ULTRACAM – AN ULTRA-FAST TRIPLE-BEAM CCD CAMERA

Vik Dhillon and Tom Marsh

28 February 1999

1 Introduction

The history of observational astronomy shows that major advances in the science almost always result when a new area of observational parameter space, such as wavelength coverage, angular resolution, sensitivity or sky coverage, becomes accessible for exploration. Such advances are usually associated with the development of new instrumentation which exploits the latest available technology. It has been the UK's ability to exploit such technology which has placed it at the leading edge of astronomy today.

This proposal is concerned with the construction and operation of just such an instrument – a new ultra-fast, triple-beam camera using the latest in CCD technology to study one of the few remaining unexplored regions of observational parameter space – high temporal resolution. The proposed camera, which we shall hereafter refer to as ULTRACAM, will be commissioned at the GHRIL focus of the 4.2-m William Herschel Telescope (WHT) and then used there and on other telescopes (including the 2.5-m Isaac Newton Telescope, the 2-m Liverpool Telescope, the 1.9-m Radcliffe Telescope and the 9.2-m South African Large Telescope) to study interacting binary stars. The instrument has been designed with versatility and portability in mind, and so we anticipate strong demand from astronomers working in different areas of astrophysics who want to use ULTRACAM on these and other telescopes.

ULTRACAM will combine the best features of CCD and photoelectric photometers, which are:

1. High efficiency over a wide wavelength range (c.f. photoelectric photometers).
2. High-speed read-out (c.f. other CCD-based systems, many of which take up to a minute to read out a small window).
3. Software apertures to cut down on the sky background (c.f. the large, fixed apertures of photoelectric photometers).
4. Simultaneous measurement of sky and comparison stars (c.f. all but the most complex photelectric photometers).
5. Simultaneous measurement of 3 bands, essential for all non-periodic variablity studies and for precise colours.
6. Multiplex advantage of measuring in 3 bands simultaneously, together with reduced filter changing overheads.

When La Palma decided to make the JKT a single-instrument telescope in 1996 in order to reduce running costs, UK astronomers lost access to the last of their research-class photoelectric photometers (the Peoples' Photometer and the Multi-Purpose Fotometer) and hence lost the capability to do high-speed, multi-waveband optical photometry. For all the advantages of the CCD cameras which replaced these instruments, they have never been able to match the speed of photoelectric photometers – until now. ULTRACAM will employ 3 frame-transfer CCD chips to image a field of up to 5 arcminutes with essentially no dead-time (even at >100 Hz frame rates) simultaneously in three different colours. ULTRACAM will be a unique instrument – to the best of our knowledge, no other such instrument either exists or is on the drawing board anywhere in the world. As well as giving back to UK astronomers a long-awaited and far-enhanced capability to do high-speed, multi-wavelength optical photometry, ULTRACAM will open up a completely new region of observational parameter space thanks to the greater sensitivity and versatility of CCDs (compared to photoelectric tubes) coupled to the efficiency and aperture of large, modern-day telescopes.

The applicants have a great deal of experience in the design, construction and commissioning of astronomical instrumentation (Dhillon) and in the development of software to reduce and analyse the resulting data (Marsh). The instrument will be assembled at Sheffield, where we will draw on the resources and expertise currently in place for the construction of detectors for the UK Dark Matter Collaboration. The instrument will be built in collaboration with the UKATC, who will deliver the detector system, and Southampton, who will deliver the pipeline data reduction software. The remainder of this document expands on all of the points raised in this introduction, starting with a detailed scientific justification for ULTRACAM.

2 Scientific justification

There are many projects that can benefit from ULTRACAM. In many cases, they can *only* be done with ULTRACAM. We now give descriptions of three projects that we wish to perform with ULTRACAM that will benefit from one or more of the capabilities listed in the introduction.

2.1 Eclipse mapping of accretion discs

Accretion discs are pivotal in the transport of angular momentum in a wide variety of astrophysical objects. It is through them that protostars can grow and the compact objects at the heart of many X-ray sources can be fuelled. Much of what we know about these discs has been learnt from close binary stars where the geometry and physical parameters can be measured precisely.

The most informative systems are eclipsing systems in which the mass donor star (which often contributes negligibly in terms of its own light) passes between us and the disc. The light curve in eclipse can then be used to map the brightnesses and temperatures over the disc, the accretor (if visible) and the impact region between the mass transfer stream and disc. It is now over a decade since Horne (1985, MNRAS, 213, 129) published a general approach to this analysis called "eclipse mapping". The basis of this technique is to model the disc as a grid of many pixels, and then to adjust their brightnesses in order to obtain a model light curve that fits the data. Of the many possible images that fit, the smoothest (in the sense of the image of maximum entropy) is selected. This technique has produced the only observational confirmation of one of the fundamental predictions of the theory of accretion discs – the $R^{-3/4}$ "law" for the temperature distribution of steady-state accretion discs.

Relatively few systems are bright enough for photoelectric photometry. This is particularly true of dwarf and X-ray novae in their low states. For the past few years, about one new X-ray nova, many of which harbour black-holes, has been discovered each year. It can only be a matter of time before a deeply eclipsing system is discovered, giving us the chance to see where in the system the optical light comes from in quiescence (the standard disc model and the more recent advective solutions differ sharply on this). It is a sobering thought that if such a system were to be discovered, the UK would not have access to any photometer that could obtain competitive data on it.

The two essential requirements for such work are short exposures, of order a second or two at most, and multi-band data (3 or more) in order to distinguish one type of spectrum from another (i.e. with only 2 bands you cannot distinguish a blackbody from a star; with 3 bands you can, e.g. Horne & Cook 1985, MNRAS, 214, 307). Simultantaneously recording the 3 bands is vital because accreting systems display erratic "flickering" which cannot be reliably disentangled from colour variations in sequentially acquired data.

2.2 Angular momentum losses and solar cycles in close binaries

The orbital periods of close binaries are continually evolving, due, it is thought, to angular momentum losses from gravitational wave radiation and the same magnetic stellar wind braking that slows down single stars. On top of this there are short-term fluctuations, possibly induced by changes in the structure of one of the stars caused by solar-type cycles (e.g. Skillman & Patterson 1988, AJ, 96, 976). Precise measurement of orbital periods can reveal these effects and is easiest in systems in which a white dwarf is eclipsed as the sharp ingress and egress provide excellent fiducials. With these systems we have the chance to measure angular momentum losses over a wide range of spectral types, and for instance we will be able to test the hypothesis that magnetic wind braking reduces when stars become fully convective at masses around $0.3\,M_{\odot}$.

As with eclipse mapping, exposure times must be no longer than a few seconds. Simultaneous multi-band data is required to correct for the effects of erratic variability upon eclipse timings. Such variability may be caused by the rapidly-rotating main-sequence star (e.g. V471 Tau), or by flickering in accreting systems.

There are now of order 15 detached and semi-detached systems for which the above method can be applied. However, there has been little work on most of them. For example, the system NN Ser, of particular interest as its short period implies a fully-convective M dwarf, has had no published eclipse times since those published by Wood & Marsh (1991, ApJ, 381, 551) soon after Haefner (1989, A&A, 213, L15) discovered it. At $V = 16.6$, NN Ser is getting faint for high-speed photometry, but would be a straight-forward target for ULTRACAM.

2.3 Short-period and stochastic variability

Accreting binaries show a wide range of rapid variability at optical wavelengths. These range from ultra-coherent pulsations originating from spinning accretors (e.g. AE Aqr, Her X-1), through "dwarf nova oscillations" (periods from 3 to 30 seconds) and quasi-periodic oscillations, to the erratic flickering already mentioned. Along with these comes possible variability from the rapidly-rotating mass donor. Apart from the coherent pulsations rather little is known about the above phenomena.

In recent years the discovery of black-hole transients has given us a new set of objects that show erratic variability. There is particular interest in the possibility that the central accretion flow in these objects may be quasi-spherical and "advective" (i.e. carrying the gravitational energy into the black-hole rather than radiating it, e.g. Naryan et al. 1997, ApJ, 482, 448). Examples of these systems are V616 Mon and V404 Cyg which have been seen to vary by ≈ 0.5 mags in quiescence, all the more remarkable given that their donor stars probably contribute $> 80\%$ of the flux (Pavlenko et al. 1996, MNRAS, 281, 1094). The black-hole transients are typically fainter than 18 in quiescence, which rules out anything but CCD-based photometry (see Appendix C). As the first multi-colour, high-speed CCD photometer, ULTRACAM would have a substantial advantage over any competitor instruments in this field.

In outburst, extraordinary variability has been observed: V404 Cyg varied by 0.7 mags on minute timescales and at X-ray wavelengths changed by factors of 500 in a matter of seconds (Kitamoto et al. 1989, Nature, 342, 518). These sorts of variations make possible Horne's echo mapping method (1997, Astronomical Time Series, p109). Data from ULTRACAM would be ideal for such investigations.

2.4 Other projects:

We do not have space here to describe other possible projects in any detail, and so we finish with a list of some of these to give an idea of the broad range of scientific applications of ULTRACAM:

1. Shift-and-add imaging, e.g. to obtain accurate light-curves of globular cluster X-ray sources.
2. Speckle imaging, studies of atmospheric turbulence and lunar occultations.
3. Follow-up on objects found in variability surveys (e.g. the INT+WFC variability survey). Many of these will be faint and have very short periods.
4. Extending techniques developed for the brighter eclipsing cataclysmic variables (CVs) to fainter systems. e.g. the use of the eclipse of the gas stream impact region and white dwarf for measuring precise system parameters.
5. Identification of emission line sources (e.g. globular cluster CVs) with simultaneous imaging in e.g. Hα and the nearby continuum.
6. Any projects requiring (a) time-series photometry of faint objects, (b) accurate colours or (c) efficient multi-band imaging.

3 Technical justification

3.1 Optical design

The proposed optical layout of ULTRACAM is shown in Appendix A. The Head of the Optics Group at the RGO prior to closure (Sue Worswick) has confirmed that our design is feasible and appropriate given our scientific requirements. Note that our concept is not a new one – similar three-way splitting of red, green and blue light using dichroics can be found in most TV cameras.

There are four optical units in ULTRACAM – the collimating fore-optics, the dichroic beamsplitters, the filters and the re-imaging lenses. The collimating fore-optics will collimate the light from the telescope. This will be an interchangeable lens unit which will enable ULTRACAM to work on telescopes of different focal ratio. Such a scheme will also allow a variety of different pixel scales to be produced, depending on the scientific requirements of the project. The light is then split into blue, green and red wavelengths by two dichroic beamsplitters. With modern coatings, a clean three-way split of the light in ULTRACAM is perfectly feasible. The light in each arm then passes through standard 50mm Harris filters of the type in use on La Palma. The light is still collimated at this stage and hence no change of focus is required for differing filter thicknesses. The light is then re-focussed onto the CCD chip using re-imaging optics positioned in front of the detectors.

In order to obtain maximum throughput from U-band through to Z-band, we will use custom-made, coated optics in ULTRACAM.

3.2 Mechanical design

It is premature to present a detailed mechanical design of ULTRACAM at this stage. All of the optics will be interchangeable by hand, i.e. there are no moving parts in ULTRACAM, which simplifies the design (and reduces the cost) considerably. We will build a single opto-mechanical chassis of black-anodised aluminium to house all of the optical components, with each component having adjustable mountings for easy alignment. The chassis will have three kinematic seats for the three CCD heads and an interchangeable flange at the collimator end to enable future attachment to Cassegrain turntables. Flexure is not of great concern with ULTRACAM as it is primarily an imaging photometer. The chassis will be designed in such a way that it flexes by no more than 100 microns (8 pixels) during typical operation at the Cassegrain focus of a telescope. With the above mechanical design, and employing reduced-volume (1.5 litre) CCD cryostats, ULTRACAM will be a simple, compact, lightweight and hence portable instrument.

3.3 Detector system

The specifications of the ULTRACAM detector system are listed in Appendix B. We will use three EEV 47-20 CCD chips as the detectors for ULTRACAM. These are high quality, state-of-the-art chips which have identical quantum efficiency and noise characteristics to the EEV CCDs recently commissioned on La Palma and which have been ordered for Gemini. The maximum delivery time for these chips is currently estimated to be 9 months. We will order three such chips from EEV, preferably with blue-optimised, green-optimised and red-optimised coatings, respectively. At the time of writing, it remains unclear if such coatings will be available. If they are, it should be possible to obtain quantum efficiencies of greater than 80% across the entire wavelength range spanned by a single arm of ULTRACAM, more than offsetting the losses (typically 0.5% per surface when narrow-band coated) in the optics required to split the light into three beams.

Smaller versions of these EEV frame transfer chips have already been developed by the UKATC for use in the wavefront sensor for the NAOMI adaptive optics instrument on the WHT and for the Gemini wavefront sensors. Hence much of the research and development necessary to procure, test, and assemble cameras using these chips has already been done by the UKATC, saving many years of staff effort from the cost of ULTRACAM.

The CCD chips will be either liquid-Nitrogen or closed-cycle cooled in order to minimize dark current; both systems costs similar amounts to construct and result in identical performance - the optimum solution will present itself during the detailed design phase once we are more aware of our mass and space constraints.

The CCDs will be controlled by a San Diego State University (SDSU) controller, provided by Bob Leach. These off-the-shelf controllers have recently been adopted by the ING on La Palma due to their high-speed/low-noise characteristics and due to the fact that they are rapidly becoming the astronomical standard world-wide, with more than 50 observatories now using them. A single SDSU controller will be able to read out all three CCD chips in two-channel mode, giving readout rates of 3.3 μseconds/pixel/readout channel with 6 e^-/pixel readout noise. Using a GPS receiver connected to the SDSU controller we will time-stamp all exposures with ULTRACAM to an absolute accuracy of better than 1 millisecond.

3.4 Data acquisition and reduction system

There are at least two methods of handling the data output by the SDSU controller, both of which cost similar amounts. The most obvious method is to write the data directly from the SDSU controller to a Sparcstation. The drawback with this method is that unix is not an ideal real-time operating system, which could cause problems when running at the high data rates output by ULTRACAM (see Appendix D). In addition, this route relies on the new and relatively untested SDSU-PCI interface and its associated software. The method we favour is based on industry standard VME technology, where the data passes from the SDSU controller into a VME system (which acts as a kind of real-time data buffer) and is then written directly onto the hard-disk of the Sparcstation. This method of data acquisition has been adopted by the UKATC for the Gemini telescopes (on their wavefront sensors) and for the UKIRT/Gemini instrument Michelle. By going the same route with ULTRACAM we ensure that there are few risks, the interfaces are well defined and the systems are well matured by the time construction begins on ULTRACAM. In addition, most of the low-level software we will require for ULTRACAM will be freely available from the UKATC, which will save many staff years of effort from the cost of ULTRACAM. Nevertheless, in spite of these advantages, we do intend to readdress the VME/SDSU-PCI issue during the design phase.

As described in Appendix D, the maximum data rate of ULTRACAM will be 2 Mbytes/second. This means that in the course of a 10-hour night, we could accumulate as much as 72 Gbytes of data. Such vast quantities of data demand two things – disk space sufficient to store at least one nights data and an archiving medium with sufficient capacity and speed to backup at least one nights data before the next night begins. The user-end of the data acquisition system will therefore consist of a SUN Sparcstation Ultra 10, running at 333 MHz with 512 Mbytes of memory, 109 Gigabytes of hard disk space (using a SUN Multipac composed of six 18 Gbyte discs) and a SUN Autoloading DDS3 DAT drive (which can archive up to 144 Gbytes of data on 6 DAT tapes). With this in place, it will be possible to observe for a whole night at the highest data rates without stopping and then archive the data at the end of the night.

Without some form of pipeline reduction, it will simply not be feasible to reduce the terrabytes of data produced in a typical observing run with ULTRACAM. Such a pipeline data reduction system will also allow an immediate assessment of the data, essential for the study of variable objects. As described in Appendix D, we have simulated the performance of a pipeline data reduction system for ULTRACAM and we find that, using the Sparcstation specified above, it will indeed be possible to reduce ULTRACAM data at the same rate as it is taken, thereby allowing us to build up light curves in real time.

4 Project management

The project plan detailed below has been discussed with Adrian Russell (Director, UKATC), René Rutten (Director, ING), Dave Carter (Project Scientist, Liverpool Telescope) and Dave Buckley (Project Scientist, South African Large Telescope). Please refer to Appendix E for additional information and supporting documentation.

4.1 Manpower

Dhillon, who will spend approximately half of his time on the project, will have primary responsibility for the budget and for guiding and coordinating the work of the three groups (Sheffield, Southampton and the UKATC) involved in the construction of ULTRACAM. Dhillon will also be responsible for the final assembly and integration of the instrument in the detector laboratory of the UK Dark Matter Collaboration at Sheffield. Marsh, who will spend approximately one day per week on the project, will have primary responsibility for the delivery of the pipeline data reduction software. All other manpower has been accounted for in the staff effort tables on page 2 of the RG2 form. Specifically, we will also require 1.45 staff years of effort from the UKATC for the delivery of the CCD cameras, CCD controller, VME hardware and all associated software. We will also require 0.1 staff years of effort from the UKATC for the optical design and procurement, 0.25 staff years of technical effort at Sheffield for the design and construction of the opto-mechanical chassis and 0.1 staff years of effort from the part-time Starlink manager at Sheffield for system management of the Sparcstation and peripherals. We will also take advantage of the 1 staff-year of effort available to the UK community at the UKATC for consultation, minor design problems and review attendance.

4.2 Track record

In this section we describe some of the relevant experience the applicants and the Department of Physics and Astronomy at Sheffield have in the field of astronomical instrumentation.

Dhillon – has worked for the ING on La Palma and the RGO in Cambridge for a total of 7 years, during which time he was a Support Astronomer, WHT Telescope Manager, GHRIL Manager and Instrument Scientist for the ISIS spectrograph, the UK Adaptive Optics Programme, WHIRCAM and the Peoples' Photometer. More recently, he was the Programme Scientist for the ING (with scientific responsibility for the ING instrument development and enhancements budget), Technical Secretary to the ING TAGs at PATT, a member of the design review panel for HROS on Gemini South, and is currently a member of the ING Instrument Working Group.

Marsh – was also a support astronomer at the RGO where he was responsible for the then new IPCS on the INT. Marsh has extensive experience of software development for astronomical data reduction and analysis. In particular, Marsh has developed the PAMELA data reduction package, which includes his algorithm for the optimal extraction of curved spectral orders (1989, PASP, 101, 1032), now a standard part of FIGARO. Marsh has also written the MOLLY data reduction and analysis package, which has become the most essential piece of software used by binary star researchers world-wide.

Sheffield – The Department of Physics and Astronomy at the University of Sheffield are very well equipped to build ULTRACAM, thanks to the resources and expertise built up by the UK Dark Matter Collaboration, led by Neil Spooner. They have offered us space and access to equipment in their detector laboratory at Sheffield, which will provide an ideal environment to assemble, integrate and test ULTRACAM. They have also offered us access to their technical staff, all of whom are experienced in the area of instrument design, and their associated optical, mechanical and electronics workshops. It is interesting to note that the UK Dark Matter Collaboration at Sheffield are also embarking on a project to investigate the use of ultra-fast CCDs for dark matter detection, and so we expect there to be many benefits to both groups from the above interaction.

4.3 Costs

A summary of the equipment costs is given in Section 8 of the RG2 form. A more detailed breakdown of these costs, all of which have been supplied by Derek Ives (Head of CCD group, UKATC), Sue Worswick (ex-Head of Optics, RGO) and John Kelly (Head of Mechanical Workshops, Sheffield), is given in Appendix B.

The cost of air-freighting the instrument twice to La Palma and back again is requested in Section 9 of the RG2 form. This will cover the first (engineering) commissioning run and the final (scientific) commissioning and first scheduled observation run. Observing runs beyond this will fall outside the period of the grant and funds for these will be sought from PATT. The anticipated travel and subsistence costs are summarized in Section 11 of the RG2 form. The UK travel covers the visits of Dhillon, Marsh and UKATC staff to the three UK sites (UKATC, Sheffield and Southampton) for project reviews and other essential exchanges relating to the project. The overseas travel covers the 2 commissioning runs on La Palma with 3 people per trip (Dhillon, Marsh and one person from the UKATC CCD group). These numbers have been derived from our experience of supporting commissioning runs on La Palma and are the minimum necessary for the successful commissioning of such an instrument.

4.4 Timescales and milestones

From our previous experience of the design, construction and commissioning of astronomical instrumentation, we believe that we will be able to deliver a fully-commissioned instrument within three years. We are therefore requesting funds for a 3-year period commencing October 1st, 1999. This timescale, and hence the costs, are critically dependent on resources being available at the UKATC when we require them. We therefore propose that if PPARC decide to fund ULTRACAM, it is provisional pending the receipt of a letter from Adrian Russell which commits the UKATC to deliver its share of the project. For further details, please refer to Appendix F.

We will adopt the review structure used by the Gemini Project and the UKATC to define major milestones in the project. There will be three reviews, each attended by a panel of experts from both the university community and the UKATC. The first of these will be the Conceptual Design Review, held 4 months after the commencement of the grant, at which we will present a ray-traced optical design (UKATC and Sheffield), an initial CAD opto-mechanical chassis design (Sheffield) and a block diagram of the data acquisition and pipeline reduction system (Southampton and UKATC). At this review we will identify the main risks which lie ahead and establish that we are ready to order those items with the longest delivery times. At the end of the first year (October 2000) we will hold the preliminary design review, in which a detailed design of every part of the instrument will be presented to the panel for comment. At this review we will confirm that the project is on the right track and identify any elements of the detailed design which need further work. The final review will be the Critical Design Review and will be held in March 2001. At this review we will demonstrate that all of the detailed design work has been completed and all that remains to be done is the remaining procurement, assembly, integration and testing of the instrument. We expect the first (engineering) commissioning run to take place on La Palma in Deember 2001 at the GHRIL focus of the WHT, which provides an ideal environment for the commissioning of such an instrument. The final (scientific) commissioning run and first scheduled observations will take place in GHRIL in September 2002, at the end of the grant period.

The lifetime of this instrument, both in a physical and in a scientific sense, will be many years longer than the lifetime of our proposed 3-year grant. Having commissioned ULTRACAM and completed the first observations of our key science projects, we will welcome applications through PATT from UK astronomers who wish to use ULTRACAM on any telescope, thereby extracting maximum scientific return from PPARC's financial investment.

A ULTRACAM – schematic representation

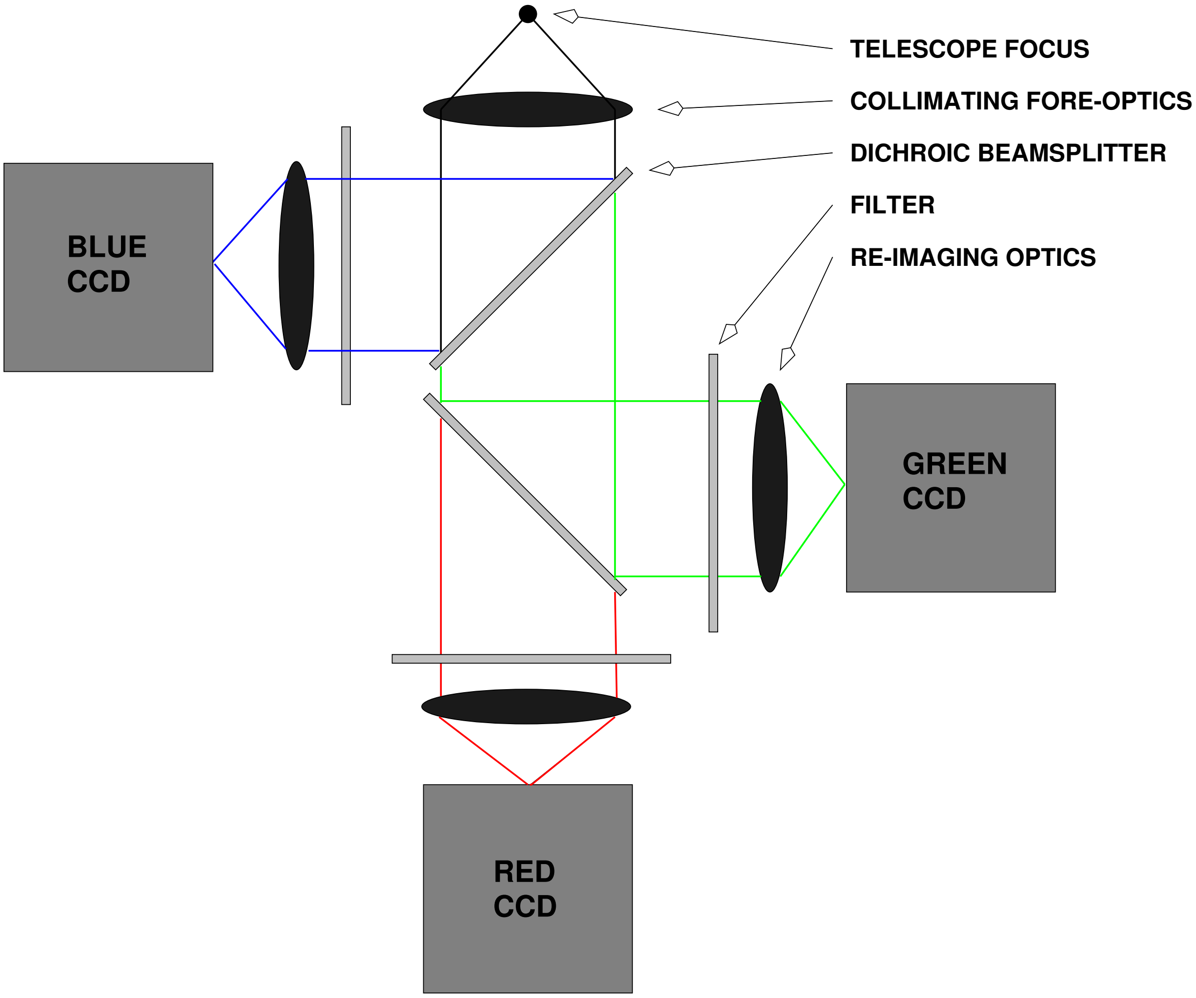

B ULTRACAM – specifications and component costs

ULTRACAM	SPECIFICATIONS	COST (incl. VAT)
DETECTOR SYSTEM		
3 × EEV 47-20 frame-transfer CCDs. Back-thinned, non-IMO, AR-coated chips. Including Oxford Instruments liquid-Nitrogen cryostats, internal wiring, faceplate, AR-coated window, pre-amps and temperature control		3×£16 k
Operating temperature	–110°C, cooled by liquid Nitrogen	
Pixel size/pixel size at GHRIL	13.5 μm/0.25″	
Imaging area/imaging area at GHRIL	1024×1024 pixels/4.25′×4.25′	
Dark current	2 e^-/hour	
Readout noise at slowest readout speed	3 e^- at 10 μseconds/pixel/channel	
Readout noise at fastest readout speed	6 e^- at 3.3 μseconds/pixel/channel	
Full well capacity	150 000 e^-	
Quantum efficiency in U/B/V/R/I/Z pass-bands	60%/83%/80%/74%/46%/20%	
DATA ACQUISITION SYSTEM		
SDSU CCD controller, including VME interface card, 12 slot chassis, 3 clock cards and 3 video cards	Data rates of up to 2 Mbytes/second	£37 k
VME hardware, including:		
VME GPS card	Absolute timing to 1 milli-second	£3.5 k
VME processor		£6.5 k
VME chassis		£2 k
VME ethernet		£1.5 k
VME memory		£2.5 k
VxWorks software		£1.5 k
SUN Ultra-10 333 MHz Sparcstation, including 512 Mbytes memory, SCSI card, SUN Multipac 109 Gbytes hard disk, 19-inch monitor and SUN autoloading 144 Gbytes DDS3 DAT drive	Data storage and archiving capacity matched to 1 night of continuous observing at the fastest data rate. Data reduction rates of at least 1 MByte/second	£18 k
OPTICAL COMPONENTS		
1 × Collimating fore-optics	2% light-losses at each of 6 surfaces	£5 k
3 × Reimaging optics	0.5% light-losses at each of 6 surfaces	£15 k
2 × Dichroic beamsplitters	5% light-losses at each of 2 surfaces	£2 k
UBVRIZ filters	10% light-losses	£3 k
MECHANICAL COMPONENTS		
Opto-mechanical chassis, including adjustable mountings for optics, kinematic seats for CCD cameras, exchangeable fore-optics unit, exchangeable mounting flange for use at Cassegrain and black anodising to minimize scattered light	Maximum flexure of 100 microns (8 pixels) from horizon to zenith when used at Cassegrain	£1 k

C ULTRACAM – limiting magnitudes

The figure below shows the limiting magnitudes for a detection at a signal-to-noise ratio of 10 as a function of exposure time on the WHT at GHRIL. Each panel corresponds to a different bandpass – U, B, V, R, I and Z. In each panel there are 6 curves. The solid curves are for ULTRACAM, with the upper curve representing dark time and the lower curve representing bright time. The dotted curves are also for ULTRACAM, but using on-chip binning of 4×4 pixels. The dashed curves are for a photoelectric photometer, again with the upper curve representing dark time and the lower curve representing bright time. Details of the signal-to-noise calculations used to make these plots and the assumptions which have gone in to them are given in the text at the end of this Appendix.

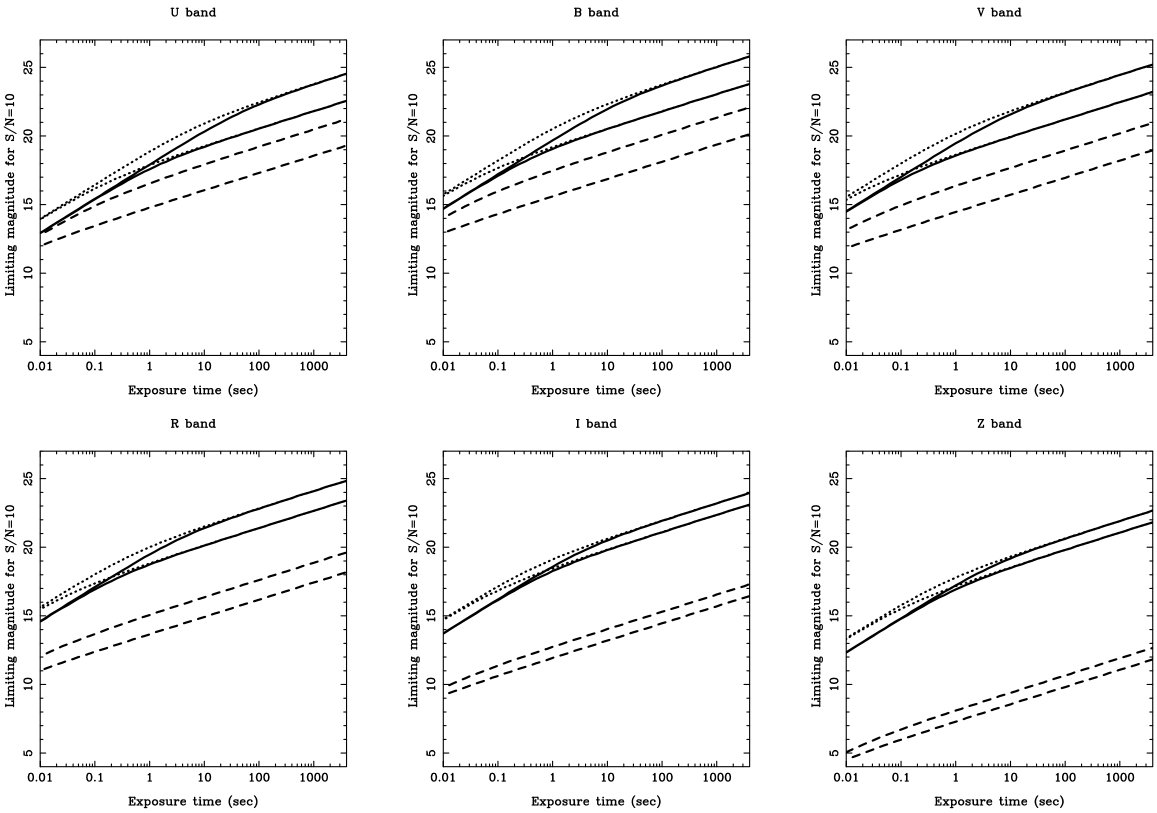

It can be seen that only in the U-band during dark time, using the shortest possible exposure times and with no on-chip binning, does the performance of a photoelectric photometer come close to matching the performance of ULTRACAM. In all other situations, ULTRACAM has a much fainter limiting magnitude than a photoelectric photometer, in some cases by as much as 8 magnitudes. Far superior limiting magnitudes over a much wider wavelength range is not the only advantage of ULTRACAM over a photoelectric photometer, of course. Another major advantage is the fact that it is possible to record sky and a comparison star simultaneously with the variable object on ULTRACAM, whereas this is impossible in all but the most complicated photoelectric photometers. Not only does this mean that ULTRACAM is capable of more accurate differential photometry of variable objects, it also means that it is possible to use ULTRACAM in non-photometric conditions, whereas a photoelectric photometer is restricted to use on (much rarer) photometric nights.

We do not, of course, only intend to use ULTRACAM on the WHT. The WHT merely provides an excellent facility, in the form of GHRIL, to commission astronomical instruments such as ULTRACAM. We intend to use ULTRACAM on telescopes of both smaller (e.g. the 2.5-m Isaac Newton Telescope, the 2-m Liverpool Telescope and the 1.9-m Radcliffe Telescope) and larger (e.g. the 9.2-m South African Large Telescope) aperture than the WHT. The figure below shows the limiting magnitudes for a detection with ULTRACAM, at a signal-to-noise ratio of 10, as a function of exposure time on standard Cassegrain 2-m, 4-m and 8-m telescopes in grey time. Each panel corresponds to a different bandpass – U, B, V, R, I and Z. In each panel there are 3 curves. The upper curve is for an 8-m telescope, the middle curve for a 4-m and the lower curve for a 2-m. Note that we have assumed that the ULTRACAM chips are unbinned in this plot. Details of the signal-to-noise calculations used to make these plots and the assumptions which have gone in to them are given in the text at the end of this Appendix.

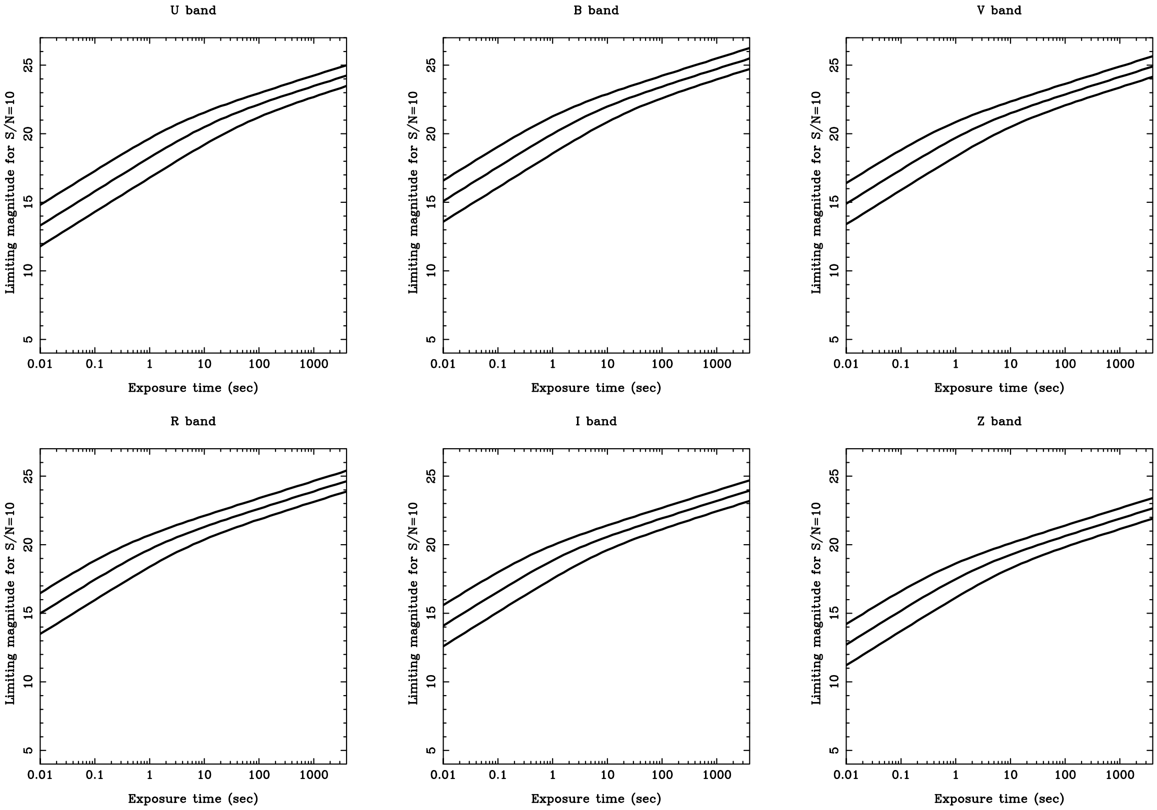

The above figure allows us to make an informed decision on which telescope to use given the scientific requirements of a particular project. For example, when studying black-hole transients, which are typically fainter than 18 in quiescence and a few magnitudes fainter in mid-eclipse, apertures of 4-m or even 8-m will be necessary to obtain time resolutions of a second. On the other hand, studies of most CVs (which are typically 14-16 in quiescence and a few magnitudes fainter in eclipse) on timescales of a second can be suitably performed on 2-m class telescopes. When looking at shorter timescale variations, such as when doing echo mapping experiments which typically require 0.1 s exposures, a 2-m will only suffice for the brightest binaries and a 4-m would in most cases be optimum.

```
DETAILS OF SIGNAL-TO-NOISE CALCULATIONS
=======================================

The number of photons detected by ULTRACAH from a point source of given apparent magnitude were calculated for each filter as follows:

                    10**(-mag/2.5)  * photons/sec/A/cm**2 giving mag = 0 at top of atmosphere
                                    * transmission of atmosphere at zenith
                                    * exposure time in sec
                                    * unobstructed area of primary mirror in cm**2 (assuming 10% central obstruction)
                                    * measured throughput of telescope mirrors
                                    * estimated throughput of ULTRACAH optics
                                    * estimated quantum efficiency of detectors in ULTRACAH
                                    * effective filter bandwidth in A

The number of photons per pixel from the sky detected by ULTRACAH was calculated in a similar manner, but using the apparent magnitude
of the sky in magnitudes/arcsec**2 and then multiplying the result by the platescale**2, where the platescale is given in arcseconds
per pixel.

The signal-to-noise ratio of the object detection was then calculated within an aperture of twice the diameter of the seeing, as
follows:

                                    pixels = pi * (seeing/platescale)**2
                           signal-to-noise = Nobj/sqrt(Nobj+pixels*(Nsky+dark+read**2))

                            where,  pixels = number of pixels in aperture
                                    seeing = seeing in arcseconds
                                    Nobj   = total number of detected photons from object in aperture
                                    Nsky   = number of photons/pixel from sky
                                    dark   = dark current in electrons/pixel/exposure
                                    read   = readout noise in electrons/pixel

Wherever possible we used values for the above parameters (i.e. extinction, sky brightnesses, etc) given by the SIGNAL program
written by Chris Benn for the ING. For all other parameters, we were careful to adopt realistic values so as to accurately
represent the expected performance of ULTRACAH. Specifically:

Parameter                         Assumed value
---------                         -------------

WHT throughput at GHRIL           15% loss at primary, secondary and tertiary
```

	25% loss at the GHRIL derotator Hence, telescope throughput = 0.46
2-m, 4-m and 8-m Cassegrain throughputs	15% loss at primary and secondary Hence, telescope throughput = 0.72
Throughput of ULTRACAM optics	12% loss at collimator (6 surfaces each with 2% loss) 10% loss at dichroic (2 surfaces each with 5% loss) 10% loss at filter 4% loss at re-imaging optics (6 surfaces each with 0.5% losses) 1% loss at cryostat window (2 surfaces each with 0.5% losses) Hence, ULTRACAM throughput = 0.69 Note that the splitting of light into 3 beams by ULTRACAM enables the use of narrow-band coatings on many of the optical surfaces in ULTRACAM. These are far more efficient than the broad-band coatings used in most common-user instrumentation, resulting in the relatively high throughput figures (given the number of optical components) for ULTRACAM.
ULTRACAM detector QE	60% in U, 83% in B, 80% in V, 74% in R, 46% in I and 20% in Z, as measured from the QE curve given for the new EEV10 chip used at the WHT. Note that it is possible that we may be able to coat the three chips with different, wavelength-optimized coatings in order to ensure that each detector has a peak QE corresponding to the arm of ULTRACAM in which it will be used. This would result in QE values of >80% in each pass-band.
Seeing	1 arcsecond. Note that the median seeing on the WHT has recently been measured by the JOSE project to be 0.7 arcseconds. The DIMM monitor at Sutherland has recently measured a median figure of around 0.8 arcseconds (Dave Buckley, SAAO, private communication).
Dark current	2 electrons/pixel/hour. Using thermo-electric cooling, this figure rises to 10\,000-60\,000 electrons/pixel/hour (depending on ambient temperature and whether water cooling is also used), which is why we have adopted liquid Nitrogen cooling for ULTRACAM.
Readout noise	4.5 electrons/pixel. At the highest readout speed of 3.3 micro-seconds/pixel/channel the readout noise would increase slightly to around 6 electrons/pixel. At a slower readout speed of around 10 micro-seconds/pixel/channel the readout noise would decrease slightly to around 3 electrons/pixel.

For the photoelectric photometer calculations a slightly different equation has been used to calculate the signal-to-noise ratios:

signal-to-noise = Nobj/sqrt(Nobj+Nsky+dark)

where, Nobj = total number of detected photons from object in aperture
Nsky = total number of detected photons from sky in aperture
dark = dark current in counts per exposure

We took the readout noise of the photoelectric photometer to be zero, the dark current to be 100 counts per second (as measured on the Peoples' Photometer on La Palma) and the aperture diameter to be 10 arcseconds (the most often used aperture with the Peoples' Photometer). Note that for accurate photometry with photoelectric photometers it is essential to use an aperture much larger than the seeing in order to avoid problems with target centering, guiding and variable seeing. The penalty for this is high sky noise, which does not affect ULTRACAM to such a great extent. The QE of a standard EMI 9658 photomultiplier tube -- as used in the Peoples' Photometer -- was assumed in our calculations -- i.e. 15% in U, 14% in B, 8% in V, 3% in R, 0.5% in I and 0.01% in Z. All other parameters were the same as used for the ULTRACAM calculations.

D ULTRACAM – data acquisition and data reduction

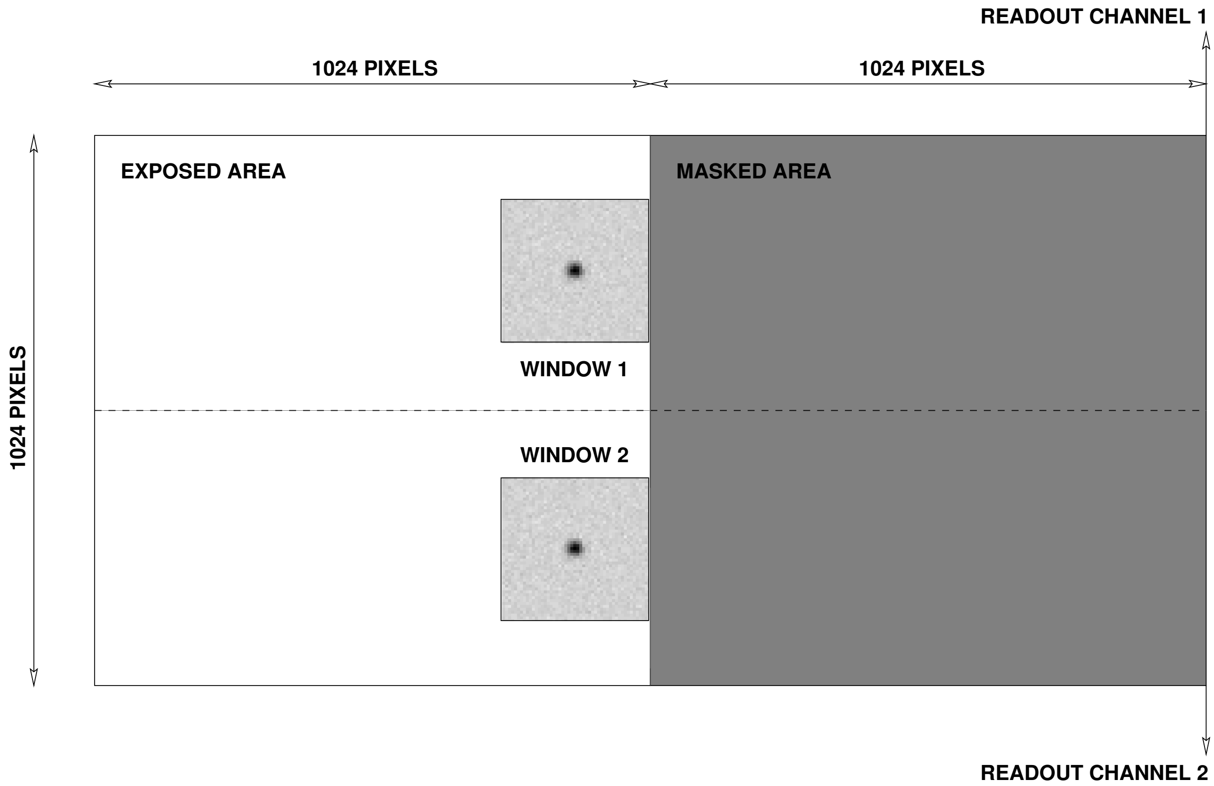

The figure above is a schematic representation of one of the ULTRACAM EEV CCD chips. Each 2048×1024 chip is split into two equal halves, each of 1024×1024 pixels. One of these halves (the left-hand part in the diagram above) is exposed to light, whereas the right-hand half is masked off. Once an exposure time of the specified length is complete, the data in the exposed area is shifted into the masked area. This shifting is very fast and takes approximately 5 μseconds per column, i.e. it takes $1024 \times 0.000005 = 0.005$ seconds for a full frame. Once the data has been transferred into the masked region, the exposed region begins accumulating data again. While this is happening, the data in the masked region is read out by the SDSU controller. This is a much slower process than the simple shifting of data and takes approximately 3.3 μseconds/pixel/readout channel. The SDSU controllers, however, can read out from two different corners of the EEV chip at the same time. Hence, the top half of the masked region in the diagram above is readout via one channel and the bottom half via a second channel. This has the effect of doubling the readout speed to 1.7 μseconds/pixel, resulting in a full frame readout time of $1024 \times 1024 \times 0.0000017 = 1.73$ seconds. This means that if the exposure time is longer than 1.73 seconds, the data in the masked region can be readout out during the next exposure, resulting in a full-frame dead-time of only 0.005 seconds, which is the time it takes to shift the data from the exposed region into the masked region. If the exposure time is shorter than 1.73 seconds, then the dead-time will be equal to 0.005 seconds plus however long it takes to read out the remaining pixels in the masked region, resulting in a maximum full-frame dead-time of only 1.735 seconds (for a 0 second exposure).

For the science projects we wish to perform, however, we will not be reading out the whole chip. Instead, we will want to read out two small windows, one containing the variable object and the other containing a comparison star to use for differential photometry. To minimize the time required to skip to these windows, we will position them as close as possible to the border between the masked and exposed regions (using the Nasmyth derotator or Cassegrain rotator), with one window in the upper half of the chip (and hence going to readout channel 1) and the other window in the lower half of the chip (and hence going to readout channel 2), as depicted in the above diagram. Given that each window will have to be used for sky estimation as well, we expect to use window sizes of around 80×80 pixels (i.e. $20'' \times 20''$). In order to cut down the readout noise, however, we will usually on-chip bin by a factor of 4, resulting in only 20×20 pixels being read out. These will first be shifted into the masked area, in a time of $20 \times 0.000005 = 0.0001$ seconds. The next exposure can then begin. In the meantime, a window

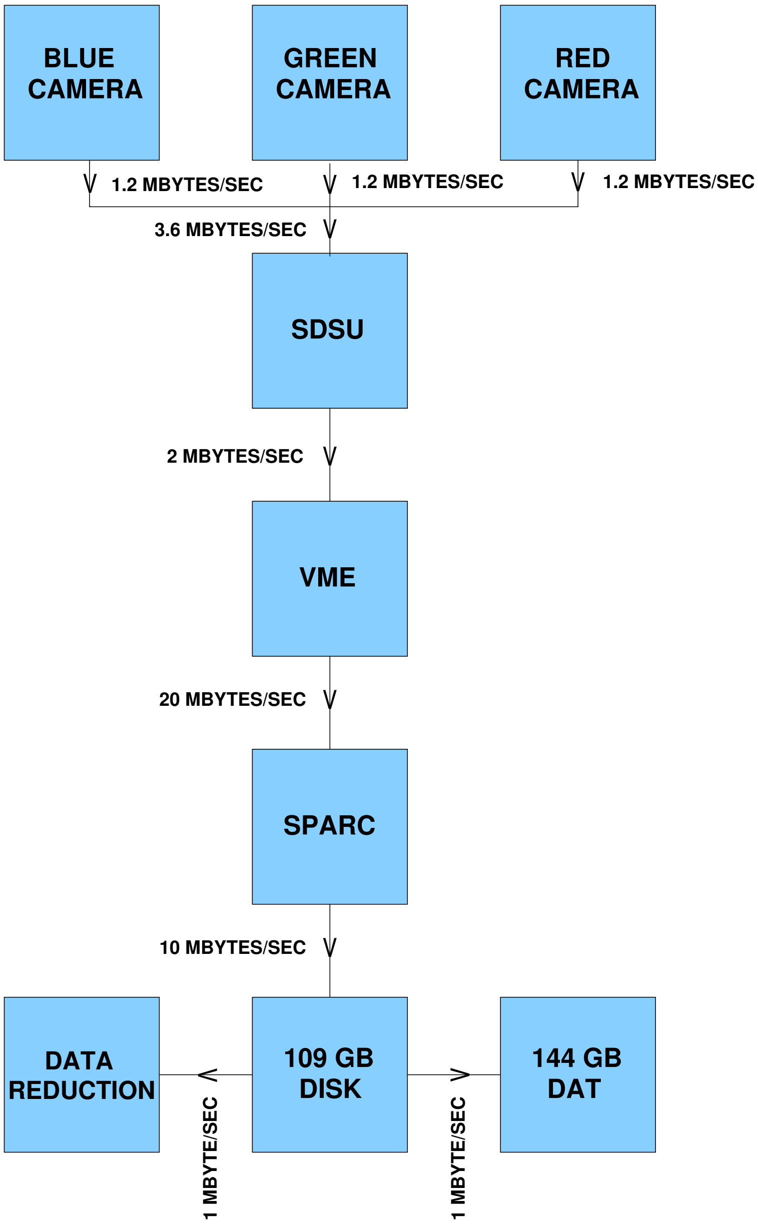

at the readout-register end of the masked region is read out, in a time of $20 \times 20 \times 0.0000033 = 0.0013$ seconds, which then makes room in the masked region for the next window, which is currently being exposed. As with the full-frame example above, if the exposure time is longer than 0.0013 seconds, it is possible to read out a window before the next exposure ends, and so the dead-time is then only the time it takes to shift the window into the masked region (0.0001 seconds). In the unlikely event that the exposure time is shorter than 0.0013 seconds, the dead-time will be equal to 0.0001 seconds plus however long it takes to read out enough pixels to make room for the window in the masked region, resulting in a maximum 20×20-window dead-time of only 0.0014 seconds (for a 0 second exposure). It should be noted here that smearing due to the finite time it takes to shift the data into the masked region is always negligible, unless the exposure time is so short that the shifting (which takes only 0.0001 seconds in the above example) becomes a significant fraction of the exposure time.

In the above examples, we are assuming that the dead-times in ULTRACAM are set solely by the pixel readout rate of the EEV chips, which is equal to 3.3 μseconds/pixel/readout channel. This is not quite true. Such a pixel rate implies that each camera will output 300 000 pixels/second/readout channel, which is equal to 1.2 Mbytes/second. The SDSU controller therefore has to handle 3.6 Mbytes/second from the three ULTRACAM cameras. In its current form, however, the SDSU controller can only handle approximately 2 Mbytes/second, although the world-wide demand for this controller means that improvements are rapid and it is hence feasible that by the time we come to buy a controller the factor of 2 increase in data rate that we require will be available. Assuming for now that our maximum data rate is 2 Mbytes/second, this effectively lowers the pixel readout rate of the EEV chips from 3.3 μseconds/pixel/readout channel to around 6 μseconds/pixel/readout channel, in which case the dead-times calculated in the paragraphs above must be suitably modified. It should be emphasized, however, that the dead-times in ULTRACAM are basically zero if the exposure time is longer than it takes to read out the data, which will almost always be the case for the key science projects we list in Section 2.

The diagram on the preceding page summarizes the data rates flowing out of each component of the ULTRACAM data acquisition system. Its purpose is to show where data-flow bottlenecks occur and how our hardware choices have minimized these. The three ULTRACAM CCDs can produce data at a maximum rate of 3.6 Mbytes/second. As discussed above, the SDSU controller can only handle a maximum of 2 Mbytes/second. Thereafter, the data is transferred to the VME system, which in turn transfers the data to the Sparcstation. These latter transfers are all over dedicated ethernet links and are capable of handling data rates in excess of 20 Mbytes/second. The Sparcstation then writes the data to the 109 Gbyte hard disk at a rate of 10 Mbytes/second. This shows that the only data-flow bottleneck occurs at the SDSU controller and hence its maximum data rate sets the maximum data rate of the whole ULTRACAM data acquisition system – 2 Mbytes/second.

Once the data is on the hard disk it can then be archived onto DAT at a maximum rate of 1 MByte/second, i.e. it will take twice as long to archive the data to DAT as it will to take the data with ULTRACAM when running at its maximum rate. This means that we will have to archive day and night to keep up with the maximum ULTRACAM data rate. This should be no problem, of course, as the DAT drive can be automatically loaded with up to 6 DDS3 tapes. In addition, it is unlikely that we will observe for a whole night with ULTRACAM running at the maximum data rate (i.e. there will be overheads due to target acquisition, calibration frames, etc).

We have emphasized in the technical justification section the importance of pipeline data reduction for ULTRACAM. In order to assess how feasible this is, we wrote two simulation programs. One program created fake (but realistic) ULTRACAM data frames – two examples of which are shown in the diagram on page 12. The other program took these fake binary data files, debiased them, flat-fielded them, searched for the star, summed the counts in an aperture centred on the star and then subtracted the sky (determined from an annulus surrounding the object aperture). The final object counts were then written to an ASCII file and checked against the (known) object counts in the fake data to ensure the photometry was correct. We set this aperture photometry program running on a set of 100 fake ULTRACAM images, each of 100×100 pixels. On a Sparcstation ULTRA 5 it took 3.6 seconds of elapsed time (with 100% of the CPU time) to reduce the 100 frames. This means that the ULTRA 5 can reduce data at a rate of $100 \times 100 \times 100 \times 2 = 2$ Mbytes every 3.6 seconds, i.e. 0.6 Mbytes/second. We learnt from Sun, however, that the new ULTRA 10 333 MHz machine is 75% faster than the ULTRA 5 we have at Sheffield, resulting in a data rate of 1 MByte/second. In conjunction with a more carefully written data reduction routine, we expect to be able to reduce data at a speed approaching 2 Mbytes/second. Also, it is unlikely that we would spend a whole night observing at the very highest data rate, as there will be inevitable overheads to do with object acquisition, calibration frames, or simply longer exposure times. The simulations we have performed therefore make us confident that it is feasible to deliver pipeline data reduction software with ULTRACAM which can produce real-time light-curves.

In summary, we have shown in this Appendix that:

- For most science projects, the dead-time of ULTRACAM is essentially zero, resulting in frame rates given by the reciprocal of the exposure time.
- The maximum date rate of ULTRACAM is 2 Mbytes/second.
- It will be feasible to archive data taken at the maximum data rate to disk and tape in real time.
- It will be feasible to reduce data taken at the maximum data rate in real time.

E ULTRACAM – letters of support from observatory directors

We have designed ULTRACAM as a standalone, versatile and portable instrument which can be used on almost any telescope. We have, however, targetted a set of 5 telescopes at which we initially intend to use ULTRACAM. These telescopes will give us the range of aperture in each hemisphere that we require to complete the key science projects listed in Section 2. For each of these 5 telescopes, we have obtained a letter of support from the relevant observatory director or project scientist stating that they are keen for their facility to support ULTRACAM as a visiting instrument. The 5 telescopes are:

1. The 9.2-m South African Large Telescope at Sutherland, South Africa. Please refer to the attached letter from Dave Buckley, SALT Project Scientist.
2. The 4.2-m William Herschel Telescope on La Palma. Please refer to the attached letter from René Rutten, Director ING.
3. The 2.5-m Isaac Newton Telescope on La Palma. Please refer to the attached letter from René Rutten, Director ING.
4. The 2.0-m Liverpool Telescope on La Palma. Please refer to the attached letter from Dave Carter, Project Scientist for the Liverpool Telescope.
5. The 1.9-m Radcliffe Telescope at Sutherland, South Africa. Please refer to the attached letter from Bob Stobie, Director SAAO.

Following discussions with René Rutten, Director ING, we have selected the GHRIL focus on the WHT from the above list of 5 telescopes to commission ULTRACAM. There are a number of reasons for this decision:

- GHRIL is isolated from the rest of the telescope and hence we will be able to set up ULTRACAM many days before the commissioning run. If we commission at a Cassegrain focus, this will not be possible as the focus will probably be in use until the morning of our run, giving us much less time to set up.
- GHRIL contains a large optical bench to which it will be a simple matter to mount ULTRACAM. If we commission at a Cassegrain focus, we will have to contend with the problem of mounting the instrument and its electronics to the rotator in addition to the other (inevitable) problems we will face during this period.
- GHRIL provides an ideal working environment. For example, the telescope can be controlled from within GHRIL and the optical bench is at a comfortable height, providing easy access to the instrument for adjustments whilst observing. If the instrument is mounted at Cassegrain, adjustments will be more time-consuming and difficult to make.

F ULTRACAM and the UKATC

This is the second time we have submitted ULTRACAM to RAP for consideration. The RAP did not award ULTRACAM a grade during its last meeting, but instead recommended we resubmit the proposal with some additional information, most of which we have provided in this resubmitted version. In fact, the only item of additional information which we have been unable to provide is a letter from Adrian Russell, Director UKATC, which commits the UKATC to provide its contribution for the cost and in the time specified in this proposal. Note that the costs and times given in this proposal do actually originate from the UKATC, and are in all likelihood deliverable, but have not been subject to the scrutiny they would give such figures if they were the subject of a formal contract.

We have had detailed discussions with Adrian Russell on this point. He is understandably unwilling to provide such a letter, which amounts to a contract between the applicants/PPARC and the UKATC. He will only write such a letter once we have a much more detailed design and plan, which will involve substantial additional effort by us and his staff, including a number of visits to Edinburgh. Given the current workload of his staff, it is unrealistic to expect this to occur before the summer. If we had waited for this letter, then, we would have been forced to resubmit on October 1st and we (and the UKATC) would have put substantial additional effort into a proposal which might then be rejected for a completely unrelated reason (e.g. quality of science).

We have therefore decided to resubmit this proposal on March 1st, without a letter of support from Adrian Russell. We propose that the RAP assess our proposal this round on the assumption that the costs and timescales given are deliverable by the UKATC and then award (or reject) the proposal accordingly. If the RAP decide to award the grant, then it would be provisional pending the receipt of a signed letter from Adrian Russell. We would then go about obtaining this letter by formulating a more detailed design and plan with the UKATC staff over the summer. The advantages of this route for us (and the UKATC) are that:

- We and the UKATC will not then be wasting our time formulating a more detailed design.
- We will be able to start building ULTRACAM up to a year earlier.

The RAP may wish to consider whether the above scenario is legitimate and then process our application accordingly.

G Publications in refereed journals since 1996

1. *The eclipsing dwarf nova HS1804+6753*
   Billington I, **Marsh T R**, **Dhillon V S**, 1996, MNRAS, 278, 673

2. *An early-time infrared and optical study of the Type Ia supernovae SN 1994D and 1991T*
   Meikle W P S, **Dhillon V S** co-author et al, 1996, MNRAS, 281, 263

3. *The EUV transient RE J1255+266*
   Watson M G, **Marsh T R**, **Dhillon V S** co-author et al, 1996, MNRAS, 281, 1016

4. *V795 Her: an SW Sex star in the period gap?*
   Casares J, Martinez-Pais I, **Marsh T R**, Charles P A, Lazaro C, 1996, MNRAS, 278, 219

5. *A detached white dwarf/M dwarf binary with an orbital period of* 2.47 *h*
   Marsh T R, Duck S R, 1996, MNRAS, 278, 565

6. *OY Car during normal outburst: Balmer emission from the red star and the gas stream*
   Harlaftis E T, **Marsh T R**, 1996, A&A, 308, 97

7. *Superhumps and ultraviolet superdips: HST observations of OY Car.*
   Billington I, **Marsh T R** et al, 1996, MNRAS, 279, 1274

8. *Slingshot Prominences during Dwarf Nova Outbursts?*
   Steeghs D, Horne K, **Marsh T R**, Donati J F, 1996, MNRAS, 281, 626

9. *Stroboscopic Doppler Tomography of FO Aqr*
   Marsh T R, Duck S R, 1996, New Ast, 1, 97

10. *Infrared spectroscopy of cataclysmic variables – II. Intermediate polars*
    Dhillon V S, **Marsh T R**, Duck S R, Rosen S R, 1997, MNRAS, 285, 95

11. *Infrared spectroscopy of the long-period polar RX J0515.6+0105*
    Harrop-Allin M K, Cropper M, Potter S B, **Dhillon V S**, Howell S B, 1997, MNRAS, 288, 1033

12. *Infrared and optical spectroscopy of Type Ia supernovae in the nebular phase*
    Bowers E J C, **Dhillon V S** co-author et al, 1997, MNRAS, 290, 663

13. *On the nature of SW Sex*
    Dhillon V S, **Marsh T R**, Jones D H P, 1997, MNRAS, 291, 694

14. *The changing face of Betelgeuse*
    Wilson R W, **Dhillon V S**, Haniff C A, 1997, MNRAS, 291, 819

15. *Supersonic line broadening in accretion discs*
    Marsh T R, **Dhillon V S**, 1997, MNRAS, 292, 385

16. *A detached double degenerate with a* 1.4 *hr orbital period*
    Moran C, **Marsh T R**, Bragaglia A, 1997, MNRAS, 288, 538

17. *A radial velocity study of the giant companion to the central star of Abell 35*
    Gatti A A, Drew J E, Lumsden S, **Marsh T R**, Moran C, Stetson P, 1997, MNRAS, 291, 773

18. *Time-resolved spectroscopy of AL Comae Berenices*
    Howell S B, Hauschildt P, **Dhillon V S**, 1998, ApJ, 494, L223

19. *The mass of the white dwarf in the old nova BT Mon*
    Smith D A, **Dhillon V S**, **Marsh T R**, 1998, MNRAS, 296, 465

20. *The secondary star in cataclysmic variables*
    Smith D A, **Dhillon V S**, 1998, 301, 767

21. *Observations of the polar ST Leonis Minoris during an extreme low state: Identification of the secondary star*
    Ciardi D R, Howell S B, **Dhillon V S**, Wagner R M, Hauschildt P H, Allard F, 1998, PASP, 110, 1007

22. *Projected rotational velocities of WD1614+136 and WD1353+409 – implications for the rate of galactic Type Ia supernovae*
    Maxted P F L, **Marsh T R**, 1998, MNRAS, 296, L34

23. *A new magnetic white dwarf: PG2329+267*
    Moran C, **Marsh T R**, **Dhillon V S**, 1998, MNRAS, 299, 218

24. *Spectroscopic evidence for two-pole accretion in the intermediate polar RXJ0558+5353*
    Still M D, Duck S R, **Marsh T R**, 1998, MNRAS, 299, 759

25. *A double-frequency dwarf nova oscillation in OY Car*
    Marsh T R, Horne K, 1998, MNRAS, 299, 921

26. *The mass and radius of the M-dwarf companion to GD448*
    Maxted P F L, **Marsh T R**, Moran C, **Dhillon V S**, Hilditch R W, 1998, MNRAS, 300, 1225

27. *The separation of the stars in the binary nucleus of the planetary nebula Abell 35*
    Gatti A A, Drew J E, Oudmaijer R D, **Marsh T R**, Lynas-Gray A E, 1998, MNRAS, 301, L33

28. *The secondary star in ST Leonis Minoris: Implications for evolutionary effects in cataclysmic variables*
    Howell S B, **Dhillon V S**, Ciardi D R, Skidmore W, 1999, ApJ, in press

29. *The orbital period and Doppler maps of the dwarf nova HL Canis Majoris*
    Still M D, Steeghs D, **Dhillon V S**, Buckley D A H, 1999, MNRAS, in press

30. *The orbital parameters of three new subdwarf B binaries*
    Moran C, Maxted P F L, **Marsh T R**, Saffer R, Livio M, 1999, MNRAS, in press

31. *Kinematics of the helium accretor GP Com*
    Marsh T R, 1999, MNRAS, in press

32. *The fraction of double-degenerates among DA white dwarfs*
    Maxted P F L, **Marsh T R**, 1999, MNRAS, in press