

ULTRACAM: an ultrafast, triple-beam CCD camera for high-speed astrophysics

V. S. Dhillon,^{1★} T. R. Marsh,^{2★} M. J. Stevenson,¹ D. C. Atkinson,³ P. Kerry,¹
 P. T. Peacocke,³ A. J. A. Vick,³ S. M. Beard,³ D. J. Ives,³ D. W. Lunney,³
 S. A. McLay,³ C. J. Tierney,³ J. Kelly,¹ S. P. Littlefair,¹ R. Nicholson,¹
 R. Pashley,¹ E. T. Harlaftis⁴ and K. O’Brien⁵

¹*Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH*

²*Department of Physics, University of Warwick, Coventry CV4 7AL*

³*UK Astronomy Technology Centre, Royal Observatory Edinburgh, Blackford Hill, Edinburgh EH9 3HJ*

⁴*Institute of Space Applications and Remote Sensing, National Observatory of Athens, I. Metaxa and B. Paulou, Lofos Koufou, Palaia Penteli, Athens 152 36, Greece*

⁵*European Southern Observatory, Alonso de Cordova 3107, Vitacura, Casilla 19001, Santiago 19, Chile*

Accepted 2007 April 16. Received 2007 April 16; in original form 2007 March 22

ABSTRACT

ULTRACAM is a portable, high-speed imaging photometer designed to study faint astronomical objects at high temporal resolutions. ULTRACAM employs two dichroic beamsplitters and three frame-transfer CCD cameras to provide three-colour optical imaging at frame rates of up to 500 Hz. The instrument has been mounted on both the 4.2-m William Herschel Telescope on La Palma and the 8.2-m Very Large Telescope in Chile, and has been used to study white dwarfs, brown dwarfs, pulsars, black hole/neutron star X-ray binaries, gamma-ray bursts, cataclysmic variables, eclipsing binary stars, extrasolar planets, flare stars, ultracompact binaries, active galactic nuclei, asteroseismology and occultations by Solar System objects (Titan, Pluto and Kuiper Belt objects). In this paper we describe the scientific motivation behind ULTRACAM, present an outline of its design and report on its measured performance.

Key words: instrumentation: detectors – instrumentation: photometers – techniques: photometric.

1 INTRODUCTION

CCDs revolutionized astronomy when they became available in the 1970s. CCDs are linear, stable, robust and low-power devices. They have large formats, small pixels and excellent sensitivity over a wide range of wavelengths and light levels; in fact, they are almost perfect detectors, suffering only from poor time resolution and readout noise compared to the photon-counting detectors that they replaced. These limitations of CCDs are inherent to their architecture, in which the photogenerated electrons must first be extracted (or *clocked*) from the detection site and then digitized.¹

There are ways in which the readout noise of CCDs can be eliminated – for example, by the use of electron-multiplying devices

(Mackay et al. 2001). There are also ways in which CCDs can be made to read out faster. First, the clocking rate can be increased and the digitization time decreased, but only if the CCD is of sufficient quality to allow these adjustments to be made without an unacceptable increase in charge-transfer inefficiency and readout noise. Secondly, the duty cycle of the exposures can be increased; for example, *frame-transfer* CCDs provide a storage area into which photogenerated charge can be clocked. This charge is then digitized whilst the next exposure is taking place and, because digitization generally takes much longer than clocking, the dead time between exposures is significantly reduced. Thirdly, the data acquisition hardware and software can be designed in such a way that the rate at which it is able to archive the data (e.g. to a hard disc) is always greater than the rate at which data are digitized by the CCD – a situation we shall refer to as *detector-limited* operation.

We have employed all three of the techniques described above to harness the greater sensitivity and versatility of CCDs (compared to photon-counting detectors) for high-speed optical photometry. The resulting instrument, known as ULTRACAM (for ULTRA-fast CAMera), was commissioned on the 4.2-m William Herschel

*E-mail: vik.dhillon@shef.ac.uk (VSD); t.marsh@warwick.ac.uk (TRM)

¹ For the purposes of this paper, the word *digitization* shall refer to both the process of determining the charge content of a pixel via correlated double sampling and the subsequent digitization of the charge using an analogue-to-digital converter.

Telescope (WHT) on La Palma on 2002 May 16, and on the 8.2-m Very Large Telescope (VLT) in Chile on 2005 May 4. To date, no detailed description of the instrument has appeared in the referred astronomical literature. In this paper, therefore, we describe the scientific motivation behind ULTRACAM, present an outline of its design and report on its measured performance.

2 SCIENTIFIC MOTIVATION

It is widely recognized that temporal resolution is a relatively unexplored region of observational parameter space (see e.g. Dravins 1994), particularly in the optical part of the spectrum. There is a simple reason for this: CCDs are the dominant detector on all of the largest ground-based telescopes and the resources of the observatories have understandably been targeted towards increasing the area and sensitivity of their detectors rather than prioritizing high-speed readout. As a result, it has been difficult for astronomers to achieve frame rates of higher than ~ 1 frame per minute using CCDs on large optical telescopes, although there are a few notable exceptions (see Table 1 for details).

Aside from the serendipitous value of exploring a new region of observational parameter space, it is the study of compact objects such as white dwarfs, neutron stars and stellar mass black holes which benefits the most from high-speed observations. This is because the dynamical time-scales of compact objects range from seconds in white dwarfs to milliseconds in neutron stars and black holes, which means that the rotation and pulsation of these objects or material in close proximity to them (e.g. in an accretion disc) tends to occur on time-scales of milliseconds to seconds. Other areas of astrophysics which benefit from observations obtained on these time-scales are studies of eclipses, transits and occultations, where increased time resolution can be used to (indirectly) give an increased spatial resolution.

Studying the above science requires a photometer with the following capabilities.

(i) *Short exposure times (from milliseconds to seconds)*. There is little point in going much faster than milliseconds, as the gravity of compact objects does not allow for bulk motions on time-scales below this. There are exceptions to this rule, however; for example, microsecond time resolution would allow the study of magnetic

instabilities in accreting systems, and nanosecond time resolution would open up the field of quantum optics (see Dravins 1994).

(ii) *Negligible dead time between exposures*. It takes a finite amount of time to archive a data frame and start the next exposure. To preserve time resolution, it is essential that this dead time is a small fraction of the exposure time.

(iii) *Multichannels (three or more) covering a wide wavelength range*. Ideally, one would obtain spectra at high temporal resolution in order to fully characterize the source of variability (e.g. to determine its temperature). The faintness of most compact objects precludes spectroscopy, unfortunately. At a minimum, therefore, at least three different pass-bands covering as wide a portion of the optical spectrum as possible need to be observed, as this would allow a blackbody spectrum to be distinguished from a stellar spectrum. It is particularly important that one of the three channels is sensitive to the far blue (i.e. between approximately 3000–4000 Å), as the flickering and oscillations observed in many accreting binaries is much more prominent at these wavelengths (see e.g. Marsh & Horne 1998).

(iv) *Simultaneous measurement of the different wavelength bands*. The requirement for multichannel photometry must be tied to a requirement that each channel is observed simultaneously. A single-channel instrument with a filter wheel, for example, could obtain data in more than one colour by changing filters between each exposure or obtaining a full cycle in one filter and then a second cycle in another filter. The first technique results in poor time resolution and both the first and second techniques are observationally inefficient (relative to an instrument which can record multichannel data simultaneously). In addition, the non-simultaneity of both techniques makes them unsuitable for the study of colour variations which occur on time-scales shorter than the duration of observations in a single filter.

(v) *Imaging capability*. A photometer with an imaging capability is essential to observe variable sources. This is because, unless the atmospheric conditions are perfectly photometric, it is necessary to simultaneously measure the target, comparison stars and sky background so that any variability observed can be unambiguously assigned to the correct source. Although non-imaging photometers based on multiple photomultiplier tubes can work around this, they still suffer from their dependency on a fixed aperture which is larger

Table 1. Instruments/modes for high-speed astrophysics with CCDs – this is not intended to be an exhaustive list, but an indication of the variety of ways in which CCDs can be used for high-speed optical observations. Those listed with a dagger (\dagger) symbol are not continuous modes and data taking has to be periodically stopped for archiving. The frame rates listed are estimates of the maximum values possible with each mode. Abbreviations: F-T – frame-transfer; MCP – microchannel plate; EMCCD – electron-multiplying CCD.

Instrument/mode	Telescope	Detector	Frame rate (Hz)	Reference
Low-smear drift mode \dagger	WHT	CCD	4	Rutten et al. (1997)
Time-series mode \dagger	AAT	CCD	100	Stathakis & Johnston (2002)
Phase-binning mode	Hale	CCD	1000	Kern (2002), Kern & Martin (2002)
Stroboscopic mode	GHO	CCD	7.5	Čadež et al. (2001), Kotar, Vidrih & Čadež (2003)
Freerun mode	Mayall	CCD + MCP	24 000	Fordham et al. (2000, 2002)
Continuous-clocking mode	Keck-II	CCD	14	O'Brien et al. (2001); Skidmore et al. (2003)
High time resolution mode \dagger	VLT	CCD	833	Cumani & Mantel (2001), O'Brien (2007)
ULTRACAM	WHT/VLT	CCD	500	This paper
Frame-storage mode	Lick 1-m	F-T CCD	0.1	Stover & Allen (1987)
UCT photometer	SAAO	F-T CCD	1	O'Donoghue (1995), Woudt & Warner (2001)
Acquisition camera	Gemini South	F-T CCD	7	Hynes et al. (2003)
JOSE camera \dagger	NOT	F-T CCD	150	St-Jacques et al. (1997), Baldwin et al. (2001)
SALTICAM	SALT	F-T CCD	10	O'Donoghue et al. (2006)
LuckyCam	NOT	EMCCD	30	Law, Hodgkin & Mackay (2006)

than the seeing disc and which therefore increases the sky noise substantially. This is not a problem with imaging photometers, in which the signal-to-noise ratio of the extracted object counts can be maximized using optimal photometry techniques (e.g. Naylor 1998).

(vi) *High efficiency and portability.* General studies of variability, that is, the study of any object which eclipses, transits, occults, flickers, flares, pulsates, oscillates, erupts, outbursts or explodes, involves observing objects which span a huge range in brightness. The faintest (e.g. the pulsars) are extremely dim and observing them at high temporal resolution is a photon-starved application on even the largest aperture telescopes currently available. Others (e.g. cataclysmic variables) are relatively bright and can be observed at sub-second time resolution at excellent signal-to-noise ratio on only a 2-m class telescope. Any instrument designed specifically for high-speed observations must therefore be portable, so that an appropriate aperture telescope can be selected for the project at hand, and highly efficient at recording incident photons in order to combat the short exposure times required.

3 DESIGN

With the exception of ULTRACAM, none of the instruments listed in Table 1 meet all six of the scientific requirements given above. In this section, we present an outline of ULTRACAM's design.

3.1 Optics

The starting point for the ULTRACAM optical design was the required field of view. The field of view has to be large enough to give a significant probability of finding a comparison star of comparable brightness to our brightest targets. The probability of finding a comparison star of a given magnitude depends on the search radius and the galactic latitude of the star. Using the star counts listed by Simons (1995), the probability of finding a star of magnitude $R = 12$ at a galactic latitude of 30° (the all-sky average) is 80 per cent if the search radius is 5 arcmin.² Most of the target types discussed in Section 2 are significantly fainter than $R = 12$, so a field of view of 5 arcmin virtually guarantees the presence of a suitable comparison star.³

The next consideration in the ULTRACAM optical design was the required pixel scale. Arguably the best compromise in terms of maximizing spatial resolution whilst minimizing the contributions of readout noise and intrapixel quantum efficiency variations (e.g. Jordan, Deltorn & Oates 1993) is to use a pixel scale which optimally samples the median seeing at the telescope. On the WHT, the median seeing ranges from 0.55 to 0.73 arcsec depending on the method of measurement (Wilson et al. 1999), so a pixel scale of 0.3 arcsec provides a good compromise. The E2V 47-20 CCDs used in ULTRACAM (see Section 3.2) were selected because their 1024 pixels on a side delivers a field of view of 5 arcmin at a scale of 0.3 arcsec pixel⁻¹, matching the requirements listed above perfectly.

There are a total of 30 optical elements in ULTRACAM, all of which have a broad-band antireflection coating with transmission

shown in Fig. 1. The maximum number of optical elements encountered by a photon passing through ULTRACAM is 14 (see Fig. 2). Light from the telescope is first collimated and then split into three beams by two dichroic beamsplitters. Each beam is then re-imaged by a camera on to a detector, passing through a filter and CCD window along the way. Each of these components is discussed in turn below.

3.1.1 Collimator

The collimator feeds collimated light to the re-imaging cameras, partially corrects for the aberrations in the telescope it is to be mounted on, and produces an image of the telescope pupil sufficiently remote to accommodate the dichroics between the last collimator element and the cameras. As the first optical element, the collimator also has to have a high transmission in the scientifically important $\sim 3000\text{--}4000 \text{ \AA}$ region (see Section 2) without introducing significant chromatic aberration. The chosen glasses were N-PSK3, CaF₂ and LLF1.

The collimator shown in Fig. 2 is used when mounting ULTRACAM on the $f/11$ Cassegrain focus of the WHT. It can be replaced with another collimator designed to accommodate the different optical characteristics of other telescopes. We already have a second collimator for use with the $f/8$ Cassegrain focus of the 2.3-m Aristarchos Telescope in Greece, and a third for use at the $f/15$ Nasmyth focus of the VLT. Each of these collimators have identical mounting plates on their lens barrels, so it is a simple matter to switch between them on the opto-mechanical chassis (Section 3.3). Note that the rest of the ULTRACAM optics remain unchanged when moving between different telescopes.

3.1.2 Dichroics

Two dichroic beamsplitters divide the light from the collimator into three different beams, which shall hereafter be referred to as the 'red', 'green' and 'blue' channels. Each dichroic consists of a UV-grade fused silica substrate with a long-wave pass (LWP) coating that reflects incident light with wavelengths shorter than the cut-point whilst transmitting longer wavelengths. Such LWP dichroics generally have a higher throughput than short-wave pass dichroics and their throughput has been further enhanced by coating the back surface of each dichroic with a broad-band antireflection coating. The ULTRACAM dichroics were manufactured by CVI Technical Optics, Isle of Man.

To keep the cost of the custom optics as low as possible, whilst still retaining maximum throughput and image quality, it was decided to dedicate the blue channel of ULTRACAM to the scientifically essential Sloan Digital Sky Survey (SDSS) u' band. As a result, the 50 per cent cut-point of the first dichroic was set to 3870 \AA (see Fig. 1). A similar argument forced us to dedicate the green channel of ULTRACAM to the SDSS g' band, giving a 50 per cent cut-point for the second dichroic of 5580 \AA . In both cases, the cut-points are measured using incident light which is randomly polarized.

It should be noted that the ULTRACAM dichroics operate in a collimated beam. This has the great advantage that ghosts produced by reflections off the back surface of the dichroics are in focus (but slightly aberrated) and fall more-or-less on top of the primary image. This means that the light in the ghosts, which is typically less than 0.001 per cent of the primary image, is included in the light of the

² An online comparison star probability calculator can be found at <http://www.shef.ac.uk/physics/people/vdhillon/ultracam>.

³ The study of transiting extrasolar planets is an exception to this rule. The brightest of these have magnitudes of $R \sim 8$ and require fields of view of approximately 30 arcmin to provide an 80 per cent chance of finding a comparison star of similar brightness.

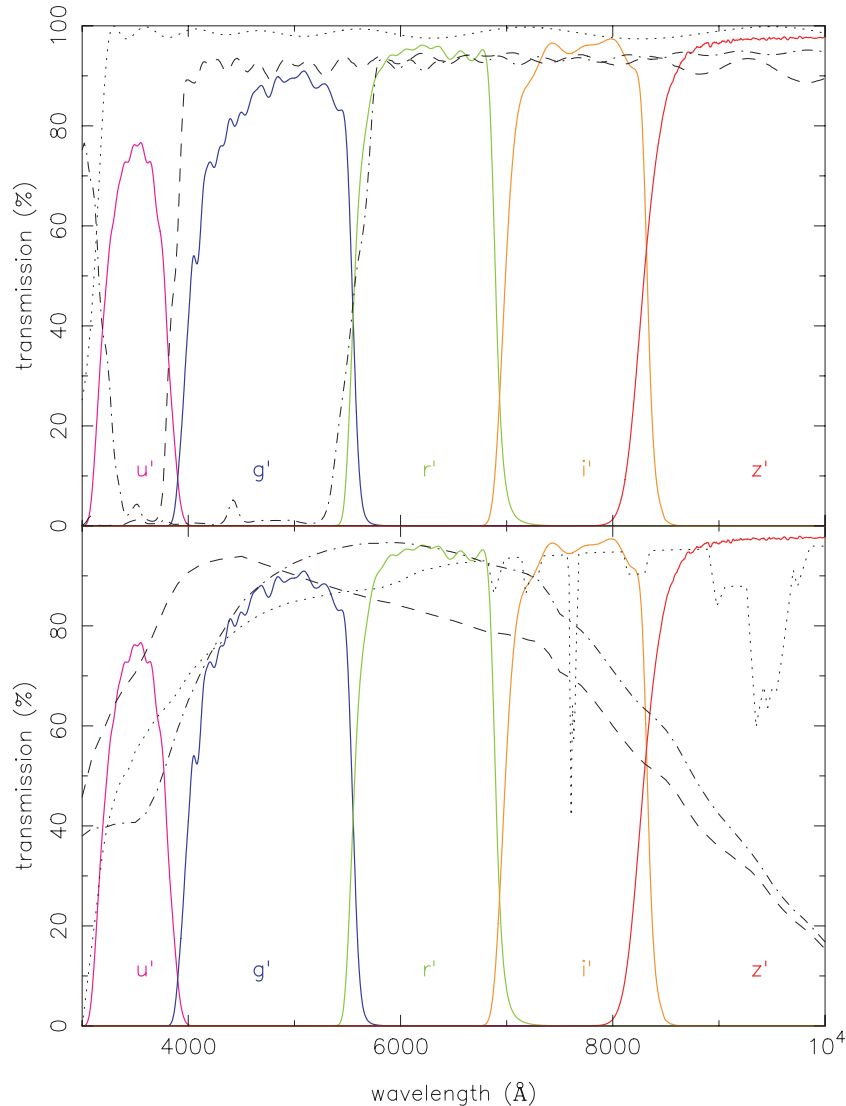


Figure 1. Top: Transmission profiles of the ULTRACAM SDSS filter set (solid lines), the antireflection coating used on the ULTRACAM lenses (dotted line), and the two dichroics (dashed line and dash-dotted line). Bottom: Transmission profiles of the ULTRACAM SDSS filter set (solid lines) and the atmosphere for unit airmass (dotted line). Also shown are the quantum efficiency curves of the blue and green CCDs (dashed line), and the red CCD (dash-dotted line).

target when performing photometry, thereby maximizing the signal-to-noise ratio of the observations and removing a potential source of systematic error.

3.1.3 Re-imaging cameras

The three re-imaging cameras are based upon a four-element double-Gauss type lens, with two glass types per camera to minimize costs (N-SK16 and LLF1 in the blue camera, N-LAK10 and SF2 in the green camera, N-LAK22 and N-SF1 in the red camera). The glass thickness has also been kept to a minimum, particularly in the blue arm, to keep the throughput as high as possible. The cameras have been optimized for an infinite object distance, with a field angle dictated by the required field size. The re-imaged pupil is located between the first and second pairs of elements (shown by the horizontal/vertical bars near the centre of each camera in Fig. 2), where the camera stop would be positioned. The back focal length of the cameras is sufficient to allow space for the filter and CCD window.

The cameras and collimators used in ULTRACAM were manufactured by Specac Ltd., UK.

3.1.4 Filters and windows

The SDSS photometric system (Fukugita et al. 1996) was adopted as the primary filter set for ULTRACAM. This system is becoming increasingly prevalent in astrophysics and consists of five colour bands (u' , g' , r' , i' and z') that divide the entire range from the atmospheric cut-off at ~ 3000 Å to the sensitivity limit of CCDs at $\sim 11\,000$ Å into five essentially non-overlapping pass-bands. The fact that SDSS filters show only negligible overlap between their pass-bands is one of the main reasons they have been adopted for use in ULTRACAM, as the dichroic cut-points then have only a negligible effect on the shape of the filter responses. The alternative Johnson–Morgan–Cousins system (U , B , V , R_C , I_C) suffers from overlapping pass-bands which would be substantially altered when used in conjunction with the ULTRACAM dichroics. The thinned CCDs used




Figure 2. Ray-trace through the ULTRACAM optics, showing the major optical components: the collimator, dichroics, cameras, filters and detector windows. The diagram is to scale – the largest lens is in the collimator and has a diameter of 120 mm.

in ULTRACAM would also suffer from fringing when used with R_c due to the extended red tail of this filter’s bandpass, but this is eliminated by the sharp red cut-off in SDSS r' . The SDSS filters used in ULTRACAM were procured from Asahi Spectra Ltd., Tokyo.

As well as SDSS filters and clear filters (for maximum throughput in each channel), ULTRACAM also has a growing set of narrow-band filters, including C III/N III + He II (central wavelength, $\lambda_c = 4662 \text{ \AA}$; FWHM, $\Delta\lambda = 108 \text{ \AA}$), Na I ($\lambda_c = 5912 \text{ \AA}$, $\Delta\lambda = 312 \text{ \AA}$), H α ($\lambda_c = 6560 \text{ \AA}$, $\Delta\lambda = 100 \text{ \AA}$), red continuum ($\lambda_c = 6010 \text{ \AA}$, $\Delta\lambda = 118 \text{ \AA}$) and blue continuum ($\lambda_c = 5150 \text{ \AA}$, $\Delta\lambda = 200 \text{ \AA}$). All ULTRACAM filters are $50 \times 50 \text{ mm}^2$ and approximately 5 mm thick, but have been designed to have identical optical thicknesses so that their differing refractive indices are compensated by slightly different thicknesses, making the filters interchangeable without having to significantly refocus the instrument. Clearly, it is only possible to use these filters in combinations compatible with the cut-points of the two dichroics, such as $u' g' r'$, $u' g' i'$, $u' g' z'$, $u' g'$ clear, $u' \text{C III/N III} + \text{He II}$ red continuum, $u' g' \text{Na I}$ or u' blue continuum H α .

The final optical element encountered by a photon in ULTRACAM is the CCD window. This allows light to fall on the chip whilst retaining the vacuum seal of the CCD head (see Section 3.2). The windows are plane, parallel discs made of UV grade fused silica and coated with the same antireflection coating as the lenses (see Fig. 1).

3.2 Detectors

ULTRACAM uses three E2V 47-20 CCDs as its detectors. These are frame-transfer chips with imaging areas of $1024 \times 1024 \text{ pixel}^2$ and storage areas of $1024 \times 1033 \text{ pixel}^2$, each pixel of which is 13 \mu m on a side. To improve quantum efficiency, the chips are thinned, back-illuminated and antireflection coated with E2V’s enhanced broadband astronomy coating (in the case of the blue and green chips) and standard mid-band coating (in the case of the red chip) – see Fig. 1. The chips have a single serial register which is split into two halves, thereby doubling the frame rate, and each of these two channels

has a very low-noise amplifier at its end, delivering readout noise of only $\sim 3.5e^-$ at pixel rates of $11.2 \text{ \mu s pixel}^{-1}$ and $\sim 5e^-$ at $5.6 \text{ \mu s pixel}^{-1}$. The devices used in ULTRACAM are grade 0 devices, that is, they are of the highest cosmetic quality available. The full well capacity of these devices is $\sim 100\,000e^-$, and the devices are hence operated at approximately unity gain with the 16-bit analogue-to-digital converters in the CCD controller (see Section 3.4).

In order for dark current to be a negligible noise source, it must be significantly lower than the number of photons received from the sky. In the worst case, in which an ULTRACAM observation is being performed in the u' band when the Moon is new, on a dark site (such as La Palma) and on a small-aperture telescope (such as a 1-m), the number of photons from the sky incident on ULTRACAM would be only $\sim 0.3e^- \text{ pixel}^{-1} \text{ s}^{-1}$. In this most pessimistic scenario, therefore, the dark current must remain below $\sim 0.1e^- \text{ pixel}^{-1} \text{ s}^{-1}$ to be a negligible noise source. Fortunately, the E2V 47-20 CCDs can be run in inversion mode (e.g. McLean 1997), which at a chip temperature of 233 K delivers a dark current of only $\sim 0.05e^- \text{ pixel}^{-1} \text{ s}^{-1}$, beating our requirement by a factor of 2. This is a relatively high temperature, so the chips can be cooled by a three-stage thermoelectric cooler (TEC) utilizing the Peltier effect. To maintain cooling stability, and hence dark current stability, the hot side of the TEC is itself maintained at a constant temperature of 283 K using a recirculating water chiller. Note that this chiller is also used to cool the CCD controller (see Section 3.4), which is located at the base of the instrument and can become extremely hot during operation.

The great advantage of using thermoelectric cooling as opposed to liquid nitrogen, for example, is that the resulting CCD heads are very small and lightweight. With three such heads in ULTRACAM, this has dramatically decreased both the mass and volume of the instrument. To reduce conductive heating and to prevent condensation on the CCDs whilst observing, the heads have been designed to hold a vacuum of below 10^{-3} torr for several weeks, and we also blow dry air or nitrogen gas across the front of the head to prevent condensation from forming on the CCD window.

3.3 Mechanics

Small-aperture telescopes generally impose much smaller size and mass constraints on instruments than large-aperture telescopes. Given the requirement for portability detailed in Section 2, we therefore designed the mechanical structure of ULTRACAM assuming the space envelope and mass limits of a typical 2-m telescope (e.g. Aristarchos).

The mechanical structure of ULTRACAM is described in detail by Stevenson (2004). Briefly, it is required to: (i) provide a stable platform on which to mount the optics, CCD heads and CCD controller; (ii) allow easy access to the optics and CCD heads for alignment, filter changes and vacuum pumping; (iii) provide alignment mechanisms for the optics and CCD heads; (iv) exhibit low thermal expansion, as all three cameras must retain parfocality; (v) exhibit low flexure (less than 1 pixel, i.e. $13\ \mu\text{m}$) at any orientation, so stars do not drift out of the small windows defined on the three chips; (vi) be electrically and thermally isolated from the telescope in order to reduce pickup noise and prevent the water chiller from attempting to cool the entire telescope structure.

The mechanical structure chosen for ULTRACAM is a double octopod, as shown in Fig. 3, which provides a rigid and open framework meeting all of the requirements described above. The Serrurier trusses of the double octopod are made of carbon fibre, which offers similar structural strength to steel but at a fivefold reduction in mass and a low thermal expansion coefficient. All of the remaining parts of the mechanical chassis are of aluminium, giving a total mass for

the instrument of only 82 kg and an overall length of just 792 mm (including the CCD controller and WHT collimator).

ULTRACAM requires a mounting collar to interface it to a telescope focus. The collar, which is constructed of steel, places ULTRACAM at the correct distance from the telescope focal plane. A layer of G10/40 isolation material is placed between the collar and the top plate of ULTRACAM to provide thermal and electrical isolation from the telescope. The mounting collar also contains a motorized focal-plane mask. This is an aluminium blade which can be moved in the focal plane to prevent light from falling on regions of the CCD chip outside the user-defined windows typically used for observing. Without this mask, the light from bright stars falling on the active area of the chip above the CCD windows would cause vertical streaks in the windows (see fig. 1 of Dhillon et al. (2005) for an example). The mask also prevents photons from the sky from contaminating the background in drift-mode windows (see Section 3.7 and Stevenson 2004, for details).

3.4 Data acquisition system

In this section we provide a brief overview of the ULTRACAM data acquisition system. A much more detailed description can be found in the paper by Beard et al. (2002).

3.4.1 Hardware

Fig. 4 shows the data acquisition hardware used in ULTRACAM. Data from the three CCD chips are read out by a San Diego State




Figure 3. Top left-hand panel: CAD image of the ULTRACAM opto-mechanical chassis, highlighting some of the components described in the text. Top right-hand panel: Photograph of ULTRACAM in the test focal station of the WHT (courtesy Sue Worswick). Bottom left-hand panel: Photograph of ULTRACAM mounted on the Cassegrain focus of the WHT. Bottom right-hand panel: Photograph of ULTRACAM mounted on the Visitor Focus of the VLT (Nasmyth on Melipal).




Figure 4. Schematic showing the principal hardware components of the ULTRACAM data acquisition system. The connections between the hardware components and their locations at the telescope are also indicated.

University (SDSU) Generation III CCD controller (Leach, Beale & Eriksen 1998; Leach & Low 2000). The SDSU controller, which is in wide use at many of the world’s major ground-based telescopes, was adopted by the ULTRACAM project due to its user programmability, fast readout, low noise and ability to operate several multiple-channel CCDs simultaneously. The SDSU controller is hosted by a rack-mounted dual-processor PC running Linux patched with real-time application interface (RTAI) extensions. RTAI is a real-time version of Linux which enables a user to have complete control over when a particular task is carried out by a processor, irrespective of what else might be running within the multitasking environment. For example, in a normal Linux system undergoing heavy input/output activity (e.g. due to the running of a data reduction task on a large file on disc), the delay between requesting a command to be performed and its execution can be hundreds of milliseconds in the worst case. With RTAI, this delay is reduced to below $\sim 10 \mu\text{s}$. In ULTRACAM, RTAI is used to provide strict control over one of the two processors, so as to obtain accurate time-stamps from the Global Positioning System (GPS) antenna located outside the dome and connected to the PC via a serial port (see Section 3.5 for details).

The instrument control PC communicates with the SDSU controller via a peripheral component interconnect (PCI) card and two 250 MHz optical fibres. As well as communicating through the fibres, the SDSU controller also has the ability to interrupt the PC using its parallel port interrupt line, which is required to perform accurate time-stamping (see Section 3.5). Data from the CCDs are passed from the SDSU PCI card to the PC memory by direct memory access, from where the data are written to a high-capacity SCSI disc

array. All of the real work in reading out the CCDs is performed by the SDSU controller; the PCI card merely forwards commands and data between the instrument control PC and the SDSU controller.

3.4.2 Software

The SDSU controller and PCI card both have onboard digital signal processors (DSPs) which can be programmed by downloading assembler code from the instrument control PC. An ULTRACAM user wishing to take a sequence of windowed images, for example, would load the relevant DSP application on to the SDSU controller (to control the CCDs) and PCI card (to handle the data). The user can also modify certain parameters, such as the exposure time or binning factors (see Section 3.7 for a full list), by writing the new values directly to the DSP’s memory.

All communication within the ULTRACAM system, including the loading of DSP applications on the SDSU controller, is via Extensible Markup Language (XML) documents transmitted using the Hyper Text Transfer Protocol (HTTP) protocol. This is an international communications standard, making the ULTRACAM data acquisition system highly portable and enabling users to operate the instrument using any interface able to send XML documents via HTTP (e.g. a web browser, perl scripts, java).

3.5 Time-stamping

Stamping CCD frames with start times accurate to a small fraction of the typical exposure time is a key requirement for any astronomical instrument. In ULTRACAM, with frame rates of up to 500 Hz, this requirement is particularly severe, as it demands time-stamping accurate to a fraction of a millisecond. Without this level of accuracy, for example, it would be impossible to determine the rate of decrease in the orbital periods of interacting binary stars as measured from eclipse timings (Brinkworth et al. 2006), or compare the pulse arrival times in the optical and X-ray light from anomalous X-ray pulsars (Dhillon et al. 2005).

Whenever an exposure is started, the SDSU controller sends an interrupt to the instrument control PC which, thanks to the use of RTAI, *instantaneously* (i.e. within $\sim 10 \mu\text{s}$) writes the current time to a first-in first-out (FIFO) buffer. A description of how the current time is determined is given below. The data handling software then reads the time-stamp from the FIFO and writes it to the header of the next buffer of raw data written to the PC memory. In this way, the time-stamps and raw data always remain synchronized. Moreover, as the SDSU controller reads out all three chips simultaneously, the red, green and blue channels of ULTRACAM also remain perfectly synchronized.

The clock provided on most PC motherboards is not able to provide the current time with sufficient accuracy for our purposes, as it typically drifts by milliseconds per second. The solution adopted by ULTRACAM is to use a PCI-CTRO5 9513-based counter/timer board, manufactured by Measurement Computing Corporation, in conjunction with a Trimble Acutime 2000 smart GPS antenna. Every 10 s, the GPS antenna reports UTC to an accuracy of 50 nanoseconds. At the same time, the number of ticks reported by the counter board is recorded. At a later instant, when a time-stamp is requested, the system records the new value reported by the counter board, calculates the number of ticks that have passed since the last GPS update, multiplies this by the duration of a tick (which we have accurately measured in the laboratory) and adds the resulting interval to the previous UTC value reported by the GPS. Since the counter

board ticks at 1 MHz, to an accuracy of 100 ppm, this is a much more reliable method of time-stamping than using the system clock of the PC.

Laboratory measurements indicate that the time-stamping in ULTRACAM has a relative (i.e. frame-to-frame) accuracy of $\sim 50 \mu\text{s}$. The absolute timing accuracy of ULTRACAM has been verified to an accuracy of $\sim 1 \text{ ms}$ by comparing observations of the Crab pulsar with the ephemeris of Lyne, Jordan & Roberts (2005) (see Stevenson 2004, for details).

3.6 Pipeline data reduction system

ULTRACAM can generate up to 1 MB of data per second. In the course of a typical night, therefore, it is possible to accumulate up to 50 GB of data, and up to 0.5 TB of data in the course of a typical observing run. To handle these high data rates, ULTRACAM has a dedicated pipeline data reduction system,⁴ written in C++, which runs on a Linux PC or Mac located in the telescope control room and connected to the instrument control PC via a dedicated 100BaseT local area network (see Fig. 4).

Data from a run on an object with ULTRACAM are stored in two files, one an XML file containing a description of the data format, and the other a single, large unformatted binary file containing all the raw data and time-stamps. This latter file may contain millions of individual data frames, each with its own time-stamp, from each of the three ULTRACAM CCDs. The data reduction pipeline grabs these frames from the SCSI disc array by sending HTTP requests to a file server running on the instrument control PC.

The ULTRACAM data reduction pipeline has been designed to serve two apparently conflicting purposes. Whilst observing, it acts as a quick-look data reduction facility, with the ability to display images and generate light curves in real time, even when running at the highest data rates of up to 1 MB per second and at the highest frame rates of up to 500 Hz. After observing, the pipeline acts as a fully featured photometry reduction package, including optimal extraction (Naylor 1998). To enable quick-look reduction whilst observing, the pipeline keeps many of its parameters hidden to the user and allows the few remaining parameters to be quickly skipped over to generate images and light curves in as short a time as possible. Conversely, when carefully reducing the data after a run, every single parameter can be tweaked in order to maximize the signal-to-noise ratio of the final data.

3.7 Readout modes

Fig. 5 can be used to illustrate how the frame-transfer chips in ULTRACAM are able to deliver high frame rates. A typical observation with ULTRACAM consists of the observation of the target star in one user-defined CCD window and a nearby non-variable comparison star in another, that is, window pair 1 in Fig. 5. Once the exposure is complete, the image area is shifted into the storage area, a process known as *vertical clocking*. This is very rapid, taking only $23.3 \mu\text{s}$ per row and hence $\sim 24 \text{ ms}$ to shift the entire 1024 rows into the storage area. As soon as the image area is shifted in this way, the next exposure begins. Whilst exposing, the previous image in the storage area is shifted on to the serial register and then undergoes *horizontal clocking* to one of two readout ports, where it is digitized. In other words, the previous frame is being read out whilst the next

frame is exposing, thereby reducing the dead time to the time it takes to shift the image into the storage area, that is, 24 ms.

ULTRACAM has no shutter – the fast shifting of data from the image area to the storage area acts like an electronic shutter, and is far faster than conventional mechanical shutters. This does cause some problems, however, such as vertical trails of starlight from bright stars, but these can be overcome in some situations by the use of a focal-plane mask (see Section 3.3).

Setting an exposure time with ULTRACAM is a more difficult concept than in a conventional non-frame-transfer camera. This is because ULTRACAM attempts to frame as fast as it possibly can, that is, it will shift the image area into the storage area as soon as there is room in the storage area to do so. Hence, the fastest exposure time is given by the fastest time it takes to clear sufficient room in the storage area, which in turn depends on the number, location, size and binning factors of the windows in the image area, as well as the vertical clocking and digitization times, all of which are variables in the ULTRACAM data acquisition system. To obtain an arbitrarily long exposure time with ULTRACAM, therefore, an *exposure delay* is added prior to the vertical clocking to allow photons to accumulate in the image area for the required amount of time. Conversely, to obtain an arbitrarily short exposure time with ULTRACAM, it is necessary to set the exposure delay to zero and adjust the window, binning and digitization parameters so that the system can frame at the required rate. As it takes 24 ms to vertically clock the entire image area into the storage area, this provides a hard limit to the maximum frame rate of $\sim 40 \text{ Hz}$, with a duty cycle (given by the exposure time divided by the sum of the exposure and dead times) of less than 3 per cent. Adopting a more acceptable duty cycle of 75 per cent results in a useable limit of only $\sim 10 \text{ Hz}$.

Clearly, an alternative readout strategy is required in order to push beyond the $\sim 10 \text{ Hz}$ frame-rate barrier and approach the desired kHz frame rates required to study the most rapid variability. For this purpose, we have developed *drift mode*, which is described in detail in Appendix A. Briefly, in drift mode the windows are positioned on the border between the image and storage areas and, instead of vertically clocking the entire image area into the storage area, only the window is clocked into the (top of) the storage area. A number of such windows are hence present in the storage area at any one time. This dramatically reduces the dead time, as now the frame rate is limited to the time it takes to clock only a small window into the storage area. For example, in the case of two windows of size $24 \times 24 \text{ pixel}^2$ and binned 4×4 , it is possible to achieve a frame rate of $\sim 500 \text{ Hz}$ with a duty cycle of ~ 75 per cent. This is currently the highest frame rate that ULTRACAM has achieved on sky whilst observing a science target. It is worth noting that at these high frame rates it is the speed at which charge can be shifted along the serial register on the ULTRACAM CCDs (currently $0.48 \mu\text{s pixel}^{-1}$), rather than the digitization time, that limits the frame rate. With larger windows and hence lower frame rates, the reverse is true.

Drift mode only offers the possibility of two windows and should only be used when frame rates in excess of 10 Hz are required. This is because the windows in drift mode spend longer on the chip and hence accumulate more dark current and, without the use of the focal plane mask, more sky photons. For frame rates of less than 10 Hz, ULTRACAM offers normal two-window, four-window, six-window and full-frame readout modes (the latter reading out in approximately 3 s with only 24 ms dead time). An example of a full-frame image taken with ULTRACAM is shown in Fig. 6.

For some applications, for example when taking flat fields or observing bright standard stars, it is desirable to use a full frame or

⁴ Available for download at <http://deneb.astro.warwick.ac.uk/phsaap/software/ultracam/html/index.html>.




Figure 5. Pictorial representation of an ULTRACAM CCD. The location of the image area, storage area and serial register are shown, along with the vertical and horizontal clocking directions. It can be seen that there are two readout channels, separated by the vertical dashed line, with the left-hand half of the chip read out via channel 1, and the right-hand half via channel 2. In two-window mode, only window pair 1 is present. In four-window mode, window pairs 1 and 2 are present, and in six-window mode all three window pairs are present. The chip can also be read out in full-frame mode, in which case the entire image area is selected (see Fig. 6). The window parameters and pixel rates defined in Appendix A are also shown.

two large windows and yet have short exposure times. This is not possible with the modes described above. To enable exposure times of arbitrarily short length, therefore, ULTRACAM also offers the so-called *clear* modes. These modes, which are available in full-frame and two-window formats, clear the chip prior to exposing for the requested amount of time. This means that any charge which has built up in the image area whilst the previous exposure is reading out is discarded. The disadvantage of this mode is that the duty cycle becomes very poor (3 per cent in the case of a 0.1-s exposure time in full-frame-clear mode).

A more detailed description of ULTRACAM’s readout modes can be found in Appendix A.

4 PERFORMANCE

During the design phase of the project, a set of functional and performance requirements for ULTRACAM were established against which the instrument was tested during the commissioning phase. A detailed description of how ULTRACAM performed when measured against these requirements is given by Stevenson (2004). Other measures of the performance of an astronomical instrument include how well the project was managed, how reliable the instrument is in operation, how much telescope time it wins, and how much science it enables. In what follows, we attempt to summarize the performance of ULTRACAM measured against these (very different) metrics.




Figure 6. Full-frame ULTRACAM image of the Crab nebula (M1) in, from left- to right-hand side, i' , g' and u' . The image was taken using the WHT on 2002 September 11 and represents a total exposure time in each filter of 413 s. The Crab pulsar is at pixel position (540, 525), just to the right-hand side of centre. The resulting light curve of the Crab pulsar was used to calibrate the ULTRACAM GPS time-stamping system (Stevenson 2004). The field of view is 5 arcmin, with north approximately at the top. The vertical line represents the border between the two readout channels on each chip, and the blank rows/columns at the top/right-hand side of the image indicate the underscan/overscan regions.

Table 2. Breakdown of the percentage of time spent observing different classes of astronomical object with ULTRACAM on the WHT and VLT. The right-hand column provides references to some of the ULTRACAM papers published in each area. The entry for cataclysmic variables includes generic studies of accretion discs around white dwarfs (6.6 per cent) and pulsating white dwarfs in binaries (3.3 per cent). The entry for X-ray binaries is composed of systems with black hole primaries (7.9 per cent) and neutron star primaries (9.5 per cent).

Target	Time (per cent)	References
Cataclysmic variables	25.3	Littlefair et al. (2006a), Feline et al. (2004)
X-ray binaries	17.4	Muñoz-Darias et al. (2006), Shahbaz et al. (2005)
sdB stars/asteroseismology	11.2	Aerts et al. (2006), Jeffery et al. (2004)
Eclipsing, detached white dwarf/red dwarf binaries	10.4	Brinkworth et al. (2006), Maxted et al. (2004)
Occultations by Titan, Pluto and Kuiper Belt objects	10	Fitzsimmons et al. (2007), Roques et al. (2006)
Ultracompact binaries	6.6	Barros et al. (2007), Marsh et al. (2006)
Pulsars	5	Dhillon, Marsh & Littlefair (2006), Dhillon et al. (2005)
Extrasolar planet transits	3.7	
Flare stars	3.3	Mathioudakis et al. (2006)
Isolated brown dwarfs	2.9	Littlefair et al. (2006b)
Isolated white dwarfs	1.7	Silvotti et al. (2006)
Gamma-ray bursts	1.7	Vreeswijk et al. (2005)
Active galactic nuclei	0.8	

(i) ULTRACAM was delivered on budget ($\sim\pounds 300\,000$) and 3 months ahead of the 3-yr schedule set to design, build and commission it.

(ii) ULTRACAM is an extremely reliable instrument, as it has no moving parts. To date, ULTRACAM has suffered from <2 per cent technical downtime, and this has nearly all been due to unavoidable hardware failures (such as hard discs and power supplies).

(iii) ULTRACAM has met nearly all of its performance requirements, including throughput (50 per cent in g'), image quality (seeing limited at 0.6 arcsec/0.3 arcsec on WHT/VLT), flexure ($10\ \mu\text{m}$), timing accuracy ($50\ \mu\text{s}$) and detector-limited data throughput. Using high-quality optics, detectors and coatings, ULTRACAM has achieved photometric zero points on the VLT of approximately $u' = 26.0$, $g' = 28.1$, $r' = 27.6$, $i' = 27.7$ and $z' = 26.7$.

(iv) The frame rates achieved on sky have ranged from 0.05 to 500 Hz on targets ranging in brightness from 8–26 mag. It can be seen that the maximum frame rate is a factor of 2 below the required 1 kHz value set at the start of the project. This is not a serious drawback, however, as the science performed by ULTRA-

CAM to date (see Table 2) has almost all concentrated on the 0.1–1 s regime.

(v) The ULTRACAM project team has adopted a policy of open access, offering the instrument to any astronomer on a shared-risks, collaborative basis. As a result, ULTRACAM has been awarded a great deal of competitively allocated observing time: 89 nights on the WHT and 39 nights on the VLT in the 5 yr since first light. A breakdown of how these 128 nights have been spent on different types of object is given in Table 2.

(vi) To date, a total of ~ 27 refereed papers have been published which are wholly or partly dependent on ULTRACAM data. Some examples of the research done with ULTRACAM are given in Table 2.

(vii) Any instrument which is not continually enhanced and upgraded will eventually cease to be competitive. For this reason, we have always pursued a vigorous enhancements programme with ULTRACAM, focused on the three most important areas for high time resolution astrophysics: increasing the maximum frame rate, reducing the readout noise and minimizing downtime/inefficiency.

The performance of the instrument has improved by at least a factor of 2 in these three areas in the 5 yr since first light.

5 CONCLUSIONS

We have described the scientific requirements, design and performance of ULTRACAM, and listed examples of the varied science that has been performed with the instrument. As for the future, we intend to continue to operate ULTRACAM for approximately one run per year on both the VLT and the WHT.⁵ We have also recently commissioned a spectroscopic version of ULTRACAM, known as ULTRASPEC (Dhillon et al. 2007). This instrument uses an electron-multiplying, frame-transfer CCD and the ULTRACAM data acquisition system to deliver high-speed, zero readout-noise spectroscopy with the EFOSC2 spectrograph on the ESO 3.6-m telescope at La Silla. In the longer term, we are exploring the possibility of building ULTRACAM-II, which would be dedicated for use in Chile and which would have four channels, a larger field of view and, if available, use multichannel, frame-transfer electron-multiplying CCDs.

ACKNOWLEDGMENTS

There are many people who have contributed to the success of ULTRACAM but who have not been included as co-authors on this paper. We would like to thank all of them for their dedication to the project, particularly the staff of the Central Mechanical Workshops at the University of Sheffield, the UK Astronomy Technology Centre in Edinburgh, the Isaac Newton Group of Telescopes on La Palma and the European Southern Observatory at Paranal. We would also like to thank PPARC for providing the funding to build, operate and exploit ULTRACAM.

We dedicate this paper to the memory of our friend and colleague, Emiliós Harlaftis, whose enthusiasm for the project and desire to use ULTRACAM on the new 2.3-m Aristarchos Telescope in Greece significantly influenced the design of the instrument.

REFERENCES

- Aerts C., Jeffery C. S., Fontaine G., Dhillon V. S., Marsh T. R., Groot P., 2006, *MNRAS*, 367, 1317
- Baldwin J. E., Tubbs R. N., Cox G. C., Mackay C. D., Wilson R. W., Andersen M. I., 2001, *A&A*, 368, L1
- Barros S. C. C., Marsh T. R., Dhillon V. S., Groot P. J., Nelemans G., Roelofs G., Steeghs D., Wheatley P. J., 2007, *MNRAS*, 374, 1334
- Beard S. M., Vick A. J. A., Atkinson D., Dhillon V. S., Marsh T. R., McLay S., Stevenson M. J., Tierney C., 2002, *Proc. SPIE*, 4848, 218
- Brinkworth C. S., Marsh T. R., Dhillon V. S., Knigge C., 2006, *MNRAS*, 365, 287
- Čadež A., Vidrih S., Galičič M., Carramiñana A., 2001, *A&A*, 366, 930
- Cumani C., Mantel K.-H., 2001, *Exp. Astron.*, 11, 145
- Dhillon V. S., Marsh T. R., Hulleman F., van Kerkwijk M. H., Shearer A., Littlefair S. P., Gavriil F. P., Kaspi V. M., 2005, *MNRAS*, 363, 609
- Dhillon V. S., Marsh T. R., Littlefair S. P., 2006, *MNRAS*, 372, 209
- Dhillon V. S., Marsh T. R., Copperwheat C., Bezawada N., Ives D. J., Vick A. J., O'Brien K., 2007, *ESO Messenger*, 127, 41
- Dravins D., 1994, *ESO Messenger*, 78, 9
- Feline W. J., Dhillon V. S., Marsh T. R., Brinkworth C. S., 2004, *MNRAS*, 355, 1
- Fitzsimmons A. et al., 2007, *A&A*, submitted
- 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
- Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, *AJ*, 111, 1748
- Hynes R. I., Charles P. A., Casares J., Haswell C. A., Zurita C., Shahbaz T., 2003, *MNRAS*, 340, 447
- Jeffery C. S., Dhillon V. S., Marsh T. R., Ramachandran B., 2004, *MNRAS*, 352, 699
- Jorden P., Deltorn J.-M., Oates P., 1993, *Gemini*, 41, 1
- Kern B., 2002, PhD thesis, California Institute of Technology
- Kern B., Martin C., 2002, *Nat*, 417, 527
- Kotar J., Vidrih S., Čadež A., 2003, *Rev. Sci. Instrum.*, 74, 3111, preprint (astro-ph/0303368)
- Law N. M., Hodgkin S. T., Mackay C. D., 2006, *MNRAS*, 368, 1917
- Leach R. W., Low F. J., 2000, *Proc. SPIE*, 4008, 337
- Leach R. W., Beale F. L., Eriksen J. E., 1998, *Proc. SPIE*, 3355, 512
- Littlefair S. P., Dhillon V. S., Marsh T. R., Gänsicke B. T., Southworth J., Watson C. A., 2006a, *Sci*, 314, 1578
- Littlefair S. P., Dhillon V. S., Marsh T. R., Shahbaz T., Martín E. L., 2006b, *MNRAS*, 370, 1208
- Lyne A. G., Jordan C. A., Roberts M. E., 2005, *Monthly ephemeris, Jodrell Bank Crab Pulsar Timing Results. Jodrell Bank Observatory, University of Manchester*
- Mackay C. D., Tubbs R. N., Bell R., Burt D., Moody I., 2001, *Proc. SPIE*, 4306, 289
- Marsh T. R., Horne K., 1998, *MNRAS*, 299, 921
- Marsh T. R. et al., 2006, in Napiwotzki R., Burleigh M. R., eds, *ASP Conf. Ser. Vol. 372, Proc 15th European Workshop on White Dwarfs. Astron. Soc. Pac.*, San Francisco, in press (astro-ph/0610414)
- Mathioudakis M., Bloomfield D. S., Jess D. B., Dhillon V. S., Marsh T. R., 2006, *A&A*, 456, 323
- Mated P. F. L., Marsh T. R., Morales-Rueda L., Barstow M. A., Dobbie P. D., Schreiber M. R., Dhillon V. S., Brinkworth C. S., 2004, *MNRAS*, 355, 1143
- McLean I. S., 1997, *Electronic Imaging in Astronomy: Detectors and Instrumentation*. Wiley, Chichester
- Muñoz-Darias T., Martínez-Pais I. G., Casares J., Marsh T. R., Cornelisse R., Steeghs D., Dhillon V. S., Charles P. A., 2006, *Adv. Space Res.*, 38, 2762, preprint (astro-ph/0507066)
- Naylor T., 1998, *MNRAS*, 296, 339
- O'Brien K., 2007, in Phelan D., Ryan O., Shearer A., eds, *High Time Resolution Astrophysics, Astronomy and Space Science Library*. Springer-Verlag, Berlin, in press
- O'Brien K., Horne K., Boroson B., Still M., Gomer R., Oke J. B., Boyd P., Vrtilsek S. D., 2001, *MNRAS*, 326, 1067
- O'Donoghue D., 1995, *Balt. Astron.*, 4, 517
- O'Donoghue D. et al., 2006, *MNRAS*, 372, 151
- Roques F. et al., 2006, *AJ*, 132, 819
- 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*
- Shahbaz T., Dhillon V. S., Marsh T. R., Casares J., Zurita C., Charles P. A., Haswell C. A., Hynes R. I., 2005, *MNRAS*, 362, 975
- Silvotti R., Pavlov M., Fontaine G., Marsh T. R., Dhillon V. S., 2006, *Mem. Soc. Astron. Ital.*, 77, 486
- Simons D. A., 1995, *Technical note no. 30. Gemini Observatory*
- Skidmore W., O'Brien K., Horne K., Gomer R., Oke J. B., Pearson K. J., 2003, *MNRAS*, 338, 1057
- St-Jacques D., Cox G. C., Baldwin J. E., Mackay C. D., Waldram E. M., Wilson R. W., 1997, *MNRAS*, 290, 66
- Stathakis R. A., Johnston H. M., 2002, *User Manual, RGO Spectrograph. Anglo-Australian Observatory*
- Stevenson M. J., 2004, PhD thesis, Univ. Sheffield
- Stover R. J., Allen S. L., 1987, *PASP*, 99, 877
- Vreeswijk P., Marsh T., Littlefair S., Dhillon V., O'Brien K., 2005, *GRB Coordinates Network*, 3445, 1

⁵ Anyone interested in using ULTRACAM on a shared-risks, collaborative basis are encouraged to contact one of the first two authors.

Wilson R. W., O'Mahony N., Packham C., Azzaro M., 1999, MNRAS, 309, 379

Woudt P. A., Warner B., 2001, MNRAS, 328, 159

APPENDIX A: MODELLING CCD FRAME RATES

In this section we expand on the brief description of ULTRACAM's readout modes given in Section 3.7 and provide an algorithm to compute ULTRACAM exposure times and dead times (and hence frame rates). The algorithm can straightforwardly be adapted to predict the performance of any frame transfer CCD.⁶

Experiments show that the algorithm is accurate to a few per cent, due mainly to uncertainties in the precise values of the input parameters such as the vertical clocking time, horizontal clocking time and digitization time. Note that we ignore the negligible dead time induced by the time it takes to execute the DSP code controlling the readout process, and the negligible influence of the various dark reference columns and rows on the calculated frame rates. For simplicity, we also ignore the negligible delay induced by the extra one or two pixels which must be added or subtracted from some of the equations given below to allow for the fact that, for example, the number of pixels in a window is given by the difference between its end and start pixel positions + 1.

Referring to Fig. 5, we define the following symbols.

1. n_{row} : The total number of rows in the image area of the CCD. In ULTRACAM, this figure is 1024. For simplicity, we assume that the storage area has the same number of rows; in reality, the storage area of the CCDs used in ULTRACAM has nine more rows than the image area.

2. n_{col} : The number of columns in the image area. The storage area has the same number of columns. The number of pixels in the serial register is also equal to this number. In ULTRACAM, this figure is 1024.

3. $n_{1x}, n_{1y}, n_{2x}, n_{2y}, n_{3x}, n_{3y}$: The x and y sizes of each window in pairs 1, 2 and 3, respectively.

4. $x_{1L}, x_{1R}, x_{2L}, x_{2R}, y_{3L}, y_{3R}$: The x position in the left-hand (L) and right-hand (R) readout channels of the start of window pairs 1, 2 and 3, respectively.

5. y_1, y_2, y_3 : The y position of the bottom of window pairs 1, 2 and 3, respectively.

6. b_x, b_y : The binning factors in the x (horizontal) and y (vertical) directions.

7. n_{win} : In drift mode, the number of windows in the stack in the storage area.

8. t_{vclock} : The time taken to vertically clock one row. In ULTRACAM, this has the same value (23.3 μs per row) in both the image and storage areas, but this need not necessarily be the case.

9. t_{hclock} : The time taken to horizontally clock one pixel along the serial register. In ULTRACAM, this is currently limited to 0.48 $\mu\text{s pixel}^{-1}$.

10. t_{video} : The time taken to determine the charge content of a pixel (via correlated double sampling) and perform the analogue-to-digital conversion. Two speeds are currently used in ULTRACAM: slow (11.2 $\mu\text{s pixel}^{-1}$) and fast (5.6 $\mu\text{s pixel}^{-1}$).

11. t_{delay} : The exposure delay, a user-defined pause in the readout process which increases the time spent accumulating photons in the image area.

12. t_{pipe} : An additional exposure delay used in drift mode which ensures that the exposure times and the gaps between exposures are uniform.

13. t_{inv} : The inversion time. This is the time taken for the chip to come out of inversion mode in preparation for the readout process. In ULTRACAM, this takes 110 μs .

14. t_{clear} : The time taken to clear the chip, that is, to vertically clock the entire image and storage areas and dump the charge.

15. t_{frame} : The frame-transfer time, that is, the time taken to transfer the entire image area into the storage area.

16. t_{y1}, t_{y2}, t_{y3} : The time taken to vertically shift window pairs 1, 2 and 3, respectively, in the storage area in order to place them adjacent to the serial register.

17. $t_{\text{line}}, t_{\text{line1}}, t_{\text{line2}}, t_{\text{line3}}$: For full-frame, window pairs 1, 2 and 3, respectively, the time taken to shift one row of the storage area into the serial register, shift it along the serial register, determine the charge content of each pixel and then perform the analogue-to-digital conversion.

18. t_{cycle} : The cycle time, defined as the total amount of time it takes to expose and then read out an image.

19. t_{exp} : The exposure time, defined as the time spent accumulating photons during which no clocking is performed in the image area.

20. t_{dead} : The dead time, defined as the time during which the CCD is not accumulating photons.

21. ν_{frame} : The frame rate, defined as $\nu_{\text{frame}} = 1/t_{\text{cycle}}$.

A1 Full-frame mode

ULTRACAM can be read out in full-frame mode with and without an overscan/underscan and with and without clearing the chip prior to each exposure. Fig. 6 shows an example full-frame image taken by ULTRACAM, in which the overscan/underscan regions can be seen at the right-hand side and top of the image. For the sake of simplicity, we shall not consider the overscan/underscan case here; this effectively adds an additional eight rows and 56 columns to the chip.

We first consider the case where the chip is cleared prior to each exposure. The time taken to clear the chip is given by the time it takes to vertically clock both the image and storage areas:

$$t_{\text{clear}} = 2n_{\text{row}}t_{\text{vclock}}. \quad (\text{A1})$$

Next, move the image into the storage area:

$$t_{\text{frame}} = n_{\text{row}}t_{\text{vclock}}. \quad (\text{A2})$$

The time it then takes to shift one row of the storage area into the serial register, shift it along the entire serial register, determine the charge content of each pixel and then digitize the data is given by

$$t_{\text{line}} = (b_y t_{\text{vclock}}) + \left(\frac{n_{\text{col}}}{2} t_{\text{hclock}} \right) + \left(\frac{n_{\text{col}}}{2b_x} t_{\text{video}} \right), \quad (\text{A3})$$

where the factors of 2 are due to the fact that the ULTRACAM CCDs are two-channel devices. The total time taken to read out one exposure is hence the sum of the following:

$$t_{\text{cycle}} = t_{\text{clear}} + t_{\text{delay}} + t_{\text{inv}} + t_{\text{frame}} + \left(\frac{n_{\text{row}}}{b_y} t_{\text{line}} \right). \quad (\text{A4})$$

Because the chip is cleared prior to each exposure, any charge accumulated in the image area during the previous exposure is lost.

⁶ An online ULTRACAM frame-rate calculator can be found at <http://www.shef.ac.uk/physics/people/vdhillon/ultracam>.

The exposure time is then simply

$$t_{\text{exp}} = t_{\text{delay}}. \quad (\text{A5})$$

Hence it is possible to obtain arbitrarily short exposure times in clear mode.

Instead, if the chip is not cleared prior to each exposure, the total time taken to read out one exposure is

$$t_{\text{cycle}} = t_{\text{delay}} + t_{\text{inv}} + t_{\text{frame}} + \left(\frac{n_{\text{row}}}{b_y} t_{\text{line}} \right), \quad (\text{A6})$$

that is, there is no t_{clear} term (which makes only a negligible difference to the cycle time). Since all charge that accumulates in the image area whilst the previous exposure is reading out is recorded, the exposure time is given by

$$t_{\text{exp}} = t_{\text{cycle}} - t_{\text{frame}}. \quad (\text{A7})$$

Hence the shortest possible exposure time obtainable in no-clear mode is given by the time it takes to read the entire chip out. In ULTRACAM, this can be up to 6 s (when not binning and using the slowest value for t_{video}).

The dead times of both the clear and no-clear modes are

$$t_{\text{dead}} = t_{\text{cycle}} - t_{\text{exp}}. \quad (\text{A8})$$

In the case of the no-clear mode, this means that $t_{\text{dead}} = t_{\text{frame}}$. This is negligible, since t_{frame} is only 24 ms in ULTRACAM. In the case of the clear mode, however, $t_{\text{dead}} = t_{\text{cycle}} - t_{\text{delay}} \simeq (n_{\text{row}}/b_y)t_{\text{line}}$, which means that the dead time is dependent on the time it takes to read the entire chip out (i.e. up to 6 s).

A2 Two-windowed mode

ULTRACAM is most often used for point-source photometry, in which case it is usually not necessary to read out the entire array. Instead, windows around the target and comparison stars are defined, as shown in Fig. 5, substantially increasing the frame rate and reducing the amount of data accumulated.

It is possible to define either two, four or six windows with ULTRACAM, with one window of each pair occupying the left-hand side of the chip and the other window occupying the right-hand side. Additionally, for each pair, the two windows must have the same sizes, vertical start positions and binning factors, and they must not overlap with each other. These rules significantly simplify the data acquisition software and yet still give full flexibility when observing, as any two stars can be located in two windows simply by adjusting the telescope position, the instrument rotator angle and the horizontal start positions and sizes of the windows.

In two-windowed mode, it is possible to read out the chip both with and without clearing. We first consider the case where the chip is cleared prior to each exposure. The time taken to clear the chip is given by the time it takes to vertically clock the image and storage areas:

$$t_{\text{clear}} = 2n_{\text{row}}t_{\text{vclock}}. \quad (\text{A9})$$

Next, move the windows into the storage area:

$$t_{\text{frame}} = n_{\text{row}}t_{\text{vclock}}. \quad (\text{A10})$$

At this stage, the windows have the same vertical position in the storage area as they had in the image area. Hence the next step is to vertically shift the windows to place them adjacent to the serial register:

$$t_{y1} = y_1 t_{\text{vclock}}. \quad (\text{A11})$$

To simplify the data acquisition software, and ensure true simultaneity between the three ULTRACAM chips, each of the windows is forced to have the same pixel position on the red, green and blue CCDs. Furthermore, the data acquisition system expects data from the left-hand channel of each chip to be processed at the same time as data from the right-hand channel, which effectively means that a window in the left-hand channel must lie the same number of pixels from the centre line of the chip as a window in the right-hand channel. In practice, such a rule would be very restrictive during target acquisition, so instead a *differential shift* is performed. On entering the serial register, each row of the window closest to the centre line of the chip is shifted by $|x_{1L} - (n_{\text{col}} - x_{1R} - n_{1x})|$ pixels until it lies the same distance from the centre line as the window in the other channel.

The time it takes to shift one row of the storage area into the serial register, perform the differential shift, horizontally shift the window to the output, determine the charge content of each pixel and then digitize the data is given by

$$t_{\text{line1}} = (b_y t_{\text{vclock}}) + (n_{\text{hclock1}} t_{\text{hclock}}) + \left(\frac{n_{1x}}{b_x} t_{\text{video}} \right). \quad (\text{A12})$$

Note that the factor of 2 present in equation (A3) is absent from the equation for t_{line1} above due to the fact that one window is read out by one of the two channels and the second window of the pair is read out via the other channel. n_{hclock1} is the total number of horizontal clocks required and is given by the sum of the differential shift, the window size, and the number of rows between the output and whichever window is closest to it. Once the serial register has been clocked n_{hclock1} times, both windows will have been read out. The charge remaining in the serial register is then dumped, rather than clocked out. This takes only $8t_{\text{hclock}}$ and therefore significantly reduces the readout time in the case where the windows are defined to lie close to the vertical edges of the detector. Strictly speaking, therefore, one should also add 8 to n_{hclock1} .

The total time taken to read out one exposure is hence the sum of the following:

$$t_{\text{cycle}} = t_{\text{clear}} + t_{\text{delay}} + t_{\text{inv}} + t_{\text{frame}} + t_{y1} + \left(\frac{n_{1y}}{b_y} t_{\text{line1}} \right). \quad (\text{A13})$$

Because the chip is cleared prior to each exposure, any charge accumulated in the windows in the image area during the previous exposure is lost. The exposure time is then simply

$$t_{\text{exp}} = t_{\text{delay}}. \quad (\text{A14})$$

Hence, it is possible to obtain arbitrarily short exposure times in clear mode.

Instead, if the chip is not cleared prior to each exposure, the total time taken to read out one exposure is

$$t_{\text{cycle}} = t_{\text{delay}} + t_{\text{inv}} + t_{\text{frame}} + t_{y1} + \left(\frac{n_{1y}}{b_y} t_{\text{line1}} \right), \quad (\text{A15})$$

that is, there is no t_{clear} term (which makes only a negligible difference to the cycle time). Since all charge that accumulates in the image area whilst the previous exposure is reading out is recorded, the exposure time is given by

$$t_{\text{exp}} = t_{\text{cycle}} - t_{\text{frame}}. \quad (\text{A16})$$

Hence the shortest possible exposure time obtainable in no-clear mode is given by the time it takes to read the windows out. This can be reduced by using smaller windows, binning and/or ensuring the windows are positioned so as to minimize n_{hclock1} and t_{y1} .

The dead times of both the clear and no-clear modes are

$$t_{\text{dead}} = t_{\text{cycle}} - t_{\text{exp}}. \quad (\text{A17})$$

In the case of the no-clear mode, this means that $t_{\text{dead}} = t_{\text{frame}}$. This is negligible, since t_{frame} is only 24 ms in ULTRACAM. In the case of the clear mode, however, $t_{\text{dead}} = t_{\text{cycle}} - t_{\text{delay}} \simeq (n_{1y}/b_y)t_{\text{line1}}$, which means that the dead time is dependent on the time it takes to read the windows out.

A3 Four and six-windowed mode

The four and six-windowed modes are only currently available in no-clear mode. For the sake of brevity, we shall only outline the four-windowed mode here – it is a simple matter to extend the algorithm to six windows.

To read out, the windows are first moved into the storage area:

$$t_{\text{frame}} = n_{\text{row}} t_{\text{vlock}}. \quad (\text{A18})$$

The next step is to vertically shift the first pair of windows to place them adjacent to the serial register:

$$t_{y1} = y_1 t_{\text{vlock}}. \quad (\text{A19})$$

Once this pair has been read out (see below), the second pair of windows must be vertically shifted to the serial register:

$$t_{y2} = (y_2 - y_1 - n_{1y}) t_{\text{vlock}}. \quad (\text{A20})$$

The differential shift for the first pair of windows is given in Section A2. In the same way, the differential shift for the second pair of windows is given by $|x_{2L} - (n_{\text{col}} - x_{2R} - n_{2x})|$ pixels.

For the first window pair, the time it takes to shift one row of the storage area into the serial register, perform the differential shift, horizontally shift the window to the output, determine the charge content of each pixel and then digitize the data is denoted by t_{line1} and is given by equation (A12). Similarly, for the second window pair, t_{line2} is given by

$$t_{\text{line2}} = (b_y t_{\text{vlock}}) + (n_{\text{hclock2}} t_{\text{hclock}}) + \left(\frac{n_{2x}}{b_x} t_{\text{video}} \right), \quad (\text{A21})$$

where n_{hclock2} is the total number of horizontal clocks required for the second window pair and is given by the sum of the differential shift, the window size, and the number of rows between the output and whichever window of the second pair is closest to it.

The total time taken to read out one exposure is hence the sum of the following:

$$t_{\text{cycle}} = t_{\text{delay}} + t_{\text{inv}} + t_{\text{frame}} + t_{y1} + t_{y2} + \left(\frac{n_{1y}}{b_y} t_{\text{line1}} \right) + \left(\frac{n_{2y}}{b_y} t_{\text{line2}} \right). \quad (\text{A22})$$

Since there is no clear, all of the charge that accumulates in the image area whilst the previous exposure is reading out is recorded. The exposure time is therefore given by

$$t_{\text{exp}} = t_{\text{cycle}} - t_{\text{frame}}. \quad (\text{A23})$$

Hence the shortest possible exposure time obtainable in no-clear mode is given by the time it takes to read the windows out. This can be reduced by using smaller windows, binning and/or ensuring the windows are positioned so as to minimize n_{hclock1} , n_{hclock2} , t_{y1} and t_{y2} .

The dead time is

$$t_{\text{dead}} = t_{\text{cycle}} - t_{\text{exp}}, \quad (\text{A24})$$

which means that $t_{\text{dead}} = t_{\text{frame}}$. This is negligible, since t_{frame} is only 24 ms in ULTRACAM.

A4 Drift mode

Drift mode is the most complex of the ULTRACAM readout modes and is used when frame rates in excess of ~ 10 Hz are required (see Section 3.7). Due to its complexity, drift mode is only available in two-windowed mode without clearing.

The readout sequence in drift mode is shown pictorially in Fig. A1. It can be seen that the basic difference between drift mode and the standard two-windowed mode is in the frame-transfer process. In two-windowed mode, the entire image area is clocked into the storage area, resulting in an irreducible minimum dead time of 24 ms. In drift mode, the two windows are positioned on the border between the storage and image areas, and only the rows containing the windows are clocked into the storage area on completion of an exposure. Since each row takes only 23.3 μs to vertically clock, a small window of, say, 24 rows takes only ~ 0.5 ms to move into the storage area, reducing the dead time by almost a factor of 50.

Fig. A1 shows that, by moving the exposed windows from the bottom of the image area into the top of the storage area, a corresponding region at the bottom of the storage area must be moved off the chip. For this to occur rapidly, the charge must be dumped at the serial register rather than clocked along it and digitized. Once the dump is complete, a previously exposed window will lie at the bottom of the storage area and will then be read out whilst another window is exposing in the image area. In this way it can be seen that a stack of exposed windows are accumulated in the storage area, with each window separated by a gap equal to or greater than n_{1y} .

As shown below, the exposure time of a window at the bottom of the image area is dependent on how long it takes to read out the window at the bottom of the storage area. Unfortunately, the CCD used in ULTRACAM has 1033 rows in its storage area, which is a prime number. This means that one of the gaps between the windows in the stack must be slightly larger than n_{1y} , and hence takes a slightly longer time to read out. The resulting variable exposure times and variable intervals between exposures would have a detrimental effect on periodicity analyses, and hence a correction is introduced. Specifically, a *pipeline delay*, t_{pipe} , is added to each window bar the one with the larger gap, thereby equalizing the exposure times of each window in the stack and the time intervals between them (see Fig. A1).

The number of windows, n_{win} , in the stack in the storage area depends on the number of rows in the windows and the size of the storage area, as follows:

$$n_{\text{win}} = \left\lfloor \left(\frac{n_{\text{row}}}{n_{1y}} + 1 \right) / 2 \right\rfloor, \quad (\text{A25})$$

where $n_{\text{row}} = 1033$ for the storage area of the ULTRACAM CCDs, and the outer brackets on the right-hand side indicate that only the integer part of the result is recorded.

The pipeline delay can now be determined by calculating the extra number of rows present in the largest gap between the windows in the stack:

$$t_{\text{pipe}} = (n_{\text{row}} - [(2n_{\text{win}} - 1)n_{1y}]) t_{\text{vlock}}. \quad (\text{A26})$$

With the above defined, we can now model the frame rate in drift mode. On completion of a user-defined exposure delay, t_{delay} , the first step in the drift-mode readout process is to move the window from the image area to the top of the storage area:

$$t_{\text{frame}} = (n_{1y} + y_1) t_{\text{vlock}}. \quad (\text{A27})$$

This will cause a previously exposed window close to the bottom of the storage area to reach the serial register. The differential shift




Figure A1. Pictorial representation of the readout sequence in drift mode with $n_{win} = 3$. Exposed windows form a vertical stack in the storage area. The storage area has 1033 rows, implying that the vertical gaps between the windows can never all be the same. To maintain uniform exposure times and intervals between exposures, therefore, a pipeline delay is added to sequences 1 and 2. On completion of sequence 3, the cycle begins again by returning to sequence 1.

required to ensure that the two windows lie the same distance from the centre line of the chip is given in Section A2. The time it then takes to shift one row of the storage area into the serial register, perform the differential shift, horizontally shift the windows to the output, determine the charge content of each pixel and then digitize is given by:

$$t_{line1} = (b_y t_{vclock}) + (n_{hclock1} t_{hclock}) + \left(\frac{n_{lx}}{b_x} t_{video} \right), \quad (\text{A28})$$

where $n_{hclock1}$ is the total number of horizontal clocks required and is given by the sum of the differential shift, the window size and

the number of rows between the output and whichever window is closest to it.

The total time taken to read out one exposure is hence the sum of the following:

$$t_{cycle} = t_{delay} + t_{pipe} + t_{inv} + t_{frame} + \left(\frac{n_{ly}}{b_y} t_{line1} \right). \quad (\text{A29})$$

Since there is no clear, the exposure time is given by

$$t_{exp} = t_{cycle} - t_{frame}. \quad (\text{A30})$$

Hence the shortest possible exposure time obtainable is given by the time it takes to read the windows out. This can be reduced by

using smaller windows, binning and ensuring the windows are positioned and sized so as to minimize n_{hclock1} and t_{pipe} . This latter fact can be particularly important. For example, selecting windows of $n_{1y} = 24$ results in $n_{\text{win}} = 22$ and $t_{\text{pipe}} = t_{\text{velock}}$. Selecting windows of $n_{1y} = 23$, however, results in $n_{\text{win}} = 22$ but $t_{\text{pipe}} = 44t_{\text{velock}}$, that is, an additional exposure time of 1 ms. In this way, it can be seen that there are a series of ‘special’ values of n_{1y} in drift mode which satisfy, for example, $t_{\text{pipe}} < 14t_{\text{velock}}$ (i.e. an additional exposure time of < 0.33 ms). These special values are: $n_{1y} = 8, 10, 13, 18, 21, 24, 31, 38, 41, 49, 54, 60, 68, 79, 93, 114, 147, 206, 344$.

The dead time is given by

$$t_{\text{dead}} = t_{\text{cycle}} - t_{\text{exp}}, \quad (\text{A31})$$

as for the rest of the time the windows are integrating on the sky. This means that $t_{\text{dead}} = t_{\text{frame}}$. Since t_{frame} depends on n_{1y} in drift

mode, as opposed to n_{row} in the other ULTRACAM readout modes, the dead time is minimized.

It should be noted that GPS time-stamping in drift mode is slightly more complicated than in the non-drift modes. As described in Section 3.5, whenever an exposure is started the SDSU controller sends an interrupt to the instrument control PC which immediately writes the current time to the next frame of data received by the PC. In the case of drift mode, however, the next frame of data is not the most recently exposed window, but the window which happened to be at the bottom of the stack in the storage area. There is hence a loss of synchronization between time-stamp and data frame. This is easily rectified by the pipeline data reduction system, however, as it is a simple matter to calculate which window each time-stamp should be attached to using equation (A25).

This paper has been typeset from a \TeX/L\TeX file prepared by the author.