 EI}{L^2}, \tag{2.8}$$

where I is the moment of inertia of a tube with inside diameter d_I and outside diameter d_o is given by:

$$I = \frac{\pi\left(d_o^4 - d_I^4\right)}{64}. \tag{2.9}$$

The example in figure 2.5.3 is for one truss section, but the equations can be applied to an arrangement of trusses such as those in ULTRACAM. The worst orientation for ULTRACAM would be horizontal, as shown in figure 2.5.4, since in this position the instrument load would be at its maximum moment about the telescope interface. In this orientation, truss sections in the vertical plane bear the majority of the instrument load whereas those in the horizontal plane support almost none of the instrument load. Therefore, the instrument load is assumed to be shared by just those trusses in the vertical plane.

To simplify the flexure analysis each half of the double octopod is treated separately. The triangular sections used for the analysis of each half are shown in figure 2.5.4. Section 1 is the telescope interface and is loaded with the whole instrument weight. Section 2 is the lower half of the double octopod and is loaded with just a CCD head and the SDSU controller (which mounts to the underside of the bottom plate – see figure 2.5.19). In reality both sections are rotated 45 degrees in the horizontal plane with respect to each other so only one section is in its worst orientation at any one time. To produce the worst flexure section 1 should be in the worst orientation because it bears the whole instrument load and the triangular section has the smallest base length, c, which is a dominant factor in equation 2.7. Section 2 will have its load shared by four truss sections, whereas the load on section 1 will be shared by just two truss sections.

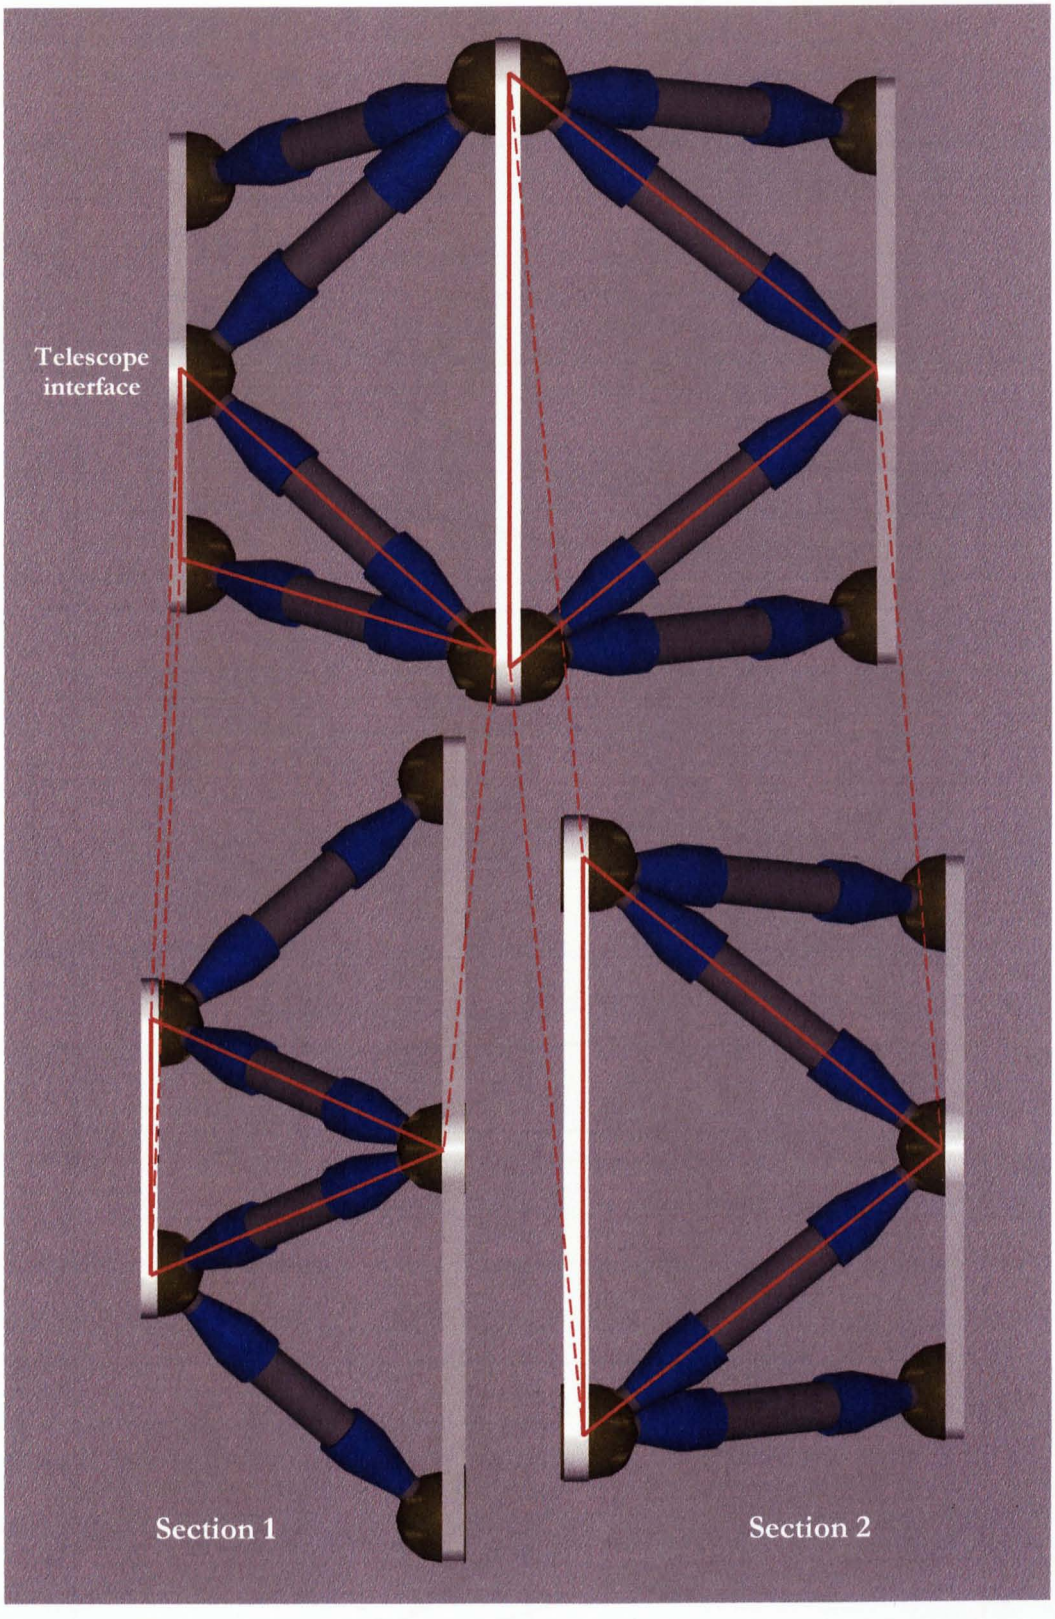

Figure 2.5.4: The flexure analysis of ULTRACAM simplifies when considering the double octopod horizontally and as two separate structures.

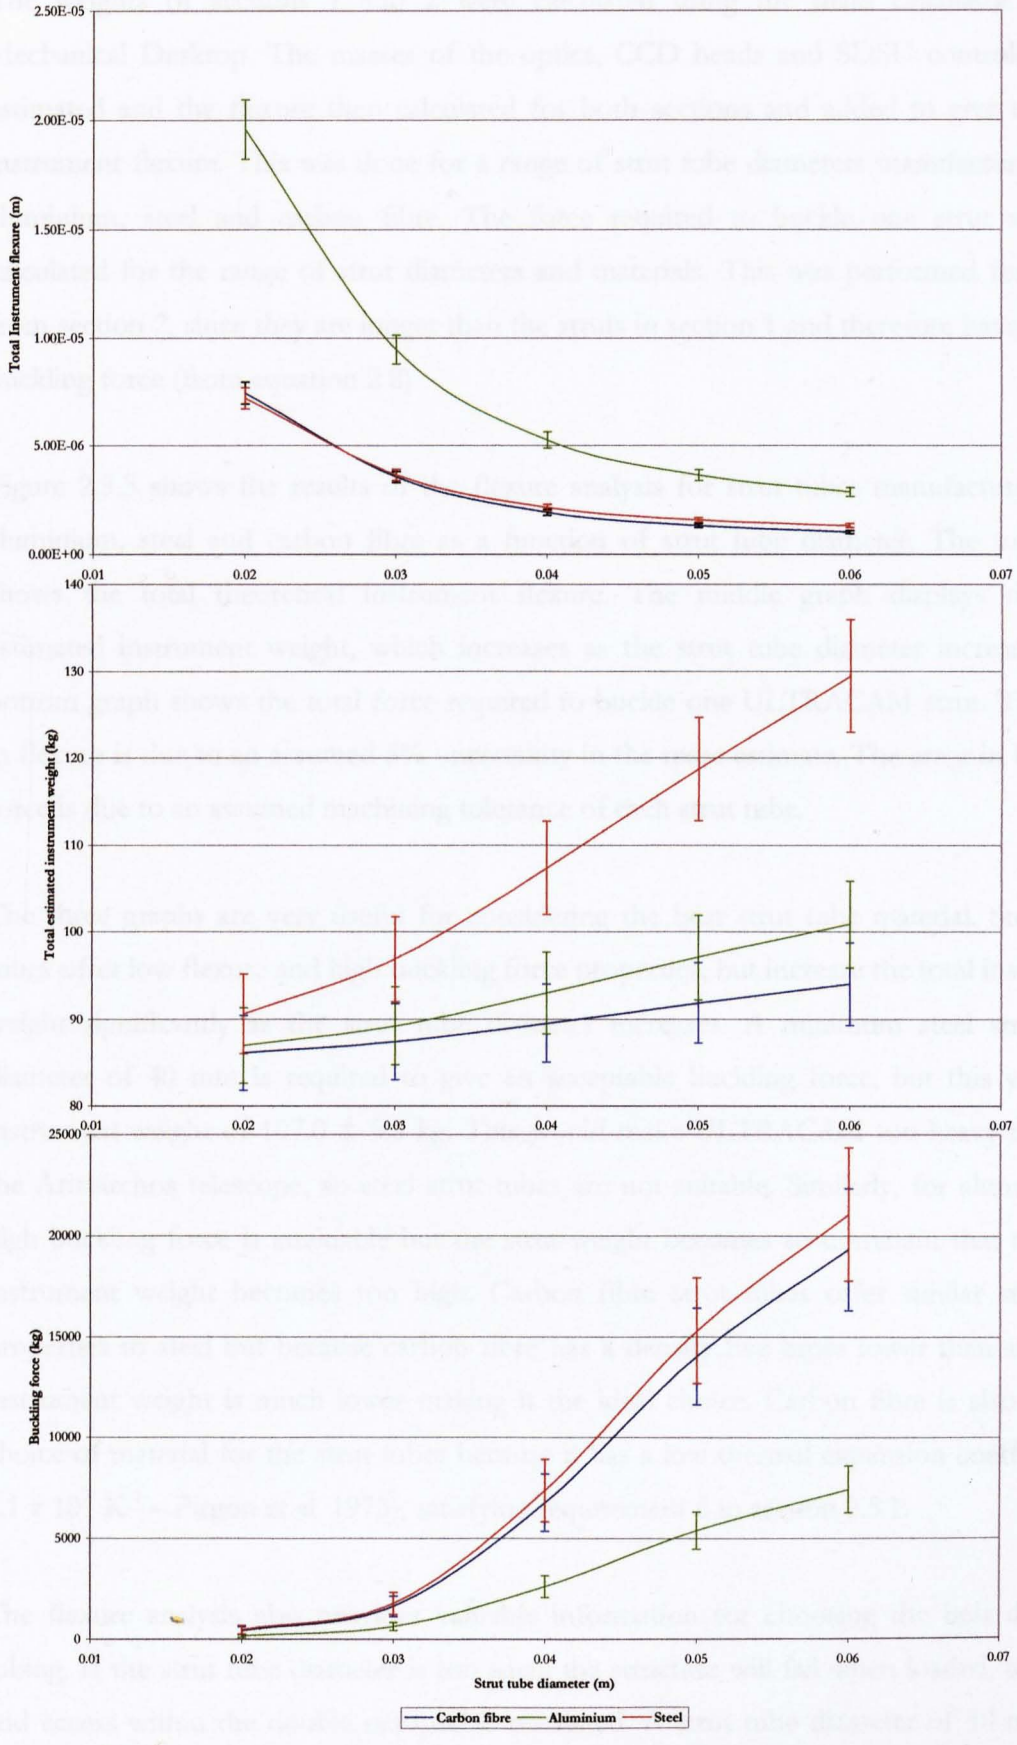

Figure 2.5.5: Theoretical models of flexure (top), total instrument weight as a function of strut tube diameter (middle), and force required to buckle an ULTRACAM strut (bottom).

The weights of sections 1 and 2 were calculated using the mass calculator tool in Mechanical Desktop. The masses of the optics, CCD heads and SDSU controller were estimated and the flexure then calculated for both sections and added to give the total instrument flexure. This was done for a range of strut tube diameters manufactured from aluminium, steel and carbon fibre. The force required to buckle one strut was also calculated for the range of strut diameters and materials. This was performed for a strut from section 2, since they are longer than the struts in section 1 and therefore have a lower buckling force (from equation 2.8)

Figure 2.5.5 shows the results of the flexure analysis for strut tubes manufactured from aluminium, steel and carbon fibre as a function of strut tube diameter. The top graph shows the total theoretical instrument flexure. The middle graph displays the total estimated instrument weight, which increases as the strut tube diameter increases. The bottom graph shows the total force required to buckle one ULTRACAM strut. The error in flexure is due to an assumed 5% uncertainty in the mass estimate. The error in buckling force is due to an assumed machining tolerance of each strut tube.

The three graphs are very useful for considering the best strut tube material. Steel strut tubes offer low flexure and high buckling force properties, but increase the total instrument weight significantly as the strut tube diameter increases. A minimum steel strut tube diameter of 40 mm is required to give an acceptable buckling force, but this yields an instrument weight of 107.0 ± 5.3 kg. This would make ULTRACAM too heavy to fit on the Aristarchos telescope, so steel strut tubes are not suitable. Similarly, for aluminium a high buckling force is attainable but the strut weight becomes so dominant that the total instrument weight becomes too high. Carbon fibre strut tubes offer similar structural properties to steel but because carbon fibre has a density five times lower than steel, the instrument weight is much lower making it the ideal choice. Carbon fibre is also a good choice of material for the strut tubes because it has a low thermal expansion coefficient (-7.1×10^{-7} K^{-1} – Pirgon et al. 1973), satisfying requirement 6 in section 2.5.1.

The flexure analysis also provides valuable information for choosing the best diameter tubing. If the strut tube diameter is too small the structure will fail when loaded, too large and access within the double octopod is restricted. A strut tube diameter of 40 mm was chosen since the theoretical flexure is calculated to be **$(1.94 \pm 0.14) \times 10^{-6}$ m** (c.f. pixel size

of 13 x 10^{-6} m), the total instrument weight is **89.5 ± 4.5 kg** and the minimum buckling force is **6763 ± 1413 kg**. Increasing the strut tube diameter beyond 40 mm provides little improvement in flexure and only increases the instrument weight. Taking the lower limit of buckling force for a carbon fibre strut of 40 mm diameter and then 50% of this value to compensate for the composite strut construction gives a buckling force of **2675 kg**. This is approximately 26 times the total estimated instrument weight and provides a reasonable safety margin.

2.5.4 Optics hull

The optics hull integrates all of the optics and is mounted on the double octopod. A rigid structure is required that forms a sealed unit to protect the optics from dust and other contaminants as well as preventing accidental movement of the aligned optics. The design of the optics hull was driven by the alignment procedure, which therefore had first to be determined.

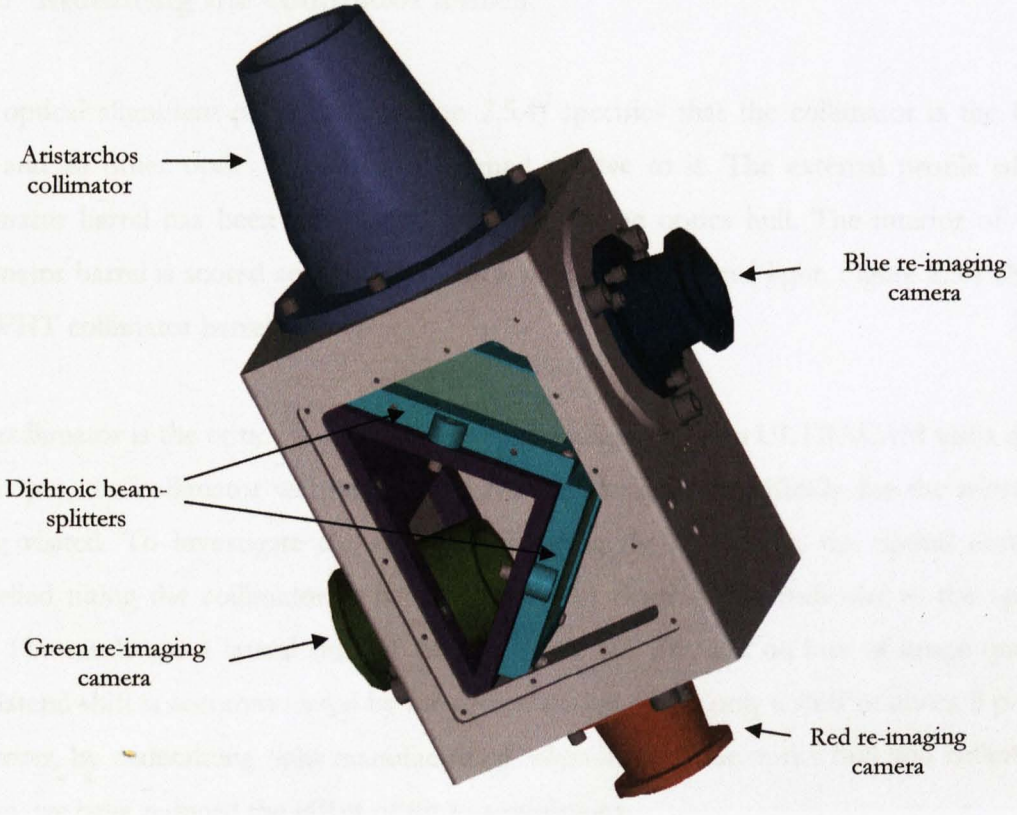

Figure 2.5.6: Optics hull with the Aristarchos collimator fitted and large access panel removed. The collimator is the first unit fixed. Both dichroic beam-splitters fit inside the hull and the blue, green and red re-imaging cameras are fixed to the outside.

The optical alignment procedure is as follows:

1. Fix collimator and keep as the reference optics
2. Align dichroic beam-splitter 1
3. Align dichroic beam-splitter 2
4. Align and focus all three re-imaging cameras, no order of preference
5. Align and focus the CCD heads, no order of preference

The optics hull is a box constructed using aluminium box section with plates at the top and bottom, as shown in figure 2.5.6. The top of the optics hull is designed to accommodate the collimating fore-optics for each telescope. Inside the hull are both dichroic beam-splitters, which can be accessed through a large removable panel. The three re-imaging cameras fix to the outside of the hull by adjustable mounts. The interior of the hull is bead blasted and the whole assembly anodised black to reduce scattered light.

2.5.5 Mounting the collimator lenses

The optical alignment procedure (section 2.5.4) specifies that the collimator is the fixed unit and all other optics elements are aligned relative to it. The external profile of the collimator barrel has been designed to integrate to the optics hull. The interior of each collimator barrel is scored and anodised black to reduce scattered light. Figure 2.5.7 shows the WHT collimator barrel.

The collimator is the optics to which everything is aligned. When ULTRACAM visits other telescopes, the collimator will be changed for one designed specifically for the telescope being visited. To investigate the effect of changing the collimator, the optical designer modelled tilting the collimator in two planes by 0.1 degrees perpendicular to the optical axis. The result was a lateral shift of approximately 100 μm and no loss of image quality. The lateral shift is accommodated by the optical design and is only a shift of about 8 pixels. However, by maintaining tight manufacturing tolerances on the optics hull and collimator design, we have reduced the effect of tilt to a minimum.

Figure 2.5.7: The WHT collimator optics barrel.

2.5.6 Dichroic beam-splitter mount

Each dichroic substrate is mounted into a frame to make handling easier and allow positional adjustments during the optical alignment. Figure 2.5.8 shows a dichroic substrate frame with retaining clamps. The dichroic substrates are very flat pieces of glass and need to remain flat since warping of the substrate due to poor mounts will affect image quality. Great care has been taken during the design of the dichroic frame to ensure the substrate is not put under stress by its mount. Each dichroic frame is machined from a single block of aluminium tool plate and then heat treated to relax the stresses introduced by removing large sections of material.

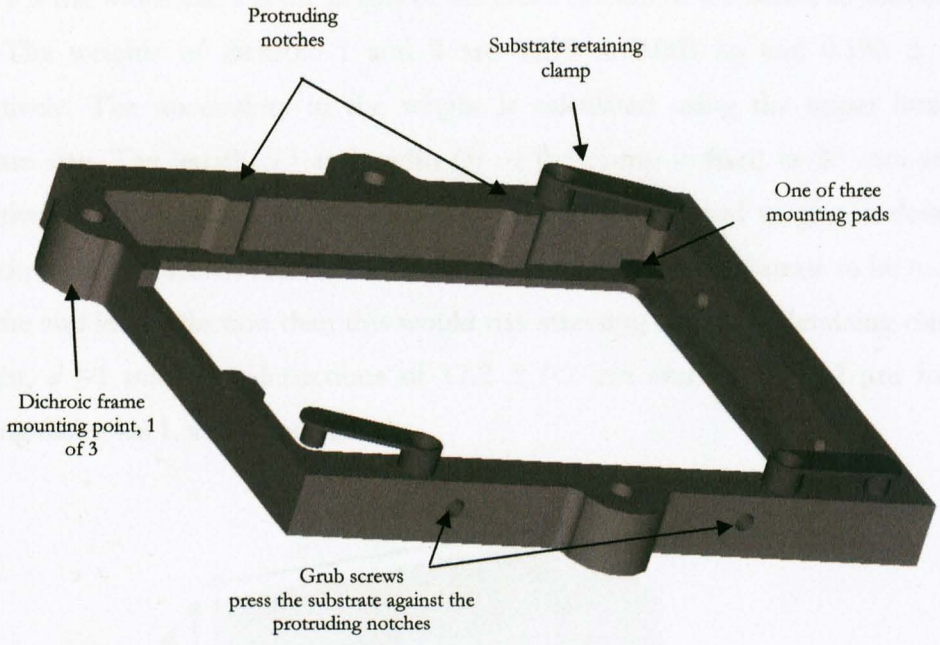

Figure 2.5.8: Dichroic substrate frame. The dichroic beam-splitter sits in the frame and is retained laterally by grub screws and vertically by three clamps.

The substrate sits in the dichroic frame on three shim pads which allow for tilt, rotation and focus adjustment. The shim pads are made from peelable shim so the thickness of each pad can be adjusted to remove any warping in the dichroic frame. Directly above each shim pad is a retaining clamp that holds the substrate in the frame. It is important to retain the substrate directly above the shim pads so as not to warp the glass. The retaining clamps were designed to not stress the glass, but to provide enough retaining force. For this part of the design I considered the retaining clamps as a bending beam supported at one end. The deflection of the beam is given by (Roark, 1975):

$$y = \frac{Wl^3}{3EI}, \tag{2.10}$$

where W is the load, l is the length of the bending beam, E is the modulus of elasticity (73.1 GPa for aluminium) and I is the area moment of inertia about the centroidal axis of the beam cross-section given by (Roark, 1975):

$$I = \frac{1}{12} bd^3, \tag{2.11}$$

where b is the width and d is the height of the cross-section of the beam, as shown in figure 2.5.9. The weights of dichroic 1 and 2 are 0.279 ± 0.003 kg and 0.183 ± 0.003 kg respectively. The uncertainty in the weight is calculated using the upper limit on the substrate size. The length (l) and width (b) of the clamp is fixed as 30 mm and 8 mm respectively. The height, d of the retaining clamps was defined to give a deflection of approximately 10 microns since more than this would cause the substrate to be too loose in its frame and less deflection than this would risk stressing the glass. Retaining clamps with a height, d =1 mm give deflections of 17.2 ± 0.2 μm and 12.3 ± 0.2 μm for clamps retaining dichroics 1, and 2 respectively.

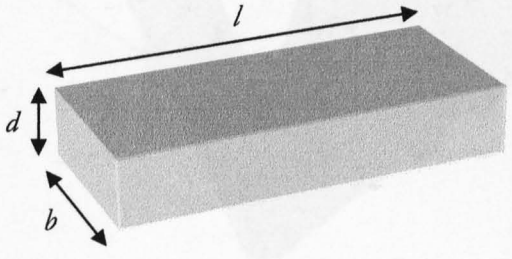

Figure 2.5.9: Diagram of the cross-section dimensions for the calculation of area moment of inertia about the centroidal axis of the beam cross-section

Lateral movement of the substrate is fixed by special grub screws that press two sides of the substrate to the frame. The special grub screws have a rubber tip which contacts the substrate and prevents damage to the glass. The rubber tip of the grub screw also accommodates any thermal expansion and contraction of the frame and substrate. The other two sides of the substrate are seated against protruding notches located directly opposite each grub screw (see figure 2.5.8). The protruding notches reduce the effects of any distortions in the frame edge by making contact with the substrate at two relatively small surface area points.

Each dichroic frame is fixed to the dichroic mount (shown in figure 2.5.10) by three bolts, which are adjusted by placing shims under each mount point (called a three-point mount). There is no lateral adjustment because the dichroic clear apertures are over-sized by 5 mm. The dichroic mount fits inside the optics hull and is mounted using another three-point mount using shim washers to provide lateral adjustment. The dichroic frames and mount are bead blasted and anodised black to reduce scattered light

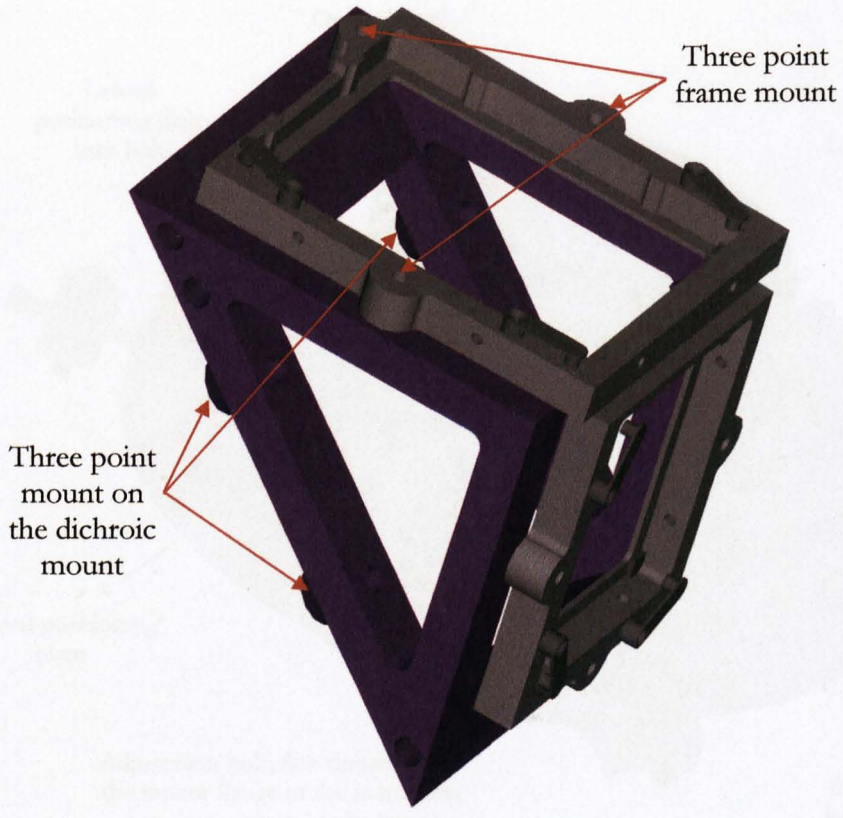

Figure 2.5.10: Dichroic mount including both dichroic frames.

2.5.7 Mounting of the re-imaging cameras

The lenses of each re-imaging camera are mounted into aluminium optics barrels which are anodised black and scored on the inside. Each camera requires tilt, rotation, focus and lateral adjustment to optically align them. Each camera fits to the outside of the optics hull by the type of mount shown in figure 2.5.11.

A lateral positioning plate is fixed to the outside of the optics hull by three bolts. The bolt holes are oversized, allowing the plate to be moved laterally by three lateral positioning blocks on the hull surface. These positioning blocks are simply aluminium blocks which temporarily fix to the optics hull surface and have a single fine threaded (M6 x 0.25) bolt lying parallel to the hull surface which pushes the lateral positioning plate. All three positioning blocks contact the lateral plate at the same time and are each incrementally adjusted to position the plate accurately.

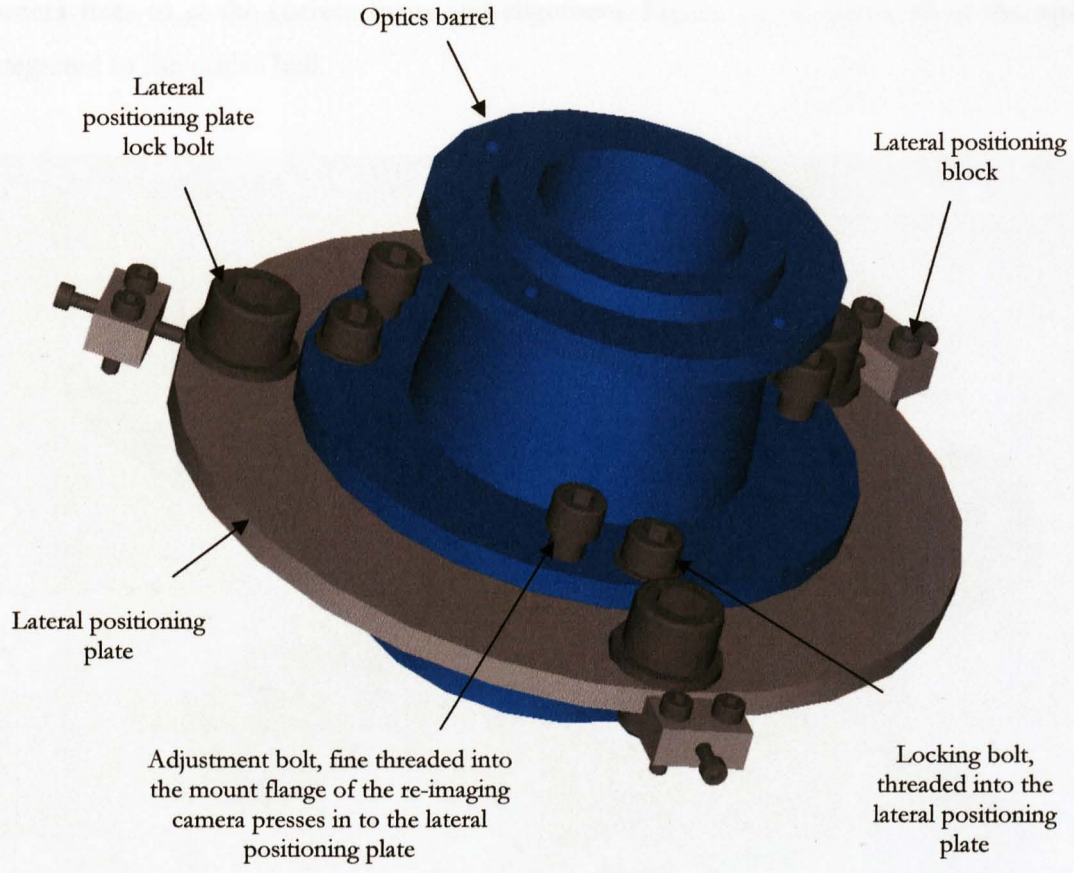

Figure 2.5.11: Re-imaging camera mount. The lateral positioning plate bolts to the optics hull and aligns the re-imaging camera to the correct lateral position. A three-point mount on the re-imaging camera controls tilt, rotation and focus to accurately align the camera.

The re-imaging camera has two external flanges, one is the barrel mounting flange and the other is the filter holder flange. The barrel mounting flange forms the interface with the lateral positioning plate. Here another three-point mount is used, but it is self adjustable and does not require shims. Each point of this three-point mount has two bolts, one that is threaded through the re-imaging camera mount flange and pressing against the lateral positioning plate, and another that passes through the re-imaging camera mount flange and is threaded into the lateral positioning plate. The former bolt positions the camera away from the lateral plate and the other locks the camera in position. Using this arrangement each mount point can be adjusted to change tilt and all three can be adjusted equally to change focus.

When in the correct position the separation of the re-imaging camera mount flange from the optics hull surface was measured in several places around the circumference. The lateral positioning plate was then replaced with a specially machined plate that the re-imaging

camera fixes to at the correct focus and alignment. Figure 2.5.12 shows all of the optics integrated to the optics hull.

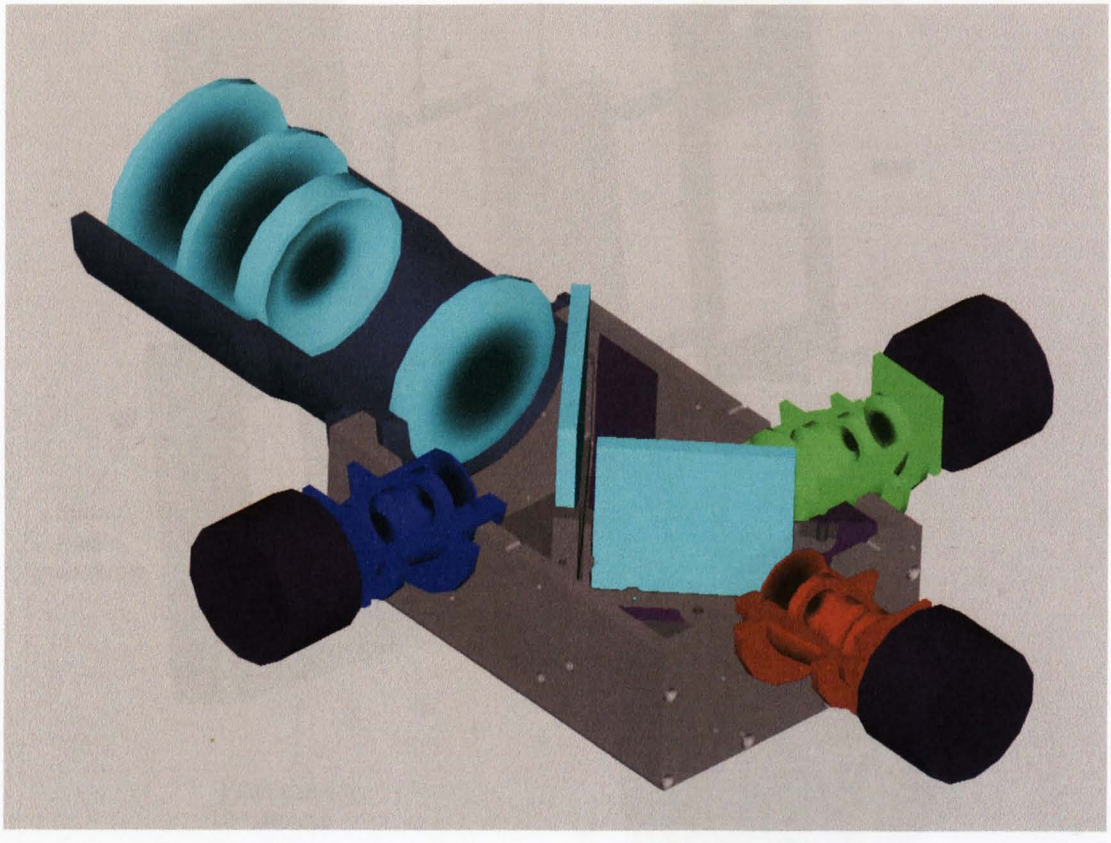

Figure 2.5.12: Cutaway image of the ULTRACAM optics hull. The collimating fore-optics, dichroic beam-splitters, re-imaging cameras and filters are shown. The CCDs are represented by the purple cylinders.

2.5.8 Mounting of the filters

An interchangeable cartridge system has been adopted to hold the filters, as shown in figure 2.5.13. The cartridges slide into a cartridge holder which bolts onto the end of each re-imaging camera. There is very little space between the end of each re-imaging camera and its respective CCD head, so the filter holders have to be as slim as possible. Each cartridge is designed to hold one 50 mm diameter/square filter of up to 5 mm thickness, but there is the option to use filters up to 10 mm thick. Such a filter will however, require a cartridge with enough depth to hold it (these have not been manufactured yet). The only problem with using filters of a different thickness is that the position of the focal plane will alter.

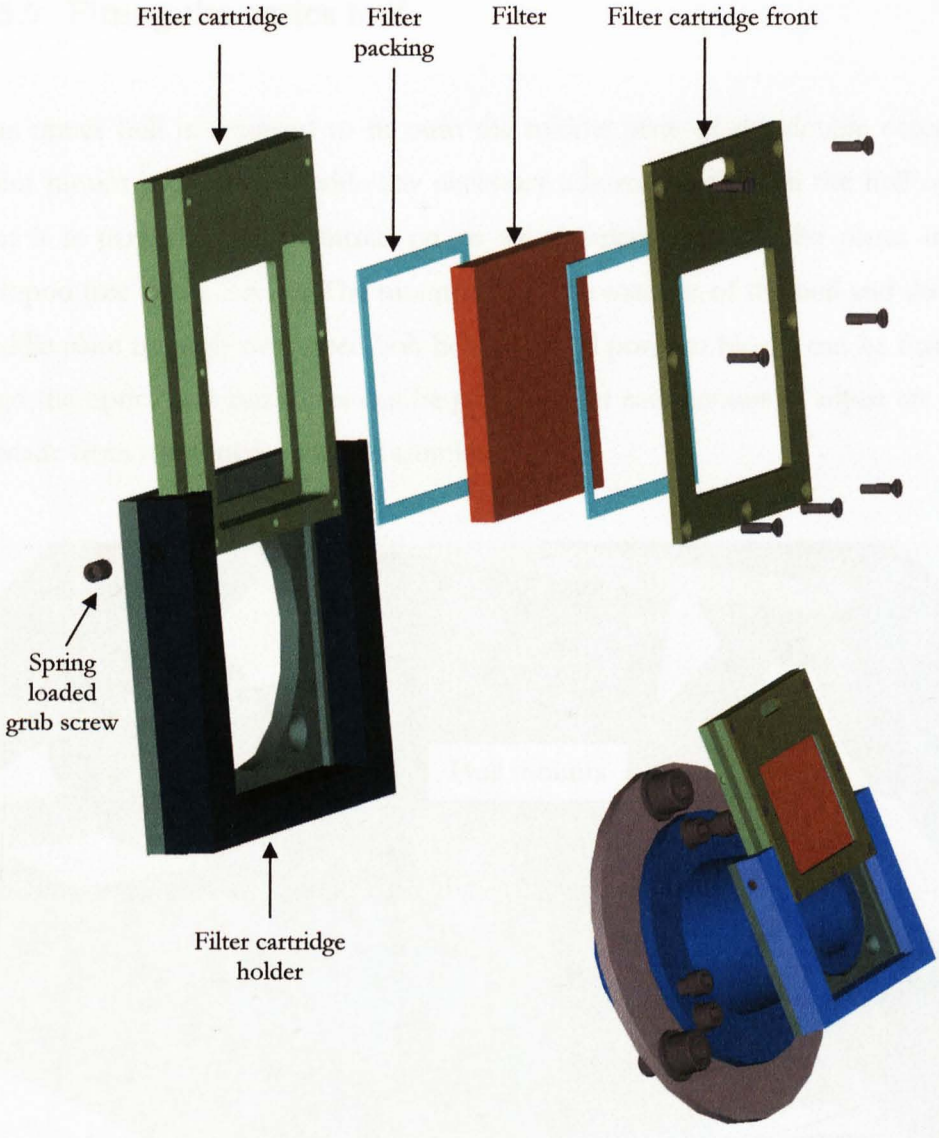

Figure 2.5.13: The interchangeable filter system. Filters are pre-mounted into cartridges which slide into a holder that bolts to the end of the re-imaging cameras.

The filters have to return to the same position each time they are inserted and they must not be tilted with respect to the optical axis. The filter cartridges slide into the filter holders on PTFE rails which provide a smooth sliding surface. A spring-loaded ball-bearing in the holder makes contact with a cut-out in the cartridge to prevent it falling out if tilted upside down. A foam baffle fits between each filter cartridge and CCD head window to prevent light leaks. It remains in place by simply expanding against the CCD head and filter cartridge. A steady flow of dry air or nitrogen is pumped into the baffle cavity using hypodermic needles to prevent condensation forming on the CCD head window.

2.5.9 Fitting the optics hull

The optics hull is designed to fit onto the middle plate of the double octopod. A three point mount is used to provide any necessary adjustment to align the hull optical axis so that it is parallel to and centred on an axis perpendicular to the plates in the double octopod (see figure 2.5.14). The mounts fix to the exterior of the hull and then bolt to the middle plate through over sized bolt holes. Lateral position blocks can be fitted to laterally align the optics hull and shims can be placed under each mount to adjust tilt. Each mount is made from machined, anodised aluminium.

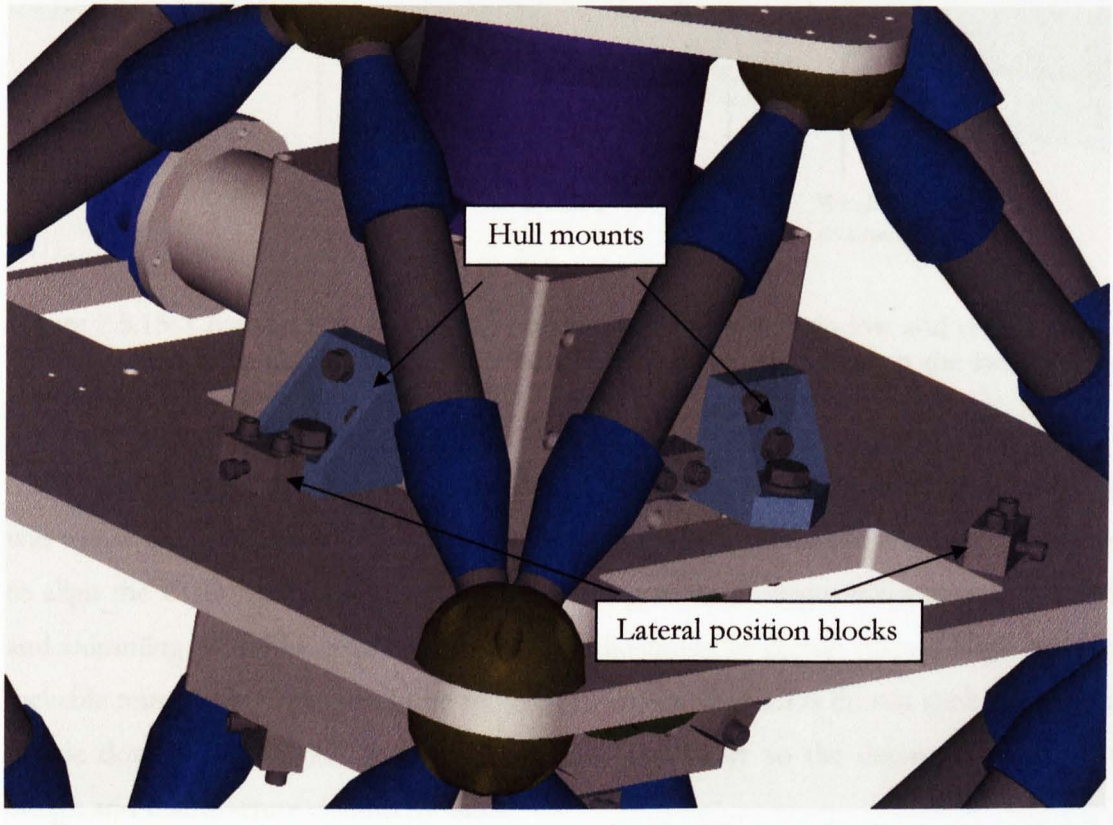

Figure 2.5.14: Optics hull mounted on the middle plate of the double octopod.

2.5.10 Mounting of the CCD heads

The CCD heads were designed at the UKATC. The CCD head lids however, were my responsibility, since they provide the best surface for mounting the CCD head. The CCD head lid incorporates the window which allows illumination of the CCD chip. The CCD head is a vacuum vessel so the lid-box and lid-window interfaces have to provide a vacuum

seal. O-rings are used at both interfaces, but because the separation between the CCD head and the re-imaging cameras is small, a low profile o-ring seal has been designed for the window. Figure 2.5.15 shows the CCD head window integrated with the lid. The lids were manufactured at Sheffield and fitted at the UKATC. They were machine finished and not anodised, since this can cause outgassing in a vacuum.

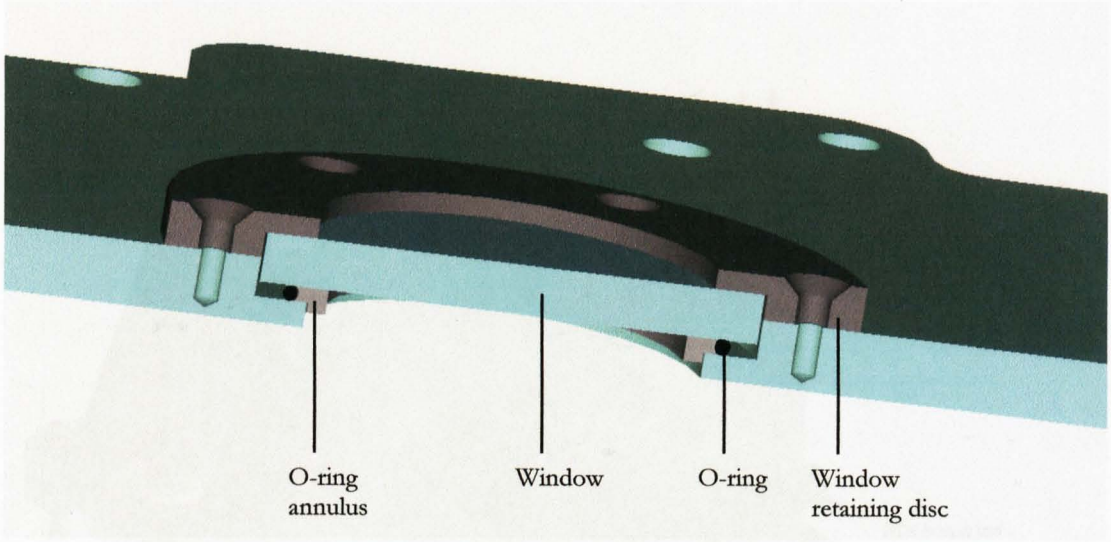

Figure 2.5.15: Cross-section of the CCD-head lid showing the window and o-ring vacuum seal. This type of o-ring seal has a low profile, which is essential given the lack of space between the CCD and filter.

The CCD heads fix to the double octopod facing each re-imaging camera. The initial idea was to use an arrangement of linear translation stages to produce the necessary adjustment to align the CCD heads. However, such an arrangement of stages takes up a lot of space and shimming would be required to produce adjustment in certain planes. High precision lockable translation stages would be required to ensure the CCDs do not shift if orientated upside down. Such stages are expensive, heavy and bulky so the decision was made to design and manufacture custom mounts.

The blue and green CCD heads use similar mounts because the CCD chips are held perpendicular to the middle plate. The red CCD chip lies parallel to the bottom plate and therefore has a different type of mount. The blue and green CCD-head mounts are L-shaped brackets that form a three-point mount with tabs extending from the CCD head lid, as shown in figure 2.5.16. The three-point mount operates in exactly the same way as with the re-imaging cameras. A fine-threaded bolt is used for adjustment and a standard-threaded bolt locks the position. The arrangement of three pairs of these bolts allows tilt and focus adjustment of the CCD head. The three point mount aligns the CCD to be

perpendicular to the optical axis and adjusts focus. Centring on the optical axis is achieved by positioning the CCD mount. Lateral positioning blocks shift the CCD in the x-direction and shimming under the three CCD mount bolts adjusts the y-direction and can correct for rotation. The combination of the three-point mount and positional adjustment of the CCD mount gives adjustment in all planes which can be locked when aligned.

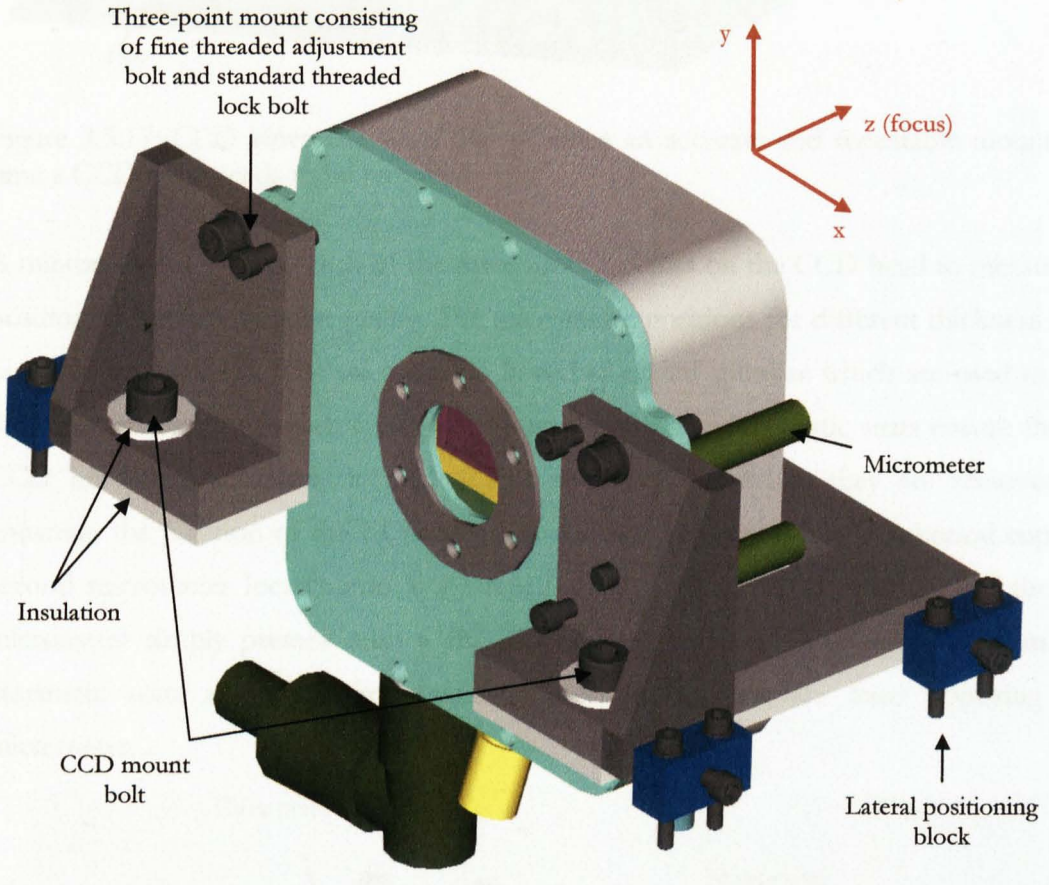

Figure 2.5.16: Green CCD head mount with CCD head fitted. The blue camera mount is very similar to this mount.

Each CCD head is electrically and thermally isolated from the opto-mechanical chassis. Thermal isolation is required to prevent unnecessary cooling of the double octopod by the cooling system. Electrical isolation is necessary to ensure that there is only one route to earth, eliminating 'grounding loops' and hence reducing detector noise. The isolation material used is G10/40 and is a glass fibre composite. There is a layer positioned between the CCD mount and the octopod. Bushes and washers were also manufactured from the same material to shield the CCD mount bolts.

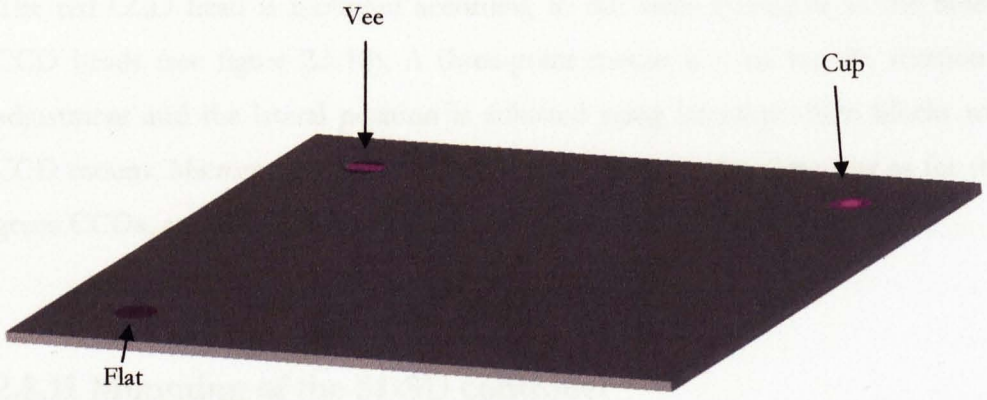

Figure 2.5.17: CCD kinematic seats that produce an accurate and repeatable mount each time a CCD head needs to be removed.

A micrometer is fitted to each of the three mount-points on the CCD head to measure the position of optimum image quality. The micrometer positions for different thickness filters can then be recorded. The micrometers have ball-ended spindles which are used to form the kinematic seats for each CCD (see figure 2.5.17). The kinematic seats ensure that the CCD heads are accurately located in the same position when they are removed. To constrain the position of the CCD, the first micrometer locates into a spherical cup. The second micrometer locates into a V-shaped channel constraining rotation and the final micrometer simply presses onto a flat disc defining the final degree of freedom. The kinematic seats are made from stainless steel inserts and are fixed opposing each micrometer.

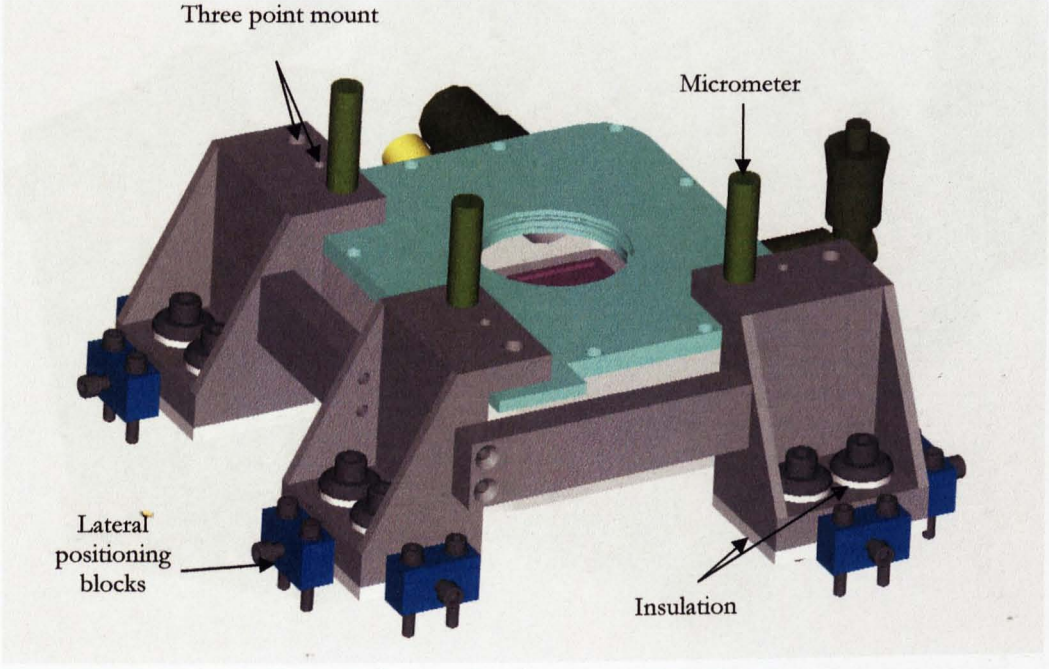

Figure 2.5.18: Red CCD head mount.

The red CCD head is mounted according to the same principles as the blue and green CCD heads (see figure 2.5.18). A three-point mount is used for tilt, rotation and focus adjustment and the lateral position is adjusted using lateral-position blocks to move the CCD mount. Micrometers are used in kinematic seats in the same way as for the blue and green CCDs, and the CCD mount is both electrically and thermally isolated.

2.5.11 Mounting of the SDSU controller

The SDSU controller needs to be close to the CCD heads to reduce pickup noise. The best available position on ULTRACAM is the base of the bottom plate. Figure 2.5.19 shows the SDSU fitted to the bottom plate of the double octopod. Blocks of aluminium have been added to the heat exchangers on the underside of the SDSU controller to allow it to act as a base so that ULTRACAM can sit upright when mounting on and dismounting from the telescope.

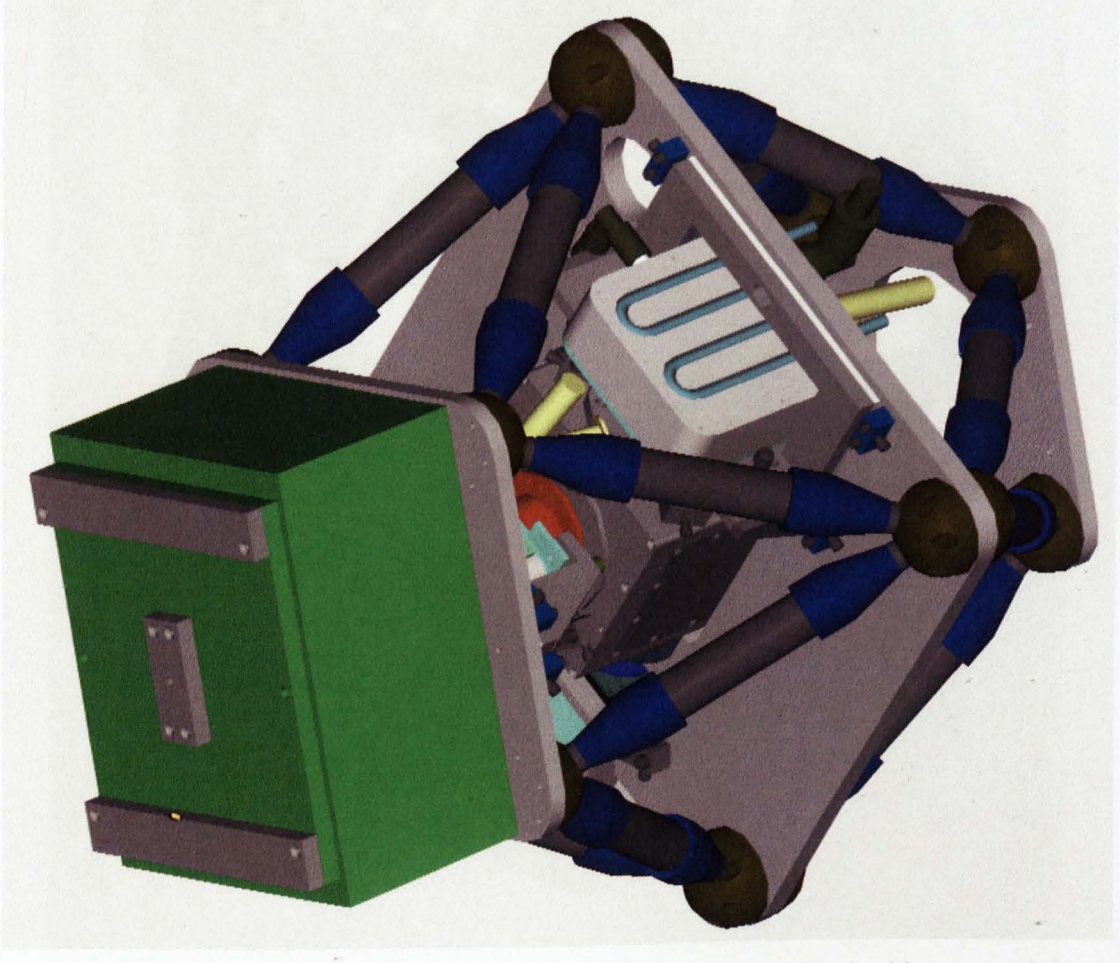

Figure 2.5.19: The SDSU controller fitted to the bottom plate of the double octopod.

2.5.12 The whole assembly

The whole ULTRACAM design with the WHT collimator fitted is shown in figure 2.5.20. The overall length of just the double octopod and SDSU controller is **792 mm** and is within the space envelope of the Aristarchos telescope. The Aristarchos telescope collimator is contained within the double octopod. The total weight of the instrument with the WHT collimator, measured by WHT engineers, is **82 kg** and is well within the mass envelope of the Aristarchos telescope.

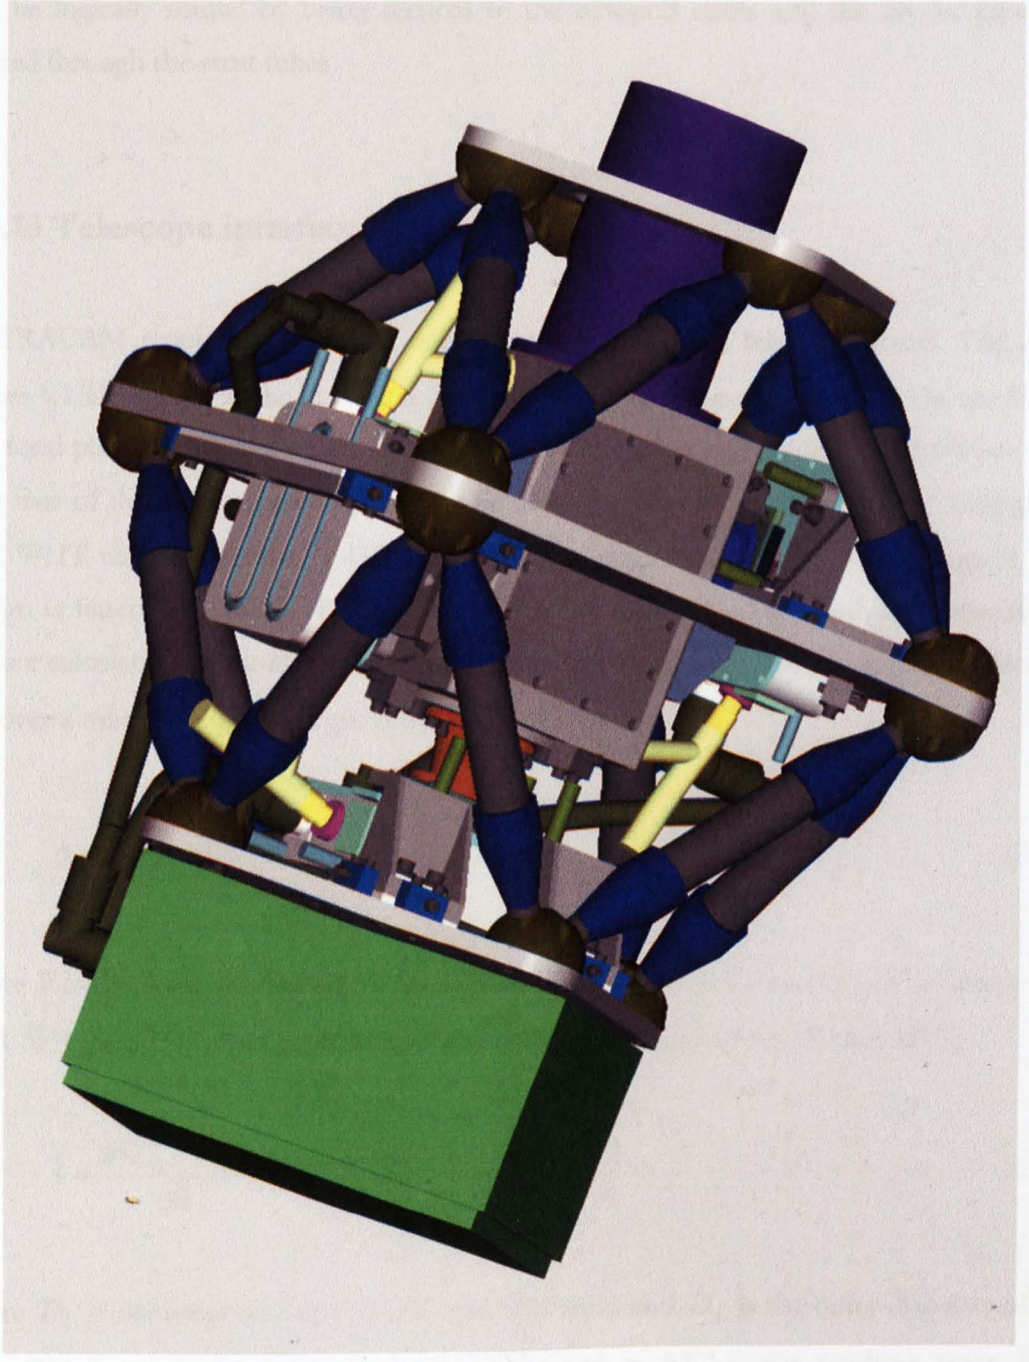

Figure 2.5.20: The whole ULTRACAM design, shown with the WHT collimator.

Figure 2.5.20 also shows the CCD-head cabling. The CAD model was used to determine the cable lengths for each CCD. Each CCD is shown with a yellow vacuum port valve which has to be fitted each time the CCD head needs to be pumped down. These are removed during normal operation but were included in the model to see if they could be fitted without removing the heads from the double octopod. The double octopod is an ideal structure for routing cables and pipes. For example, the water cooling pipes are heavy and can be secured to the struts to prevent stress on the pipe connectors. Electrical cables can be logically routed by being secured to the octopod struts and the dry air pipes are routed through the strut tubes.

2.5.13 Telescope interface

ULTRACAM requires a mounting collar to interface it to a telescope focus. The collar places ULTRACAM at the correct distance from the telescope focal plane. On the WHT, the focal plane is 150 mm below the Cassegrain acquisition and guidance (A&G) box. The first lens of the collimator is designed to be 235 mm away from the telescope focal plane. The WHT mounting collar is 455 mm long, so a rigid structure has been designed. The design is based on a piece of steel tubing with two flanges at each end. An approximate flexure calculation was undertaken to test the rigidity of this type of design. The deflection, Δ, over a tube of length L, is given by (Roark, 1975):

$$\Delta = \frac{FL^3}{3EI}, \tag{2.12}$$

where F is the force causing the deflection (in Newton), E is the modulus of elasticity (for steel, E = 207 GPa) and I is the area of moment of inertia, given by (Roark, 1975):

$$I = \frac{\pi\left(D_2^4 - D_1^4\right)}{64}, \tag{2.13}$$

where D_1 is the inner diameter of the tube (228 mm) and D_2 is the outer diameter of the tube (253 mm). Using the mass of ULTRACAM as 82 kg, the resulting force F, is

approximately 820 N. The deflection or flexure is calculated to be (1.8 ± 0.20) x 10^{-6} m. This would not be noticed in ULTRACAM images since the pixel size is 13 microns. With reasonable confidence that the design was rigid enough, as a precautionary measure, ribs were welded between the flanges (see figure 2.5.21).

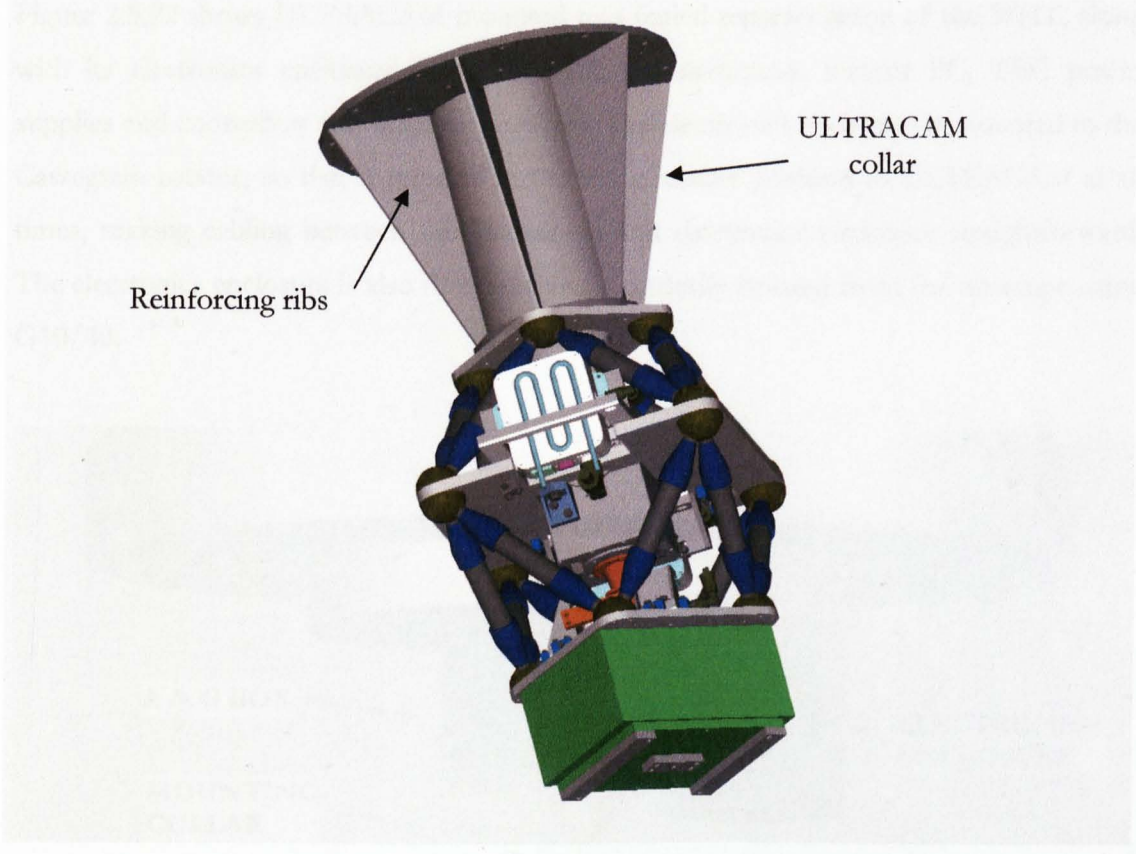

Figure 2.5.21: Model drawing of the collar designed to interface ULTRACAM to the WHT.

The optical designer modelled ULTRACAM tilted 0.2 degrees with respect to the telescope optical axis. The result was virtually no lateral shift and no noticeable degradation of image quality. A 0.5 degree tilt caused a 100 μm lateral shift and an acceptable reduction of image quality. The optical designer also modelled ULTRACAM laterally shifted by 2 mm from the telescope optical axis. This caused no image quality degradation, but did cause a lateral shift of approximately 30 pixels in the images on the CCD. These figures set the manufacturing tolerance on the collar.

There is a layer of G10/40 isolation material between the collar and the top plate of ULTRACAM. This thermally and electrically isolates ULTRACAM and the SDSU controller from the telescope. The three plates of the double octopod have been designed

and assembled to be parallel, so special attention has been given to ensuring the collar is parallel too. Simple joints have been used between the tube and flanges and these have been welded with care to ensure the steel tube did not warp during manufacture. The collar has been painted black to reduce scattered light.

Figure 2.5.22 shows ULTRACAM mounted to a scaled representation of the WHT, along with its electronics enclosure (which houses the instrument control PC, TEC power supplies and controllers and other electronics). The electronics enclosure is mounted to the Cassegrain rotator, so that it remains in the same relative position to ULTRACAM at all times, making cabling between the instrument and electronics enclosure straightforward. The electronics enclosure is also thermally and electrically isolated from the telescope using G10/40.

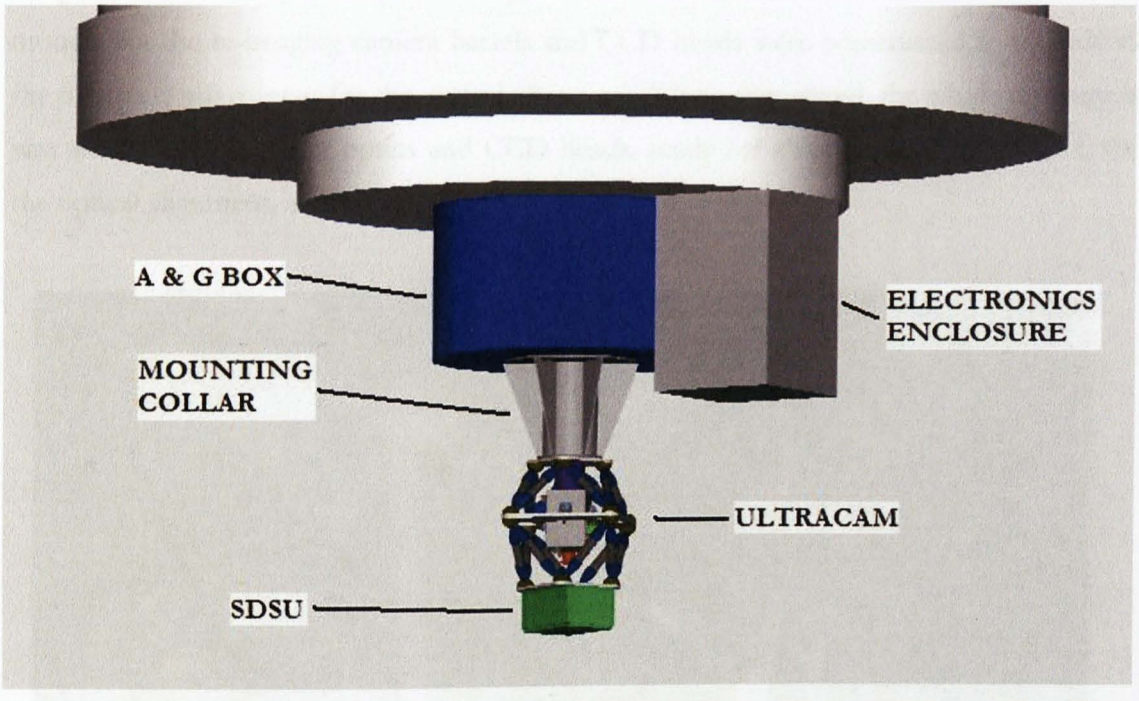

Figure 2.5.22: ULTRACAM fitted to a scaled model of the WHT. The mounting collar provides the interface between ULTRACAM and the telescope. It is designed to position ULTRACAM at the correct vertical spacing, parallel to and aligned with the telescope optical axis. Also shown is the electronics rack mounted to the Cassegrain rotator.

Chapter 3

Instrument Assembly

3.1 Double Octopod and Optics Hull

All parts for the double octopod and optics hull were manufactured at Sheffield. Custom mounts for the re-imaging camera barrels and CCD heads were constructed to provide all the necessary adjustment for the optical alignment. When completed, the whole instrument was assembled, excluding optics and CCD heads, ready for shipment to the UKATC for the optical alignment, as shown in figure 3.1.1.

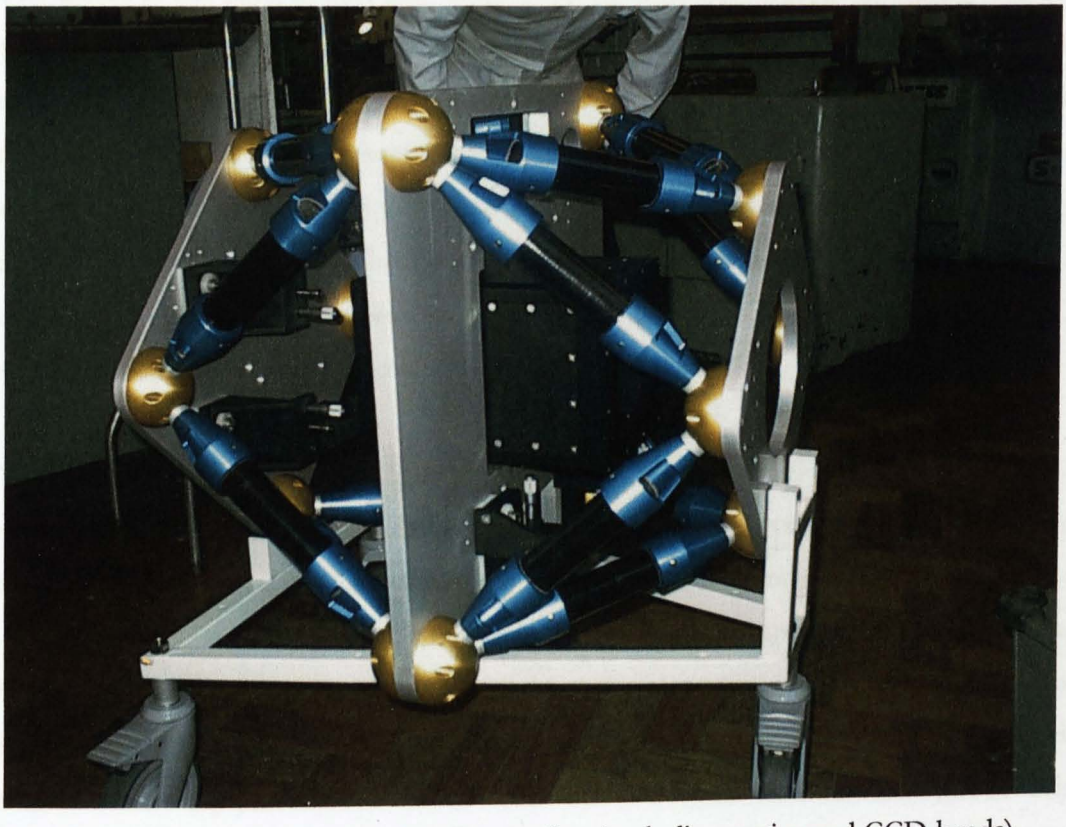

Figure 3.1.1: ULTRACAM on its handling trolley (excluding optics and CCD heads).

Also shown in figure 3.1.1 is the handling trolley which I designed to help move ULTRACAM around. The anti-static wheels of the trolley are removable and there are holes in the trolley frame that are at the correct locations so it and ULTRACAM can be securely fixed to a standard optical bench.

3.2 Optical alignment

The optical alignment began on 25th March 2002 at the UKATC and was performed by Tully Peacocke and myself. Telescope commissioning was scheduled for 16th May 2002, but ULTRACAM had to be shipped by 2nd May 2002 at the latest in order to be at the telescope in time.

The alignment was undertaken on an optics bench using an alignment telescope. An alignment telescope projects target rings and a cross-wire through the lenses to be aligned and can focus on the reflections of these from each lens surface. The alignment telescope can be laterally moved to align it to the optical axis of a lens and can then be used to align other lenses to the same optical axis. Because of the operating wavelength range of the blue optics, the alignment telescope could not detect reflections from the blue camera lenses and dichroic beam-splitter. A HeNe laser was directed through the blue optics to provide reflections from the lens surfaces that could be detected by the alignment telescope. A penta-prism was used during the alignment to align mirrors with respect to defined reference optics; a light beam incident onto a penta-prism always reflects the beam at an angle of 90 degrees with respect to the incident beam angle. The use of mirrors and reference optics will be explained later in this section.

The first stage of the alignment fixed the dichroic substrates at 90 degrees to each other in an L-shape. The frames with the dichroic beam-splitters fitted were attached to the dichroic mount. Figure 3.2.1 shows a schematic of the optics bench during the dichroic alignment. Both the alignment telescope and HeNe laser were aligned square to two mirrors opposing each dichroic and each other. The green dichroic (dichroic 2) was aligned to the alignment telescope and fixed as the reference optic. The blue dichroic (dichroic 1) was then adjusted using the laser so that reflected light from the mirror returns to its point of origin. Adjustments were made until the smallest shim could not improve the alignment leaving a

tilt in the green dichroic beam-splitter in the z-plane, measured as a 0.25 mm offset in the reflection received at the alignment telescope over an optical path of 3.6 m. The angular tilt is 0.07 mrad or 0.04 degrees.

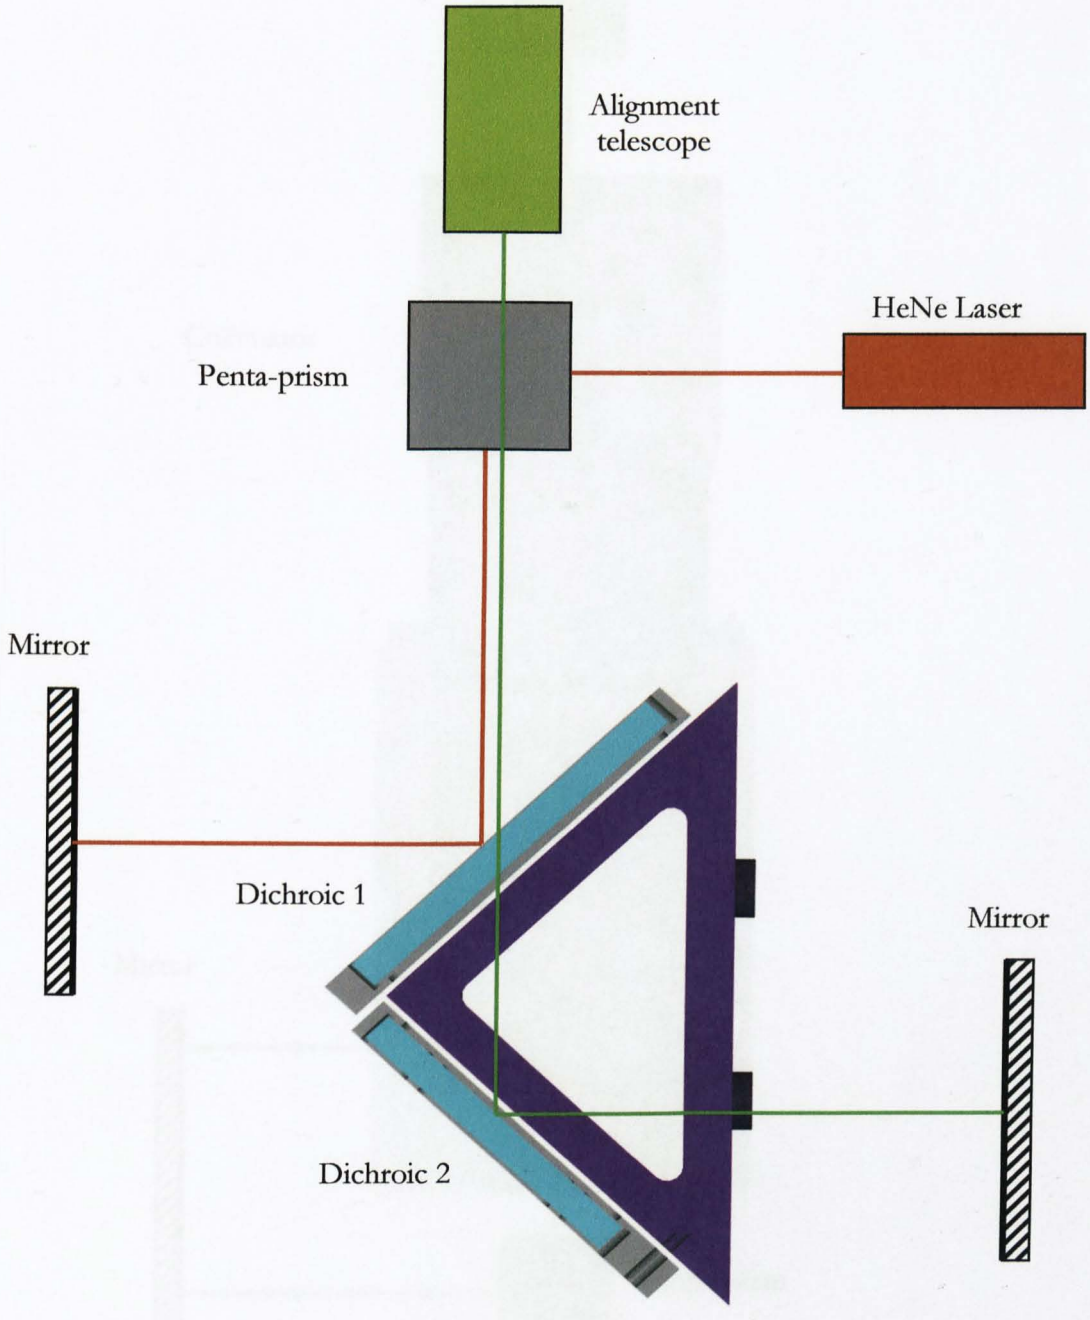

Figure 3.2.1: Schematic of the optical bench during the dichroic beam-splitter alignment. An alignment telescope was used to align the green dichroic (dichroic 2) and a HeNe laser was used to align the blue dichroic (dichroic 1). The penta-prism is used to reflect incident light at 90 degrees.

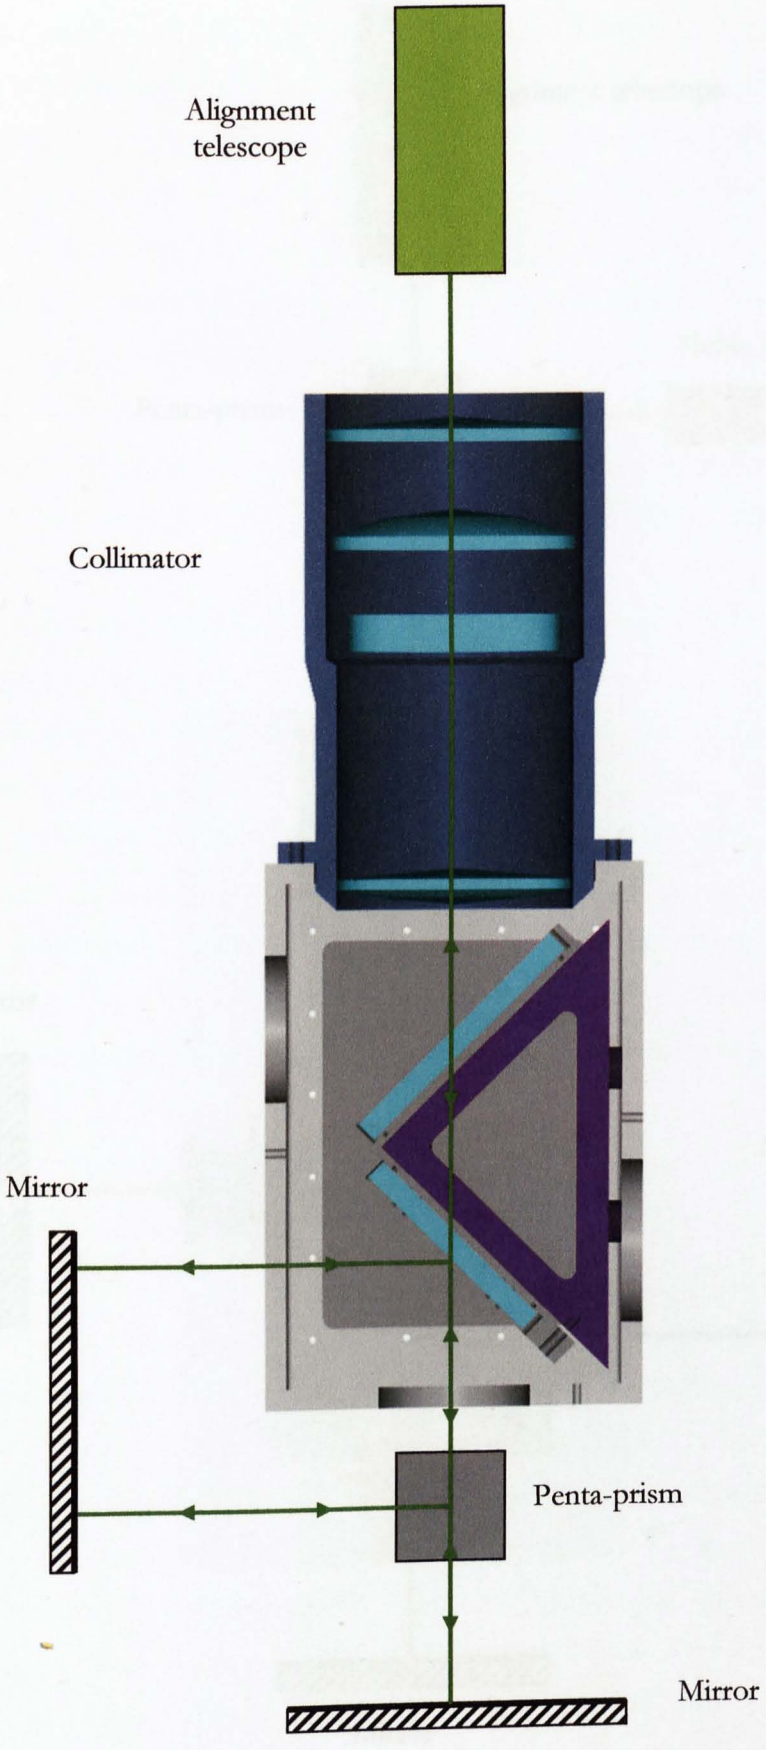

Figure 3.2.2: Schematic of the layout used to align the collimator and dichroics.

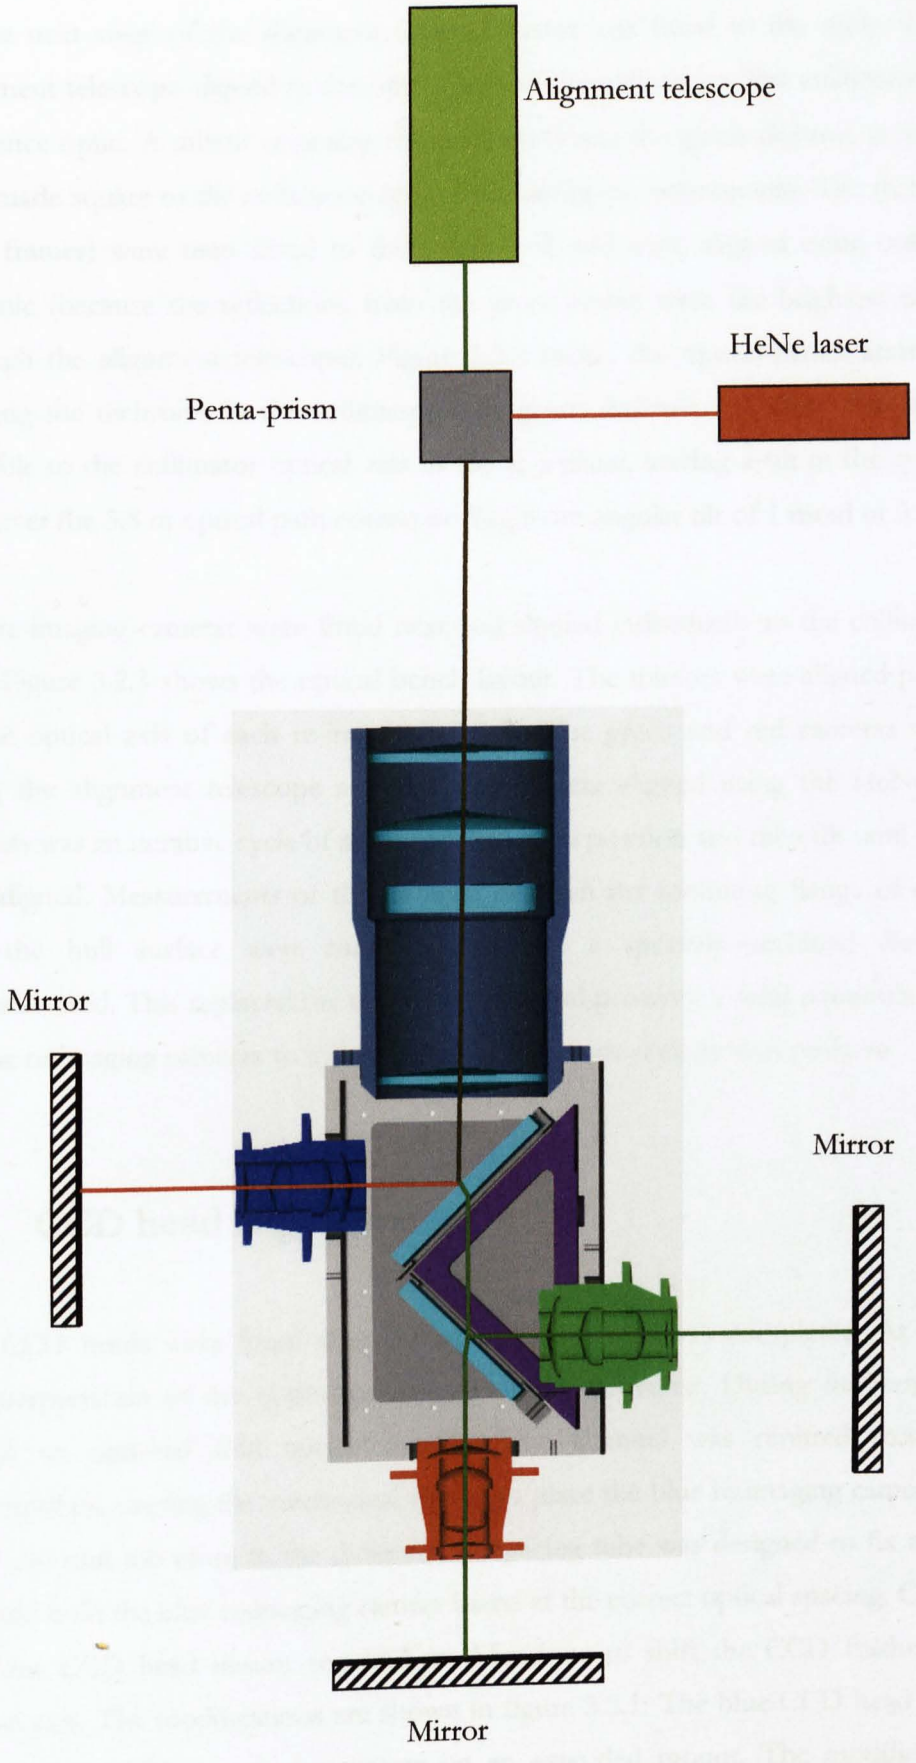

Figure 3.2.3: Schematic of the layout used to align the re-imaging cameras.

In the next stage of the alignment, the collimator was fitted to the optics hull and the alignment telescope aligned to the optical axis of the collimator. The collimator is now the reference optic. A mirror opposing the location where the green dichroic is mounted was first made square to the collimator optical axis using the penta-prism. The dichroic mount (and frames) were then fitted to the optics hull and were aligned using only the green dichroic (because the reflections from the green lenses were the brightest when viewed through the alignment telescope). Figure 3.2.2 shows the optical-bench arrangement for aligning the dichroics to the collimator. The green dichroic was aligned as accurately as possible to the collimator optical axis in the x, y plane, leaving a tilt in the z-plane of 0.5 mm over the 3.8 m optical path corresponding to an angular tilt of 1 mrad or 0.08 degrees.

The re-imaging cameras were fitted next and aligned individually to the collimator optical axis. Figure 3.2.3 shows the optical bench layout. The mirrors were aligned perpendicular to the optical axis of each re-imaging camera. The green and red cameras were aligned using the alignment telescope and the blue camera aligned using the HeNe laser. The process was an iterative cycle of adjusting the lateral position and then tilt until each camera was aligned. Measurements of the distance between the mounting flange of each camera and the hull surface were measured so that a specially-machined disc could be manufactured. This replaced the temporary disc and provides a solid continuous surface to fix the re-imaging cameras to the optics hull at the correct alignment position.

3.3 CCD head alignment

The CCD heads were fitted after the optical alignment was completed. At this stage a misinterpretation of the optical prescription was discovered. During the iterative optical design an optional fold mirror for the blue channel was omitted from the final prescriptions, causing the mechanical design to place the blue re-imaging camera and CCD head 100 mm too close to the dichroics. A spacing tube was designed to fix to the optics hull and hold the blue re-imaging camera barrel at the correct optical spacing. Consequently the blue CCD head mount required modifications to shift the CCD further down the optical axis. The modifications are shown in figure 3.3.1. The blue CCD head is now held outside the double octopod structure on an extended mount. The modifications were

quickly designed by myself and manufactured at Sheffield and ULTRACAM was returned to the UKATC for the CCD head alignment.

The CCD heads were aligned and focused by illuminating a target placed at the telescope focal-plane position and reading out the CCD to measure the lateral position and focus of the target image. Lateral alignment of the CCDs was achieved by shifting the position of the CCD mounts, whilst focus and tilt of the CCD was achieved by adjusting the three point mount on each CCD head. A temporary CCD alignment was undertaken prior to the commissioning run, since there was little time and it was clear that modifications were required to the CCD head mount designs after the observing run.

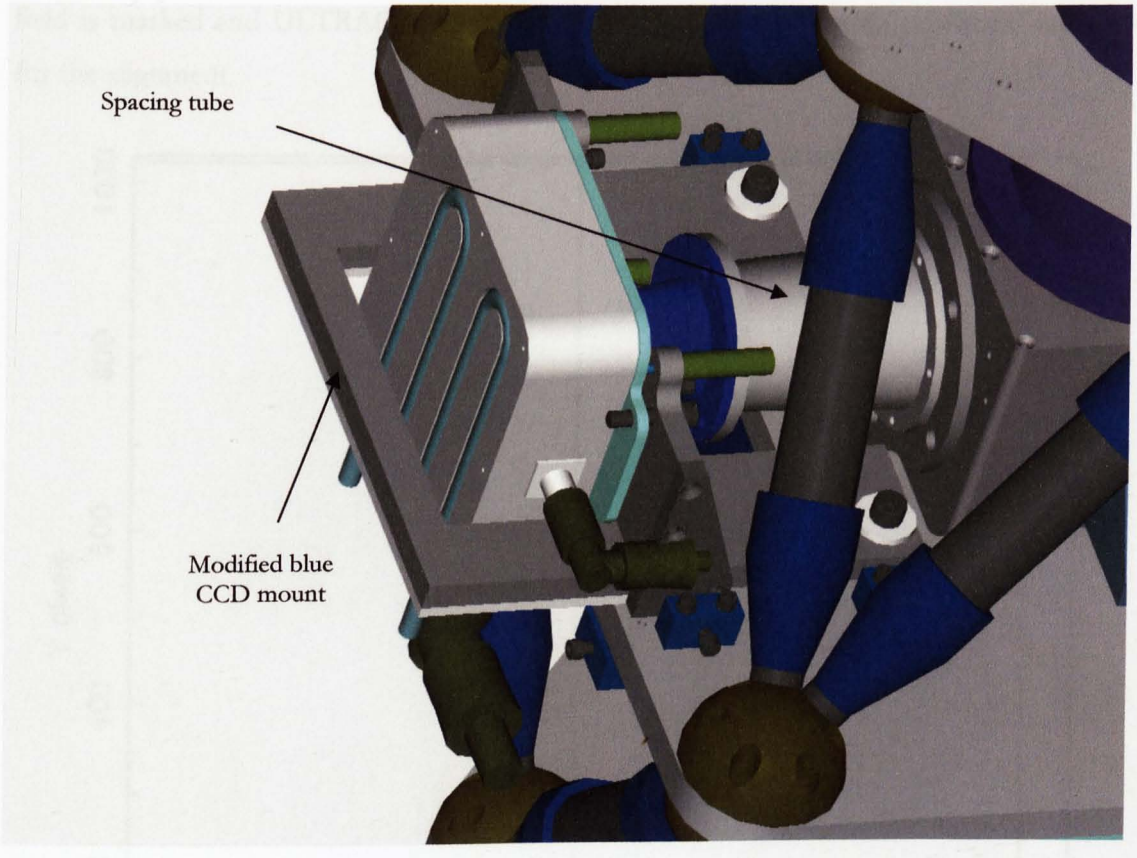

Figure 3.3.1: View of the modifications required to allow the optical alignment to proceed. A spacing tube holds the blue re-imaging camera an additional 100 mm away from the optics hull. The modified blue CCD mount also hold the CCD head further along the optical axis.

The post commissioning modifications to the CCD mounts included fitting kinematic seats, as described in section 2.5.10, so that the CCD heads could be removed and replaced in the same location each time they are removed. The red CCD mounts required

modification because the first design made it very difficult to align accurately. The new design is based on the red CCD head sitting on three pedestals forming a three point mount. In the initial design the three pedestals were not joined, making lateral alignment of the CCD very difficult. The solution, detailed in section 2.5.10 incorporates beams that join all three pedestals. All of the modifications required all three CCD heads and mounts to be removed resulting in each CCD head requiring re-alignment.

The CCD head re-alignment was done at Sheffield by me. A specially designed star field tube was manufactured which slides over the collimator and holds a disc of glass at the position of the WHT focal plane. A black acetate sheet with an array of clear holes was glued to the glass disc creating an artificial array of stars (see figure 3.3.2). The centre of the field is marked and ULTRACAM simply images the field to provide positional feedback for the alignment.

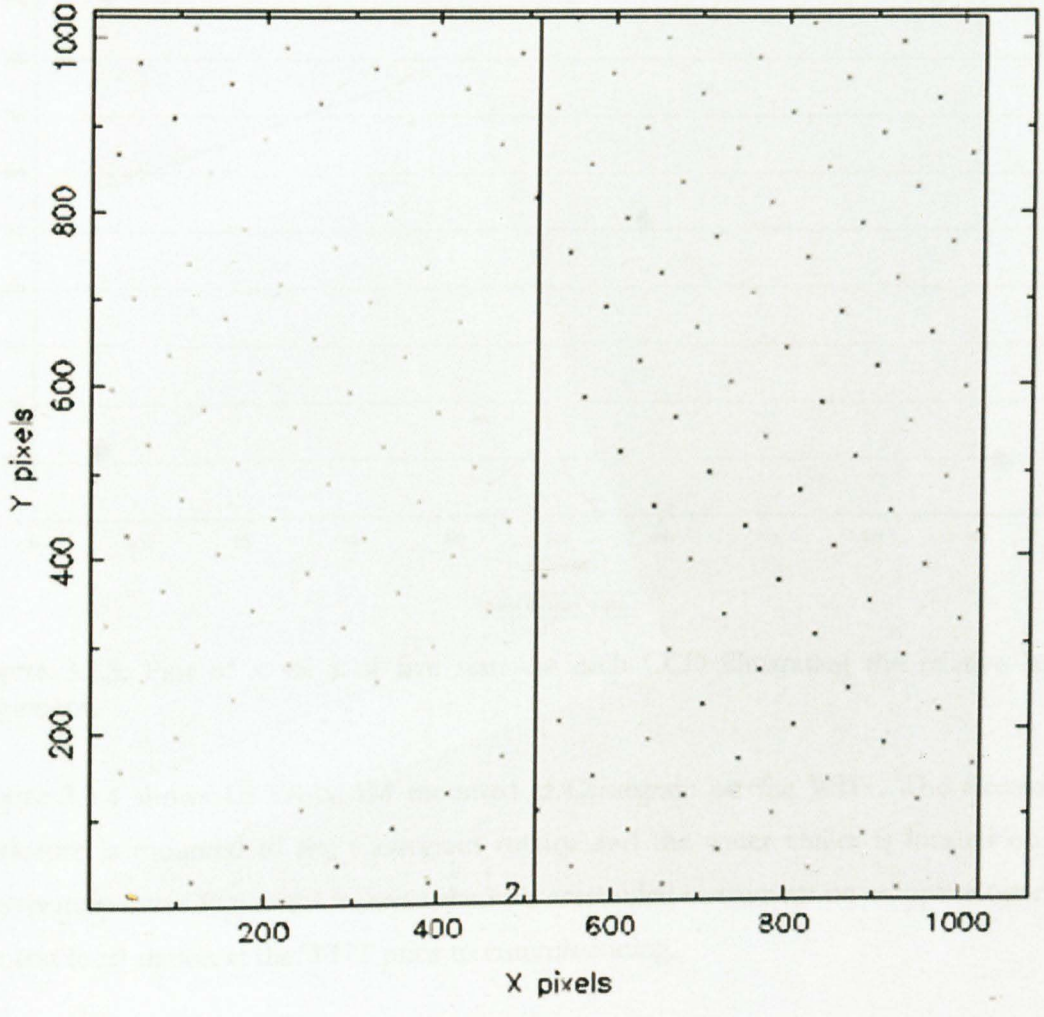

Figure 3.3.2: An ULTRACAM image with the star-field tube fitted.

Five stars were chosen, one at the centre and four at the corners of the field, so that focus could be checked over the entire array. The focus was measured by the FWHM of the stars and alignment was measured by the x,y positions of the stars. Figure 3.3.3 shows a plot of the x and y positions of five stars on each CCD after alignment of the heads. The three plots are superimposed to show relative CCD alignment. The three CCDs are not perfectly aligned in rotation about the CCD centre, the worst being the red CCD. A star appearing on the green CCD would be located 10 pixels away on the red CCD at the corners of the field. This is due to the CCDs not being perfectly aligned and the plate scale differing slightly in each band. The effect of this is that the minimum window size at the corners of the CCD field (the most commonly used location in drift mode) is limited by the pixel-to-pixel misalignment of the CCDs.

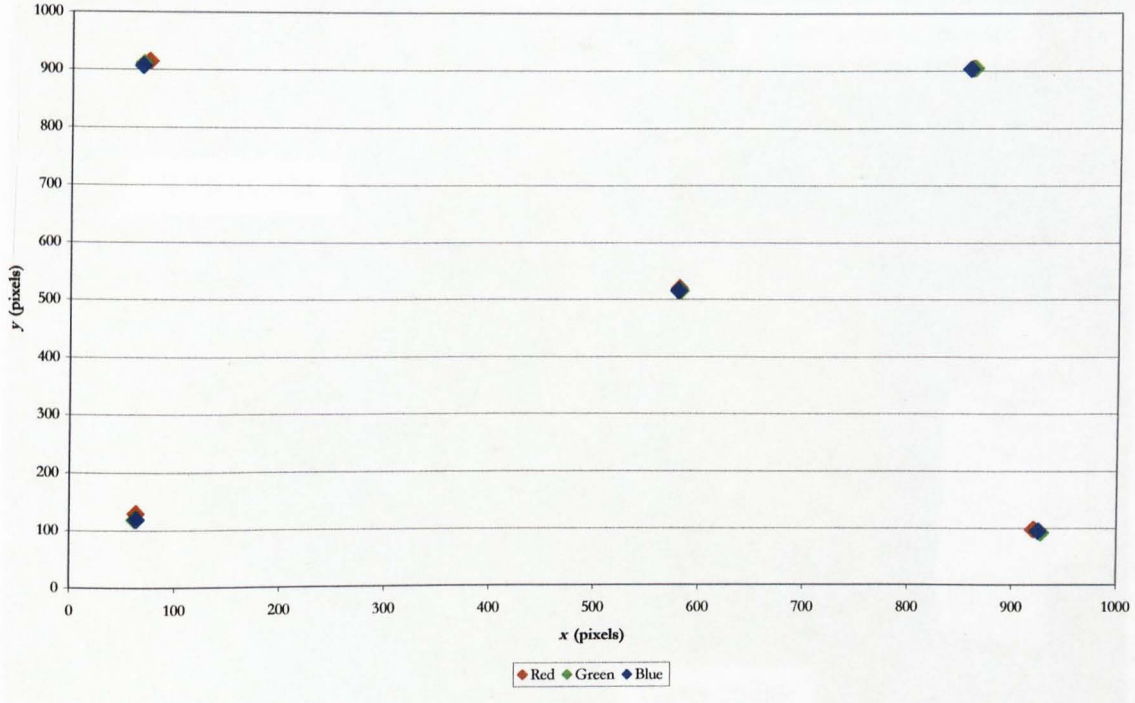

Figure 3.3.3: Plot of x vs. y of five stars on each CCD illustrating the relative lateral alignment.

Figure 3.3.4 shows ULTRACAM mounted at Cassegrain on the WHT. The electronics enclosure is mounted to the Cassegrain rotator and the water chiller is located on the observatory floor. Figure 3.3.5 shows the fully assembled instrument on an optics bench in the test focal station at the WHT prior to commissioning.

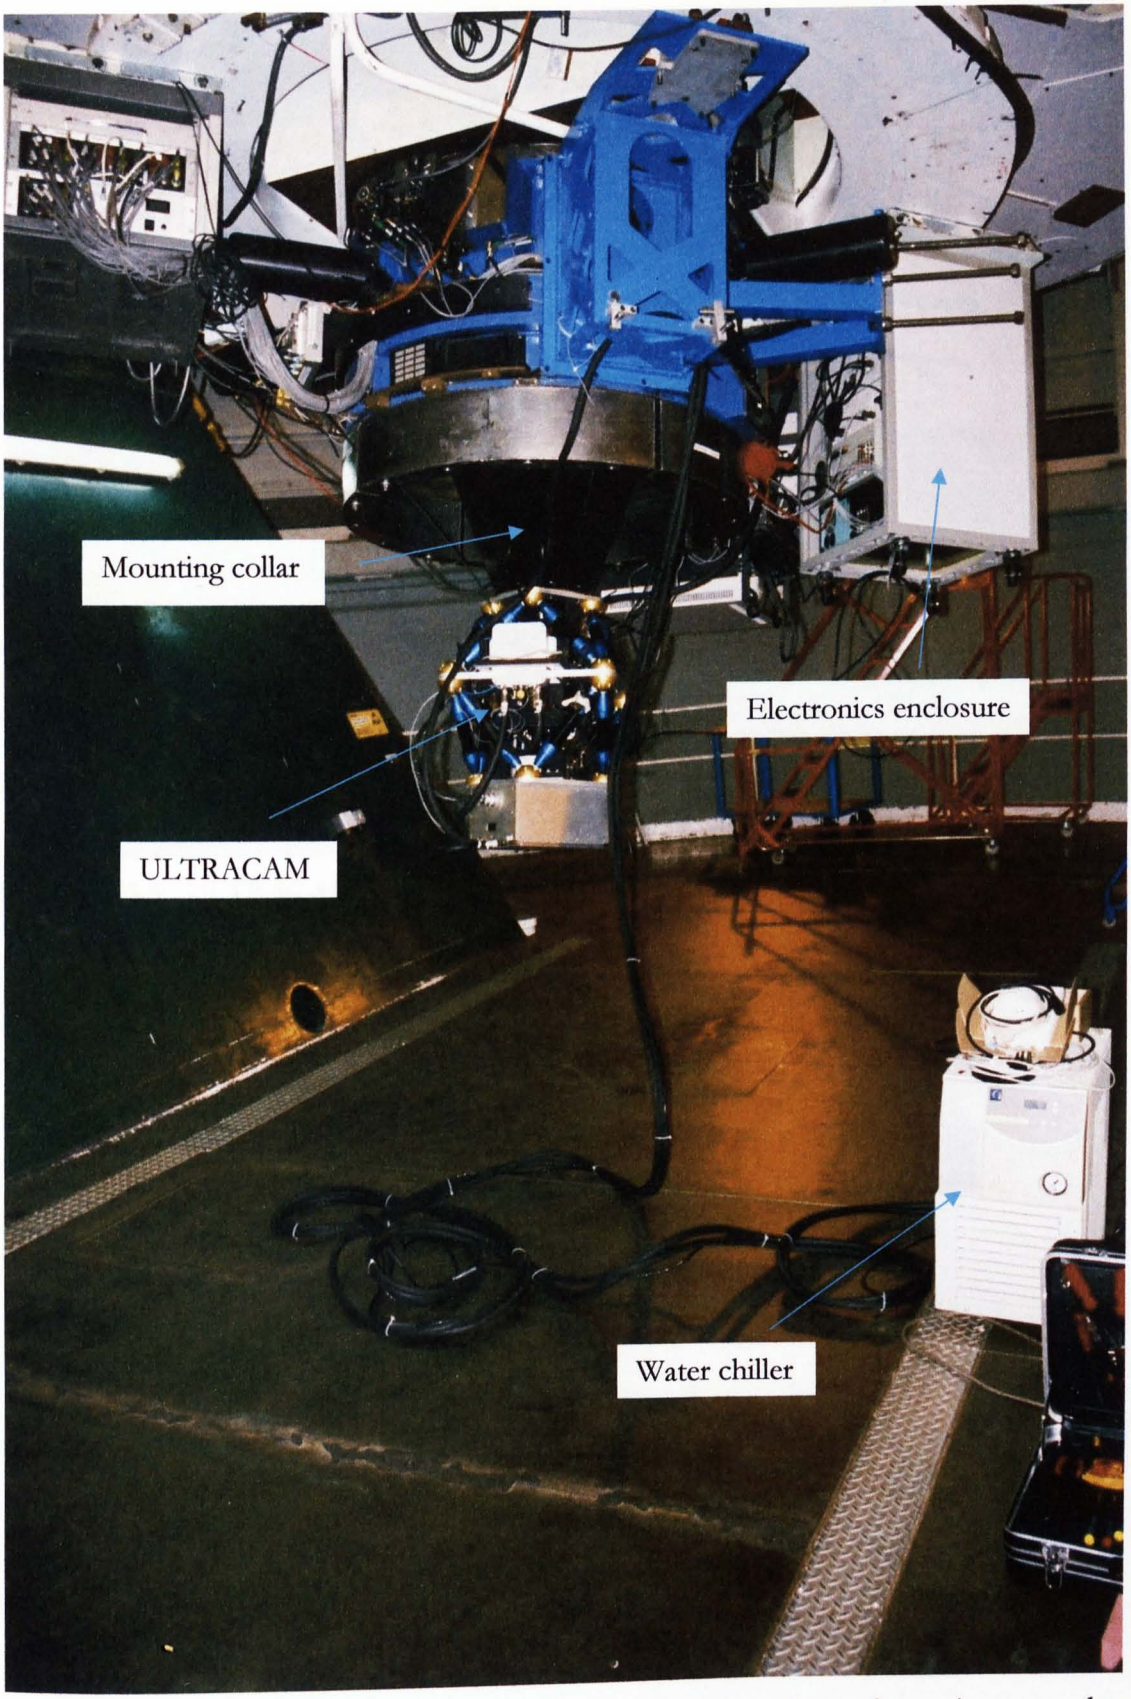

Figure 3.3.4: ULTRACAM mounted to the WHT. The electronics enclosure is mounted to the rotator and the water chiller is located on the observatory floor.

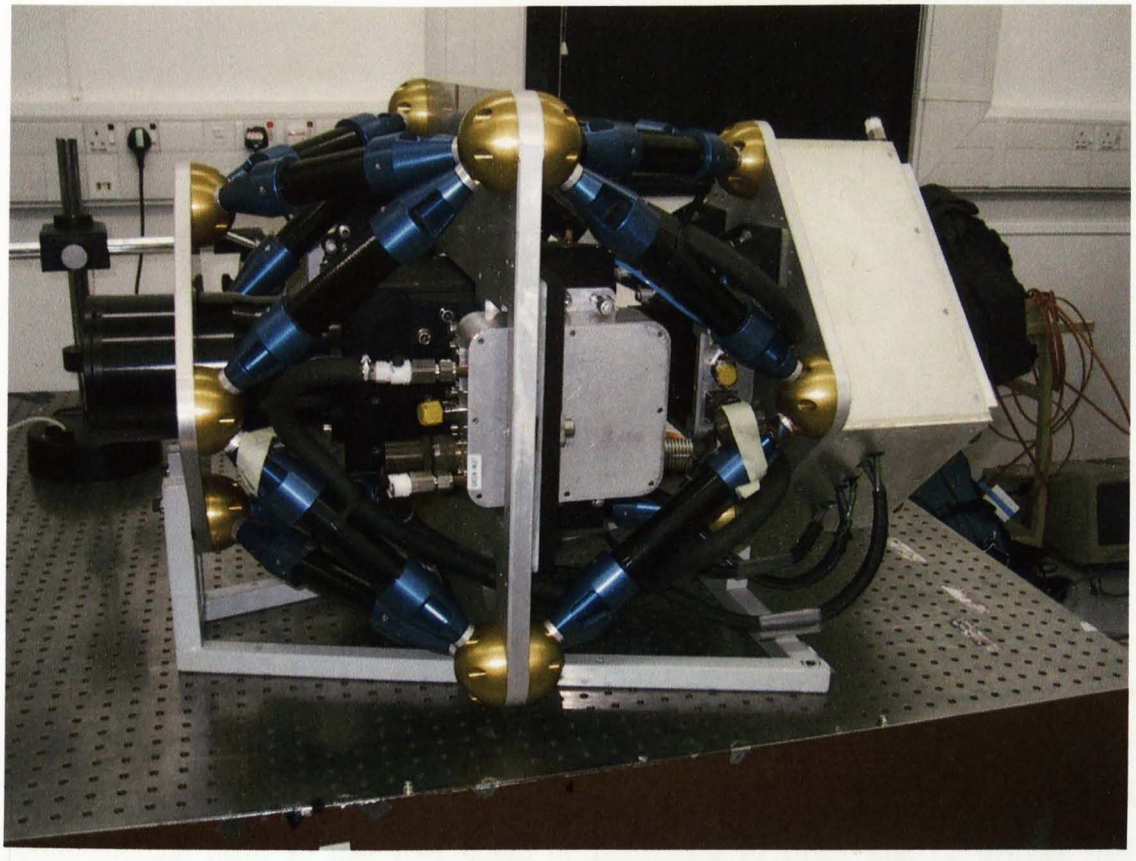

Figure 3.3.5: ULTRACAM on an optics bench at the WHT prior to commissioning (Courtesy of Sue Worswick).

Chapter 4

Commissioning

4.1 Introduction

ULTRACAM was commissioned on 16th May 2002 on the WHT, La Palma. One night was allocated to make the instrument operational, and this was immediately followed by four nights of science time. ULTRACAM arrived at the WHT 8 days prior to the commissioning night so that tests could be performed to check that no damage had occurred during shipping. This chapter will present the results of the commissioning tests which investigate how well the optics, CCDs, software and mechanical structures perform. Commissioning data was obtained not only on 16th May 2002, but also during subsequent observing runs. All the work presented in this chapter was undertaken by myself unless otherwise stated.

4.2 Optics

4.2.1 Image quality

The image quality was measured to determine whether the optics perform as predicted. Significant differences between expected and actual image quality may indicate an imperfect lens surface or misalignment. The image quality was tested by measuring the FWHM of theoretical and observed images of point sources.

The optical designer modelled a point source passing through the ULTRACAM optics and presented the output in the form of numerical arrays defining intensity in each CCD focal plane. This was done for point sources of wavelength 365, 500 and 625 nm, so that the blue, green and red re-imaging cameras have been individually modelled. The point sources originated from three locations in the field of view, as shown in figure 4.2.1.

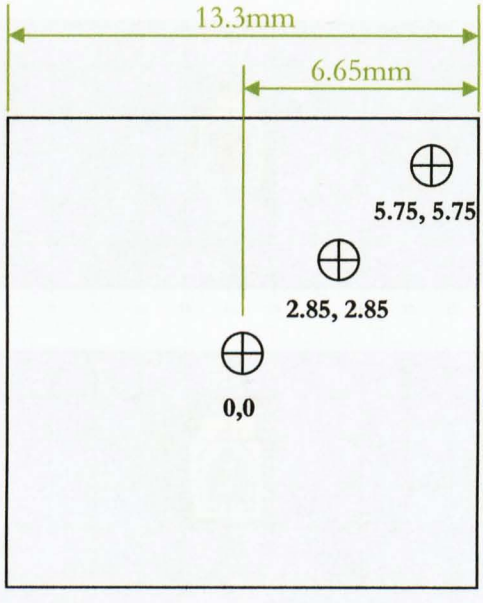

Figure 4.2.1: Diagram showing the location of the origin of three-point sources in the ULTRACAM field of view used in the image quality analysis. The locations of the point sources, (0,0), (2.85,2.85) and (5.75,5.75) are measured in millimetres from the centre of the field.

The intensity array output by the optical designer was divided into pixels which were much smaller than the ULTRACAM CCD pixels. To compare the FWHM of observed point sources, the intensities of the theoretical data pixels had to be binned to match the larger ULTRACAM CCD pixels.

The binned arrays are plotted in figure 4.2.2. The plots show the theoretical intensity of a point source re-imaged onto an ULTRACAM CCD array at three different field locations. Good image quality is indicated by intensity being concentrated in fewer pixels. The plots are a quantitive form of spot diagram, which allow FWHM to be measured. Similar to the spot diagrams in section 2.2.9, the plots in figure 4.2.2 show that the position of optimum focus is an annulus around the centre of the chip with the exception of the blue camera, where the optimum focus is at the centre of the field.

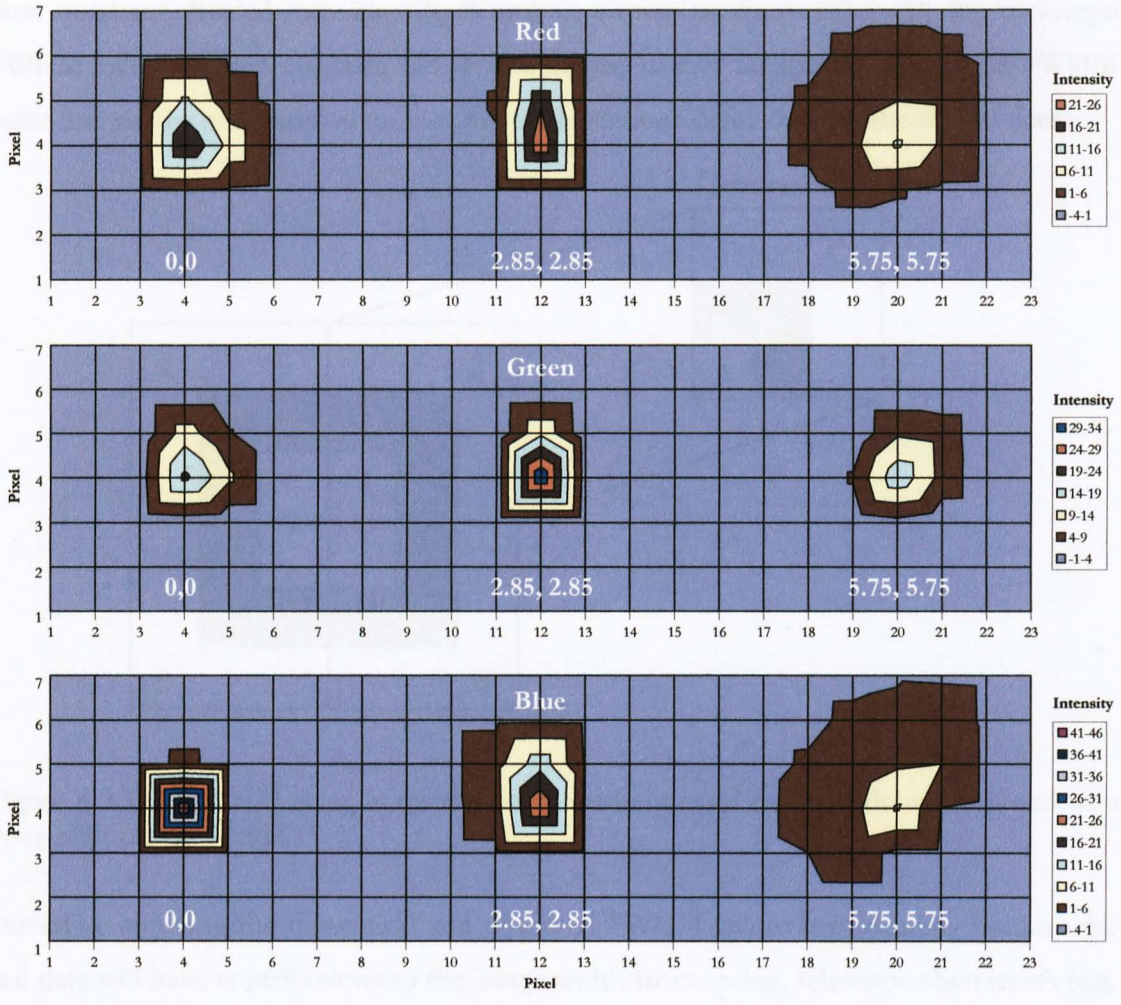

Figure 4.2.2: Theoretical images of a point source at three locations on each CCD.

The FWHM of each point source was measured in x and y and these values were then averaged and plotted as a function of position on the CCD field. The theoretical data assumes the image quality to be symmetrical about the centre of the field. What this means is that the FWHM at location 2.85, 2.85 mm is the same as at locations -2.85, 2.85 mm and 2.85, -2.85 mm and -2.85, -2.85 mm.

Data obtained in September 2002 of the crab nebula provided frames containing numerous stars. The observations were undertaken during 0.9 arcsecond seeing (as measured by the ING Differential Image Motion Monitor (DIMM)). The FWHM of as many objects as possible were measured on each CCD. There were 26, 32 and 6 stars in the *r'*, *g'* and *u'* frames, which were used for the analysis. The FWHM and associated uncertainties of these stars were measured using DIPSO, and the uncertainties were combined in quadrature. To compare the results with the theoretical data the CCD array was divided into quadrants and

each quadrant divided into bins A, B and C, shown in figure 4.2.3. All the measured FWHM values in each quadrant bin were averaged to give an approximate single FWHM value that can be compared at each of the field locations defined in the theoretical data.

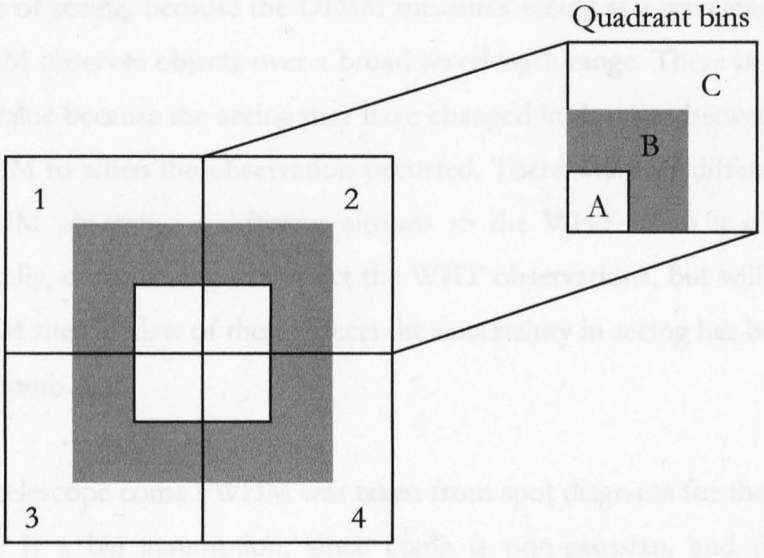

Figure 4.2.3: The CCD array is divided into quadrants and each quadrant is binned into three regions, A, B and C.

Directly comparing the theoretical and observed FWHM values is misleading because the real data will have contributions to the image width from seeing, telescope aberrations (e.g. coma), telescope focus errors and guiding errors. Seeing can be accounted for by approximating the theoretical and seeing Point Spread Function (PSF) as Gaussian, and convolving them, giving the total FWHM, σ_{Total}:

$$\sigma^2_{Total} = \sigma^2_{Theoretical} + \sigma^2_{Seeing}, \tag{4.1}$$

where $\sigma_{Theoretical}$ is the theoretical FWHM and σ_{Seeing} is the seeing FWHM. Similarly, telescope coma can be approximated and corrected in the same way giving:

$$\sigma^2_{Total} = \sigma^2_{Theoretical} + \sigma^2_{Seeing} + \sigma^2_{Coma}, \tag{4.2}$$

where σ_{Coma} is the approximated FWHM contribution from coma.

The seeing was recorded by the DIMM at a wavelength of 500 nm, one hour prior to the observing run. There are several sources of uncertainty in the seeing value measured by the DIMM to what is observed by the WHT. Firstly, there is uncertainty due to the wavelength dependence of seeing, because the DIMM measures seeing at a wavelength of 500 nm but ULTRACAM observes objects over a broad wavelength range. There is also uncertainty in the seeing value because the seeing may have changed in the time between being measured by the DIMM to when the observation occurred. There will be a difference in seeing due to the DIMM observing a different airmass to the WHT when it observes the target objects. Finally, dome seeing will effect the WHT observations, but will not be significant at the DIMM site. In view of these effects the uncertainty in seeing has been assumed to be $\pm$ 0.1 arcseconds.

The WHT telescope coma FWHM was taken from spot diagrams for the telescope (Carter, 1996). This is a big assumption, since coma is non-gaussian, and this will introduce uncertainty. It is assumed that telescope focus is optimum for all three colour bands, and that telescope tracking error will have a very small effect on FWHM since the exposure times of the crab nebula are short (15 seconds). The WHT is known to track to <1 arcsecond over 10 minutes (Benn, 1999), so contributions due to tracking are assumed to be insignificant.

Figures 4.2.4, 4.2.5 and 4.2.6 show plots of FWHM as a function of field location *x,y*. For each CCD the upper plot shows FWHM values for each field quadrant defined in figure 4.2.3 and the lower plot compares the average FWHM with the theoretical FWHM (the latter with contributions from seeing and coma added). The lower plots also show the quadrant with the lowest FWHM or best image quality in each case. This represents the best achievable image quality for each colour band. Deviations from this curve are due to the CCD not being properly aligned. Unfortunately, there is very little data for the blue optics because there were few blue stars in the Crab nebula field.

The top panel of figure 4.2.4 indicates that the red CCD is tilted with respect to the focal plane; quadrants 2 and 4 show lower FWHM than quadrants 1 and 3, showing that the tilt increases from right to left according to figure 4.2.3. The theoretical FWHM with contributions from seeing and telescope coma is closest to the best results from the observed quadrant suggesting that the red optics are operating close to theoretical

expectation. The average observed FWHM are skewed by the tilt in the CCD, which has since been adjusted to improve the red image quality.

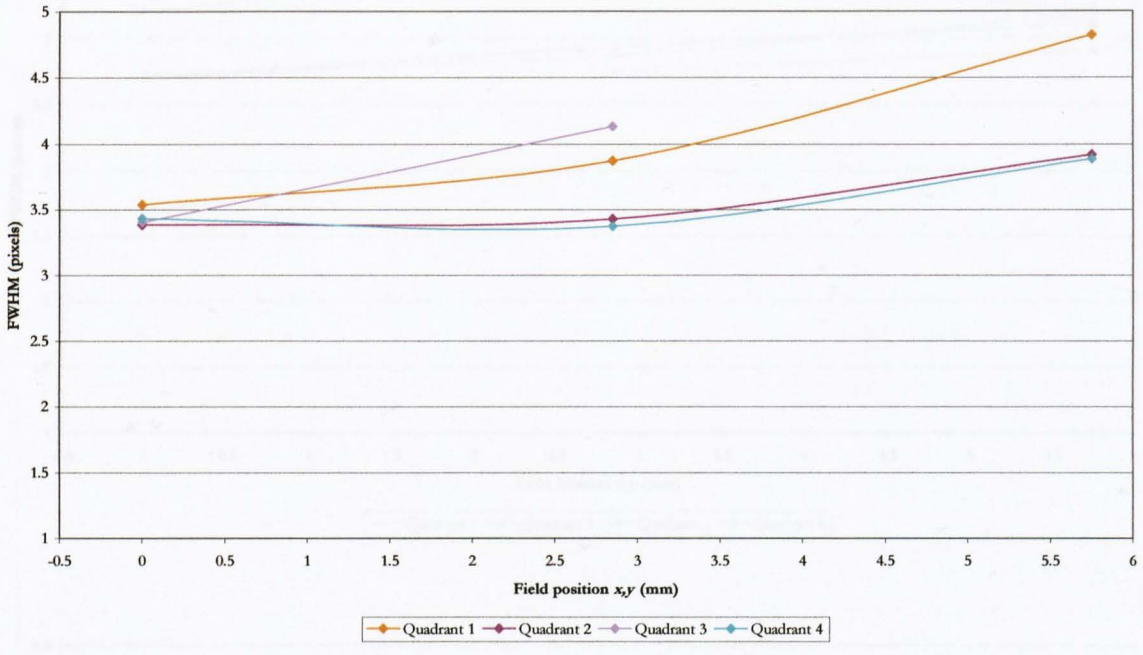

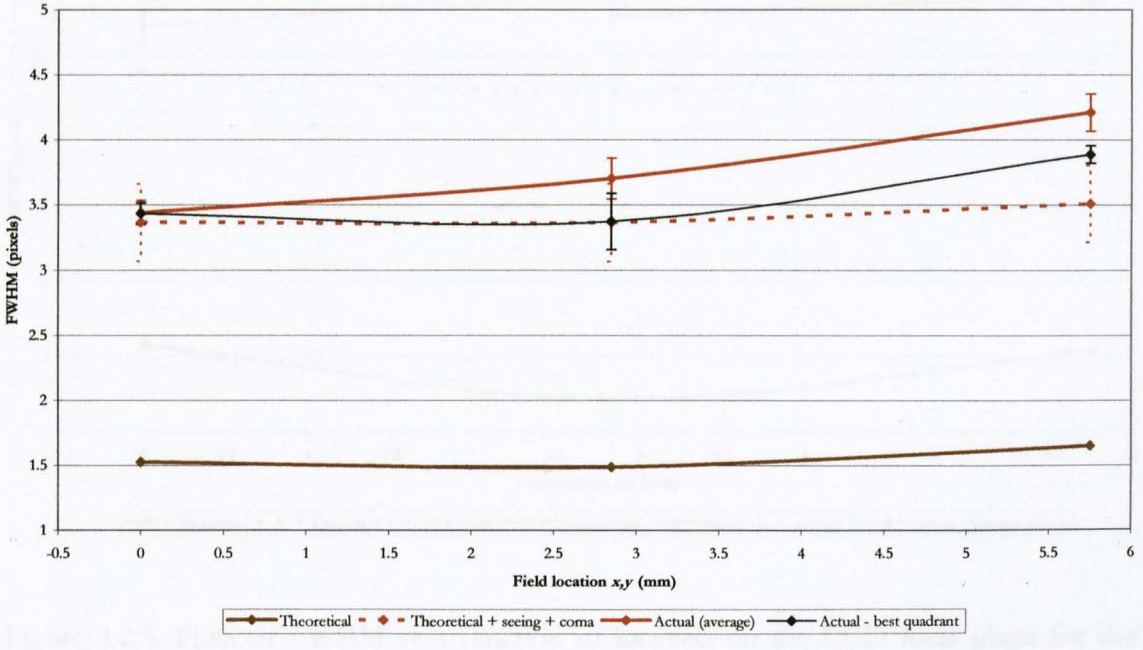

Figure 4.2.4: Plots of FWHM as a function of location on the CCD focal plane for the red optics. The top plot shows observed FWHM for each quadrant specified in figure 4.2.3. The bottom plot shows the average observed FWHM, theoretical FWHM and theoretical FWHM with seeing and coma contributions included. Also plotted on the bottom graph is the quadrant with the lowest observed FWHM, i.e. the quadrant with the best image quality.

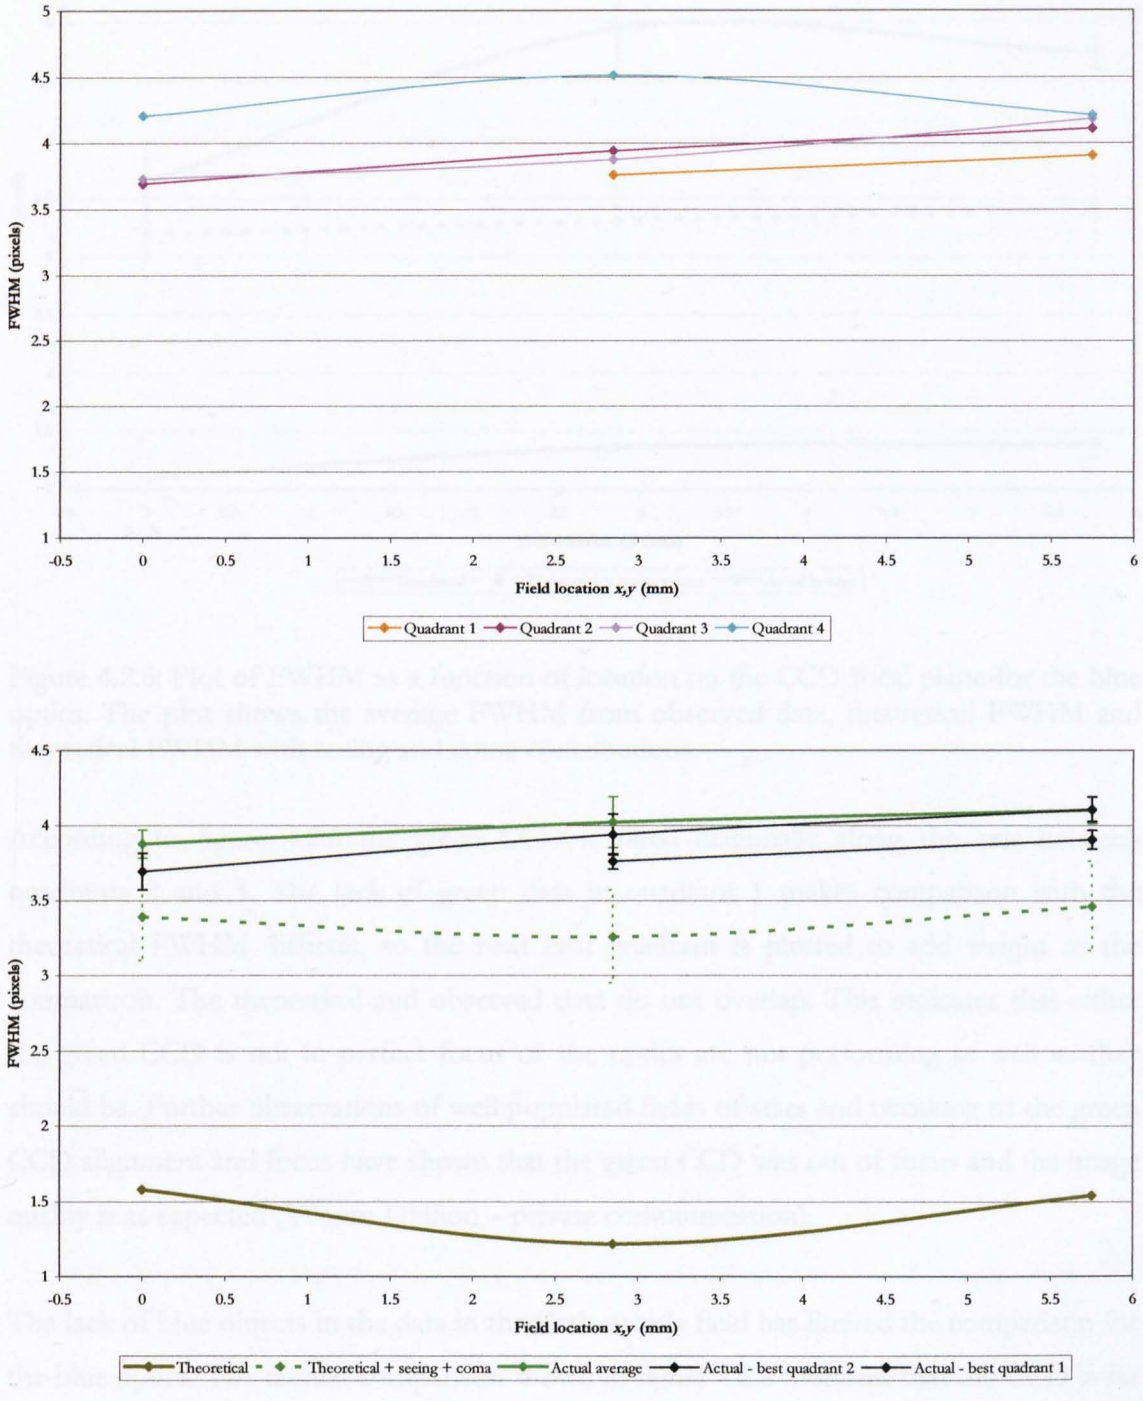

Figure 4.2.5: Plots of FWHM as a function of location on the CCD focal plane for the green optics. The top plot shows observed FWHM for each quadrant specified in figure 4.2.3. The bottom plot shows the average observed FWHM, theoretical FWHM and theoretical FWHM with seeing and coma contributions included. Also plotted on the bottom graph is the quadrant with the lowest observed FWHM which only has two data points, so the next best quadrant is also plotted.

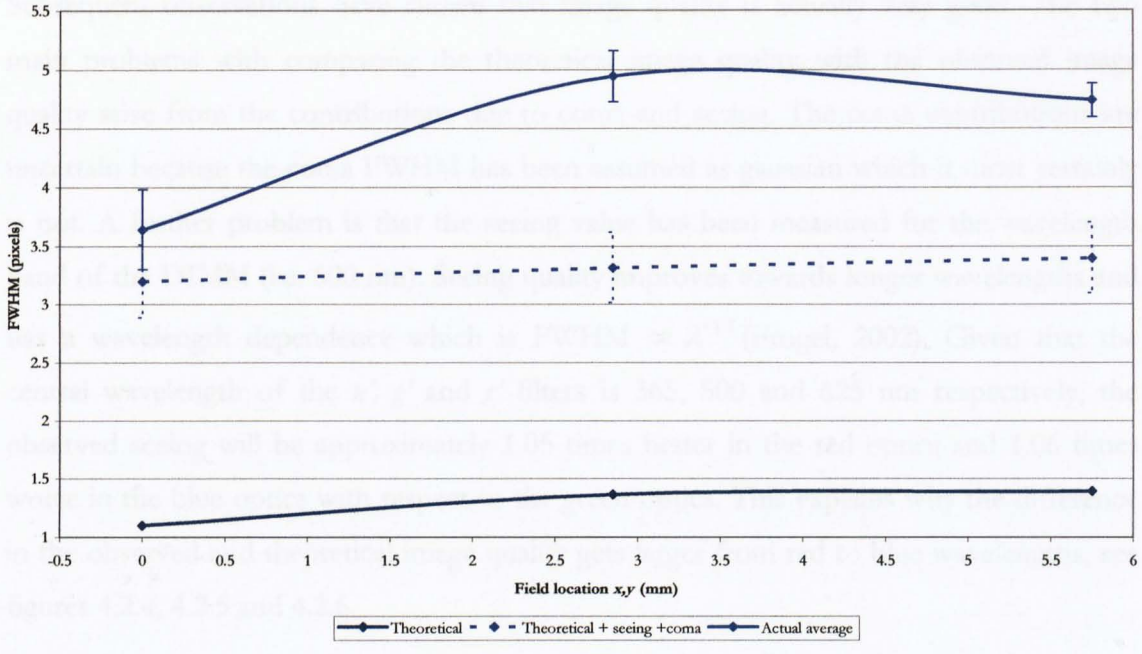

Figure 4.2.6: Plot of FWHM as a function of location on the CCD focal plane for the blue optics. The plot shows the average FWHM from observed data, theoretical FWHM and theoretical FWHM with seeing and coma contributions.

According to figure 4.2.5 the green CCD is tilted diagonally along the axis between quadrants 2 and 3. The lack of green data in quadrant 1 makes comparison with the theoretical FWHM difficult, so the next best quadrant is plotted to add weight to the comparison. The theoretical and observed data do not overlap. This indicates that either the green CCD is not in perfect focus or the optics are not performing as well as they should be. Further observations of well populated fields of stars and tweaking of the green CCD alignment and focus have shown that the green CCD was out of focus and the image quality is as expected (Vikram Dhillon – private communication).

The lack of blue objects in the data in the Crab nebula field has limited the comparison for the blue optics. The limited comparison shown in figure 4.2.6 indicates that the CCD is far from optimum focus. Observing a well populated field of blue stars and tweaking of the blue CCD alignment and focus would help determine whether the optics or the CCD alignment are affecting image quality. A subsequent observation during 0.6 arcseconds seeing measured the FWHM of the target as 0.6 arcseconds indicating that image quality is actually very good and only limited by atmospheric effects (Vikram Dhillon – private communication).

Subsequent observations have shown that image quality is actually very good. The two main problems with comparing the theoretical image quality with the observed image quality arise from the contributions due to coma and seeing. The coma contributions are uncertain because the coma FWHM has been assumed as gaussian which it most certainly is not. A further problem is that the seeing value has been measured for the wavelength band of the DIMM (i.e. 500 nm). Seeing quality improves towards longer wavelengths and has a wavelength dependence which is FWHM $\propto \lambda^{-1/5}$ (Frogel, 2002). Given that the central wavelength of the *u'*, *g'* and *r'* filters is 365, 500 and 625 nm respectively, the observed seeing will be approximately 1.05 times better in the red optics and 1.06 times worse in the blue optics with respect to the green optics. This explains why the difference in the observed and theoretical image quality gets larger from red to blue wavelengths, see figures 4.2.4, 4.2.5 and 4.2.6.

4.2.2 Throughput

Throughput is a measure of how much incident light entering an optical system emerges as detectable photons. The analysis presented in this section involves modelling the throughput and comparing this with the observed throughput calculated from observations of standard stars. Any discrepancies between modelled and measured throughput will indicate how well the ULTRACAM optics are operating. A large difference between the modelled and measured throughput could indicate, for example the AR coatings are not performing to specification, or that there is a source of vignetting in the light path.

Starting from above the atmosphere, the light from any celestial source is attenuated by the following processes on its journey to the ULTRACAM detectors:

1. **Atmospheric extinction.**
   Light from a celestial source will first be attenuated by the transmission of the atmosphere. Data for this has been obtained from Asahi Spectra Co. for airmass 1.

2. **Reflectivity of the telescope mirrors.**

   After the light from a source has passed through the atmosphere it is attenuated by reflection losses at the telescope primary and secondary mirrors, which were 83% and 88% efficient, respectively, at the WHT when the standard stars were observed (Benn, 2002). The mirror reflection losses are also adjusted for wavelength dependence using data on the wavelength dependence of aluminium reflectivity (Orion optics).

3. **ULTRACAM lens reflection losses.**

   All light entering ULTRACAM is attenuated by reflection losses at each lens surface. All lenses have an anti-reflection coating and typically reflect 2% of the incident light at each surface. Transmission curves for anti-reflection coatings applied to the glasses used in ULTRACAM have been provided by CVI Technical Optics.

4. **ULTRACAM lens transmission losses.**

   All light entering ULTRACAM is attenuated by the internal transmission of each lens. The amount of light absorbed is dependent on wavelength, the glass type and glass thickness. Data for the internal transmission of the glass types used in ULTRACAM has been obtained from the Schott AG, Mainz, Germany.

5. **Transmission/Reflectivity of the dichroic beam splitters.**

   The light now enters ULTRACAM and is attenuated by the dichroic beam-splitters. For u', the light is attenuated by the reflectivity of dichroic 1. For g', the light is attenuated by the transmission of dichroic 1 and the reflectivity of dichroic 2. For r', i' and z', the light is attenuated by the transmission of dichroics 1 and 2. Measured transmission curves for the ULTRACAM dichroic beam-splitters have been provided by CVI Technical Optics.

6. **Transmission of the filters.**

   After the dichroic beam-splitters and re-imaging cameras, the light then passes through the filters (SDSS for the model case). Measured transmission curves for the SDSS filters used in ULTRACAM have been provided by Asahi Spectra Co.

7. **CCD quantum efficiency.**

   The light finally reaches the CCD detectors. The blue and green CCDs have a 'broad band' AR coating and the red CCD has a 'mid band' AR coating to improve

their quantum efficiencies. The quantum efficiency curves for the ULTRACAM CCDs have been provided by E2V.

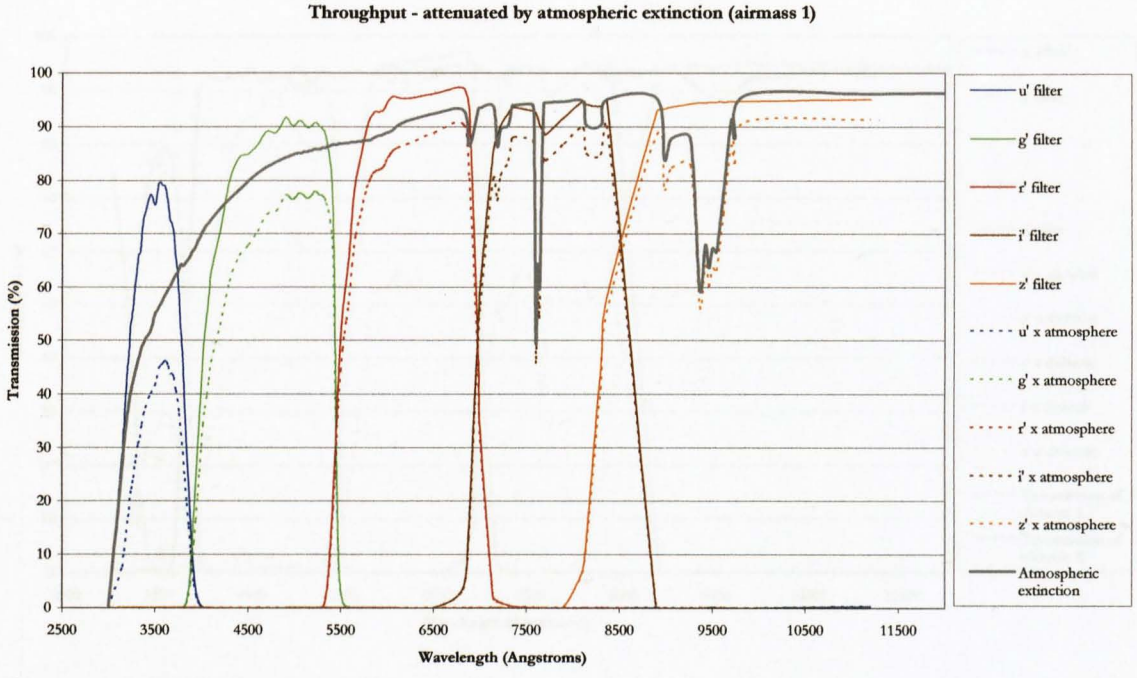

Figure 4.2.7: ULTRACAM filter profiles attenuated by atmospheric extinction at airmass 1.

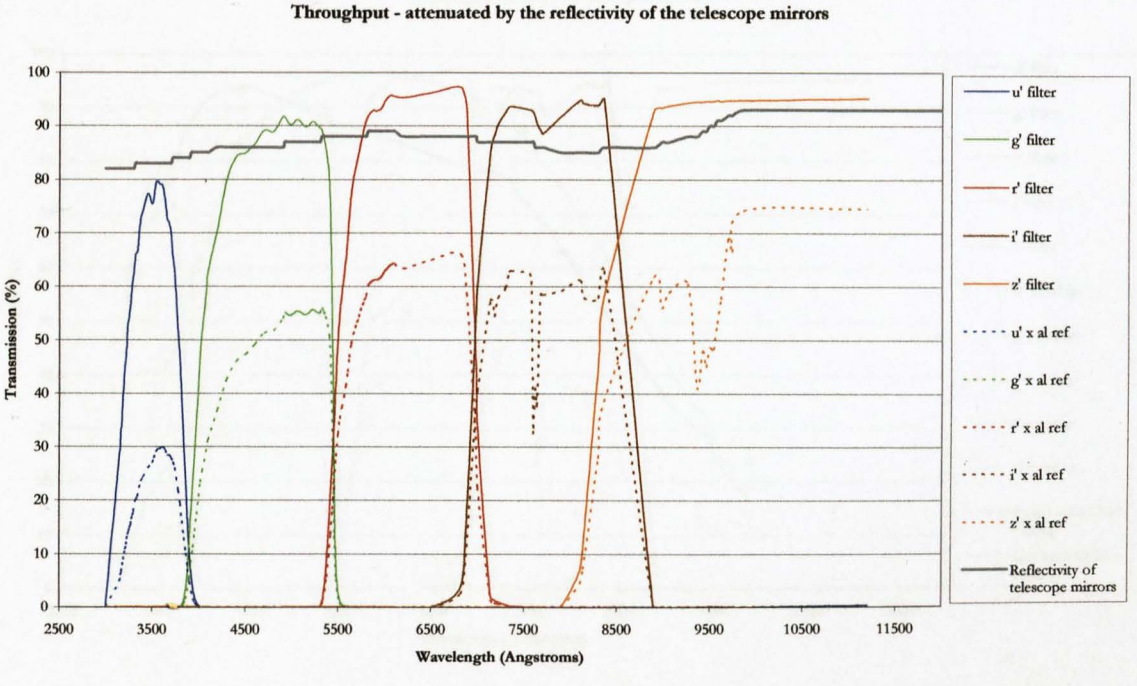

Figure 4.2.8: ULTRACAM filter profiles attenuated by atmospheric extinction and the reflectivity of the telescope mirrors.

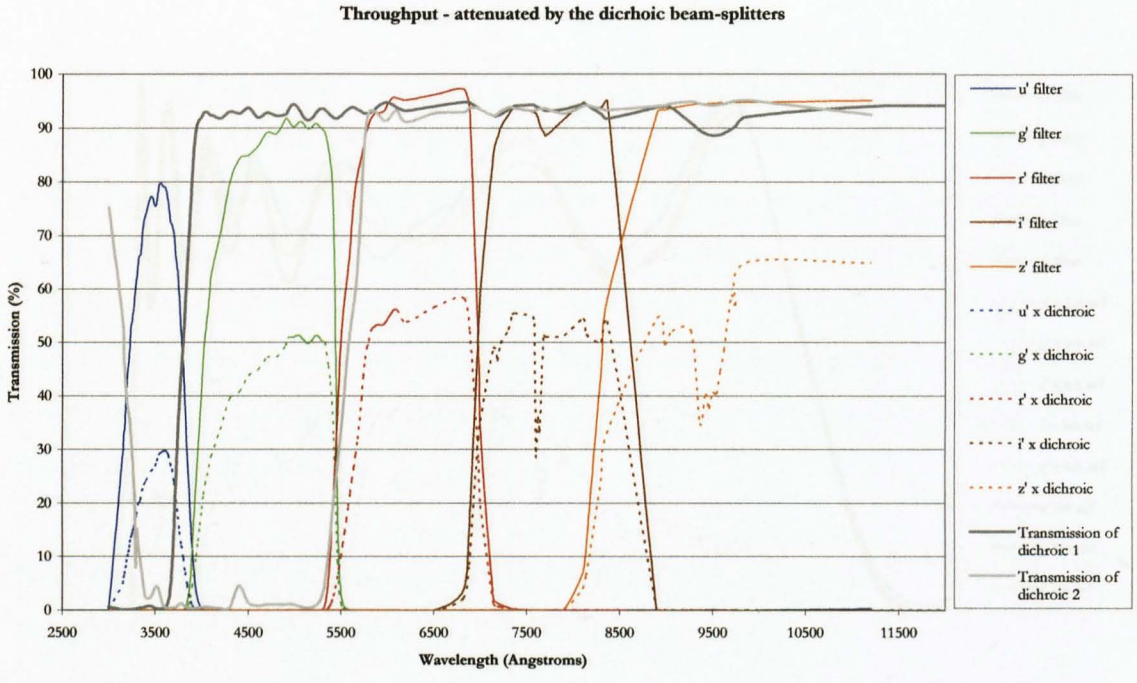

Figure 4.2.9: ULTRACAM filter profiles attenuated by atmospheric extinction, reflectivity of the telescope mirrors and the dichroic beam splitters.

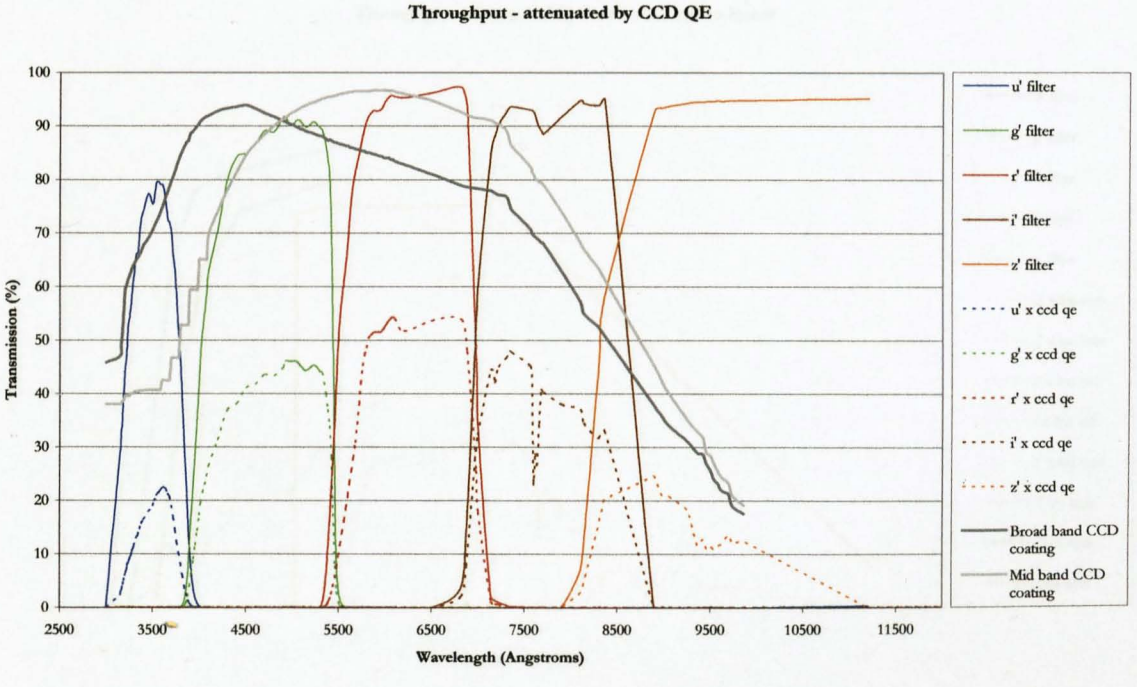

Figure 4.2.10: ULTRACAM filter profiles attenuated by atmospheric extinction, reflectivity of the telescope mirrors, dichroic beam-splitters and the CCD QE.

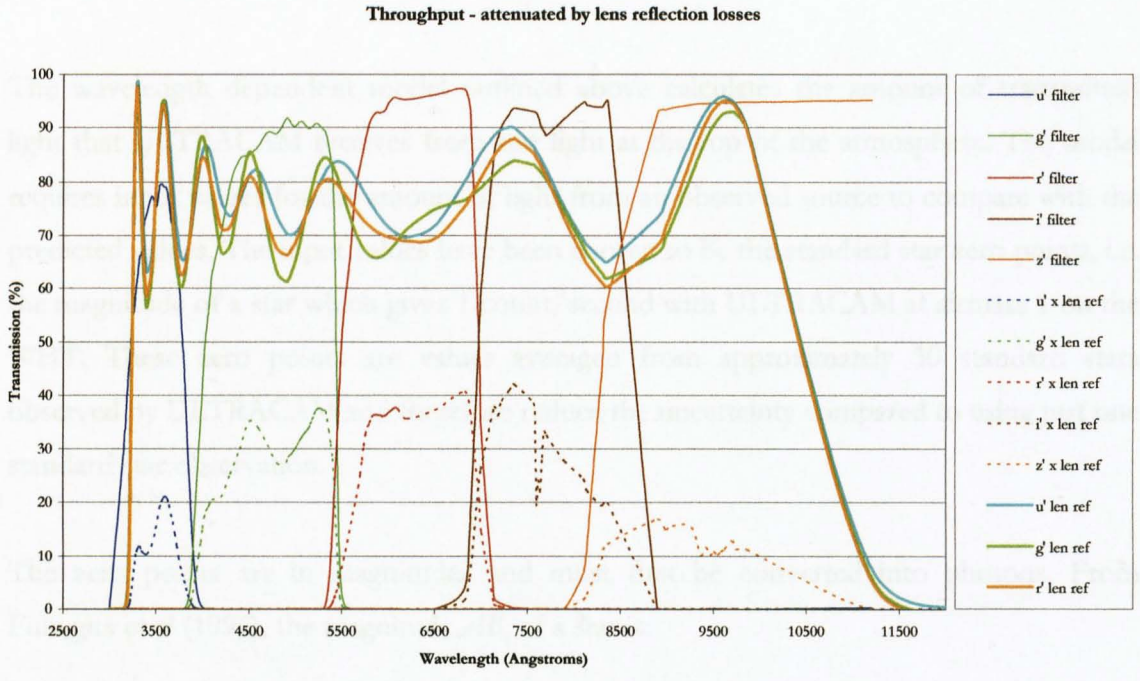

Figure 4.2.11: ULTRACAM filter profiles attenuated by atmospheric extinction, reflectivity of the telescope mirrors, dichroic beam-splitters, CCD QE and the total lens reflection losses.

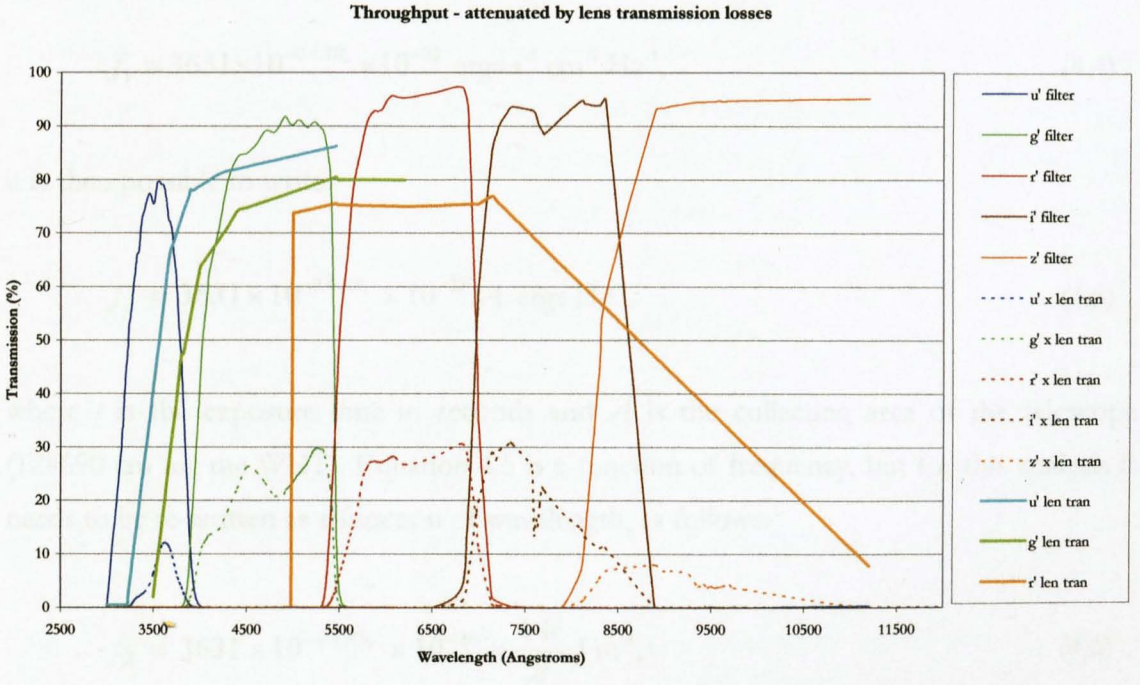

Figure 4.2.12: ULTRACAM filter profiles attenuated by atmospheric extinction, reflectivity of the telescope mirrors, dichroic beam-splitters, CCD QE, total lens reflections losses and the total internal lens transmission losses.

Figures 4.2.7 to 4.2.12 show the modelled attenuation of each stage, starting with the SDSS filter profiles.

The wavelength dependent model outlined above calculates the amount of transmitted light that ULTRACAM receives from star light at the top of the atmosphere. The model requires input values for the amount of light from an observed source to compare with the predicted values. The input values have been chosen to be the standard star zero points, i.e. the magnitude of a star which gives 1 count/second with ULTRACAM at airmass 1 on the WHT. These zero points are values averaged from approximately 30 standard stars observed by ULTRACAM and therefore reduce the uncertainty compared to using just one standard star observation.

The zero points are in magnitudes and must first be converted into photons. From Fukugita et al (1996), the magnitude AB_ν of a star is:

$$AB_\nu = -2.5 \log f_\nu - 48.60 \text{ ergs s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}, \quad (4.3)$$

where f_ν is the flux per unit frequency from the object. Rearranging gives:

$$f_\nu = 3631 \times 10^{-0.4 AB_\nu} \times 10^{-23} \text{ ergs s}^{-1} \text{ cm}^{-2} \text{ Hz}^{-1}, \quad (4.4)$$

it is then possible to write:

$$f_\nu = 3631 \times 10^{-0.4 AB_\nu} \times 10^{-23} tA \text{ ergs Hz}^{-1}, \quad (4.5)$$

where t is the exposure time in seconds and A is the collecting area of the telescope (124690 cm for the WHT). Equation 4.5 is a function of frequency, but for this analysis it needs to be re-written as a function of wavelength, as follows:

$$f_\lambda = 3631 \times 10^{-0.4 AB_\nu} \times 10^{-30} \times \frac{tAc}{\lambda^2} \text{ J m}^{-1}, \quad (4.6)$$

where a conversion from ergs to joules is also made. The number of photons from a source of flux, f_ν, is:

$$N_p = \frac{f_\lambda}{E} = \frac{f_\lambda}{hc/\lambda}, \tag{4.7}$$

where E is the energy of one photon of wavelength λ, h is the Planck constant and c is the velocity of light. Rearranging gives:

$$f_\lambda = \frac{N_p hc}{\lambda}. \tag{4.8}$$

Equation 4.8 can be substituted into equation 4.6 and re-arranged to give the total number of photons per angstrom for a given magnitude AB_ν.

$$N_p = 5.48\times10^6 \times 10^{-0.4AB_\nu} \times \frac{tA}{\lambda} \text{ photons / angstrom.} \tag{4.9}$$

Using equation 4.9 the total number of photons incident on ULTRACAM has been calculated from each zero point. This was done for wavelength intervals centred at wavelength, λ. The total number of photons per wavelength interval is then multiplied by the estimated throughput of ULTRACAM in each interval to give the number of detected electrons per wavelength interval, where the photons here have been converted to electrons due to the CCD quantum efficiency term in the throughput prediction. The throughput corrected electrons per wavelength interval are then summed to give the total number of electrons in each filter. Finally, the total number of detected electrons in each filter, are converted to counts using the measured CCD gain value (see section 4.3.1). The resulting predicted and measured counts are listed in table 4.1. The uncertainty in the measured counts arises from the uncertainty in the zero point values.

The predicted counts should be equal to the measured counts. Table 4.1 also shows the percentage error in the prediction. A negative percentage indicates that the predicted throughput is lower than the measured value, whilst a positive percentage indicates that the

predicted throughput is higher than the measured value. With the exception of *r'*, the predicted throughput always is under-estimated.

Table 4.1: Results of the ULTRACAM throughput modelling.

	u'	*g'*	*r'*	*i'*	*z'*
Zero point	24.55	26.69	26.21	26.28	25.24
CCD gain (electrons/ADU)	1.200	1.150	1.225	1.225	1.225
Predicted counts	0.846	0.849	1.139	0.808	0.667
Measured counts	1 ± 0.04	1 ± 0.03	1 ± 0.05	1 ± 0.04	1 ± 0.04
Percentage error in model	-15.4	-15.1	13.9	-19.2	-33.3

The deviation of predicted counts from measured counts is due to one or more variables in the throughput model being underestimated. There is data in the throughput model that has relatively low confidence associated with it, such as the lens transmission losses and the transmission of the atmosphere. The throughput model uses values for internal transmission for the glass types used in ULTRACAM from Schott AG. The internal transmission values given are for glasses of thickness 25 mm and 10 mm. The lenses in ULTRACAM have a central thickness range of 2.5 mm to 17.6 mm. The throughput model takes all of the lenses of the same glass type and adds their thickness and chooses the combinations of same glass types that are closest to 25 mm and 10 mm central thickness. However, these were never exact and therefore some uncertainty exists in the estimation of the lens transmission losses.

The largest uncertainty will be due to underestimating the transmission of the atmosphere. The atmosphere data used in the model has no information regarding the amount of dust and water present during the observation. This will have a significant effect on the transmission properties of the atmosphere as a function of wavelength.

Figure 4.2.13 shows the predicted throughput of the ULTRACAM optics alone. Also shown on the graph is the upper limit of throughput (lower limit for *r'*). This is the value of throughput that would be achieved if the model estimated all variables perfectly. The upper limit curves peak at 49% in *u'*, 65% in *g'*, 49%, in *r'*, 66% in *i'*, 42% in *z'* and the lower limit curves peak at 41% in *u'*, 56% in *g'*, 43 in *r'*, 54% in *i'* and 28% in *z'*. Figure 4.2.13 also shows that these values would have dropped to 19% in *u'*, 26% in *g'*, 24% in *r'*, 23% in *i'* and 15% in *z'* if no AR coatings had been used.

Figure 4.2.14 shows the predicted throughput of the optics and the CCD detectors. Finally, figure 4.2.15 shows the predicted throughput of the optics, CCD detectors and the transmission of the WHT primary and secondary mirrors; this shows the expected throughput of the whole ULTRACAM – telescope system. Figures 4.2.14 and 4.2.15 both show the upper limit throughput values.

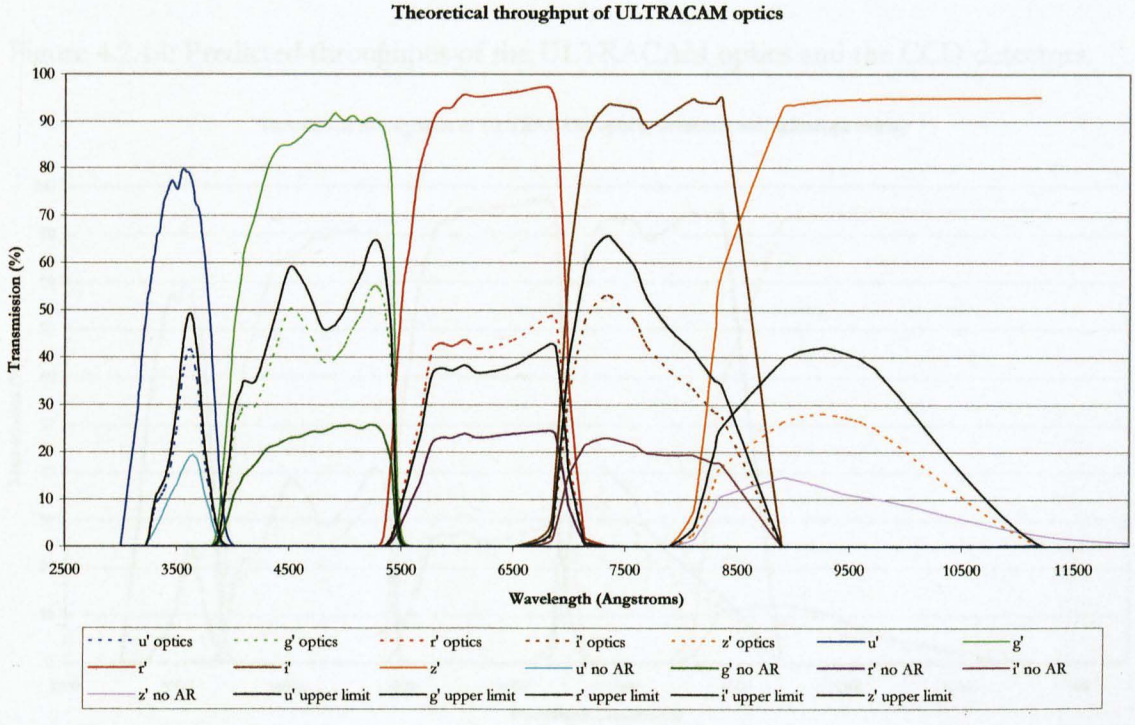

Figure 4.2.13: Predicted throughput of ULTRACAM optics only. Also shown are curves for the optics with no anti-reflection coating applied and the upper limit of throughput.

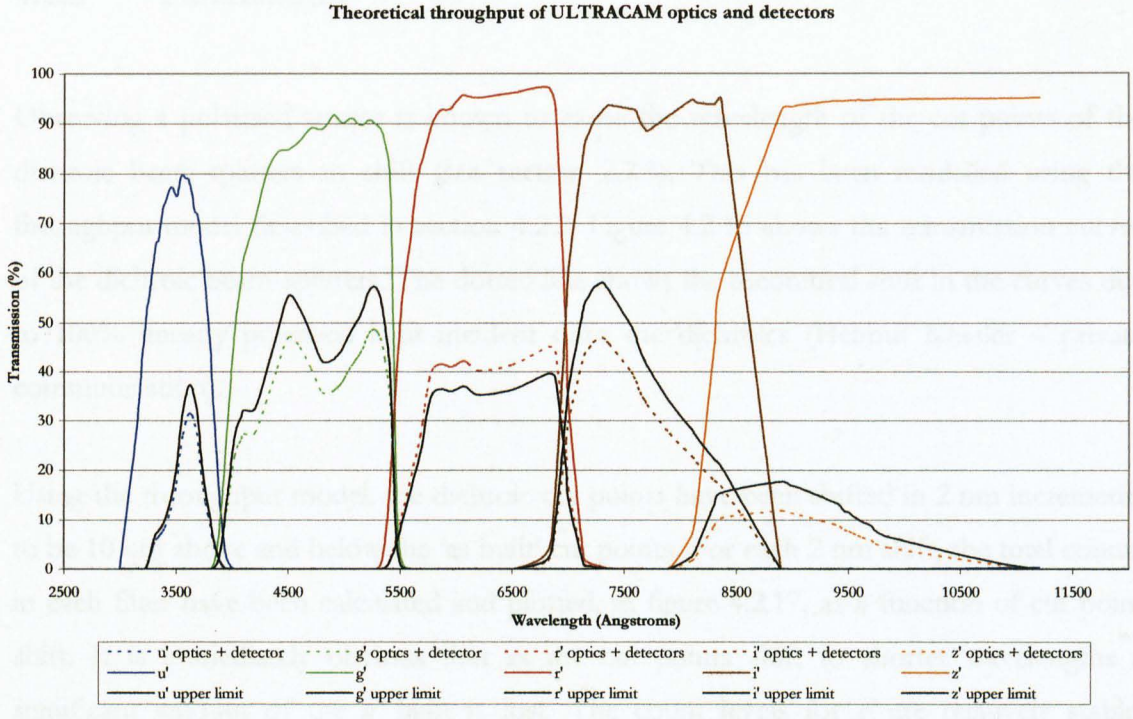

Figure 4.2.14: Predicted throughput of the ULTRACAM optics and the CCD detectors.

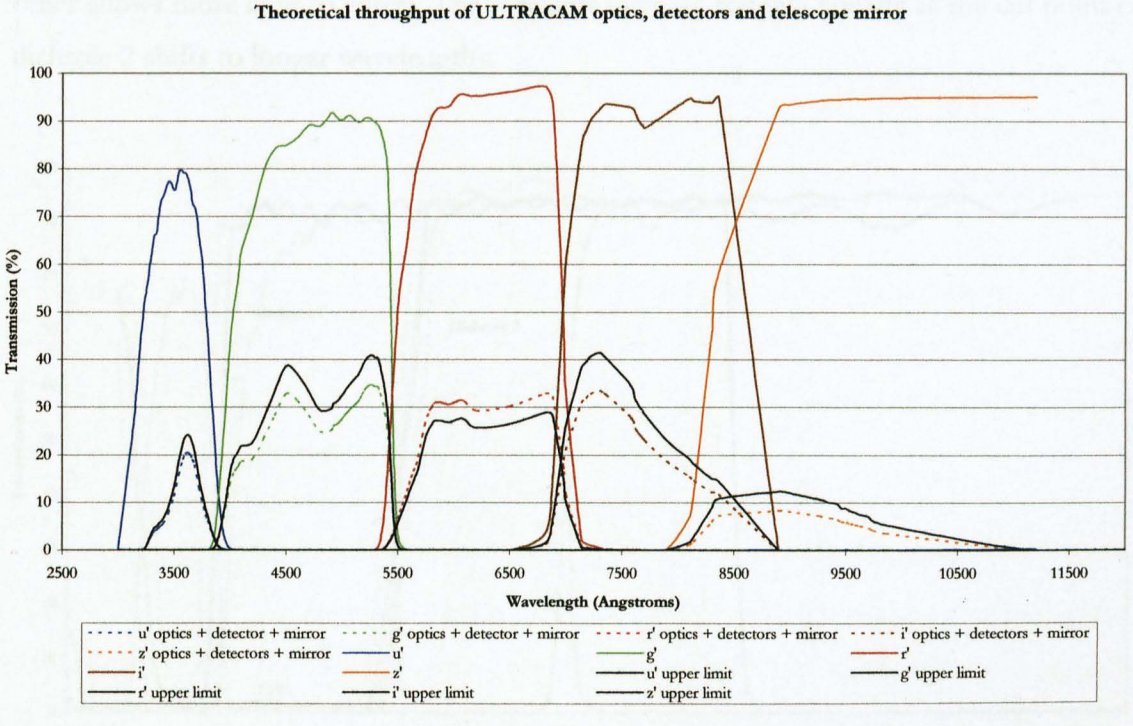

Figure 4.2.15: Predicted throughput of the ULTRACAM optics, the QE of the CCD detectors and the transmission of the WHT primary and secondary mirrors.

4.2.3 Polarisation

Observing a polarised source is known to cause the wavelength of the cut points of the dichroic beam-splitters to shift (see section 2.2.3). This has been modelled using the throughput model described in section 4.2.2. Figure 4.2.16 shows the transmission curves of the dichroic beam-splitters. The dotted line shows the theoretical shift in the curves due to 100% linearly polarised light incident onto the dichroics (Helmut Kessler – private communication).

Using the throughput model, the dichroic cut points have been shifted in 2 nm increments to be 10 nm above and below the 'as built' cut points. For each 2 nm shift, the total counts in each filter have been calculated and plotted, in figure 4.2.17, as a function of cut point shift. It is immediately obvious that as the cut points shift to shorter wavelengths a significant amount of the *u'* light is lost. The count levels for *g'* are relatively stable, however, because this band is affected by both dichroic cut-points shifting, which tend to cancel each other out (i.e. one dichroic cut removes light from the pass-band whilst the other allows more light to enter). The *r'* counts increase roughly linearly as the cut point of dichroic 2 shifts to longer wavelengths.

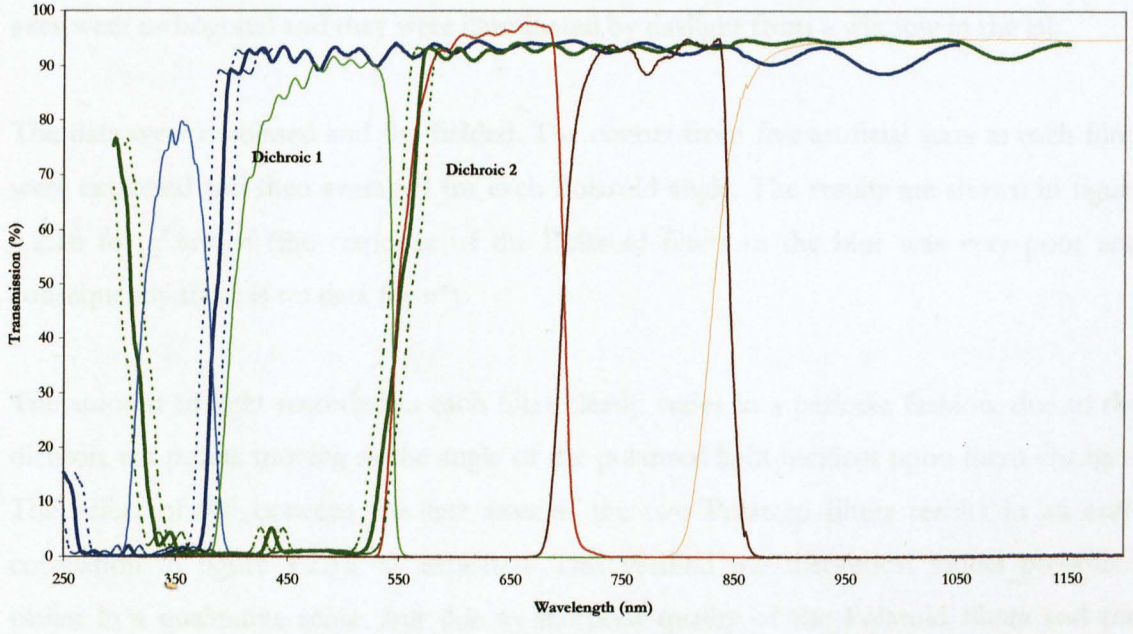

Figure 4.2.16: Transmission of the dichroic beam-splitters shown with the SDSS filter curves. The dotted line shows the theoretical shift of the dichroic cut points due to 100% linearly polarised light incident on the dichroics.

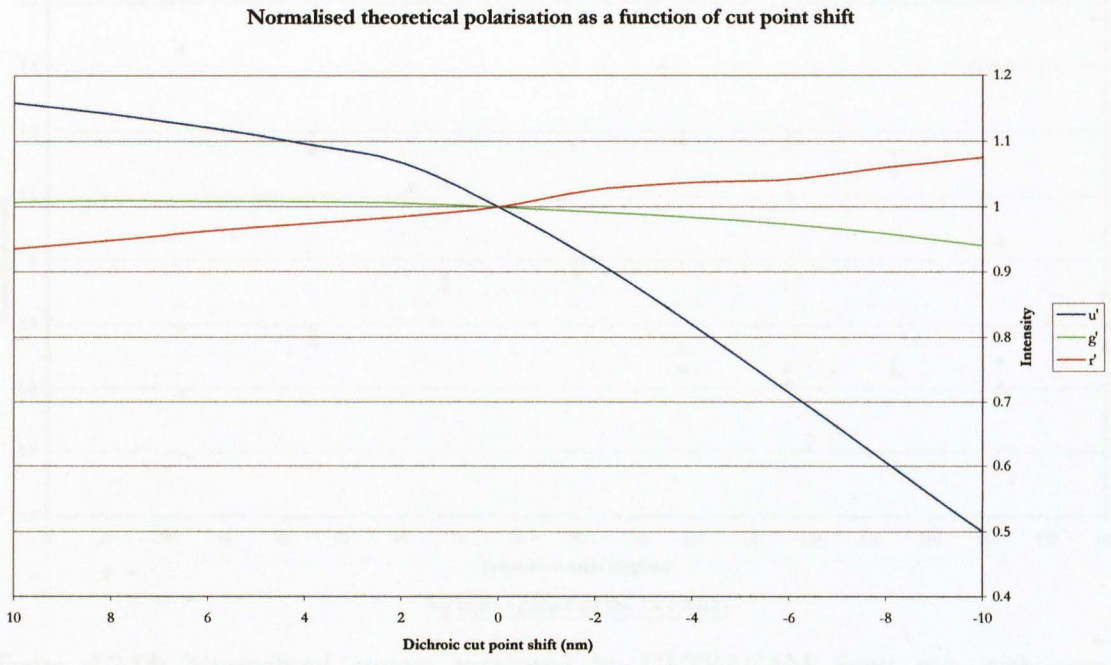

Figure 4.2.17: Normalised counts in each filter as a function of the shift in the dichroic cut-point from the 'as built' value.

An experiment in the laboratory was conducted to quantitatively test the above model. Two Polaroid filters were fixed to the artificial star-field tube (for details of the artificial star-field tube see section 3.3) and the filters were rotated through 180° in 15° increments while ULTRACAM was imaging. The Polaroid filters were positioned so that their fast axes were orthogonal and they were illuminated by daylight from a window in the lab.

The data were de-biased and flat-fielded. The counts from five artificial stars in each filter were extracted and then averaged for each Polaroid angle. The results are shown in figure 4.2.18 for *g'* and *r'* (the response of the Polaroid filters in the blue was very poor and consequently there is no data for *u'*).

The amount of light recorded in each filter clearly varies in a periodic fashion, due to the dichroic cut points moving as the angle of the polarised light incident upon them changes. The offset of 90° between the fast axes of the two Polaroid filters results in an anti-correlation in figure 4.2.18, as expected. This verified the theoretical model presented earlier in a qualitative sense, but due to the poor quality of the Polaroid filters and the changing daylight sky levels, no quantitative information can be obtained from this experiment.

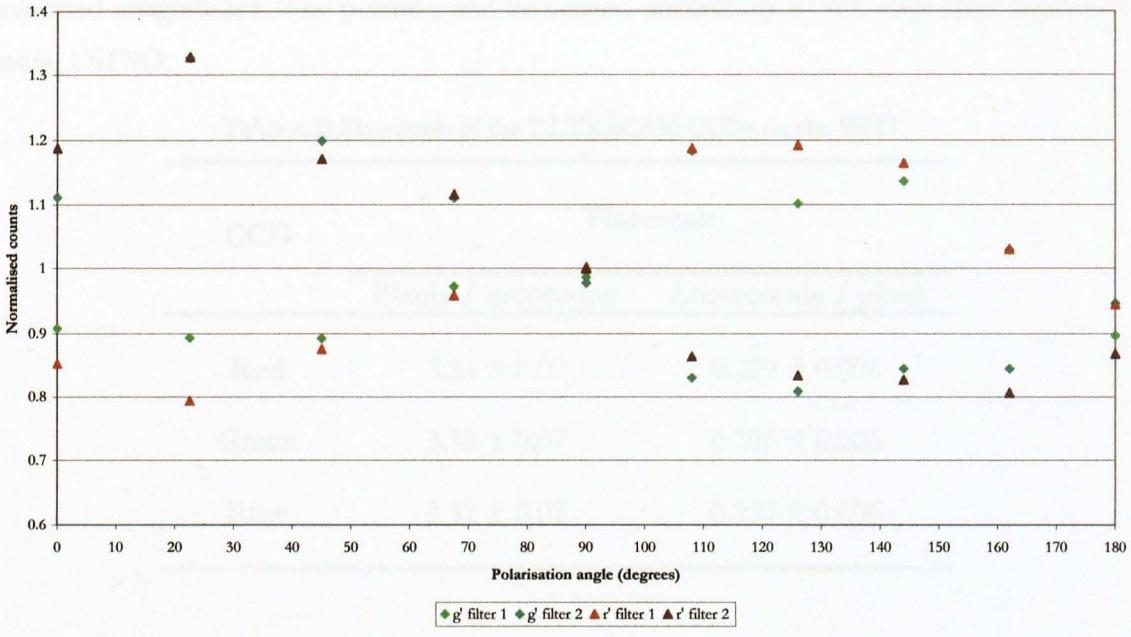

Figure 4.2.18: Normalised counts measured by ULTRACAM from two orthogonal Polaroid filters as a function of their (arbitrary) angle relative to the dichroic beam-splitters. Both *g'* and *r'* data is plotted.

To quantify the effect of polarisation, it would be better to observe a zero polarisation standard star and insert high quality Polaroid filters between the telescope and ULTRACAM. The telescope Cassegrain rotator could then be used to rotate ULTRACAM with respect to the filters and hence change the plane of polarisation of the incident light. This would be a much more controlled experiment, but unfortunately, such an experiment is not currently possible on the WHT.

4.2.4 Platescale and CCD orientation

The platescale and orientation of each ULTRACAM CCD was measured from data frames where the WHT has been moved a known offset. The WHT offsetting is accurate to < 0.2 arcseconds (Benn, 2001). Data from September 9th 2002 have a 1 arcminute shift East between runs 7 and 8. The x and y position of 3 stars were measured in run 7. The same stars were then measured again in run 8 and the linear shift between the two sets of pixel coordinates were calculated to give the platescale. Table 4.2 shows the platescale data and figure 4.2.19 shows the orientation of the CCDs on the WHT. The uncertainties quoted are due to the accuracy in measuring the star positions which is approximately 0.4 pixels for each position measurement. This will dominate the offsetting error, so they have been

assumed insignificant. The position and associated uncertainty of the stars were measured using DIPSO.

Table 4.2: Platescale of the ULTRACAM CCDs on the WHT

CCD	Platescale	
	Pixels / arcsecond	Arcseconds / pixel
Red	3.34 ± 0.07	0.299 ± 0.006
Green	3.38 ± 0.07	0.296 ± 0.006
Blue	3.37 ± 0.07	0.297 ± 0.006

We measured the angle, 164.44°, between the vertical plane of the CCD and the North position of the telescope with ROT SKY PA = 0. This was then input into the telescope control system as an 'Instrument ULTRACAM' setting, so that now, ROT SKY PA = 0 corresponds to North equals parallel to the y – plane of the CCD and East to the x – plane of the CCD, where east points towards decreasing x – pixel values.

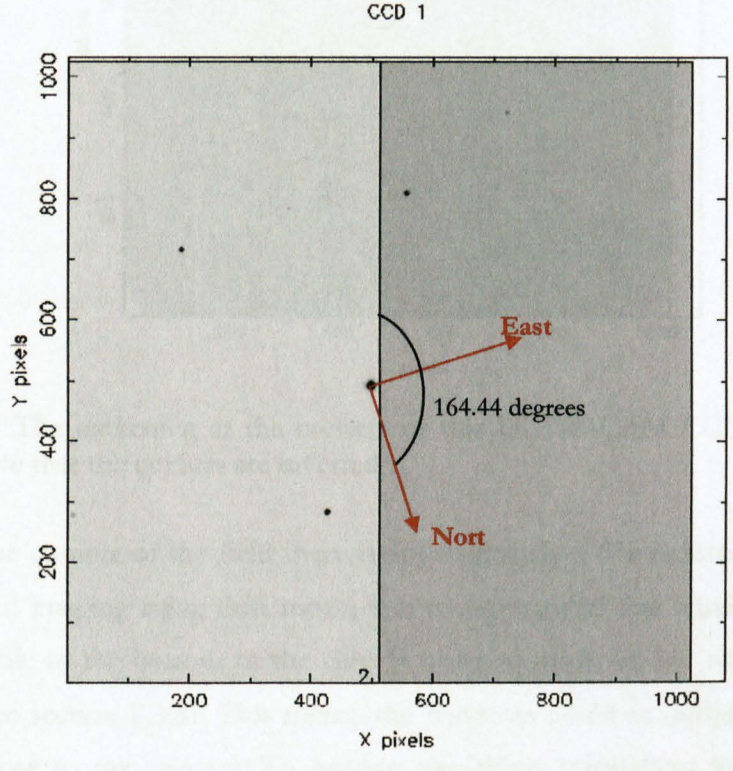

Figure 4.2.19: Orientation of the CCDs when ULTRACAM is fitted to the WHT at ROT SKY PA = 0. Note that ULTRACAM is always fitted to the telescope in the same orientation.

4.2.5 Vignetting

Vignetting is the unintended obstruction of light falling onto the focal plane. An example of vignetting in an ULTRACAM image is shown in figure 4.2.20. The corners of the field are darkened due to the size of the clear apertures in the optics.

To quantify vignetting a standard star was imaged at the field centre and the telescope then moved so that the star shifted from the centre of the field to each corner, where another image was obtained. The images were bias subtracted and flat fielded and the total counts from the star in each position were extracted. Table 4.3 lists the percentage decrease in total counts from the star at the field corners compared to the field centre for each ULTRACAM channel.

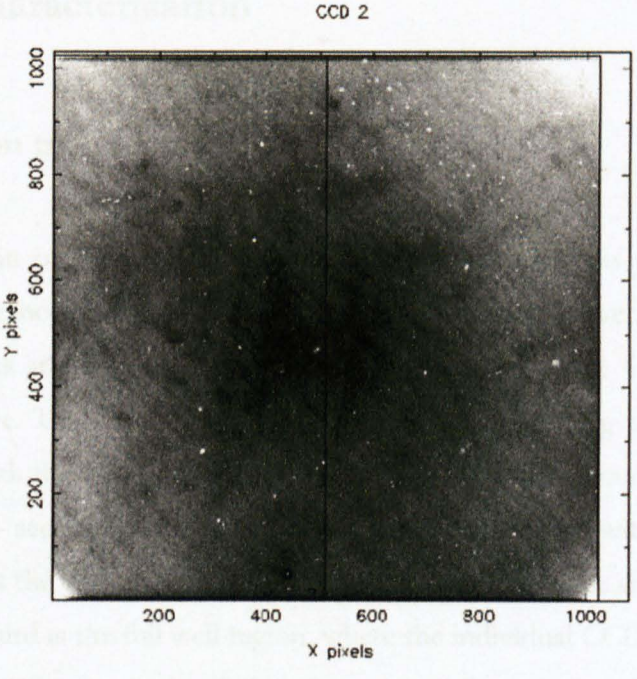

Figure 4.2.20: The darkening at the corners of this ULTRACAM CCD image illustrates vignetting (note that the colours are inverted).

At the extreme corners of the field there is approximately a 6% reduction in throughput. For high speed imaging using drift mode, it is recommended that windows are placed as close as possible to the bottom of the chip in order to minimise the number of V-clocks performed (see section 2.3.3). This means the windows could be subject to vignetting if placed too close to the corners. To prevent significant throughput loss, it is therefore recommended that the windows should be approximately 100 pixels away from the edge of the CCD.

Table 4.3: Vignetting in ULTRACAM

Result quoted as percentage throughput decrease at the field corners compared to the field centre.

CCD	Vignetting (%)
Red	6.29 ± 0.01
Green	6.12 ± 0.01
Blue	6.60 ± 0.05

4.3 CCD characterisation

4.3.1 Photon transfer curve

The characterisation of CCD gain, read noise and granularity[7] was performed using the photon transfer method (Howard, 2002). The photon transfer curve is a log – log plot of noise versus counts and has the distinctive shape shown in figure 4.3.1. There are three regions in the curve. The first is the read noise region, representing the noise floor of the CCD associated with the output amplifier (see section 1.5). The second region is shot noise (or Poisson noise – see section 1.5.2) dominated, and is associated with the random arrival times of photons at the detector; this noise is inherent to the source and does not originate in the CCD. The third is the full well region, where the individual CCD pixels are unable to hold additional charge; the peak of the photon transfer curve represents the full well capacity of the CCD.

The data for this analysis were obtained in the laboratory, where the CCDs were uniformly illuminated by a light source and the exposure time incrementally increased. Data was recorded in both fast and slow readout modes (see section 2.3). For the data reduction, each CCD readout channel was treated independently where the mean standard deviation

[7] Size of the pixel – to – pixel sensitivity variations which are normally corrected by flat fielding.

of the pixel counts were calculated for each half. These values were then averaged and plotted as log counts versus log noise for each frame.

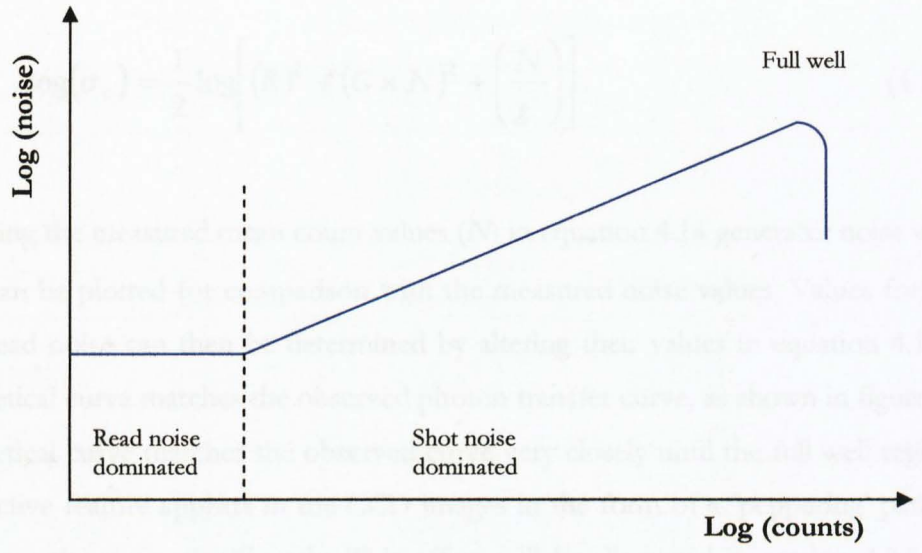

Figure 4.3.1: An illustrative photon transfer curve showing the three separate noise regions described in the text.

To quantify the CCD gain, read noise and granularity a noise model has been derived. The noise sources add in quadrature to give the total noise, σ_T, as follows:

$$\sigma_T^2 = \sigma_R^2 + \sigma_S^2 + \sigma_G^2, \tag{4.10}$$

where σ_R is the noise attributed to read noise, σ_S is the noise attributed to shot noise and σ_G is the noise attributed to granularity. We can define each noise source as:

$$\sigma_R^2 = R^2, \tag{4.11}$$

$$\sigma_S^2 = \left(\frac{\sigma_e}{g}\right)^2 = \frac{N_e}{g^2} \qquad (\sigma_e = \sqrt{N_e}), \tag{4.12}$$

$$\sigma_G^2 = (G \times N)^2, \tag{4.13}$$

where R, is the read noise in counts, σ_e is the shot noise in electrons, g is gain in electrons/ADU (ADU and counts are the same thing), N_e is mean number of electrons, G

is the grain noise with a value between 0 – 1 and N is mean number of counts. Combining equations 4.11 – 4.13 and taking logs gives:

$$\log(\sigma_T) = \frac{1}{2}\log\left[(R)^2 + (G \times N)^2 + \left(\frac{N}{g}\right)\right]. \qquad (4.14)$$

Entering the measured mean count values (N) in equation 4.14 generates noise values (σ_T) that can be plotted for comparison with the measured noise values. Values for gain, grain and read noise can then be determined by altering their values in equation 4.14 until the theoretical curve matches the observed photon transfer curve, as shown in figure 4.3.2. The theoretical curve matches the observed curve very closely until the full well region. Here a distinctive feature appears in the CCD images in the form of a 'peppering' pattern, which increases the noise significantly. This effect will be discussed in section 4.3.3. Table 4.4 quotes the results of the gain, grain and read noise measurements for each channel of each CCD. The uncertainty has been measured by fitting the model curve to the upper and lower limits of the observed data.

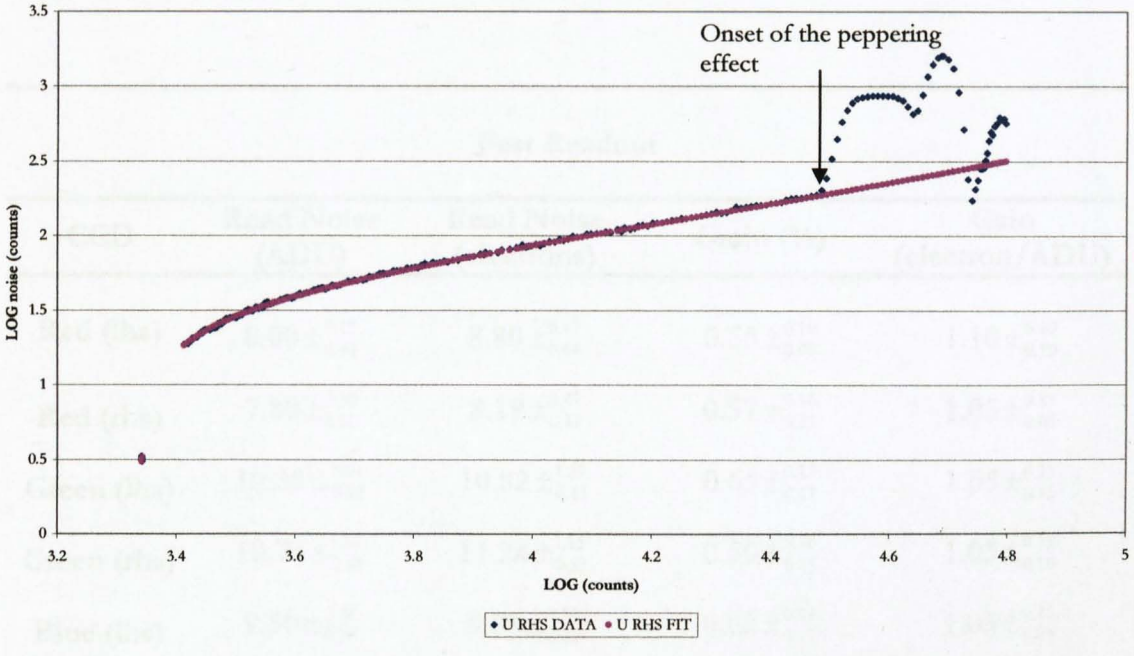

Figure 4.3.2: Photon transfer curve for the left-hand readout channel of the blue CCD. Both the measured and a theoretical fit to it are plotted. At high count levels a 'peppering' pattern appears on the CCD, resulting in the sharp increase in noise evident at the highest light levels.

Table 4.4: Photon transfer curve results

Read noise, grain and gain measured from the photon transfer curves for each CCD in slow and fast readout mode. Values are presented for the right hand side (rhs) and left hand side (lhs) of the chip.

Slow Readout

CCD	Read Noise (ADU)	Read Noise (electrons)	Grain (%)	Gain (electron/ADU)
Red (lhs)	$3.60 \pm^{0.40}_{0.55}$	$4.50 \pm^{0.50}_{0.69}$	$0.52 \pm^{0.08}_{0.09}$	$1.25 \pm^{0.30}_{0.20}$
Red (rhs)	$3.65 \pm^{0.45}_{0.35}$	$4.38 \pm^{0.42}_{0.42}$	$0.55 \pm^{0.20}_{0.05}$	$1.20 \pm^{0.30}_{0.15}$
Green (lhs)	$3.70 \pm^{0.50}_{0.60}$	$4.26 \pm^{0.57}_{0.69}$	$0.62 \pm^{0.18}_{0.17}$	$1.15 \pm^{0.10}_{0.15}$
Green (rhs)	$3.80 \pm^{0.50}_{0.50}$	$4.37 \pm^{0.58}_{0.58}$	$0.45 \pm^{0.15}_{0.15}$	$1.15 \pm^{0.10}_{0.15}$
Blue (lhs)	$3.60 \pm^{0.60}_{0.60}$	$4.32 \pm^{0.72}_{0.72}$	$0.62 \pm^{0.10}_{0.40}$	$1.20 \pm^{0.25}_{0.15}$
Blue (rhs)	$3.40 \pm^{0.30}_{0.50}$	$4.08 \pm^{0.36}_{0.60}$	$0.65 \pm^{0.25}_{0.20}$	$1.20 \pm^{0.25}_{0.10}$

Fast Readout

CCD	Read Noise (ADU)	Read Noise (electrons)	Grain (%)	Gain (electron/ADU)
Red (lhs)	$8.00 \pm^{0.15}_{0.40}$	$8.80 \pm^{0.17}_{0.44}$	$0.55 \pm^{0.10}_{0.09}$	$1.10 \pm^{0.10}_{0.10}$
Red (rhs)	$7.80 \pm^{0.30}_{0.30}$	$8.19 \pm^{0.41}_{0.41}$	$0.57 \pm^{0.16}_{0.12}$	$1.05 \pm^{0.15}_{0.05}$
Green (lhs)	$10.30 \pm^{1.20}_{0.50}$	$10.82 \pm^{1.57}_{0.11}$	$0.65 \pm^{0.15}_{0.17}$	$1.05 \pm^{0.15}_{0.15}$
Green (rhs)	$10.70 \pm^{1.10}_{0.40}$	$11.24 \pm^{1.15}_{0.53}$	$0.50 \pm^{0.20}_{0.15}$	$1.05 \pm^{0.15}_{0.10}$
Blue (lhs)	$9.50 \pm^{0.50}_{0.30}$	$9.98 \pm^{0.52}_{0.32}$	$0.62 \pm^{0.23}_{0.17}$	$1.05 \pm^{0.15}_{0.05}$
Blue (rhs)	$8.70 \pm^{0.30}_{0.60}$	$9.57 \pm^{0.33}_{0.66}$	$0.65 \pm^{0.25}_{0.19}$	$1.10 \pm^{0.10}_{0.10}$

4.3.2 Linearity

For a perfect CCD detector the conversion of photons to electrons is a linear process. In practice, CCDs are imperfect and deviate from linearity at high and low light levels (Fellers and Davidson). CCD linearity is measured by plotting a range of exposure times versus counts whilst observing a non-variable source. The data for this analysis were recorded with ULTRACAM fitted to the WHT during the daytime, with the telescope dome closed and all dome lights off. The tungsten lamp and neutral density filters in the A&G box were used to evenly illuminate the CCDs.

Data were acquired for both fast and slow readout modes and each CCD readout channel was treated independently. The clipped mean and RMS for one large window in each CCD channel were determined and plotted versus exposure time. A straight line was fitted to the linear region of the graphs and the equation of the line used to predict count values. The difference between the observed and predicted count values have been plotted in figures 4.3.3 and 4.3.4 in both fast and slow readout modes. For each CCD there is a region between approximately 3000 and 35000 counts where the deviation from linearity is within 0.1%; table 4.5 shows count levels on each CCD.

Table 4.5: Linearity results

Bias subtracted count levels above and below the minimum and maximum values will be in the non-linear region of the CCDs.

	Left-hand side of CCD		Right-hand side of CCD	
	Minimum	**Maximum**	**Minimum**	**Maximum**
Red CCD				
Fast readout speed	160	36400	750	43000
Slow readout speed	1450	42300	1500	50200
Green CCD				
Fast readout speed	2800	37500	2900	39000
Slow readout speed	1650	33400	1750	35500
Blue CCD				
Fast readout speed	3000	43000	2700	38000
Slow readout speed	1350	54500	1850	48000

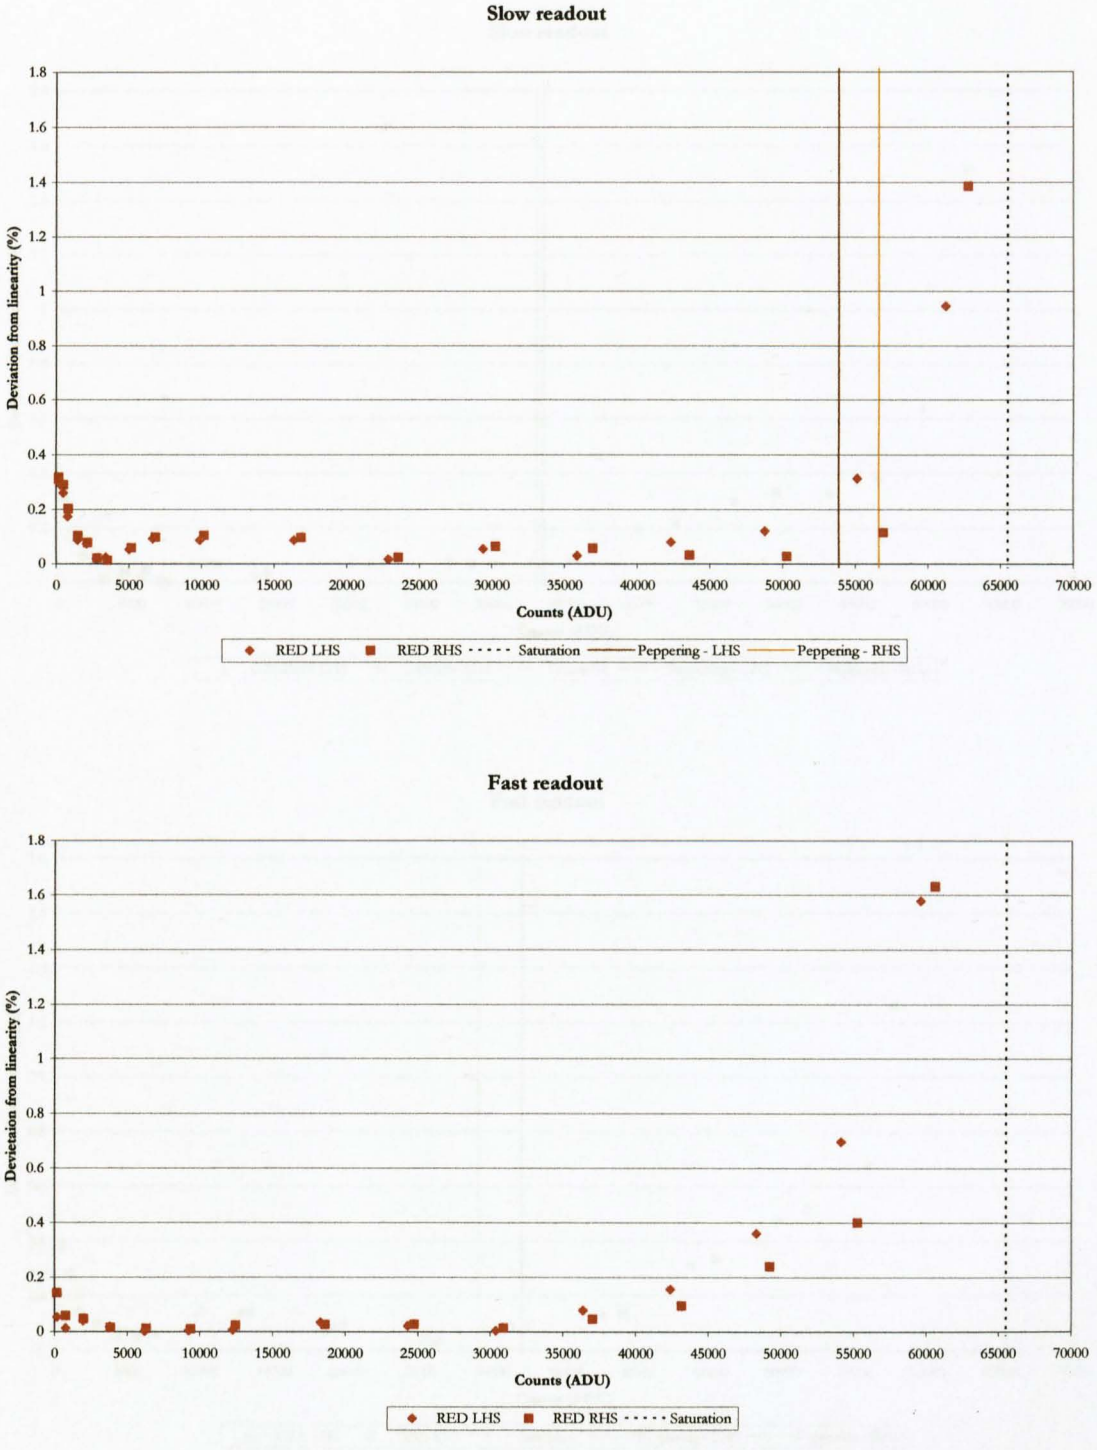

Figure 4.3.3: Graphs of percentage deviation from linearity as a function of bias subtracted count level for the red CCD. Plots for fast and slow readout are presented and each plot shows linearity data for each readout channel of the CCD. The onset of the peppering effect in each CCD half is shown in the top graph.

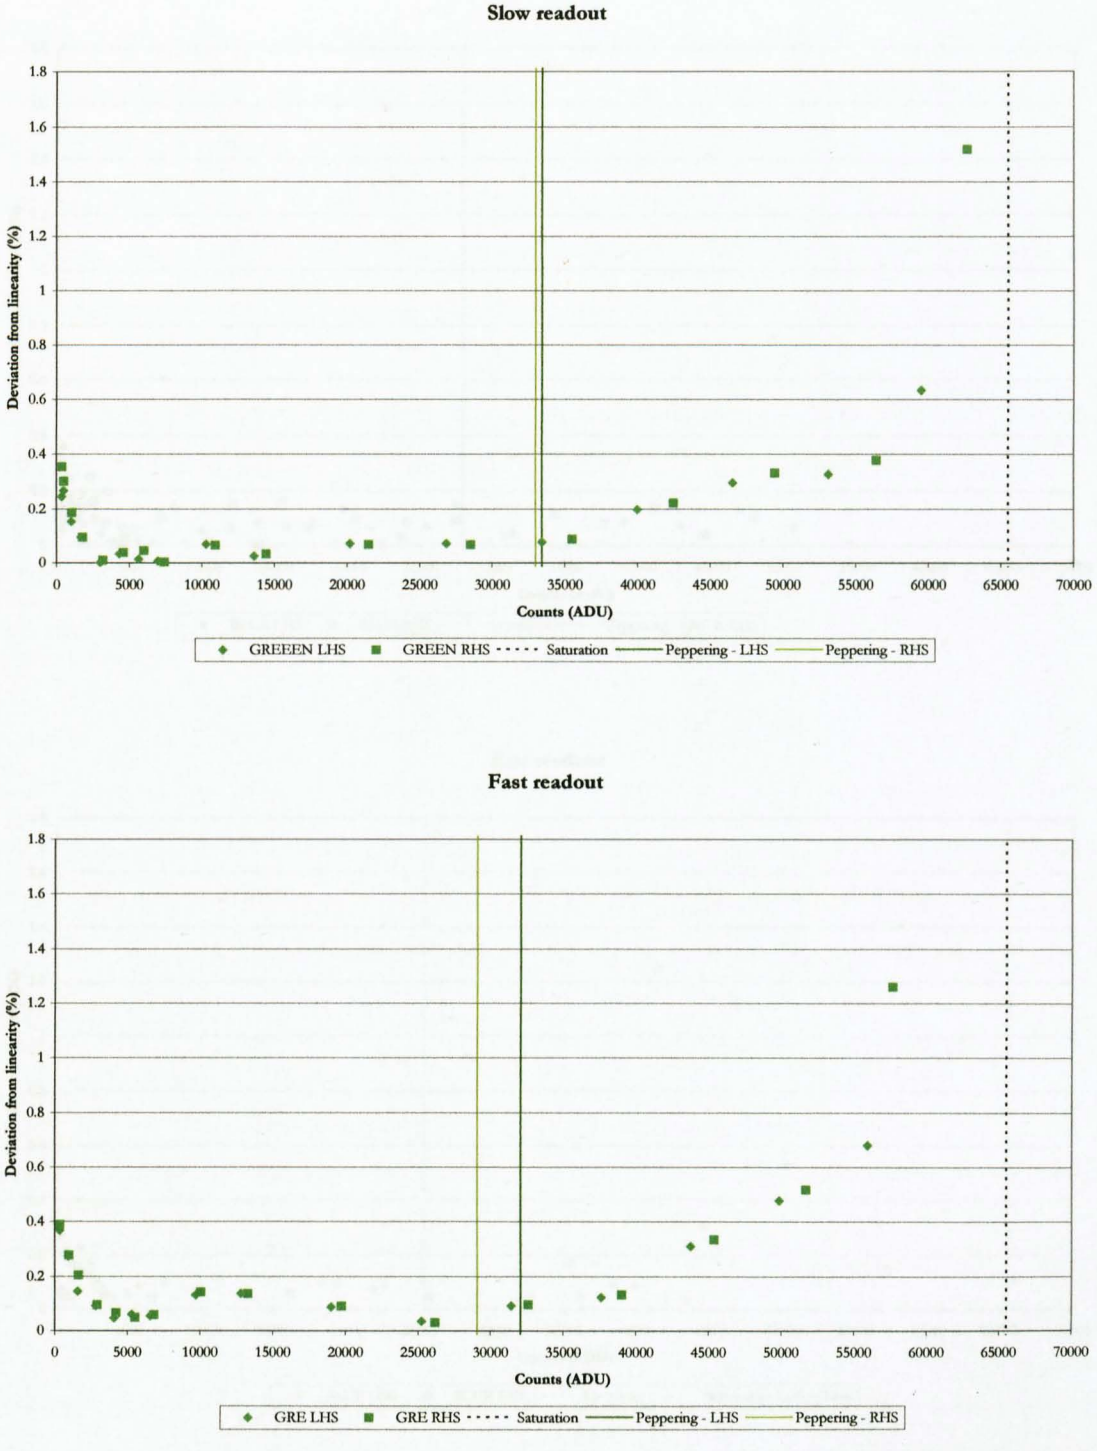

Figure 4.3.4: Graphs of percentage deviation from linearity as a function of bias subtracted count level for the green CCD. Plots for fast and slow readout are presented and each plot shows linearity data for each readout channel of the CCD. The onset of the peppering effect in each CCD half is shown in the top graph.

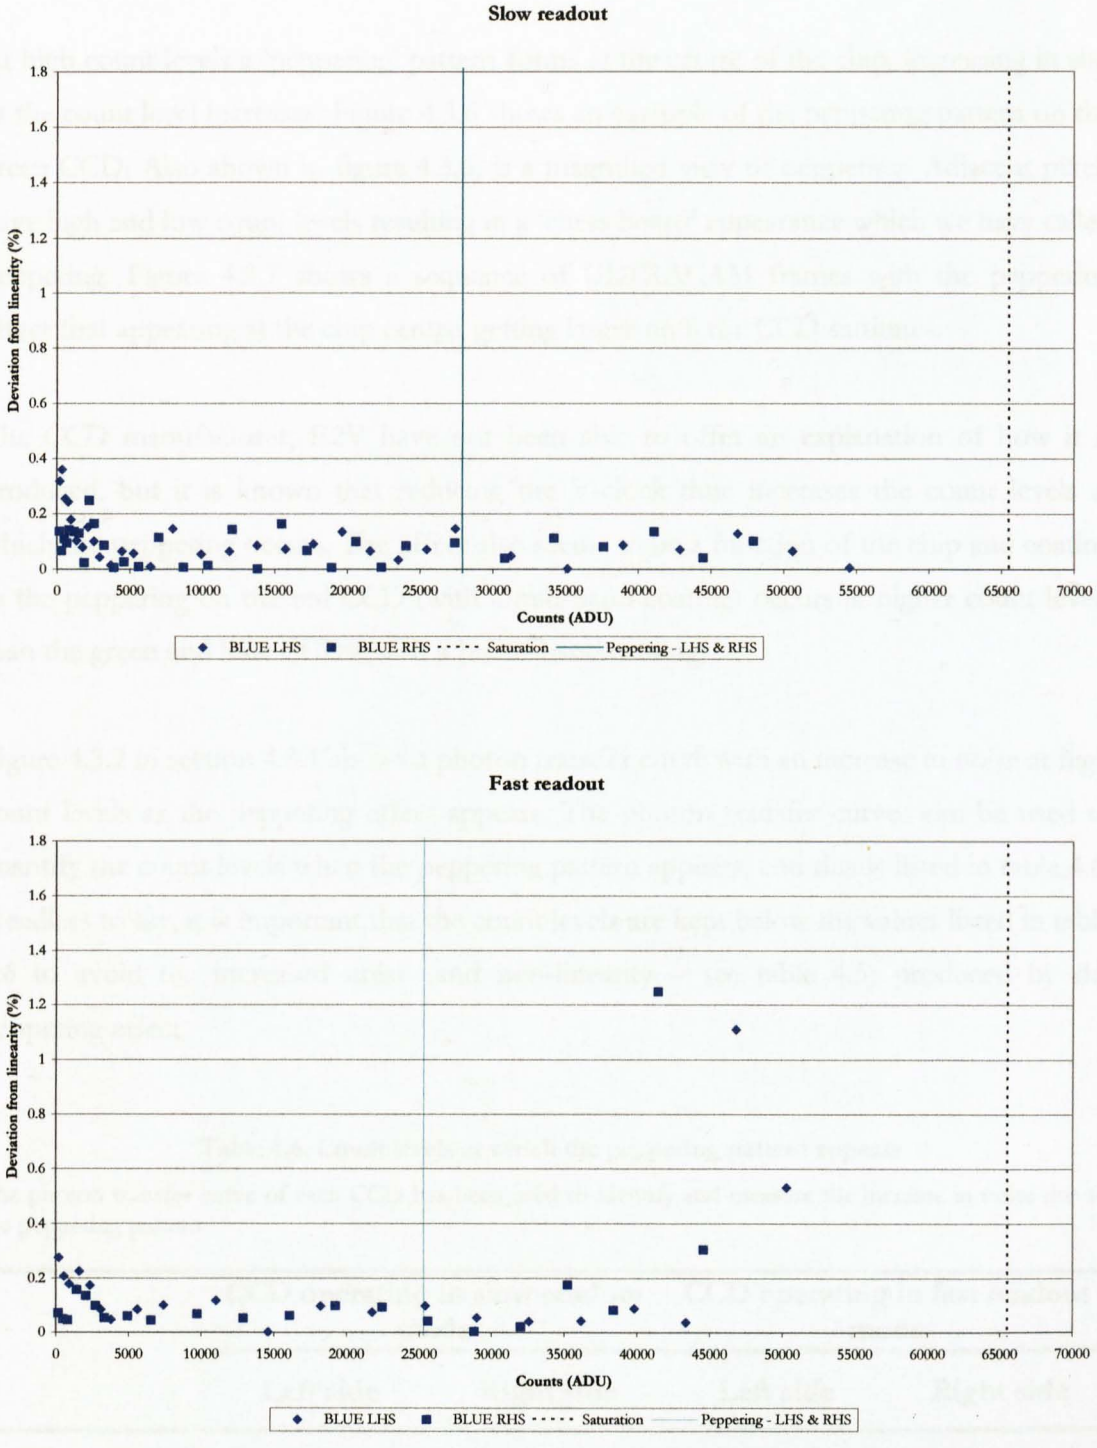

Figure 4.3.5: Graphs of percentage deviation from linearity as a function of bias subtracted count level for the blue CCD. Plots for fast and slow readout are presented and each plot shows linearity data for each readout channel of the CCD. The onset of the peppering effect in each CCD half is shown in the top graph.

4.3.3 Peppering

At high count levels a 'peppering' pattern forms at the centre of the chip, increasing in size as the count level increases. Figure 4.3.6 shows an example of the peppering pattern on the green CCD. Also shown in figure 4.3.6, is a magnified view of peppering. Adjacent pixels have high and low count levels resulting in a 'chess board' appearance which we have called peppering. Figure 4.3.7 shows a sequence of ULTRACAM frames with the peppering effect first appearing at the chip centre, getting larger until the CCD saturates.

The CCD manufacturer, E2V have not been able to offer an explanation of how it is produced, but it is known that reducing the V-clock time increases the count levels at which the peppering occurs. The effect also seems to be a function of the chip and coating as the peppering on the red CCD (with a mid-band coating) occurs at higher count levels than the green and blue CCDs (with a broad-band coating).

Figure 4.3.2 in section 4.3.1 shows a photon transfer curve with an increase in noise at high count levels as the peppering effect appears. The photon transfer curves can be used to quantify the count levels when the peppering pattern appears, and this is listed in table 4.6. Needless to say, it is important that the count levels are kept below the values listed in table 4.6 to avoid the increased noise (and non-linearity – see table 4.5) produced by the peppering effect

Table 4.6: Count levels at which the peppering pattern appears

The photon transfer curve of each CCD has been used to identify and measure the increase in noise due to the peppering pattern.

	CCD operating in slow readout mode		CCD operating in fast readout mode	
	Left side	**Right side**	**Left side**	**Right side**
Red CCD	56000	58800	No effect observed	
Green CCD	35400	35000	34600	31600
Blue CCD	30200	30200	28100	28100

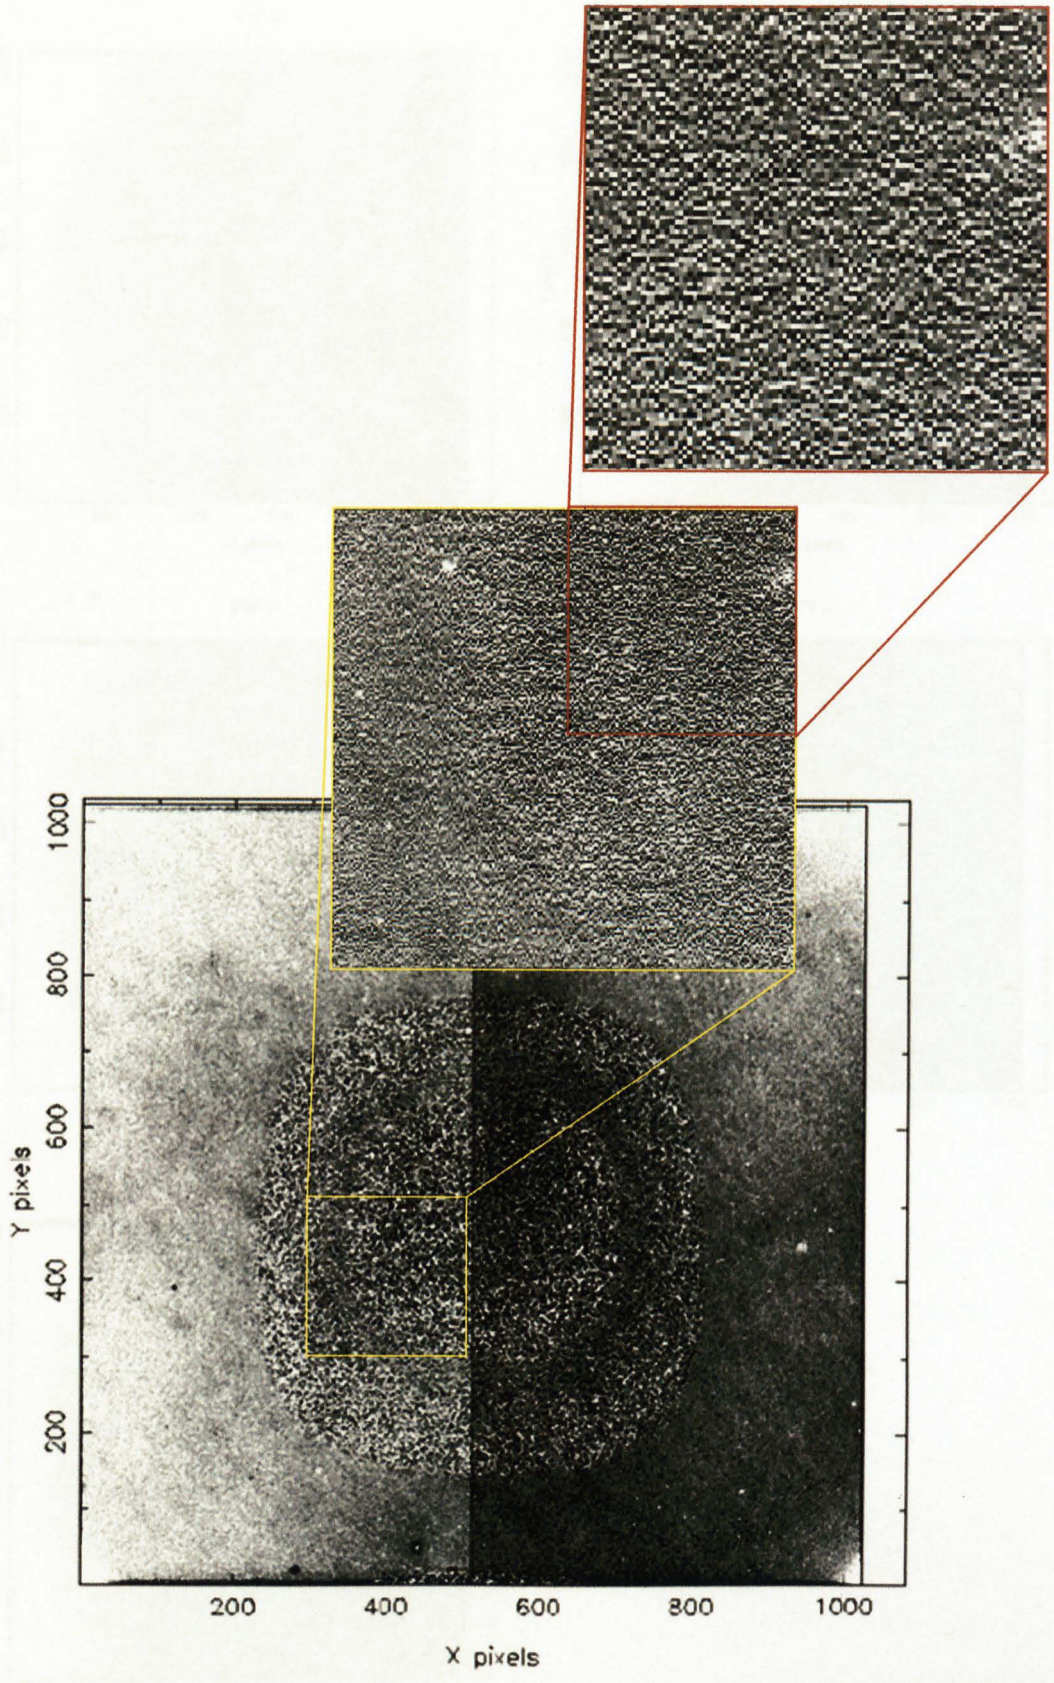

Figure 4.3.6: An example of the peppering pattern that forms on the CCDs at high count levels. The pattern first appears at the chip centre and radiates outwards as the count level increases. The example shown is the green CCD. Magnified views of peppering are also shown.

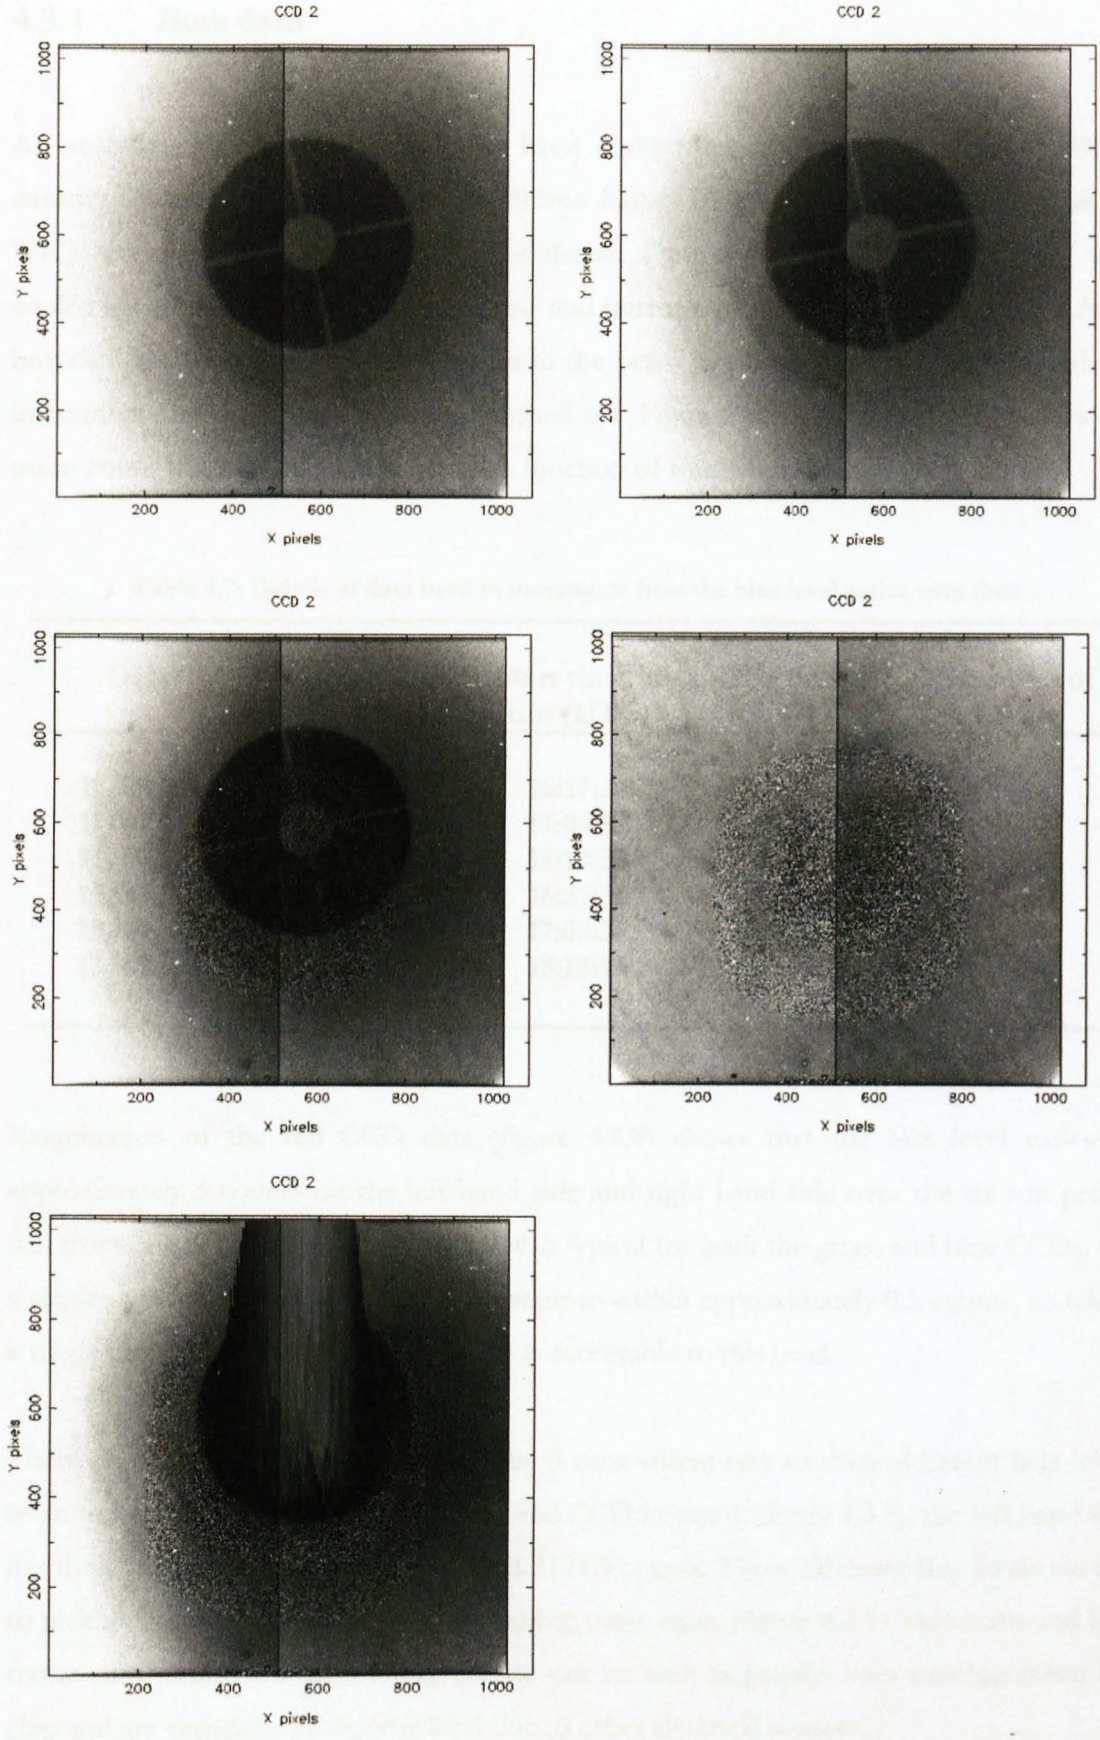

Figure 4.3.7: Example sequence of the peppering pattern appearing, starting at the top left.

4.3.4 Bias data

An analysis of ULTRACAM data has been undertaken to investigate the time varying nature of the bias level. A total of 28000 bias frames obtained in six separate runs on the WHT were analysed – see table 4.7 for details. During data collection the CCDs were cooled to -40°C and the observatory dome and mirror covers closed. In addition, the A&G box comparison lamp mirror was placed in the beam to prevent light passing through the instrument and the dome lights were turned off. Figures 4.3.8, 4.3.9 and 4.3.10 show the mean count levels in the bias frames as a function of frame number.

Table 4.7: Details of data used to investigate how the bias level varies with time

Date	Run number	Start time of run (UTC)	End time of run (UTC)	Number of frames
11/09/03	4	16:17:27	17:03:19	5442
11/09/03	5	17:04:52	17:51:14	5500
11/09/03	6	18:13:34	18:33:37	2380
13/09/03	1	16:31:35	17:02:24	3656
13/09/03	2	17:02:38	18:09:06	7884
13/09/03	3	18:18:45	18:44:09	3014

Examination of the red CCD data (figure 4.3.8) shows that the bias level varies by approximately 3 counts on the left hand side and right hand side over the six run period (i.e. from one night to the next). This is also typical for both the green and blue CCDs. On a single night, the bias level remains constant to within approximately 0.5 counts, so taking a single bias run at the start of each night is acceptable to this level.

There is a feature present in 5 out of the 6 runs where two to three different bias levels seem to be present. For example, on the red CCD in run 4 (figure 4.3.8), the left hand side has three bias levels at 2123.8, 2124.1 and 2124.5 counts. These different bias levels are due to pickup noise, which was a problem during these runs. Figure 4.3.11 shows the red bias frame with pickup features. These feature can be seen as parallel lines running down the chip and are variations in the bias level due to other electrical sources.

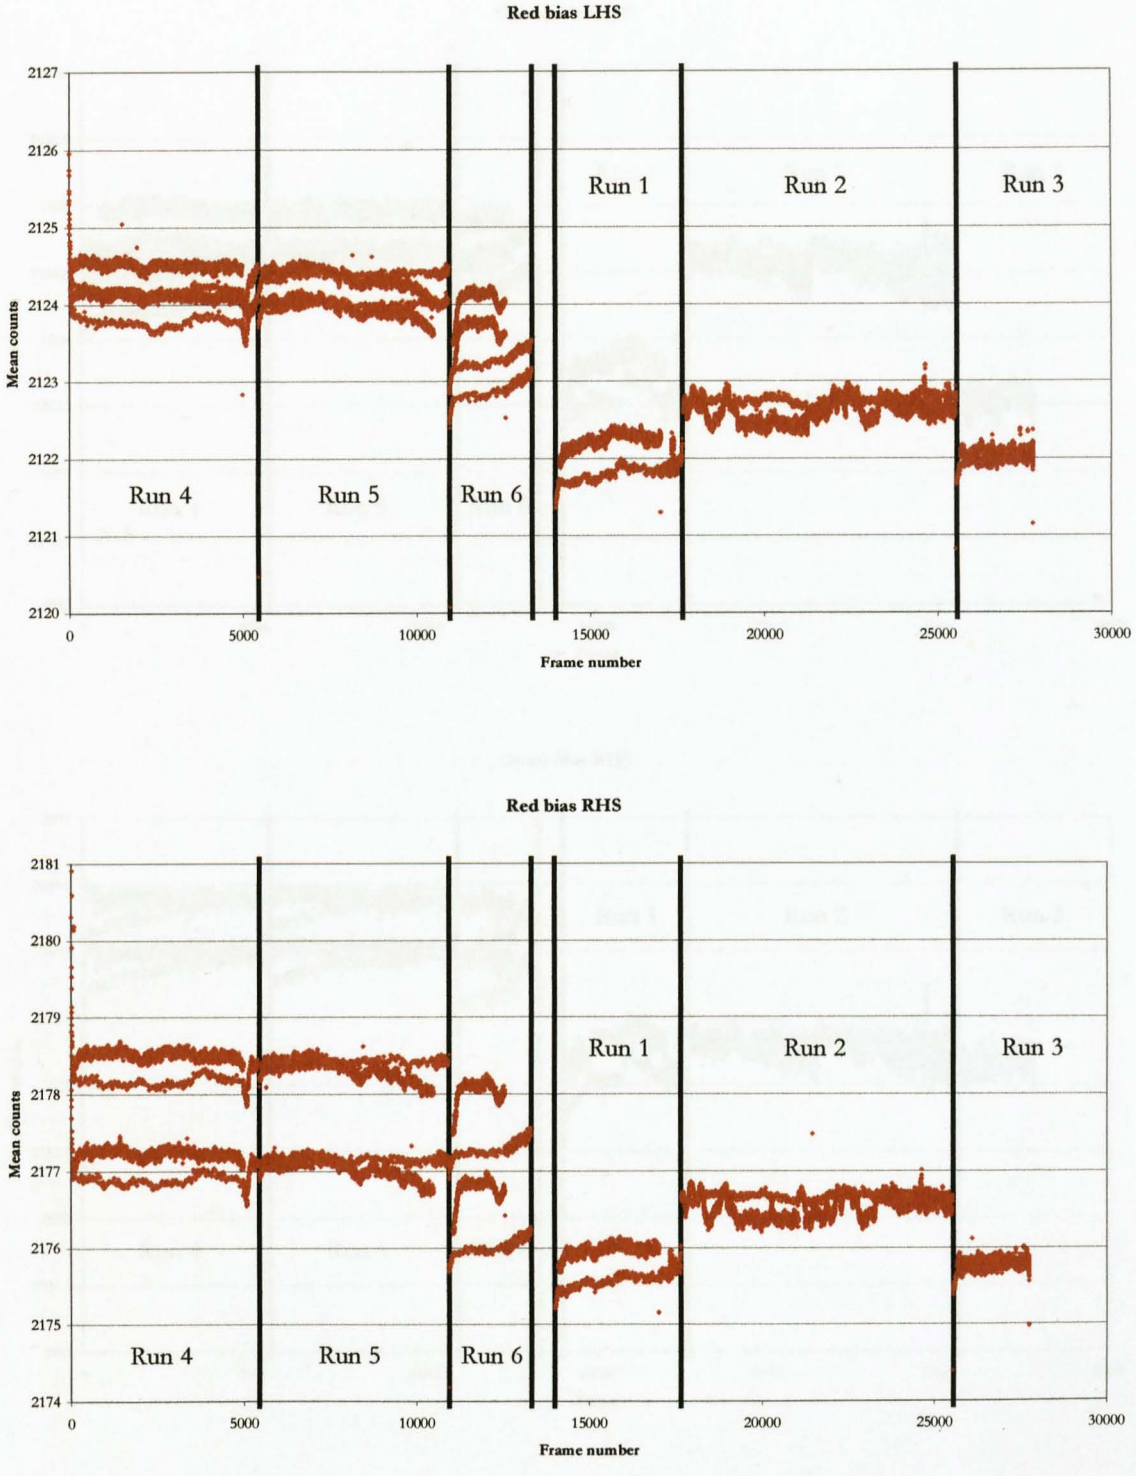

Figure 4.3.8: Plots of the bias level on the left and right channels of the red CCD as a function of frame number for 6 different runs from September 2002.

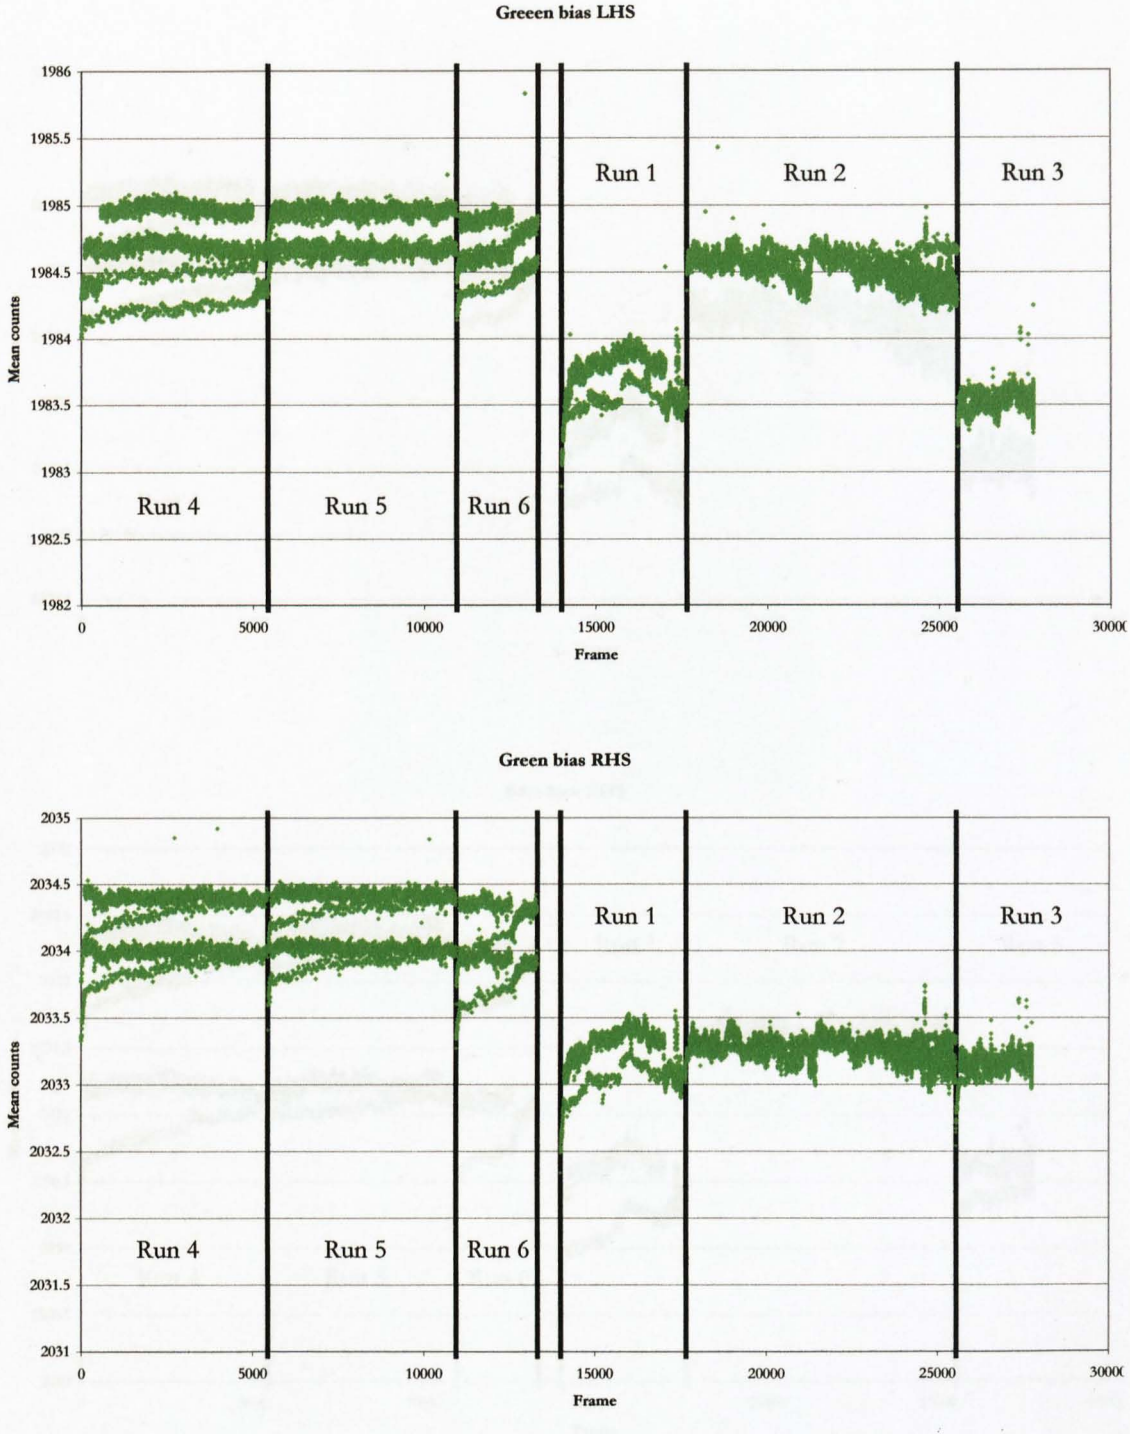

Figure 4.3.9: Plots of the bias level on the left and right channels of the green CCD as a function of frame number for 6 different runs from September 2002.

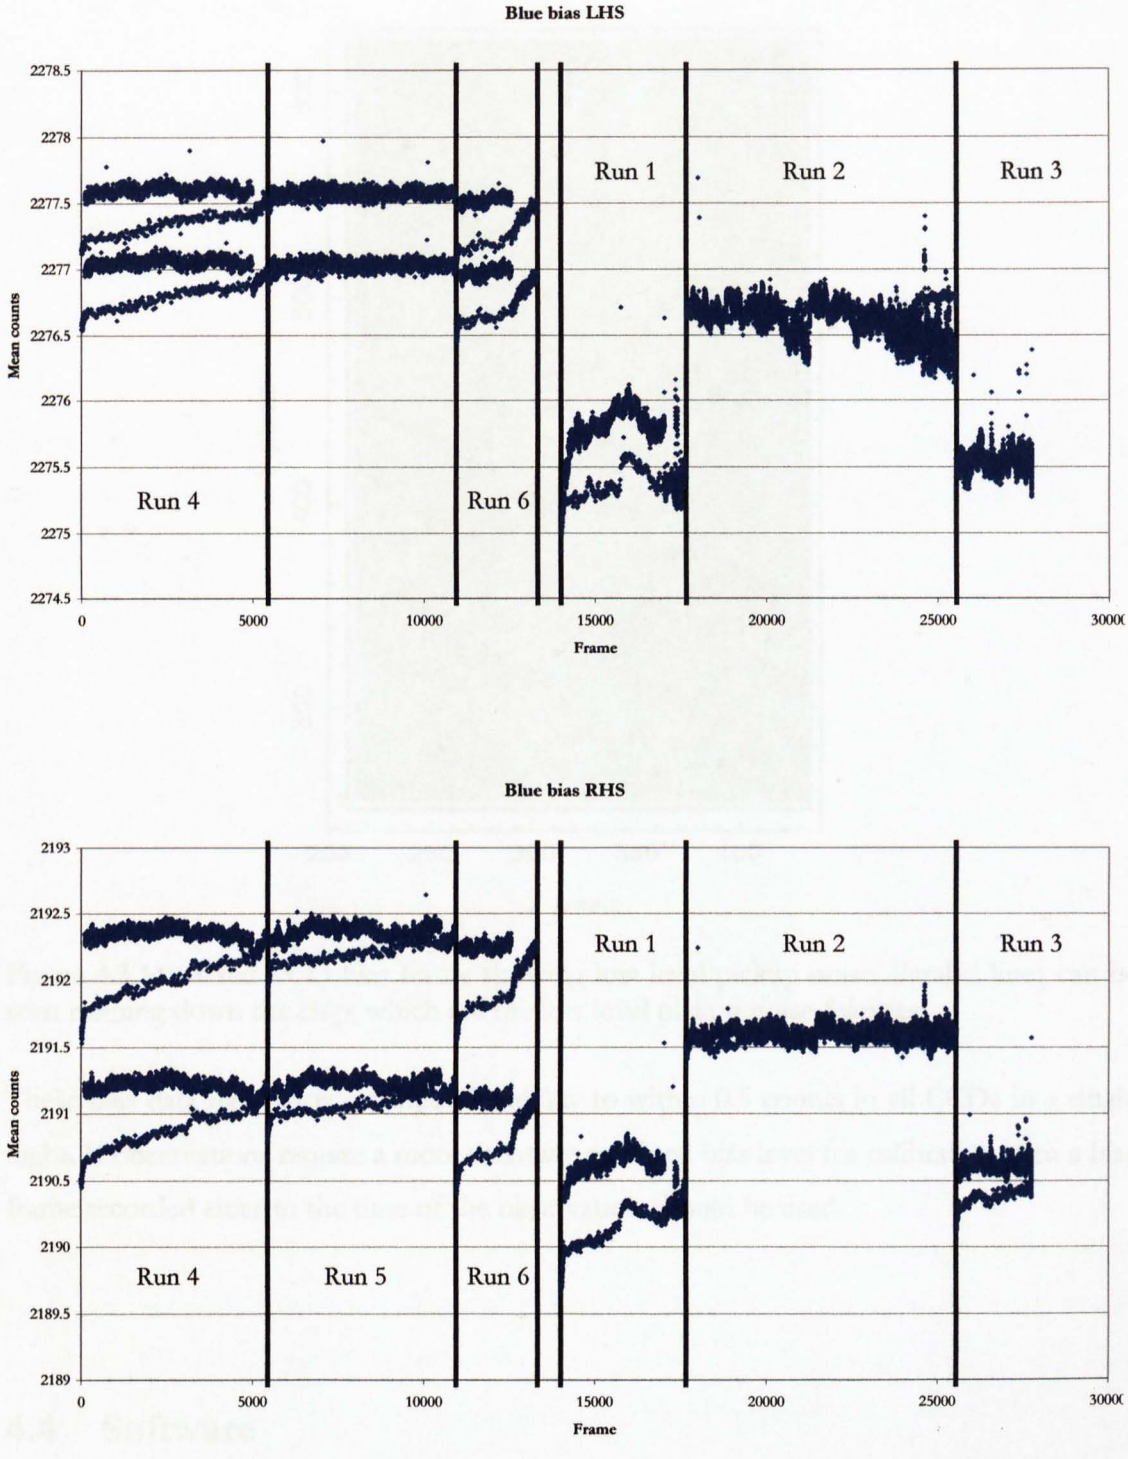

Figure 4.3.10: Plots of the bias level on the left and right channels of the blue CCD as a function of frame number for 6 different runs from September 2002.

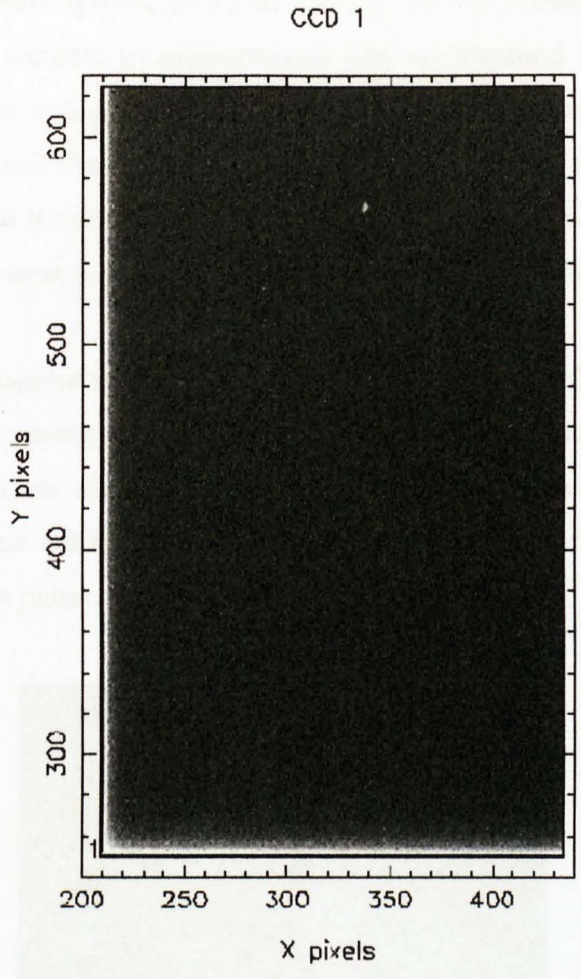

Figure 4.3.11: A red CCD bias frame showing low level pickup noise. Parallel lines can be seen running down the chip, which are the low level pickup noise features.

These bias data have shown temporal stability to within 0.5 counts in all CCDs in a single night. If observations require a more narrowly confined bias level for calibration then a bias frame recorded close to the time of the observation should be used.

4.4 Software

4.4.1 Timing accuracy

It is crucial for the science that ULTRACAM has been designed for that the absolute start-time of each frame is accurately known. Time stamping is achieved using the clock signals

from the GPS satellite system, which in the case of the Trimble Acutime 2000 used in ULTRACAM, are accurate to approximately one microsecond (Acutime 2000 reference sheet). Although the timing accuracy of the GPS system is known, it is not clear that this accuracy is maintained throughout the data acquisition system to the time-stamp that is eventually written to the data frame. To verify the timing accuracy of ULTRACAM frames it is necessary to observe a source that exhibits a periodic feature at a known absolute time.

Pulsars are ideal sources to test time-stamping accuracy. A pulsar is a rapidly-spinning neutron star with a strong magnetosphere which has a spin axis that is not aligned to the star's rotation axis, as shown in figure 4.4.1. To an observer in the line of sight, synchrotron emission arising from the magnetosphere will therefore be modulated on the orbital period of the pulsar.

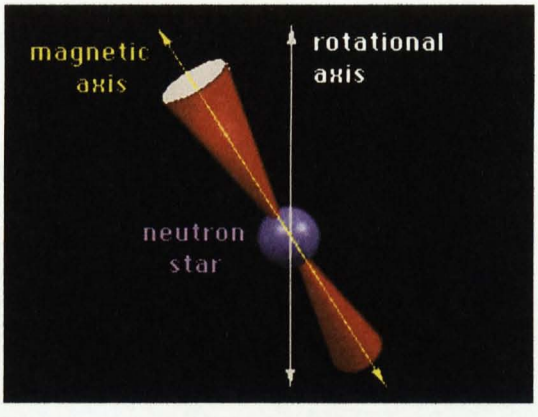

Figure 4.4.1: Diagram of a pulsar. The rotation axis and magnetic axis are not aligned, so when the neutron star spins radiation is beamed into space.

The Crab pulsar is an ideal pulsar to observe for this experiment, being bright (m_v = 16.5), having an accessible rotation period (33 milliseconds) and, more importantly, an accurate, regularly updated ephemeris (Lyne et al, 2003). Figure 4.4.2 shows the light-curve of the Crab pulsar obtained with S-Cam2 (Perryman et al. 1999) – see section 1.3.4.

The Jodrell Bank radio ephemeris accurately predicts the absolute time of arrival of every radio pulse from the Crab pulsar (Lyne et al, 2003). The optical pulse peak is known to precede the radio peak by 140μs (Sanwal et al 1999). The Crab pulsar was observed by ULTRACAM in September 2002 from the WHT, with the aim of comparing the arrival

times of the pulses measured by ULTRACAM with the Jodrell Bank ephemeris to determine the accuracy of the time stamping.

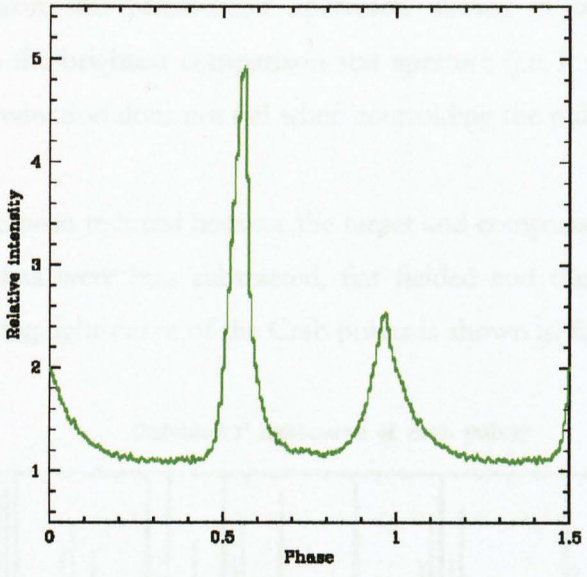

Figure 4.4.2: The light-curve of the Crab pulsar (Perryman et al. 1999).

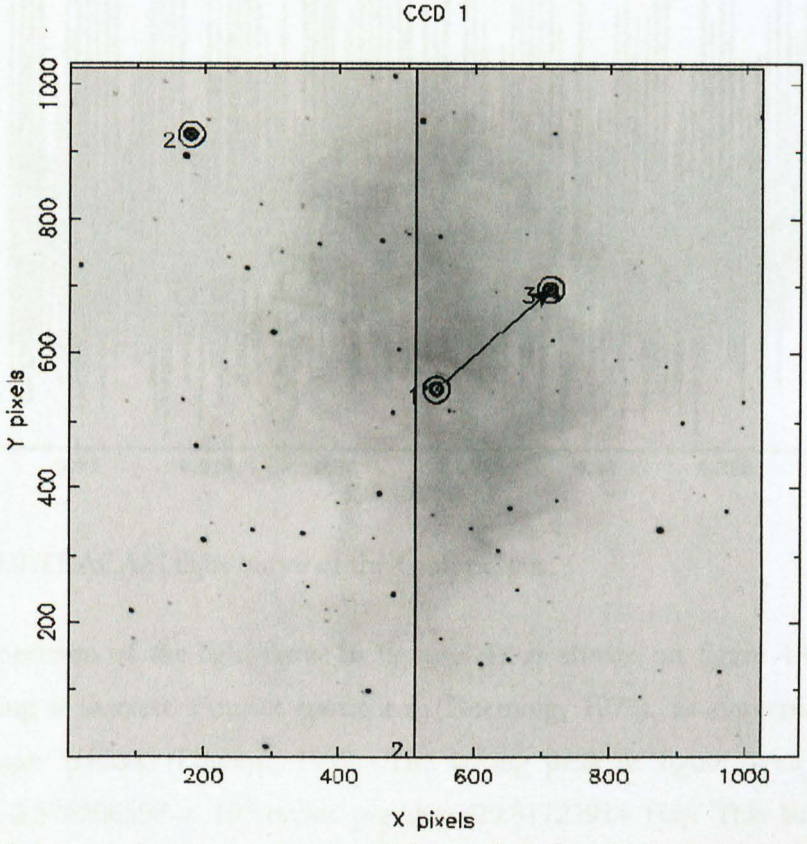

Figure 4.4.3: A 5 second exposure image of the Crab pulsar in r'. Aperture 1 is the Crab pulsar and apertures 2 and 3 are comparison stars. The arrow extending from aperture 1 to aperture 3 indicates that they are linked, i.e. the reduction software uses the position of the bright star in aperture 3 and a constant offset from it to deduce the position of the (faint) Crab pulsar in aperture 1.

A total of 422 frames of the Crab pulsar were obtained in full frame clear mode, with an exposure time of 5 milliseconds. We also obtained a 5 second exposure image to locate the target and comparison star photometry apertures, shown in figure 4.4.3. The target aperture is linked to the brightest comparison star aperture (i.e. it maintains a fixed offset from it) so that the reduction does not fail when centroiding the pulsar in quiescence.

Only the *r'* data have been reduced because the target and comparison stars are brightest in these frames. The data were bias subtracted, flat fielded and the target and sky counts extracted. The resulting light-curve of the Crab pulsar is shown in figure 4.4.4.

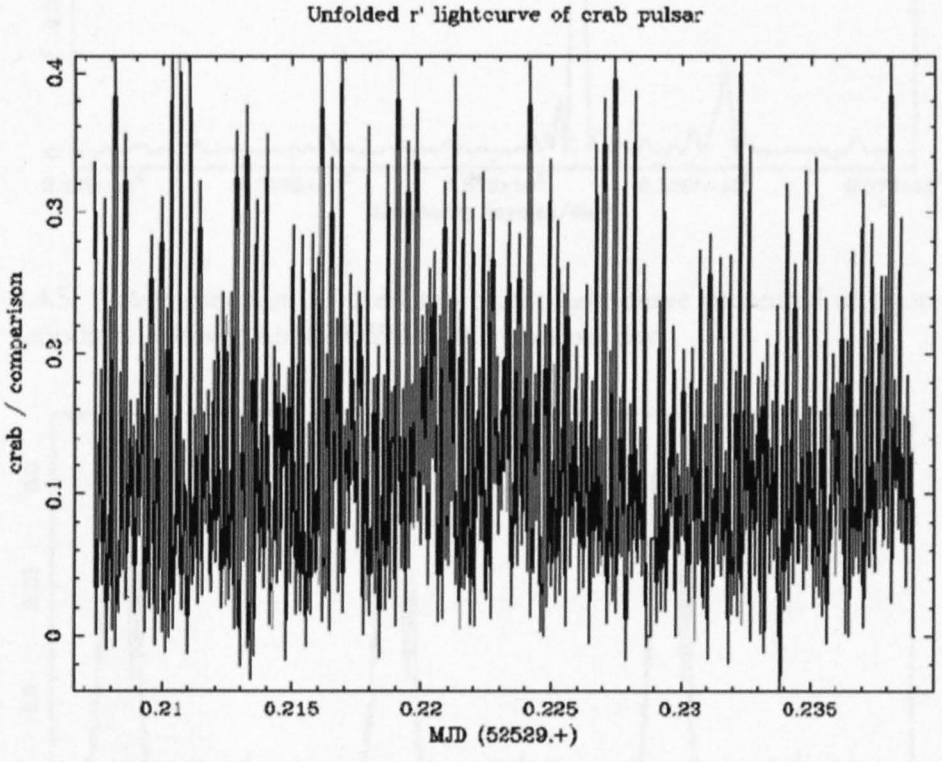

Figure 4.4.4: ULTRACAM light-curve of the Crab pulsar.

The power spectrum of the light-curve in figure 4.4.4 is shown on figure 4.4.5. This was calculated using a discrete Fourier transform (Deeming, 1975), as implemented in the Starlink package 'period' (Dhillon, 1997). The strong peak in figure 4.4.5 occurs at a frequency of 2.576208598 x 10^6 cycles per day (29.81722914 Hz). This period (i.e. 33 milliseconds) was used to fold the data into 50 phase bins. Figure 4.4.6 shows three cycles of the resulting folded, binned light-curve, which has a very similar morphology to the Crab nebula light-curve shown in figure 4.4.2 (albeit in a lower time resolution and signal-to-noise).

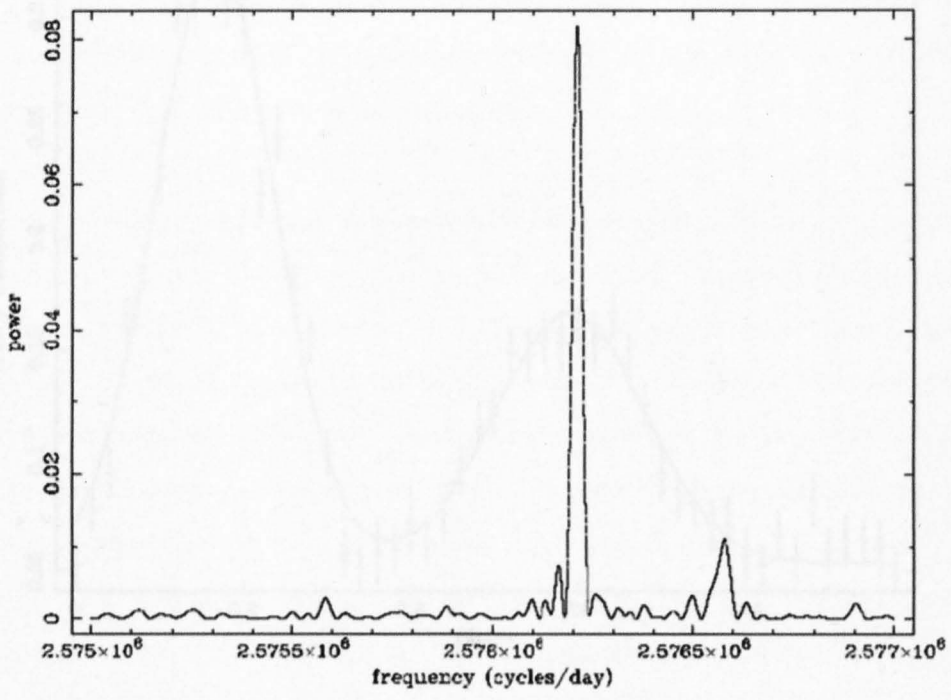

Figure 4.4.5: Power spectrum of the Crab pulsar light-curve presented in figure 4.4.4. A strong periodicity exists at 2.576208598 x 10^6 cycles per day.

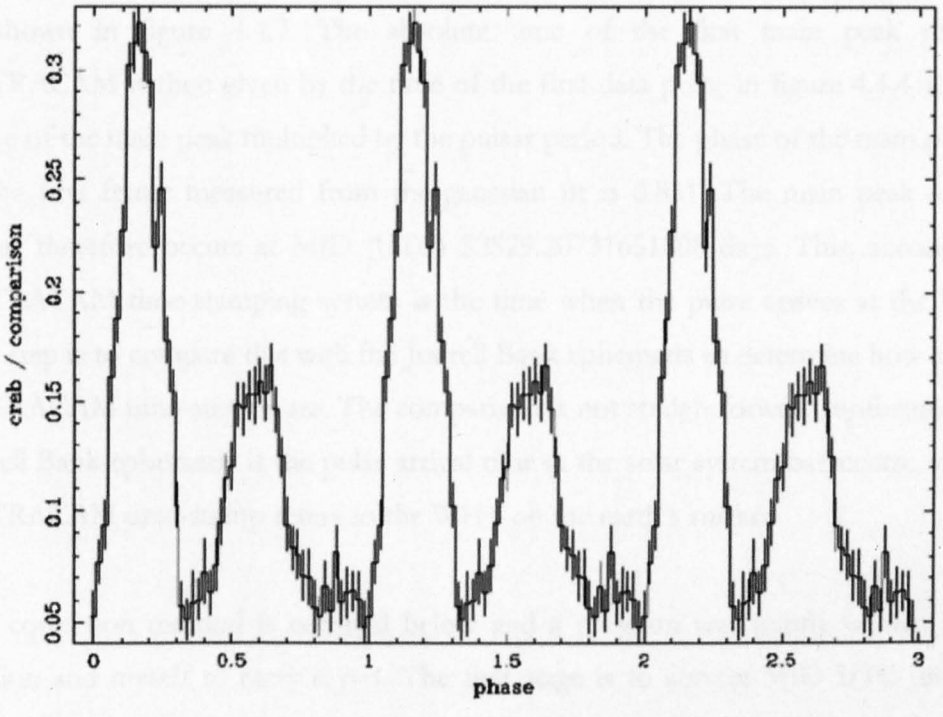

Figure 4.4.6: Phase-binned r' light-curve of the Crab pulsar, folded to show showing three complete spin cycles.

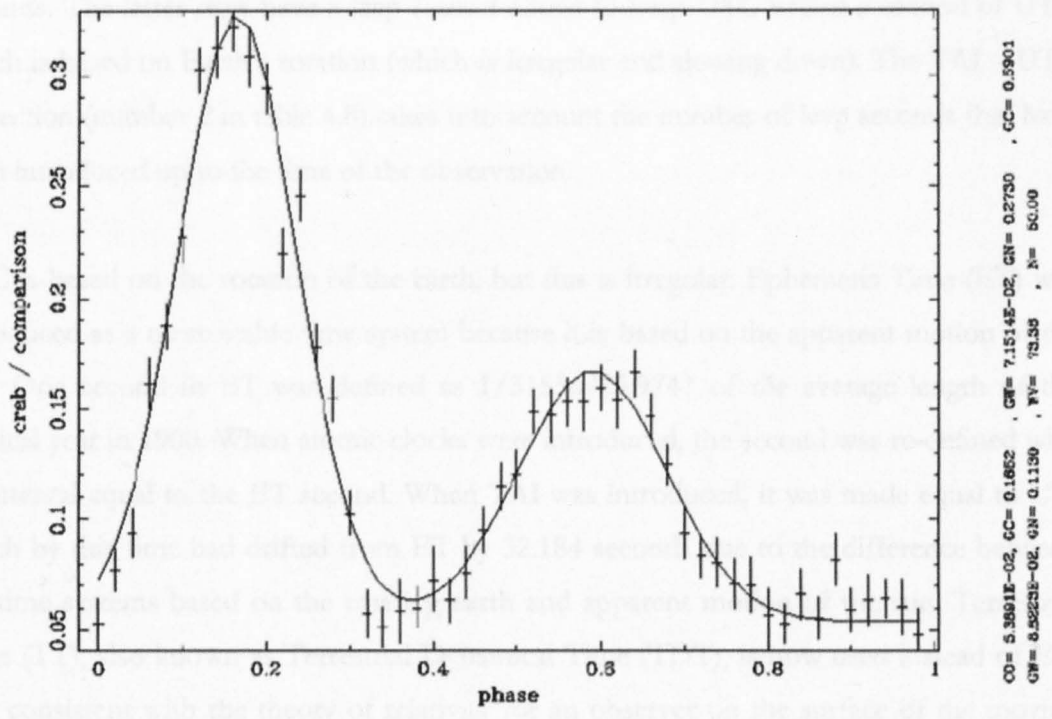

Figure 4.4.7: The Crab pulsar light-curve of figure 4.4.4 with a double-gaussian fitted to determine the phase of the main peak.

To determine the observed time of pulse maximum, we fitted a Gaussian to the main peak, as shown in figure 4.4.7. The absolute time of the first main peak recorded by ULTRACAM is then given by the time of the first data point in figure 4.4.4 added to the phase of the main peak multiplied by the pulsar period. The phase of the main peak relative to the first frame measured from the gaussian fit is 0.831. The main peak of the Crab Pulsar therefore occurs at MJD (UTC) 52529.20731651808 days. This, according to the ULTRACAM time-stamping system is the time when the pulse arrives at the WHT. The next step is to compare this with the Jodrell Bank ephemeris to determine how accurate the ULTRACAM time-stamps are. The comparison is not straightforward, unfortunately as the Jodrell Bank ephemeris is the pulse arrival time at the solar system barycentre, whereas the ULTRACAM time-stamp refers to the WHT on the earth's surface.

The correction method is outlined below and a program was jointly written by Vikram Dhillon and myself to carry it out. The first stage is to correct MJD UTC (number 1 in table 4.8) to International Atomic Time (TAI). International Atomic Time is based on the SI second, which is defined as 9192631770 cycles of radiation from the transition between two energy levels of the ground state of Caesium 133. TAI is independent of earth's rotation whereas UTC is based on the SI second and knows of days of 86400 and 86401

seconds. The latter days have a leap second added to keep UTC within a second of UT1, which is based on Earth's rotation (which is irregular and slowing down). The TAI – UTC correction (number 2 in table 4.8) takes into account the number of leap seconds that have been introduced up to the time of the observation.

UTC is based on the rotation of the earth, but this is irregular. Ephemeris Time (ET) was introduced as a more stable time system because it is based on the apparent motion of the sun. One second in ET was defined as 1/31556925.9747 of the average length of the tropical year in 1900. When atomic clocks were introduced, the second was re-defined with an interval equal to the ET second. When TAI was introduced, it was made equal to UT, which by this time had drifted from ET by 32.184 seconds due to the difference between the time systems based on the rotating earth and apparent motion of the sun. Terrestrial Time (TT), also known as Terrestrial Dynamical Time (TDT), is now used instead of ET. It is consistent with the theory of relativity for an observer on the surface of the moving earth. The TT – TAI correction (number 3 in table 4.8) is a constant offset of 32.184 seconds which is required so there is no discontinuity at the changeover to atomic clocks.

The next stage is to transform the TT to the solar system barycentre, or Barycentric Dynamical Time (TDB). When earth is at perihelion it moves faster and deeper in the sun's gravitational potential well, causing the TT to be slower than TDB. When earth is at aphelion, the opposite occurs. The TDB – TT correction (number 4 in table 4.8) is a periodic term determined by the orbital parameters of earth with respect to the solar system barycentre.

The correction so far has calculated the time when a photon from the pulsar hits the WHT on Earth as measured by a clock at the solar system barycentre. We need to know the time when the photon hits the solar system barycentre as measured by a clock at the solar system barycentre. This then removes the problem of photons arriving at different times relative to each other due to Earth moving in its orbit around the Sun. This is done by calculating the barycentric earth position/velocity vector and the geocentric position of the telescope using SLALIB routines (Wallace, 2003). The two vectors are added to give the barycentric telescope position/velocity vector. The pulsar position/velocity vector is also calculated using SLALIB routines. The velocity components of the position – velocity vectors are removed and the vectors normalised. The angle, θ, between the sun – earth

vector and the sun – pulsar vector is calculated and this is used to calculate the arrival time difference, Δt between a photon arriving at the solar system barycentre and at the WHT on the earth's surface. Figure 4.4.8 shows a schematic of the sun – earth and sun – pulsar vectors. The pulsar – earth and pulsar – sun angle, γ, is very small because the pulsar is very far away. Therefore the arrival time of a photon at the WHT and at point X on figure 4.4.8 is assumed equal, i.e. $A = B$. The arrival time difference between a photon arriving at the solar system barycentre and at the WHT on the earth's surface (number 5 in table 4.8) is added to the TDB and gives the arrival time of the pulsar peak (number 6 in table 4.8) observed by ULTRACAM at the solar system barycentre. The observed pulsar peak corrections are shown in table 4.8.

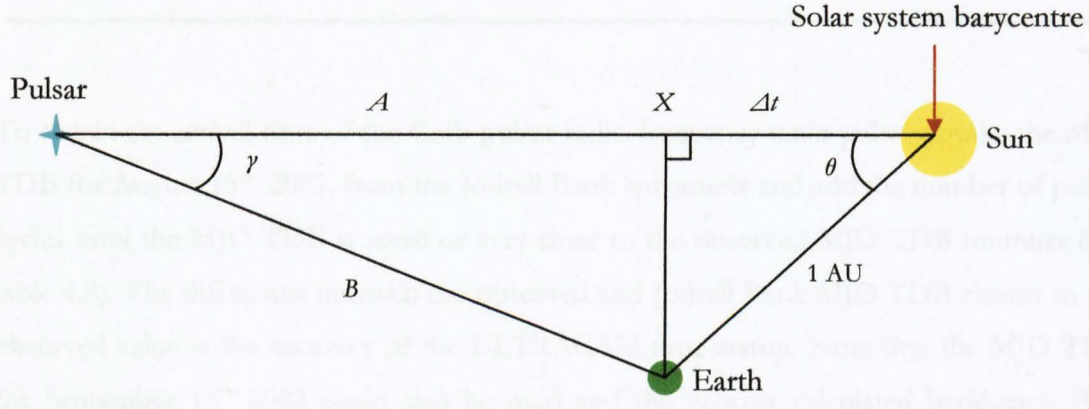

Figure 4.4.8: Diagram of the sun – earth vector and the sun – pulsar vector. The angle between the two vectors, θ, is calculated and used to determine the time difference, Δt, between a photon arriving at the solar system barycentre and at the WHT on the earth's surface.

The Jodrell Bank ephemeris gives the arrival time of the centre of the first Crab-pulsar main peak after midnight at infinite frequency (of the light the pulsar is observed in) on the 15th of every month. The pulsar frequency, ν, and the rate of change of the pulsar frequency, $\dot{\nu}$, for each month is also given in the ephemeris. The quoted arrival time accuracy in the Jodrell Bank ephemeris for August and September 2002 is 1 millisecond. This determines the accuracy to which the observed pulse can be compared with the ephemeris.

Table 4.8: Timing accuracy results

1	MJD (UTC) – time observed on the WHT	52529.207465966096
2	TAI – UTC (seconds)	32
3	TT – TAI (seconds)	32.184
4	TDB – TT (seconds)	-0.00149797986
5	Barycentre Travel time, Δt (seconds)	-44.9998293
6	MJD (TDB) – time observed at the solar system barycentre	52529.207687987771

To obtain the arrival time of the Crab-pulsar radio-frequency main pulse we take the MJD TDB for August 15th, 2002, from the Jodrell Bank ephemeris and add the number of pulsar cycles until the MJD TDB is equal or very close to the observed MJD TDB (number 6 in table 4.8). The difference between the observed and Jodrell Bank MJD TDB closest to the observed value is the accuracy of the ULTRACAM time-stamp. Note that the MJD TDB for September 15th 2002 could also be used and the process calculated backwards. The pulsar period is slowing and this affects the pulse arrival-time calculation, so the rate of change of pulsar frequency needs to be incorporated into the calculation. The total number of pulsar cycles, N_C, between the quoted Jodrell Bank MJD TDB to the observed MJD TDB is given by the following series:

$$N_C = \nu \Delta e + \dot{\nu} \frac{\Delta e^2}{2} + \ddot{\nu} \frac{\Delta e^3}{6}, \tag{4.15}$$

where $\ddot{\nu}$, is the rate of change of the rate of change of the pulsar frequency and e, is the epoch given by:

$$\Delta e = TDB_{Obs} - TDB_{Jod}, \tag{4.16}$$

where TDB_{Obs}, is the TDB recorded by ULTRACAM time-stamp and TDB_{Jod}, is the TDB calculated using the Jodrell Bank ephemeris. The rate of change of pulsar frequency, $\dot{\nu}$, is given by:

$$\dot{\nu} = \frac{d\nu}{dt} = \frac{dP}{dt}\frac{d\nu}{dP}, \tag{4.17}$$

where P, is the pulsar period given by:

$$\dot{P} = \dot{\nu}\frac{dP}{d\nu}. \tag{4.18}$$

The pulsar period P, is inversely related to the pulsar frequency ν, and can be written as:

$$P = \nu^{-1}, \tag{4.19}$$

and the derivative of P with respect to pulsar frequency is therefore:

$$\frac{dP}{d\nu} = -\nu^{-2}. \tag{4.20}$$

It is then possible to write:

$$\dot{P} = -\frac{\dot{\nu}}{\nu^2}. \tag{4.21}$$

The rate of change of the rate of change of the pulsar frequency $\ddot{\nu}$, is given by:

$$\ddot{\nu} = \frac{d\dot{\nu}}{dt} = \frac{d\dot{\nu}}{dP}\frac{dP}{dt} = \frac{d\dot{\nu}}{dP}\dot{P} \tag{4.22}$$

where $\dot{\nu}$, can be rewritten by rearranging equation 4.21 and combining with equation 4.20 to give:

$$\dot{\nu} = -\dot{P}\nu^2 = -\frac{\dot{P}}{P^2} = -\dot{P}P^{-2}. \tag{4.23}$$

By taking the derivative of equation 4.23, we can write $\ddot{\nu}$ as:

$$\frac{d\dot{\nu}}{dP} = 2\dot{P}P^{-3} - \frac{d\dot{P}}{dP}P^{-2}. \tag{4.24}$$

The last term on the right-hand side of equation 4.24 is ignored because it is very small and therefore insignificant, so:

$$\ddot{\nu} = \frac{2\dot{P}^2}{P^3} \tag{4.25}$$

Table 4.9: Timing accuracy results

The MJD (TDB) is the arrival time of the centre of the main radio pulse after midnight at the barycentre of the solar system. The observed barycentric frequency, ν, its first derivative ν - dot at the arrival time, the error in the quoted arrival time are also given. The number of pulsar cycles between the MJD (TDB) in this table and the observed MJD (TDB) is calculated and the time between this number and a whole number of cycles is determined to give the timing accuracy of the ULTRACAM time stamp.

	August 15th 2002	**September 15th 2002**
MJD (TDB) (Days)	52501.027142	52532.002024
ν (Hertz)	29.8152123201	29.8142110435
$\dot{\nu}$ (10^{-15} second^{-2})	-373842.19	-373809.73
Error in ephemeris (microsecond)	1000	1000
Number of cycles	72662861.997608870268	-7192861.019236300141
Timing accuracy (seconds)	**0.000080198313**	**0.000645205714**

The results using the August 15th and September 15th 2002 data from the Jodrell Bank ephemeris are shown in table 4.9. The results show that the accuracy of ULTRACAM's absolute timing is approximately **0.6 milliseconds**, where the observed pulse arrival times precede the Jodrell Bank timings by this amount. The accuracy of the analysis is limited, however, by the uncertainty of 1 millisecond in the ephemeris, so the absolute timing accuracy of ULTRACAM should be quoted as 1 millisecond. To reduce this uncertainty further, the Crab pulsar will have to be observed in a month that has a low arrival time uncertainty in the ephemeris (e.g. due to Jodrell Bank observing it more frequency).

4.5 Mechanical

4.5.1 Flexure

It is vital that ULTRACAM exhibits low flexure so that each CCD remains co-aligned (see section 2.5.1). The mechanical design analysis of section 2.5.3 predicts the flexure of the double octopod to be **$(1.94 \pm 0.14) \times 10^{-6}$ m**. This value is for flexure over the entire length of the double octopod and assumes rigid joints. The analysis in this section will calculate the actual instrumental flexure and determine whether it is differential between the CCDs (i.e. all three CCDs shift due to flexure in the same direction and by the same amount). This will be done by using observational data to record the position of a star on each CCD chip whilst ULTRACAM moves through different orientations.

The data used in this analysis is an observation of HS0039+4302 taken on 19th September 2002. The observation began at 23:15 UTC and ended at 06:51 UTC, giving a total time of 07:36 on the target and 7521 frames. The x and y positions of the stars on each frame were determined by Gaussian fitting. The point of least flexure is when the telescope points at the zenith, and hence the flexure has been calculated relative to this value as shown in figure 4.5.1.

Figure 4.5.1 shows that there are jumps in the flexure on all three CCDs. Examination of other data obtained in May and September 2002 also shows these jumps and verifies that

they always occur in all three arms of ULTRACAM. A distinctive feature of the jumps is that they each shift by the same number of pixels in each arm and each x and y plane on the CCDs, as shown in figures 4.5.2 and 4.5.3, which show plots of flexure in both the x and y planes of each CCD.

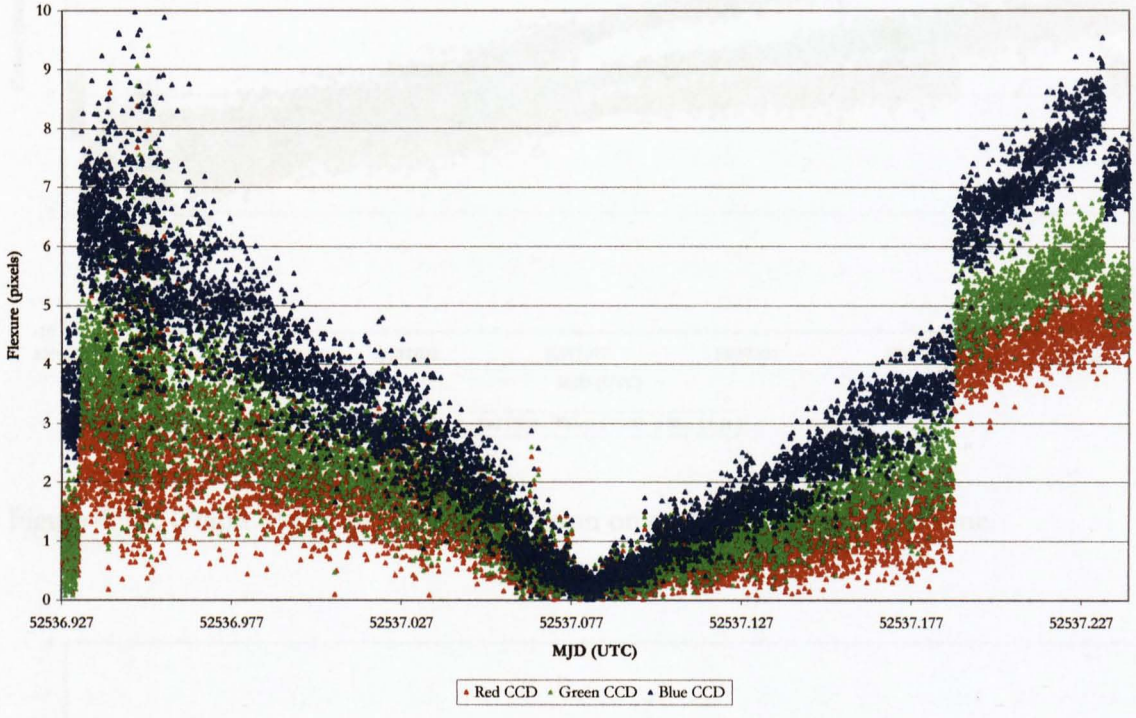

Figure 4.5.1: Instrumental flexure in each CCD plotted as a function of time. The sharp jumps in flexure are thought to be due to flexure in the autoguider probe in the Cassegrain A&G box.

A structural shift in ULTRACAM could cause the observed flexure jumps, but it would then be unlikely that the jump is the same in all three CCDs. A more likely explanation is that something has caused the telescope image to shift before it enters the ULTRACAM collimator. This could be due either to flexure in the ULTRACAM collar or in the telescope. However, the flexure in the mount collar has been calculated to be lower than **$(1.8 \pm 0.2) \times 10^{-6}$ m** (see section 2.5.13). This is an order of magnitude lower than the 2 – 3 pixel jumps (i.e. 26 – 33 x 10^{-6} m) shown in figures 4.5.2 and 4.5.3. The flexure jumps are therefore more likely to be attributable to some feature of the WHT, such as flexure in the autoguider probe; if the autoguider probe moves, the position of the guide star on the autoguider CCD moves, forcing the telescope to move to compensate for the shift. This telescope movement would cause the stars to move on the ULTRACAM CCDs.

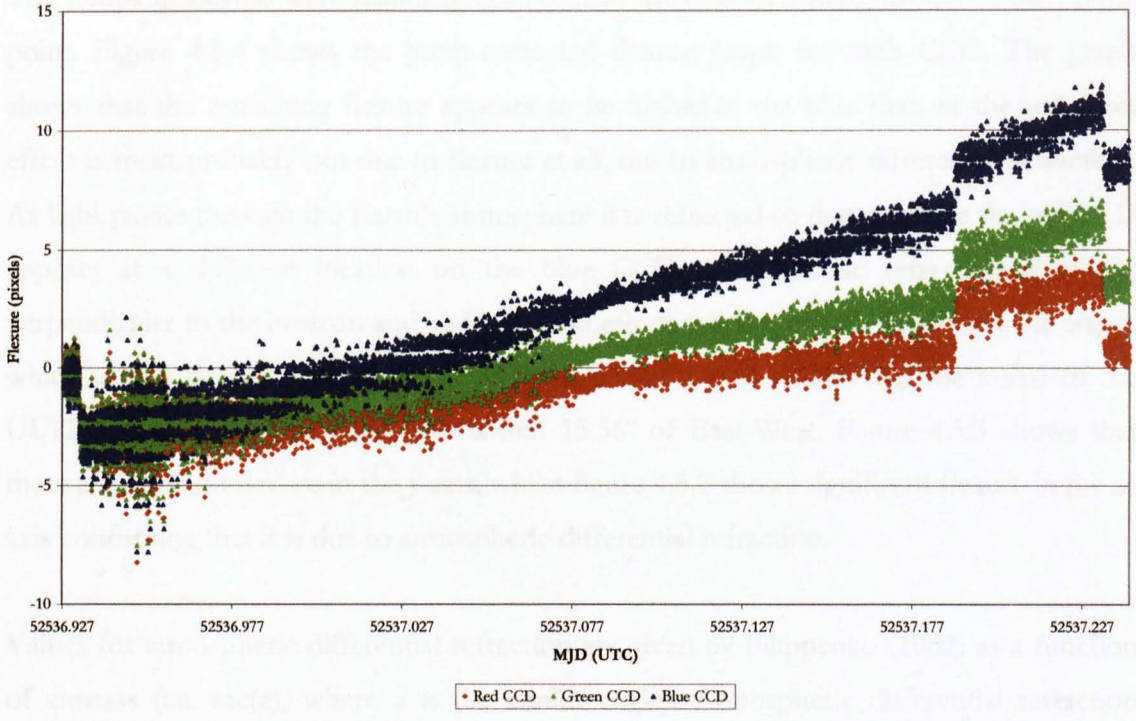

Figure 4.5.2: Graph of flexure in the *x*-direction on each CCD chip versus time.

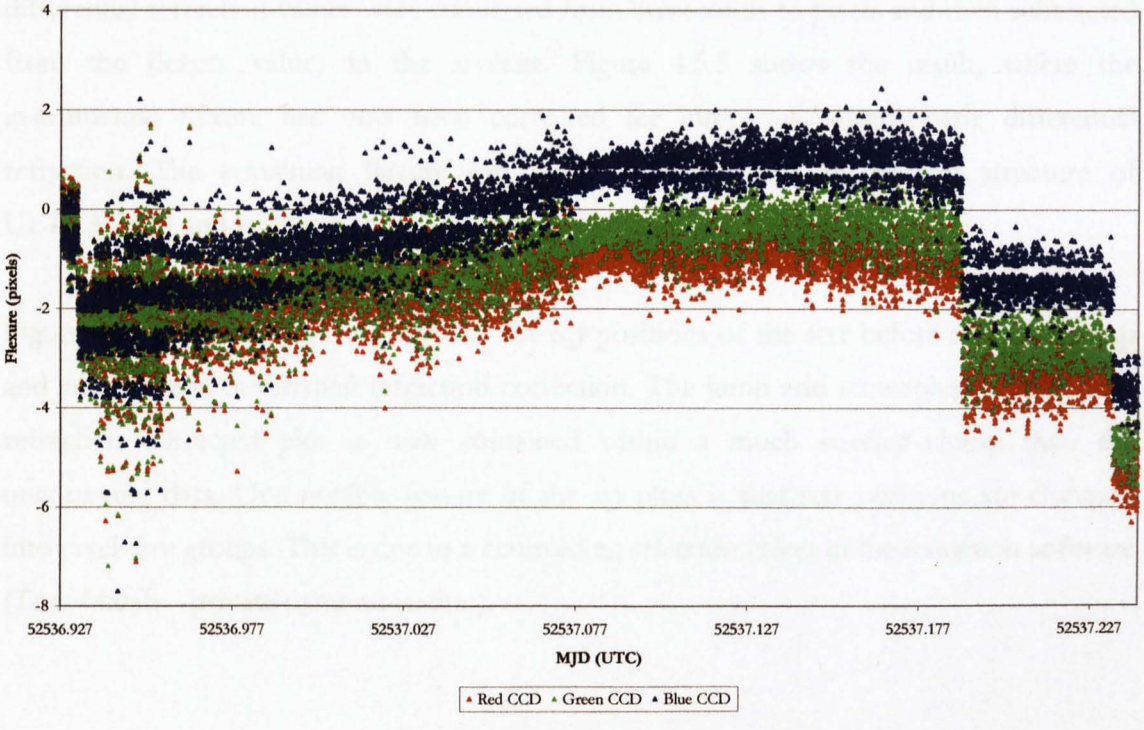

Figure 4.5.3: Graph of flexure in the y-direction on each CCD chip versus time.

The jumps in flexure were manually corrected in the data by adding offsets at each jump point. Figure 4.5.4 shows the jump-corrected flexure graph for each CCD. The graph shows that the remaining flexure appears to be higher in the blue than in the red. This effect is most probably not due to flexure at all, but to atmospheric differential refraction. As light passes through the Earth's atmosphere it is refracted so that a star on the red CCD appears at a different location on the blue CCD. Atmospheric refraction is always perpendicular to the horizon and varies in strength depending on the airmass of the target, which is also shown in figure 4.5.4. In section 4.2.4 it was shown that the x-axis of the ULTRACAM CCDs are aligned to within 15.56° of East-West. Figure 4.5.3 shows that there is almost no flexure in the y-axis, whilst figure 4.5.2 shows significant flexure in the x-axis confirming that it is due to atmospheric differential refraction.

Values for atmospheric differential refraction are given by Filippenko (1982) as a function of airmass (i.e. sec(z), where z is the zenith angle). Atmospheric differential refraction versus sec(z) was plotted for each colour and the equation of the line calculated. Values of atmospheric differential refraction were then generated for each colour band using the equation of the line by inputting airmass values. The colour-dependent atmospheric differential refraction values were converted from arcseconds to pixels and then subtracted from the flexure values in the x-plane. Figure 4.5.5 shows the result, where the instrumental flexure has now been corrected for jumps *and* atmospheric differential refraction. The remaining flexure can be attributed to the mechanical structure of ULTRACAM and is shown to be non-differential.

Figures 4.5.6 and 4.5.7 show graphs of the x,y positions of the star before and after jump and atmospheric differential refraction correction. The jump and atmospheric differential refraction corrected plot is now contained within a much smaller clump than the uncorrected data. One notable feature of the x,y plots is that star positions are clumped into pixel-size groups. This is due to a centroiding selection effect in the reduction software (Tom Marsh – private communication).

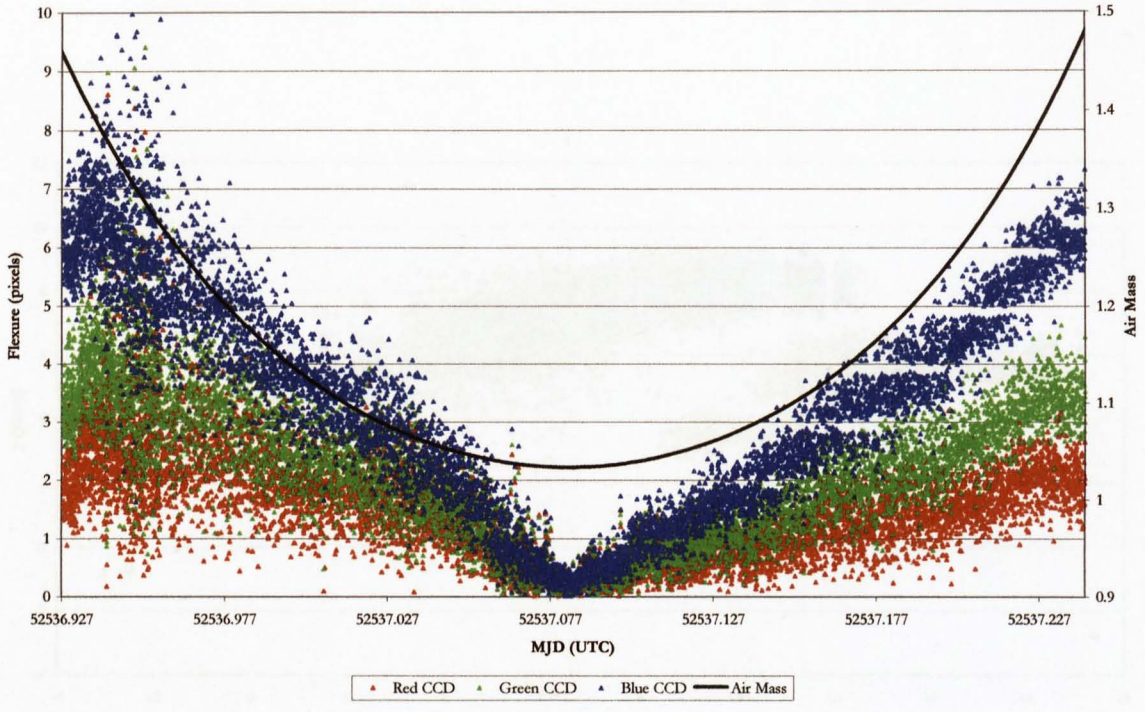

Figure 4.5.4: Graph of flexure versus time after corrections for jumps due to flexure in the autoguider probe (see text for details).

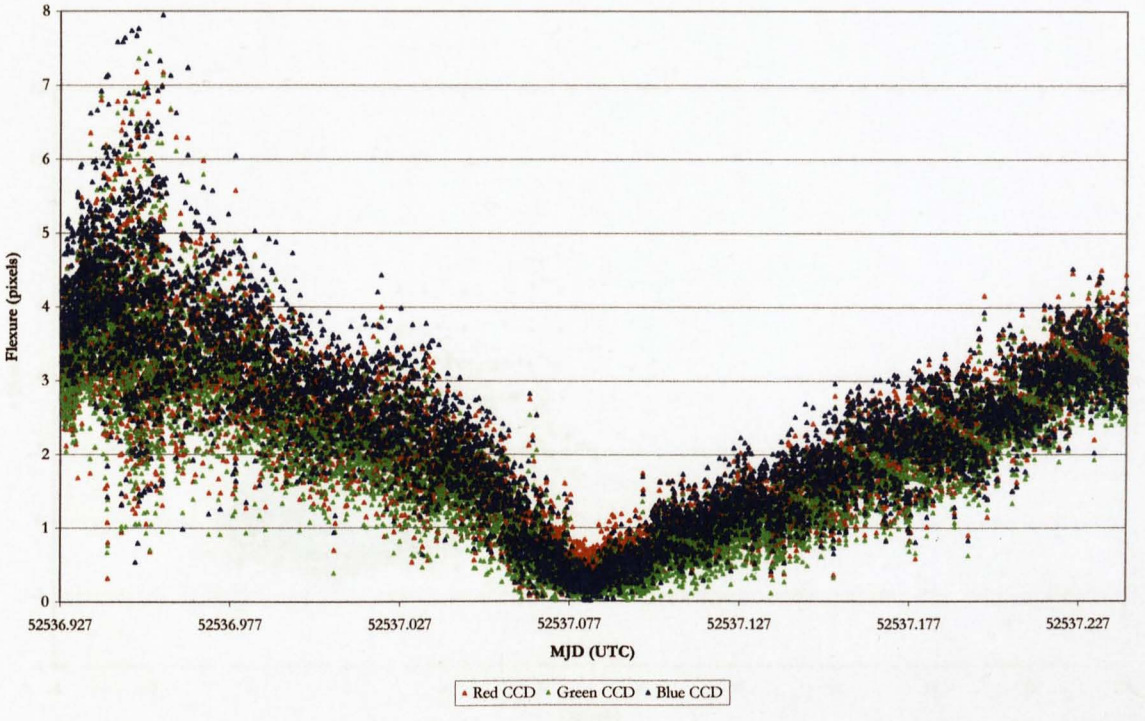

Figure 4.5.5: Graph of instrumental flexure versus time after corrections for jumps due to flexure in the autoguider probe and atmospheric differential refraction.

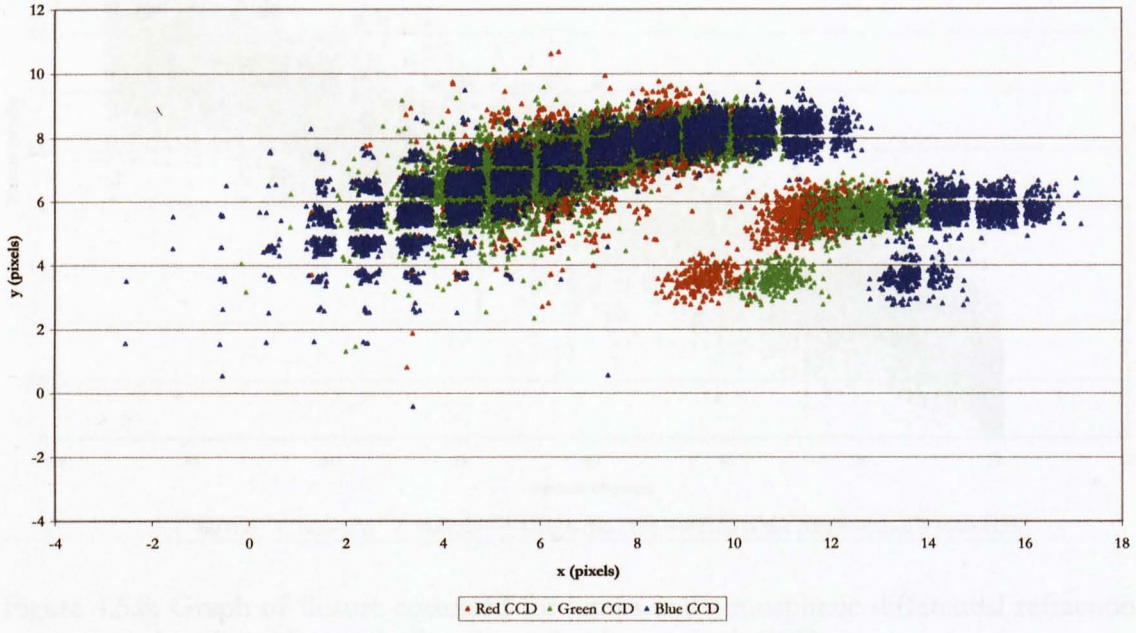

Figure 4.5.6: Graph of star *x,y* positions (no jump or atmospheric differential refraction correction).

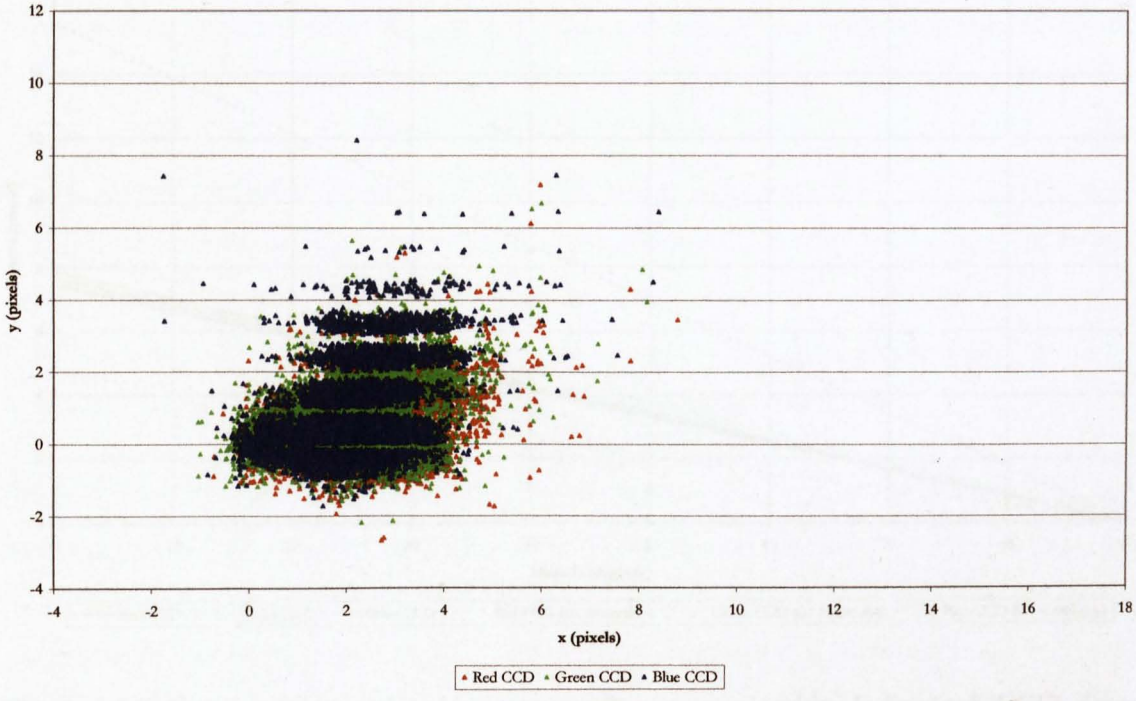

Figure 4.5.7: Graph of star *x,y* positions corrected for jumps and atmospheric differential refraction.

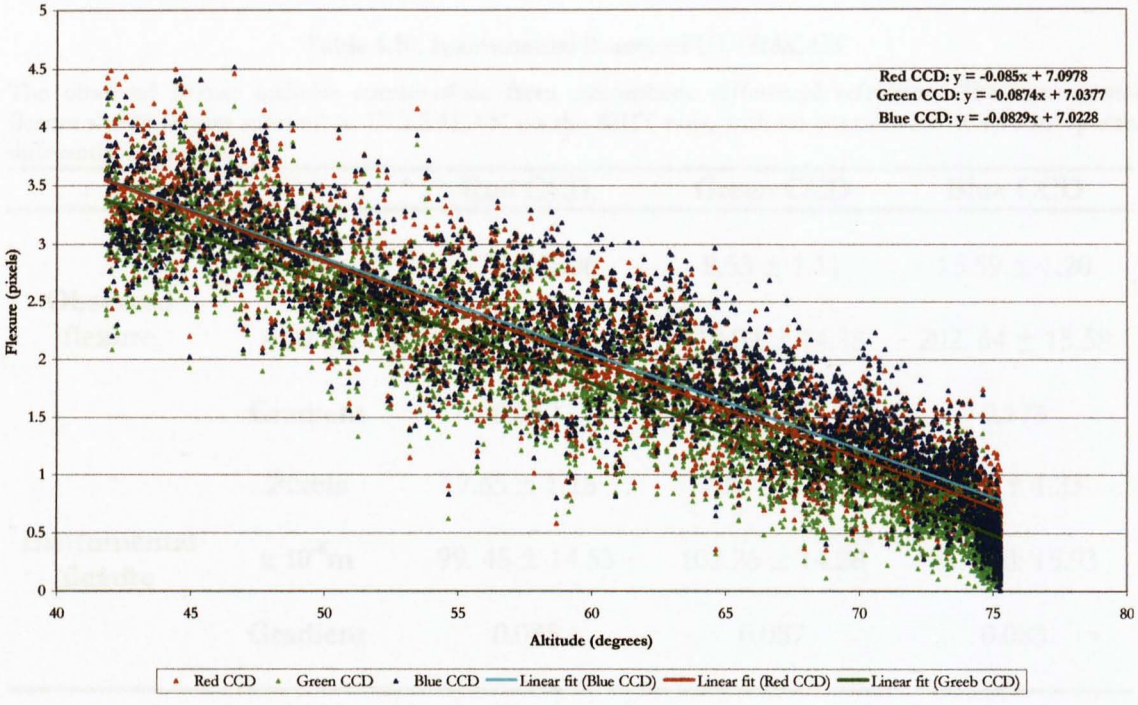

Figure 4.5.8: Graph of flexure corrected for jumps and atmospheric differential refraction versus star altitude, with straight-line fits to the data on each CCD.

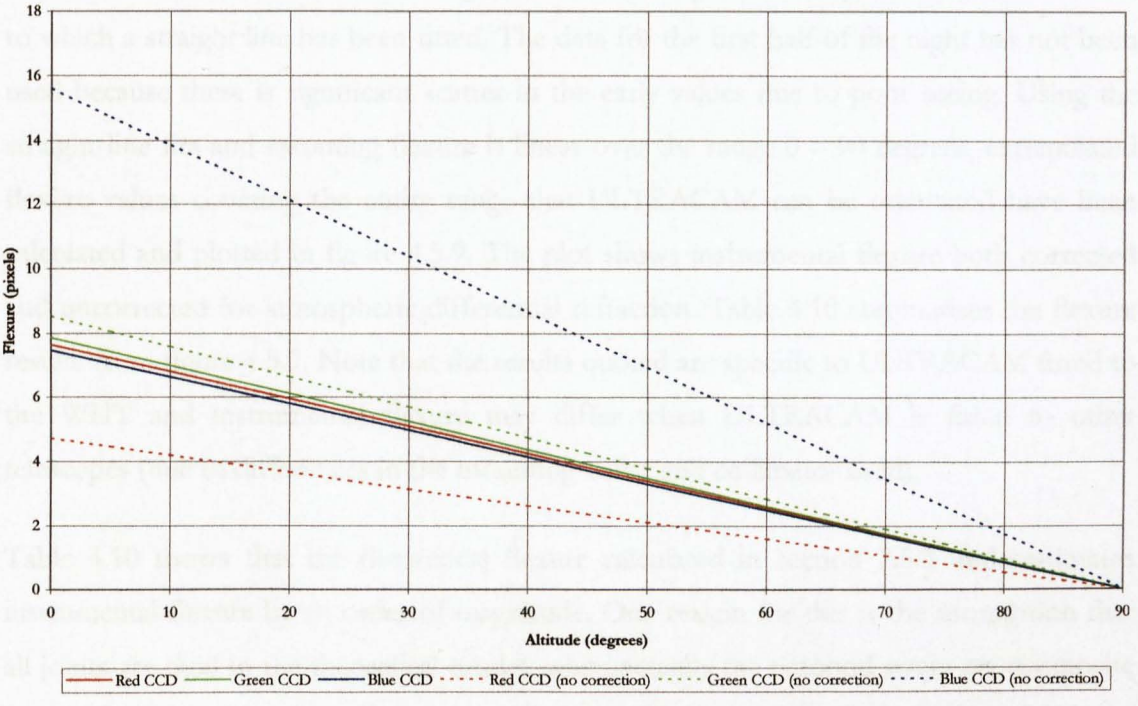

Figure 4.5.9: Graph of extrapolated instrument flexure versus altitude for each CCD. Two sets of flexure values are plotted, one set corrected and one set not corrected for atmospheric differential refraction.

Table 4.10: Instrumental flexure of ULTRACAM

The observed flexure includes contributions from atmospheric differential refraction. The instrumental flexure shows flexure inherent to ULTRACAM on the WHT only, with no contribution from atmospheric differential refraction.

		Red CCD	Green CCD	Blue CCD
Observed flexure	Pixels	4.72 ± 1.06	8.53 ± 1.11	15.59 ± 1.20
	x 10^{-6}m	61.31 ± 13.81	110.92 ± 14.38	202. 64 ± 15.58
	Gradient	0.052	0.095	0.173
Instrumental flexure	Pixels	7.65 ± 1.18	7.87 ± 1.10	7.46 ± 1.23
	x 10^{-6}m	99. 45 ± 14.53	102.26 ± 14.28	96.99 ± 15.93
	Gradient	0.085	0.087	0.083

The data can now be used to predict instrumental flexure for orientations not achieved during the HS0039+4302 observation. HS0039+4302 was observed from an altitude of 43.38°, through transit (at an altitude of 75.28°), down to an altitude of 41.80°. A plot of flexure versus altitude is shown in figure 4.5.8 for the post transit part of the observation, to which a straight line has been fitted. The data for the first half of the night has not been used because there is significant scatter in the early values due to poor seeing. Using the straight-line fits and assuming flexure is linear over the range 0 – 90 degrees, extrapolated flexure values covering the entire range that ULTRACAM can be orientated have been calculated and plotted in figure 4.5.9. The plot shows instrumental flexure both corrected and uncorrected for atmospheric differential refraction. Table 4.10 summarises the flexure results from figure 4.5.9. Note that the results quoted are specific to ULTRACAM fitted to the WHT and instrumental flexure may differ when ULTRACAM is fitted to other telescopes (due to differences in the mounting collar and collimator used).

Table 4.10 shows that the theoretical flexure calculated in section 2.5.3 underestimates instrumental flexure by an order of magnitude. One reason for this is the assumption that all joints are rigid in the theoretical model, when actually the octopod struts are composite parts which can introduce flexure at each joint. Moreover, the theoretical analysis was simplified by only including the double octopod and mounting collar; the observed values

however, measure flexure at each CCD and are therefore also sensitive to flexure in the CCD head mounts.

The figures in table 4.10 are very useful for planning observations on the WHT. The ranges of flexure relative to each CCD can be used to define CCD window sizes. The values in table 4.10 give the observed flexure in each CCD as ULTRACAM tracks a star from an altitude of 30° to the zenith. If the altitude range of an observation is known in degrees, then multiplying by the gradient in table 4.10 will give the observed differential flexure in pixels on each CCD. For example, if a star is observed through an altitude range of 22 degrees, the observed flexure would be r' = 1.14, g' = 2.09 and u' = 3.81 pixels relative to where the star had been positioned on the CCD at the start of the observation.

Chapter 5

Development work

5.1 Introduction

After the commissioning run in May 2002, there were several areas where ULTRACAM had to be enhanced before the next observing run (in September 2002). This chapter details the work undertaken to develop ULTRACAM for subsequent observations.

5.2 Clear filters

Some ULTRACAM observations will require the maximum possible throughput. This can be achieved by removing the SDSS filters, but this causes the position of the focal plane to change, so consequently the CCDs would have to be re-focussed. This is not desirable during an observing run as it would take considerable time. A suitable alternative is to replace the SDSS filters with high throughput, broad-band clear filters which have the same optical thickness (i.e. the product of the physical thickness and refractive index) as the SDSS filters.

Replacing the SDSS filters with clear filters will cause significant ghosting in the blue and green channels and very weak ghosting in the red channel. Figure 2.2.4 of the simple ghosting model predicts that primary ghosts (i.e. the first ghost image originating from the second surface of the dichroic beam-splitter) of light in the wavelength range 300 – 1100nm will be 1.8%, 1.7% and 0.13% for the blue, green and red channels respectively. The primary ghost images in the blue and green channels are so high because there are no

SDSS filters to block the ghosts. This section describes the procurement of the clear filters and the tests we performed using them in ULTRACAM.

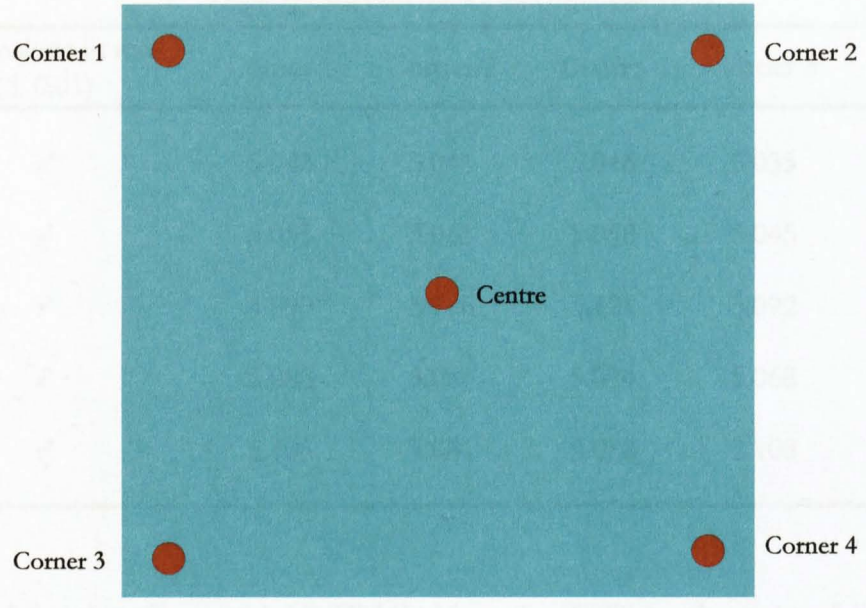

Figure 5.2.1: Filter positions where each thickness measurement was made.

The physical thickness of the SDSS filters needed to be accurately measured to calculate the required thickness of the clear filters. This was not originally known accurately and each filter was probably different. I took the filters to Stephenson Blake, a Sheffield-based company that owns a contact probe micrometer. The machine operates by moving a probe using accurate drives until it senses contact with the surface of the filters. Measurements were taken at the corners and centre of each filter, as shown in figure 5.2.1. The resulting thickness measurements (listed in table 5.1) are accurate to approximately 0.01mm, and the probe is non-destructive and therefore does not damage the filters.

The information in table 5.1 was given to Asahi Spectra Co., the same company that we procured the SDSS filters from, who used the data to calculate the thickness of UV-grade fused silica required to give the same optical thickness as the SDSS filters. As a compromise for cost, the clear filters supplied by Asahi Spectra Co. were all the same thickness based on producing the best results in the red channel. We decided to buy three of these filters, knowing that if they were used in the blue and green channels that the images would be out of focus and would contain ghosts. However, if an observation

specifically required clear filters in all three channels then the blue and green CCDs could be re-focussed for the clear filters and the target counts could be extracted.

Table 5.1: ULTRACAM SDSS filter thickness

Dimensions in mm (± 0.01)	Corner 1	Corner 2	Centre	Corner 3	Corner 4
u'	5.043	5.041	5.046	5.035	5.032
g'	5.055	5.063	5.058	5.045	5.057
r'	5.090	5.106	5.121	5.092	5.097
i'	5.088	5.080	5.094	5.068	5.081
z'	5.105	5.081	5.088	5.103	5.085

We tested the clear filters with ULTRACAM on the WHT on the night of 6th June 2003. The aim of the observation was to combine the data on a very faint and heavily reddened X-ray binary from all three CCDs to improve the signal-to-noise of the light-curve. As expected the blue CCD data is out of focus, so much so that it is impossible to determine the ghost images (figure 5.2.2). The green CCD data is in focus and the ghosting is clearly evident. There is no ghosting visible in the red CCD images and the data is in focus as expected. Figure 5.2.2 shows ULTRACAM images in each channel with clear filters fitted.

When the SDSS filters are used in ULTRACAM most of the primary ghost image is outside the pass-band of the filters and is blocked. This is also true for the light outside the pass-band that is reflected at the first surface due to the less than perfect transmission efficiency of the dichroics. Ghost images are not expected to form separate images in ULTRACAM because the dichroic beam-splitters operate in a collimated beam. However, this is not the case. The last lens in the collimator is bigger than the first lens in the re-imaging cameras and therefore the light beam within ULTRACAM converges slightly. This results in ghost images from the second surface of the dichroic beam-splitters.

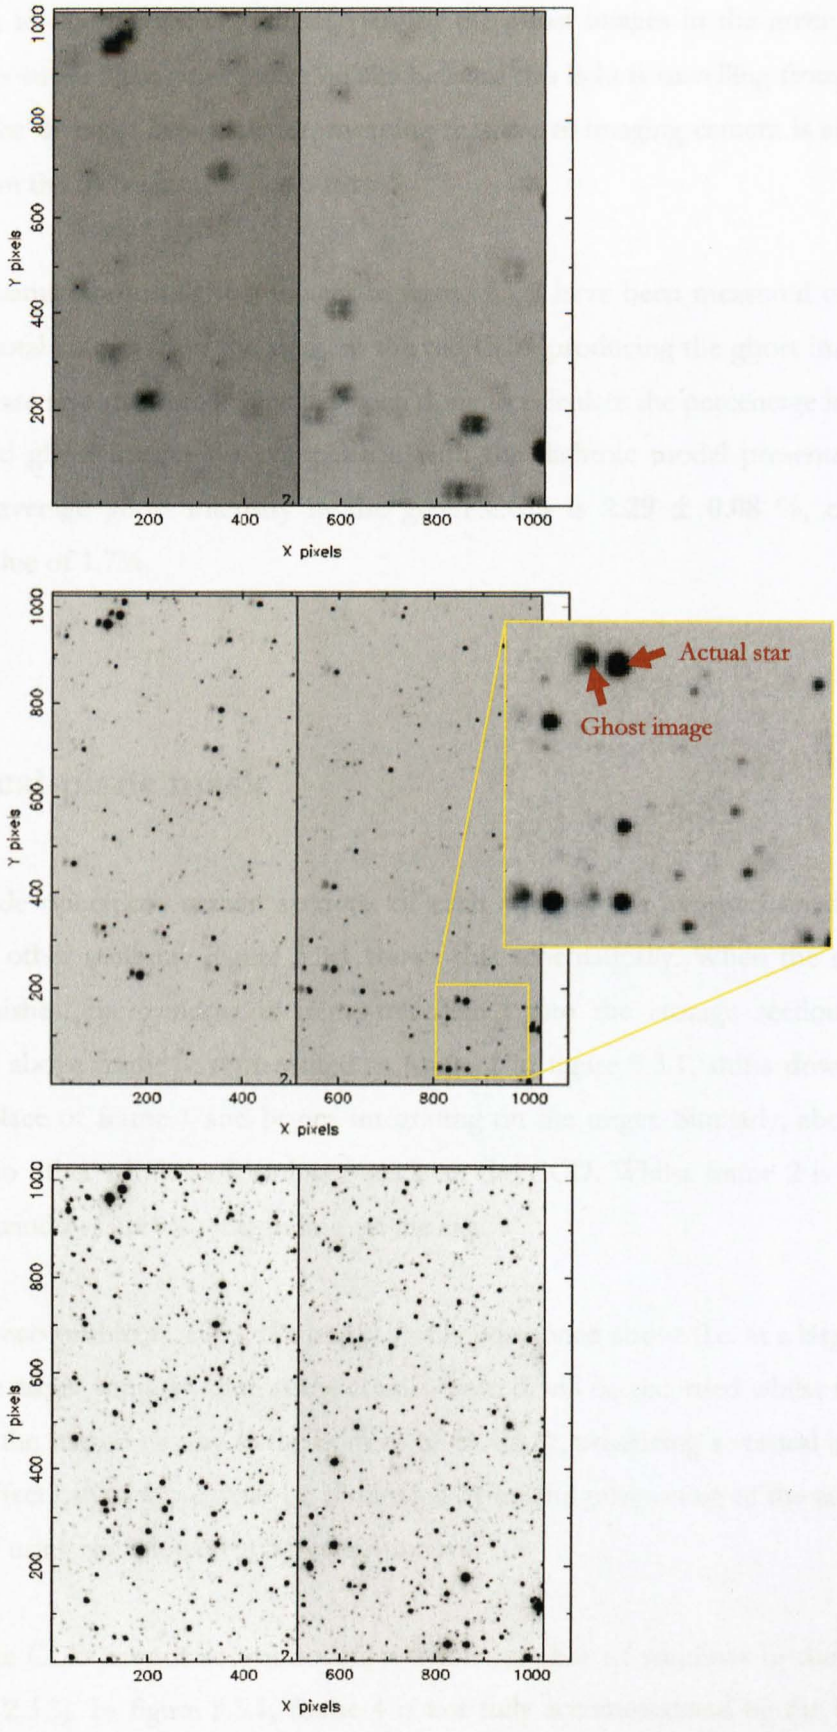

Figure 5.2.2: ULTRACAM images in the blue (top), green (middle) and red (bottom) channels with clear filters fitted.

Also shown in figure 5.2.2 is a closer view of the ghost images in the green CCD. The ghost images suffer from poor image quality because this light is travelling from the second surface of the dichroic beam-splitter, meaning that the re-imaging camera is at the wrong distance from the dichroic.

The total counts from the ghost images in figure 5.2.2 have been measured on the green CCD. The total counts from the stars on the red CCD producing the ghost images on the green CCD are also measured. This has been done to calculate the percentage intensities of the observed ghost images for comparison with the dichroic model presented in figure 2.2.4. The average ghost intensity in the green CCD is **2.29 ± 0.08** %, close to the predicted value of 1.7%.

5.3 Focal-plane mask

In drift mode operation, certain sections of each window are exposed on the chip for longer than other sections. Figure 5.3.1 shows this schematically. When the exposure of frame 1 finishes, the window is frame-transferred into the storage section. The area immediately above frame 1, represented as frame 2 in figure 5.3.1, shifts down the CCD taking the place of frame 1 and begins integrating on the target. Similarly, above frame 2 there are two other windows forming a stack on the CCD. Whilst frame 2 is integrating, these other windows are also integrating on the sky.

This causes two problems. First, if a bright star is positioned above (i.e. at a larger y – pixel number) the target window, then some counts from it will be recorded whilst the window moves over the star on its way to the bottom of the stack, producing a vertical streak in the data. This effect can in some cases be eliminated by careful positioning of the target objects on the CCD using the telescope Cassegrain rotator.

Secondly, the CCD cannot accommodate a whole number of windows in the y-direction (see section 2.3.3). In figure 5.3.1, frame 4 is not fully accommodated by the CCD array. The lower section of the window is exposed whilst the upper section is not. This results in one section of the frame being exposed for longer than the other, resulting in the bi-modal

background, shown in figure 5.2.3. The only way to reduce this effect is to define a pair of windows that stack as close to a whole number of windows on the CCD as possible.

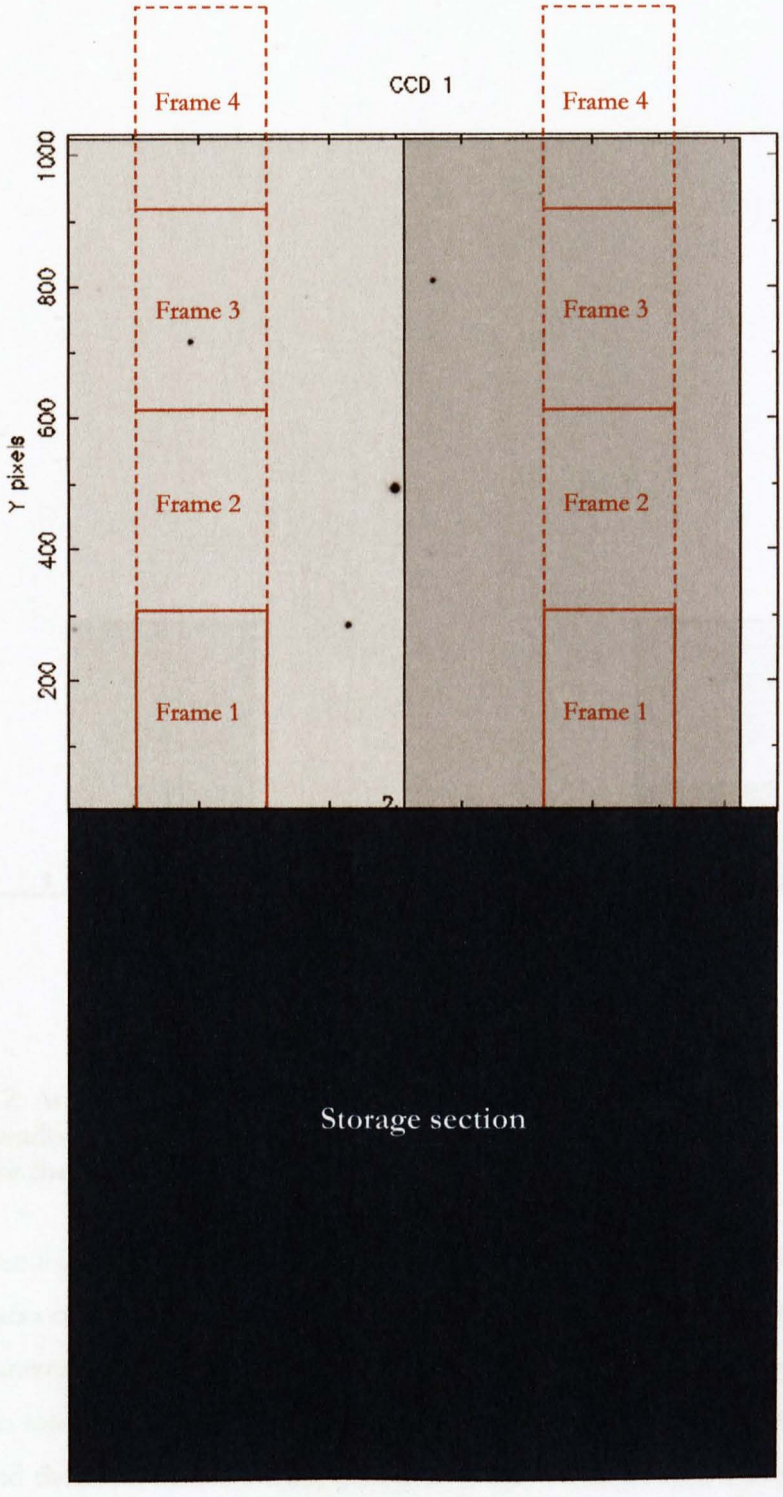

Figure 5.3.1: Schematic of drift mode operation, with windows drifting down the imaging section prior to integrating on the target in frame 1. Because a whole number of windows can never be stacked, the bottom part of frame 4 exists on the chip for longer than the top part.

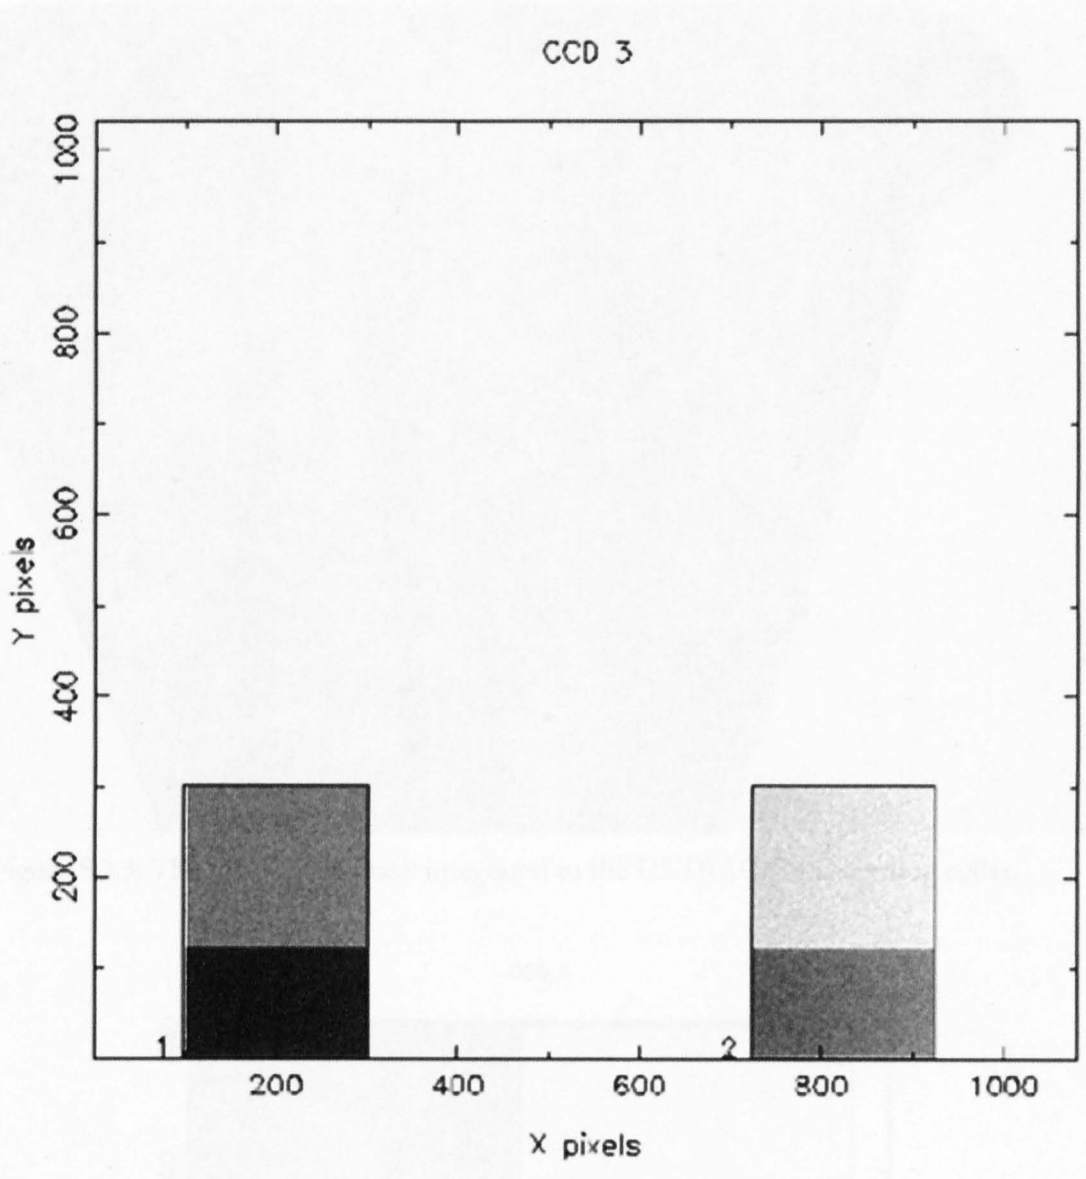

Figure 5.3.2: An ULTRACAM drift-mode image showing the bi-modal background level in a pair of windows. The lower half of the windows are at a higher count level than the top half because they have been exposed for longer.

The problems described above prompted me to develop a mask that can be inserted to cover the area of the CCD above the windows, thereby eliminating the streaks from stars in the field above the windows and the bi-modal background level. The mask has been designed to intersect the light path at the telescope focal plane, so that it is in focus when inserted and therefore forms a sharply defined edge. The focal plane drift mask integrates with the ULTRACAM mounting collar and is shown in figure 5.3.3.

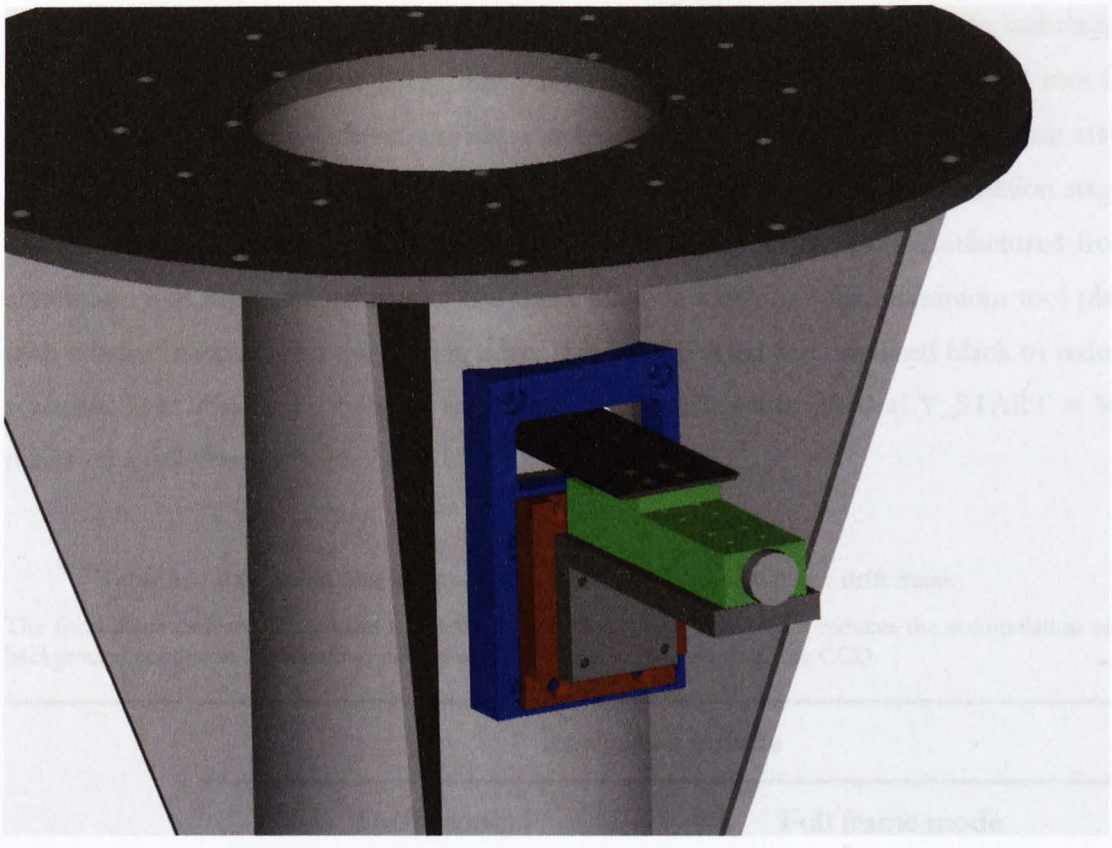

Figure 5.3.3: The focal-plane mask integrated to the ULTRACAM mounting collar.

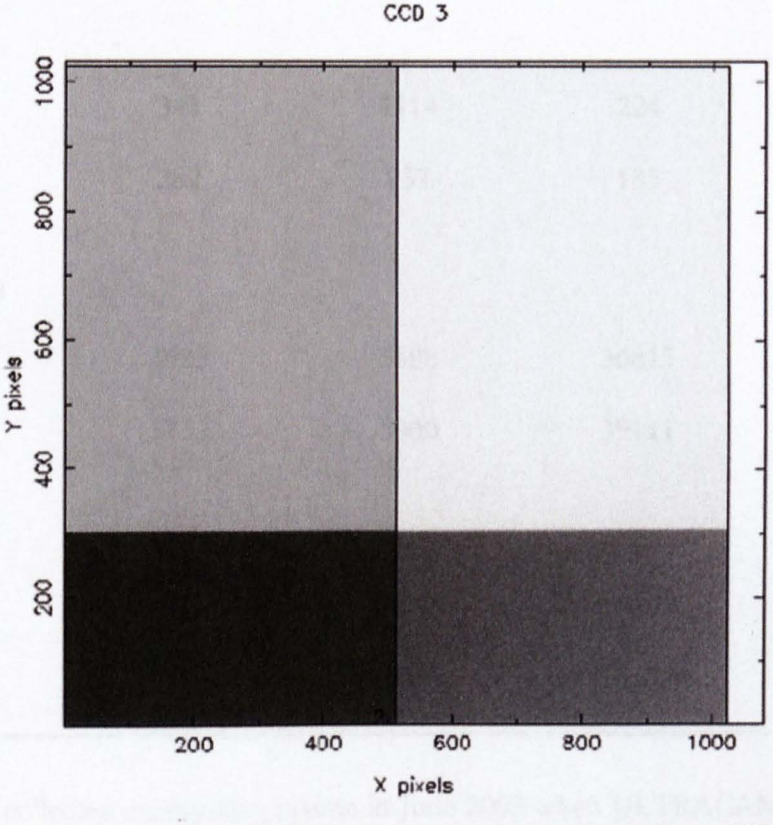

Figure 5.3.4: A full-frame ULTRACAM image with the focal plane drift mask inserted to Y_START = 300 pixels.

The focal-plane mask is inserted into the light path by a large, lockable linear translation stage (the green unit in figure 5.3.3). The scale of movement on the linear stage is 1 mm for every 16 pixels in the y – direction. The mask can focussed by a second translation stage (the red unit in figure 5.3.3). There is an L-shaped bracket to join the two translation stages and a plate to interface to the ULTRACAM mount collar. These are manufactured from aluminium and are anodised black. The mask blade is machined-flat aluminium tool plate with a bevel machined on the cutting edge. It is bead blasted and anodised black to reduce scattered light. Figure 5.3.4 shows the focal plane drift mask inserted at Y_START = 300 pixels on a full-frame image.

Table 5.2: Raw mean bias subtracted counts using the focal plane drift mask

The focal plane drift mask has been inserted and retracted to test how well it reduces the accumulation of background counts on the windows as they shift down the entire length of the CCD.

	Raw mean counts			
	Drift mode		Full frame mode	
	Mask inserted	Mask retracted	Outside of mask	Under the mask
Blue				
LHS	343	1114	224	10
RHS	282	937	185	8
Green				
LHS	1625	5388	36815	1191
RHS	1732	5800	39111	1246
Red				
LHS	5369	17995	16087	592
RHS	5470	18583	16531	590

Data were collected during the daytime in June 2003 when ULTRACAM was fitted to the WHT. Tungsten lamps and neutral density filters were used to evenly illuminate the CCDs

whilst ULTRACAM operated in drift mode with and without the focal-plane mask inserted. Data were also taken in full frame mode with the mask inserted. An analysis window was defined that eliminates the bad rows and columns in each window. The mean of each window was then calculated in drift mode with the mask retracted and inserted to Y_START = 300 pixels. This has also been done for full-frame data. Table 5.2 shows the result from the focal plane drift mask analysis.

Figure 5.3.5 shows an example drift-mode image with the focal plane mask inserted to Y_START = 300. The bi-modal background level seen in figure 5.3.2 is significantly reduced, but is just visible at a low level. This is because the lower area of the window exists on the CCD for longer than the upper half and therefore still accumulates dark current and possibly stray light.

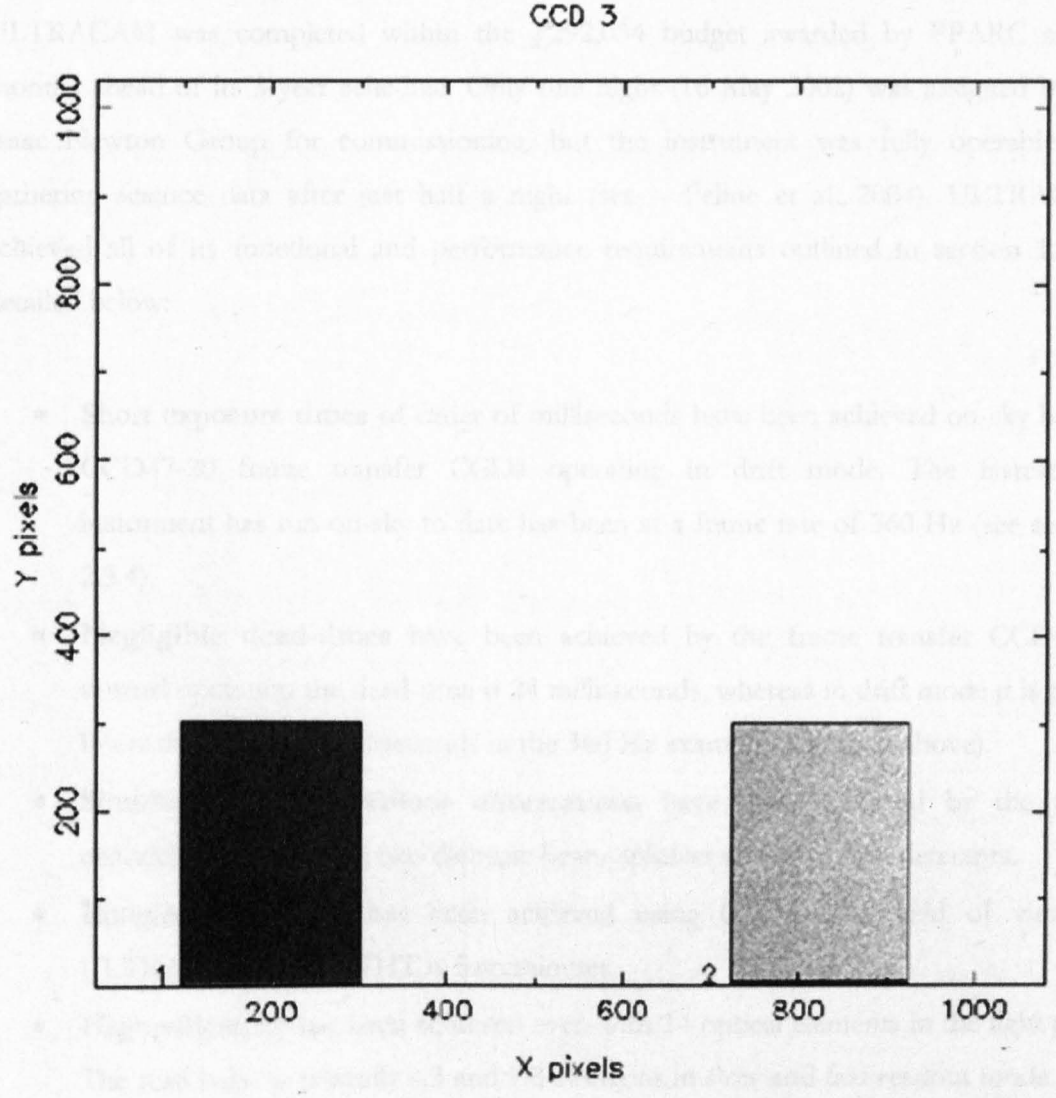

Figure 5.3.5: Example drift-mode image with the focal-plane mask inserted.

Chapter 6

Conclusions and Future Work

6.1 How successful was the project?

ULTRACAM was completed within the £292,034 budget awarded by PPARC and 3 months ahead of its 3-year schedule. Only one night (16 May 2002) was assigned by the Isaac Newton Group for commissioning, but the instrument was fully operable and gathering science data after just half a night (see – Feline et al. 2004). ULTRACAM achieved all of its functional and performance requirements outlined in section 1.4, as detailed below:

- **Short exposure times** of order of milliseconds have been achieved on-sky by the CCD47-20 frame transfer CCDs operating in drift mode. The fastest the instrument has run on-sky to date has been at a frame rate of 360 Hz (see section 2.3.4).
- **Negligible dead-times** have been achieved by the frame transfer CCDs. In normal operation the dead-time is 24 milliseconds, whereas in drift mode it is much lower than this (0.5 milliseconds in the 360 Hz example described above).
- **Simultaneous multi-colour observations** have been achieved by the three channels created by the two dichroic beam-splitters and opposing detectors.
- **Imaging capability** has been achieved using CCDs. The field of view of ULTRACAM on the WHT is 5 arcminutes.
- **High efficiency** has been achieved even with 14 optical elements in the light path. The read noise is typically 4.3 and 9.8 electrons in slow and fast readout mode. The

throughput of ULTRACAM is 49% in u', 65% in g', 49%, in r', 66% in i' and 42% in z'.

- **Portability** has been achieved by designing ULTRACAM to fit the smallest telescope size and weight envelope that it will visit; the instrument weighs only 82-kg and has an overall length of only 792 mm allowing it to mount onto the Aristarchos telescope (as well as the WHT and VLT).

The success of ULTRACAM is reflected by the science data that it has produced and the demand for the instrument. Currently 7 refereed papers have been written, along with several other publications, and these are listed in appendix 1. A total of 61 nights on the WHT have been awarded to ULTRACAM, to date (spanning the period May 2002 – August 2004). These observing proposals originate from collaborators at the Universities of Sheffield, Southampton and Warwick and from the UK and overseas astronomy community. The successful observing proposals are listed in appendix 2.

A further mark of success is that on 1st September 2004, ULTRACAM was officially approved to mount on the visitor focus of the 8.2 m VLT in Chile during period 75 (April – September 2005). This makes ULTRACAM not only the first UK instrument on the VLT, but also the first visitor instrument on the VLT.

6.2 Future work

The ULTRACAM project has been awarded an additional £67, 000 by PPARC to mount it on the Aristarchos telescope and the VLT. The money has been used to procure ULTRACAM collimators for each telescope and will fund the manufacture of ULTRACAM – telescope interface collars. A collaborative arrangement with the Aristarchos telescope will provide the instrument team an allocation of 37 nights per year on the telescope. This is on the condition, however, that ULTRACAM remains mounted on the Aristarchos telescope at all times other than when it is required on a larger telescope (e.g. WHT or VLT). This will ensure that ULTRACAM will spend the majority of its time on a telescope gathering data and not in an instrument store.

Further funding has also been awarded to the ULTRACAM project by the EU to adapt a Low-Light-Level CCD (LLLCCD) to replace one of the ULTRACAM detectors to perform spectroscopic observations. The LLLCCDs offer a reduced readout noise of less than one electron rms (Jerram et al), which makes them ideal for low light level observations, such as spectroscopy.

At slow readout speeds (e.g. kHz pixel rates, such as used in ULTRACAM) typical CCDs exhibit a readout noise of 3 electrons, whilst at fast readout speeds (MHz pixel rates) the readout noise is typically 10 or more electrons (Basden et al 2003). The readout noise originates in the output electronics used to convert the electrons to ADUs (counts). In an LLLCCD, the array and output electronics are largely the same as a conventional CCD except that the serial register is extended to apply an on-chip gain to the original signal (see figure 6.2.1). The extended readout register uses high electric fields to cause avalanche multiplication; accelerating electrons collide with other electrons in the semiconductor releasing them from their bonds. By the time the original electrons from a pixel have reached the output amplifier a shower of electrons exists. Therefore, if 1 photon is detected in a pixel, it will be amplified during the gain stage, read out and be accurately identified as signal and not noise.

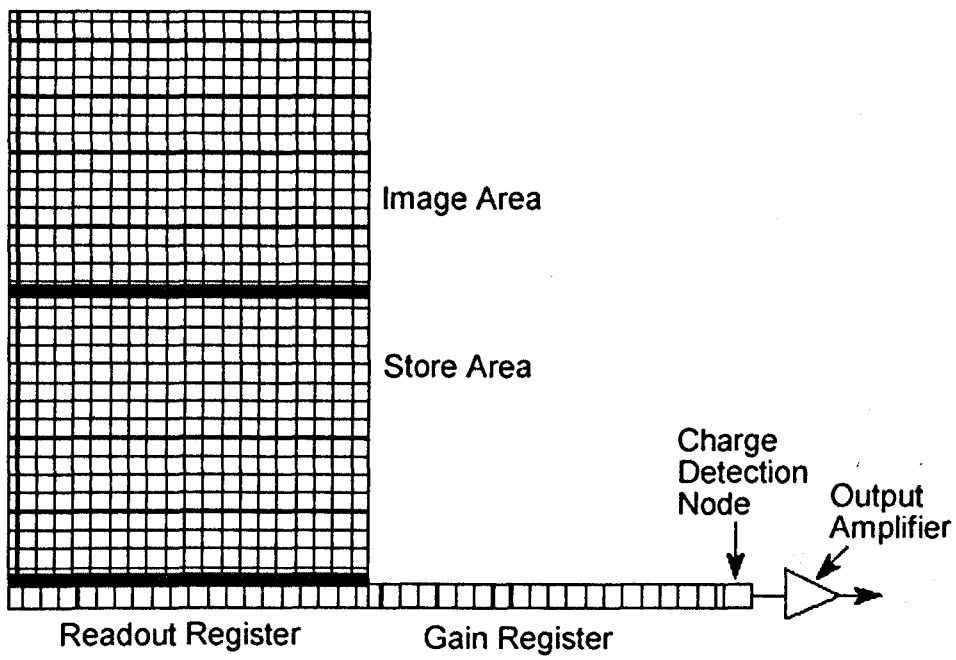

Figure 6.2.1: Layout of the frame transfer E2V LLLCCD.

The amount of gain between each pixel shift along the serial register is small, typically 1% – 2%, so there are many pixels in the extended serial register to give sufficient total gain (Basden et al. 2003), e.g. a typical mean gain of 6629 has been quoted for the CCD65 from E2V Technologies. As the gain is increased, the effective readout noise decreases. However this is limited by the dynamic range of the CCD. The full well capacity of the CCD can be exceeded by the charge built up by the gain process. This restricts the number of observed photons that the detector can be exposed to and therefore limits the brightness of the observed object and/or the integration time.

The gain process is statistical and the total number of final electrons is a probability distribution, as shown in figure 6.2.2. The probability distribution has a mean input of n photons (1 – 4 in figure 6.2.2) multiplied by the mean gain g and a variance of ng^2. There is uncertainty in determining the number of input photons from the final detected signal because of overlap between the different probability distributions in figure 6.2.2.

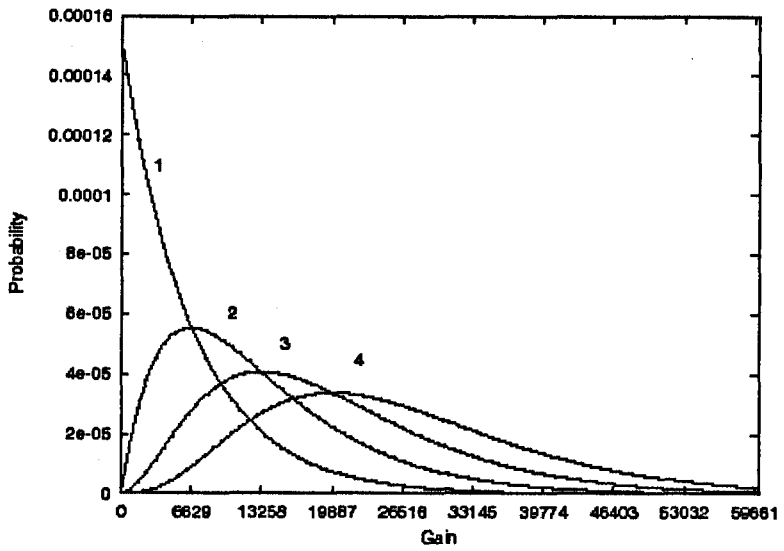

Figure 6.2.2: Probability distribution of output electrons for input photons n = 1, 2, 3, 4. (Basden et al. 2003)

Readout mode strategies for astronomical LLLCCDs have been investigated by Basden et al (2003). The recommended strategy is to have a low signal input so that one or less photons are detected per pixel per frame. Any detectable signal above the noise is assumed to originate from the target source. This is ideal for spectroscopy since light levels are usually low by the time it has passed through the grating and is incident on the detector. Other strategies for regimes where the signal input is >1 photons involve setting threshold

levels. The threshold levels are defined by estimating the most likely input number of photons for a given output. However, uncertainties in the predictions arise because the gain and mean light level is not accurately known. This is a considerable problem for photometric studies because although the readout noise during the observation is reduced, the calculated absolute photometric data will have uncertainty associated with it because the input signal is not known accurately.

Appendices

Appendix 1 – Papers resulting from ULTRACAM data

Jeffery, C. S.; Dhillon, V. S.; Marsh, T. R.; Ramachandran, B. Multicolour high-speed photometry of pulsating subdwarf B stars with ULTRACAM, 2004, MNRAS, 352, 699J.

Feline, W. J.; Dhillon, V. S.; Marsh, T. R.; Stevenson, M. J.; Watson, C. A.; Brinkworth, C. S. ULTRACAM photometry of the eclipsing cataclysmic variable OU Vir, 2004, MNRAS, 347, 1173

Shahbaz, T.; Dhillon, V. S.; Marsh, T. R.; Zurita, C.; Haswell, C. A.; Charles, P. A.; Hynes, R. I.; Casares, J. Multicolour observations of V404 Cyg with ULTRACAM, 2003, MNRAS, 346, 1116

Feline, W. J.; Dhillon, V. S.; Marsh, T. R.; Brinkworth, C. S. ULTRACAM photometry of the eclipsing cataclysmic variables XZ Eri and DV UMa, MNRAS, in press

Beard S.; Vick A.; Atkinson A.; Dhillon V. S.; Marsh T. R.; McLay S.; Stevenson M. J.; Tierney C.; The Ultracam Camera Control and Data Acquisition System", in "Advanced Telescope and Instrumentation Control Software II", Lewis, H., SPIE 4848, "Astronomical Telescopes and Instrumentation", Waikoloa 2002.

Other publications include are:

"New instruments give UK astronomers control of the time domain". RAS Press Notice PN04-03, Date: 9 February 2004. Issued by Peter Bond, RAS Press Officer (Space Science).

"Multicolour observations of V404 Cyg and J1118+480 with ULTRACAM". Poster at CfA conference on X-ray timing

"Multiwavelength observations of eclipsing polars". Cape Workshop on Magnetic CVs paper

"ULTRACAM successfully commissioned on the WHT". ING bulletin article

"New Telescope Camera Targets Fast-Spinning Hearts of Distant Galaxies". By Robert Myers, Special to SPACE.com

"100 Photographs in the Blink of an Eye". PPARC press release.
www.pparc.ac.uk/Nw/Press/ultracam.asp

Appendix 2 – List of successful ULTRACAM observing proposals

Marsh (Warwick). Stochastic variability of accreting white dwarfs

Littlefair (Sheffield). The quiescent accretion disc in the dwarf nova IP Peg

Dhillon (Sheffield). ULTRACAM observations of the transiting extrasolar planet HD209458b

Aerts (Nijmegan). Asteroseismology of the pulsating sdB star PG0014+067

Jeffery (Armagh). Mode identification from multicolour photometry of the pulsating sdB star PG0014+067

Casares (IAC).Echo tomography of flourescence lines in Sco X-1

ING. ULTRACAM service time

Marsh (Warwick). ULTRACAM observations of detached white dwarf/M dwarf binary stars

Roques (Meudon). Search for small Kuiper Belt objects by stellar occultations

ING. ULTRACAM service time

Shahbaz (IAC). High time-resolution imaging of the black hole X-ray transient J0422+32

van der Klis (Amsterdam) & 2 & Optical analogues of X-ray timing phenomena in X-ray binaries using ULTRACAM and RXTE

Dhillon. ULTRACAM observations of the transiting extrasolar planet HD209458b

Mathioudakis (QUB). High frequency oscillations in active cool stars

Marsh. Stochastic variability of accretion discs

Fitzsimmons (QUB). The 13th/14th November stellar occultation by Titan

Shahbaz (IAC). High time-resolution optical studies of quiescent X-ray transients

Marsh. ULTRACAM observations of interacting binary stars

Jeffery (Armagh). Colorimetric mode identification in pulsating subdwarf B stars

O'Brien (Amsterdam/ESO). Optical analogues of X-ray timing phenomena in X-ray binaries using ULTRACAM and RXTE

Shahbaz (IAC). High time-resolution optical studies of quiescent black hole X-ray transients

van Kerkwijk (Utrecht/Toronto). The nature of magnetars: ULTRACAM observations of the anomalous X-ray pulsar 4U 0142+61

Dhillon. ULTRACAM commissioning

Marsh. Magnetic braking and solar cycles in detached binary stars

Dhillon. Coordinated optical and X-ray observations of the eclipsing polar HU Aqr

Bibliography

Acutime 2000 reference sheet, Trimble Navigation Limited.

Atkinson D.; 2001, (CCD) Application documentation issue 1.

Atkinson D.; 2001, Ultracam System Maintenance Document (CCD) issue 1.

Baldwin J. E.; Tubbs R. N.; Cox G. C.; Mackay C. D.; Wilson R. W.; Andersen M. I.; 2001, Astronomy and Astrophysics, **368**, L1.

Basden A.G.; Haniff C. A.; Mackay C. D. 2003. Monthly Notices of the Royal Astronomical Society. **345**, 985-991.

Beard S.; 28 August 2002, Advanced Telescope and Instrumentation Control Software II, 4848-23.

Benford D. J.; Voellmer G. M.; Chervenak J. A.; Irwin K. D.; Moseley S. H; Shafer R. A.; Stacey G. J.; Staguhn J. G.; 2002; Thousand - element multiplexed superconducting bolometer arrays.

Benn C.; 1999, WHT observers guide.

Benn C.; 2001, WHT observers guide.

Benn C.; 2002, WHT observers guide.

Brown T. M.; Charbonneau D.; Gilliland R. L.; Noyes R. W.; Burrows A.; 2001, Ap J, **552**, 699 – 709.

Cabrera B.; Clarke R.; Miller A.; Nam S. W.; Romani R.; Saab T.; Young B.; 1999, Cryogenic detectors based on superconducting Transition-Edge Sensors for time-energy resolved single-photon counters and for dark matter searches, Physica B.

Čadež A.; Vidrih S.; Galičič M.; Carramiñana A.; 2001, Astronomy and Astrophysics, **366**, 930.

Carter D.; 1995, Change due to uncollimated light, Appendix D of the observers guide, ING.

Carter D.; 1996, ING observers guide.

Charbonneau D.; 2001, Scientific frontiers on research of extrasolar planets, ASP conference series.
Cheng, K. S.; Ruderman, M.; Zhang, L.; 2000,ApJ, 537, 964

Cocke, W. J.; Disney, M. J.; Taylor, D. J.; 1969, Nature, **221**, 525.

Clampin; M., 1997, Instrument Science Report, ISR 96-31

Cumani C.; Mantel K.-H.; 2001, Experimental Astronomy, **11**, 145.

CVI Laser; www.cvilaser.com/static/tech_dichroic.asp.

Daugherty, J. K.; Harding, A. K.; 1996, ApJ, 458,278

De Bruijne J. H. J.; Reynolds A. P.; Perryman M. A. C.; Favata F.; Peacock A, 2001; Preprint arXiv:astro-ph/0108232.

Deeming T.J.; 1975. Fourier analysis with unequally spaced data. Astrophys. Space Sci, **36**, 137 – 158.

Dhillon V. S.; 1997. Starlink user note 167.

Dravins D.; 1994, ESO Messenger No. 78, 9 – 19.

Dravins D.; Faria D. Nilsson B.; 2000, Avalanche Diodes as photon-counting detectors in astronomical photometry, SPIE Proc. Vol. **4008**, 298 – 307.

Elliot J. L. et al; 1998 Icarus **77**, 148-170.

ESO 0.5 m Telescope, La Silla, Photometer – characteristics web page ls.eso.org/lasilla /Telescopes/2p2T/E0p5M/Photometers/photometers.html.

Feline W. J.; Dhillon V. S.; Marsh T. R.; Stevenson M. J.; Watson C. A.; Brinkworth C. S.; ULTRACAM photometry of the eclipsing cataclysmic variable OU Vir, 2004, MNRAS, 347, 1173.

Feline W. J.; Dhillon V. S.; Marsh T. R.; Stevenson M. J.; Watson C. A.; Brinkworth C. S. Erratum: ULTRACAM photometry of the eclipsing cataclysmic variable OU Vir, 2004, MNRAS in press.

Feline W. J.; Dhillon V. S.; Marsh T. R.; Brinkworth C. S.; ULTRACAM photometry of the eclipsing cataclysmic variables XZ Eri and DV UMa, 2004, MNRAS in press.

Fellers T. J.; Davidson M. W.; Charge-Coupled Device (CCD) Linearity, National High Magnetic Field Laboratory

Fordham J. L. A.; Kawakami H.; Michel R. M.; Much R.; Robinson J. R.; 2000, MNRAS, **319**, 414.

Fordham J. L. A.; Vranesevic N.; Carramiñana A.; Michel R.; Much R.; Wehinger P.; Wyckoff S.; 2002, ApJ, **581**, 485.

Filippenko A. V.; 1982, PASP, **94**, 715.

Frogel J. A.; 2002, Image quality at selected astronomical observatories, Lawrence Berkeley National Laboratory.

Fukugita M.; Ichikawa T.; Gunn J. E.; Doi M.; Shimasaku K.; Schneider D. P.; 1996, ApJ, **111**, 1748.
Gilliland R. L.; 1992, Astronomical CCD observing and reduction techniques, ASP conference series, volume **23**.

Hewish, A.; Bell, S. J.; Pilkington, J. D. H.; Scott, P. F; Collins, R. A; 1968, Nature, **217**, 709.

Hippel T.; Winge C.; 2002, User guide, Acquisition camera detector, Gemini observatory.

Howard N. E.; 2002, Opsci Application Note OAN-006, Photon Transfer Technique.

Hynes R.; 2001, MNRAS, **323**, L26.

Jeffery C. S.; Dhillon V. S.; Marsh T. R.; Ramachandran, B.; Multicolour high-speed photometry of pulsating subdwarf B stars with ULTRACAM, 2004, MNRAS, 352, 699J.

Jerram P.; Pool P.; Bell R.; Burt D.; Bowring S.; Spencer S.; Hazelwood M.; Moody I.; Catlett N.; Heyes P.; The LLLCCD: Low Light Imaging without the need for an intensifier. Marconi Applied Technologies, Chelmsford, Essex, UK.

Janesick J., Elliott T., 1992, Astronomical CCD observing and reduction techniques, ASP conference series, volume **23**.

Kern B.; 2002, Ph.D. Thesis, California Institute of Technology.

Kotar J.; Vidrih S.; Čadež A.; 2003, Rev Sci Instrum, **74**, 3111 (astro-ph/0303368)

Koren B.; 2001, SPIE's **OE**Magazine, August, 34.

Lyne A. G.; Jordan C. A.; Roberts M. E.; 2003. Jodrell Bank Crab Pulsar Timing Results, Monthly Ephemeris.

Marsh T. R.; Horn K.; 1998, MNRAS, **299**, 921 – 928.

Marsh T. R.; 2004, Pipeline data reduction software help notes.

Mather J.; Seery B.; Stockman P.; Bely P.; Burg R.; Burt J.; Geithner P.; Greenhouse M.; Hunter D.; Isaacs J.; Ferguson H.; Long K.; Petro L.; 1999, NGST Reference Design Requirements, Recommendations And Guidelines, NGST MONOGRAPH NO. 5

McLean I. S.; 1997, Electronic imaging in astronomy, Detectors and instrumentation. Wiley-Praxis series in astronomy and astrophysics.

Middleditch, J.; Pennypacker, C.; 1985, *Nature*, **313**, 659.

Nave R.; 2004, Anti-Reflection Coatings, hyperphysics.phy-astr.gsu.edu/hbase/phyopt/antiref.html.

Nolan P. J.; Fundamentals of college physics, 2nd edition, William C. Brown Publishers, 1995.

Orion Optics. www.orionoptics.co.uk/hilux.htm
Pacini, F.; 1971, ApJ, 163,17

Pacini, F.; Salvati, M.; 1983, ApJ, 274, 369

Pacini, F.; Salvati, M.; 1987, ApJ, 321, 445
Parker Q. A., Bland-Hawthorn J., 1998, PASA, 15 (1), 33.

Peacock T.; Verhoeve P.; Rando N.; Erd C.; Bavdaz M.; Taylor B. G.; Perez D.; 1998, Astronomy and astrophysics supplement series, **127**, 497.

Perryman, M. A. C.; et al, 1999, Astron. Astrophys, **346**, L30.

Pirgon O.; Wolstenholm G.H.; Yates B.; 1973, J. Phys. D: Appl. Phys. **6** 309-321.

Rando N.; Andersson S.; Collaudin B.; Favata F.; Gondoin P.; Peacock A.; Perryman M.; Verveer J.; Goldie D. J.; 1999, First astronomical results of S-Cam, 8th International workshop on low temperature detectors.

Rapkin E.; 1997, Photomultipliers and how they operate, http:/www.intl-isotope-soc.org/tech1.htm.

Redfern M.; Shearer A.; O'Kane P.; O'Byrne C.; Wouts R.; Read P.; Carter M.; Jordan B.; Cullum M.; 1992, Gemini, newsletter of the RGO, **38,** 1.

Roark R. J.; Young W. C.; Formulas for Stress and Strain, 1975, McGraw-Hill Kogakusha, Ltd. p324.

Romani R. W.; Miller A. J.; Cabrera B.; Figuera-Feliciano F.; Nam S. W.; 1999, ApJ, **521**, L153 – L156.

Rutten R. G. M.; Gribbin F. J.; Ives D. J.; Bennett A.; Dhillon V. S.; 1997, User manual, Drift-mode CCD readout. Isaac Newton Group, La Palma.

Sanwal D.; Robinson E. L.; Stiening R. F.; AAS Meeting #193, Session 112. Pulsars II. **112.04**. (http://www.aas.org/publications/baas/v30n4/aas193/226.htm)

Schott, Mainz, Germany. Schott.com

Shearer A.; et al, 1997, *ApJ*, **487**, L181.

Shearer A.; et al, 1998, *A&A*, **335**, L21.

Shearer A.; Golden A.; 1999, ASP Conference series, Vol.

Siegmond O. H. W.; 2000, Microchannel plate imaging detector technologies for UV instruments, Space science laboratory.

Sinnott R. W.; Sky and Telescope, Feb, 1994, 91

Simons D.; Longitudinally averaged R-band field star counts across the entire sky. Gemini Technical Note, August 1995.
Smith J. A.; et al, 2002, *ApJ*, **123**, 2121.

Stathakis R. A.; Johnston H. M.; 2002, User manual, RGO spectrograph. Anglo-Australian Observatory.

Straubmeier C.; Knabach G., Schrey, F., 2001, Pre-print arXiv:astro-ph/0109181.

Wallace, P. T.; et al. 1977, Nature, **266**, 692.

Wallace P. T.; 2003. Starlink user note 67.61.

Wilson et al.; 1999, **309**, 379. MNRAS.