pt5m – a 0.5 m robotic telescope on La Palma

L. K. Hardy,[1] T. Butterley,[2] V. S. Dhillon,[1,3]★ S. P. Littlefair[1] and R. W. Wilson[2]

[1]*Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH, UK*
[2]*Centre for Advanced Instrumentation, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*
[3]*Instituto de Astrofisica de Canarias, E-38205 La Laguna, Tenerife, Spain*

ABSTRACT

pt5m is a 0.5 m robotic telescope located on the roof of the 4.2 m William Herschel Telescope (WHT) building, at the Roque de los Muchachos Observatory, La Palma. Using a five-position filter wheel and CCD detector, and bespoke control software, *pt5m* provides a high-quality robotic observing facility. The telescope first began robotic observing in 2012, and is now contributing to transient follow-up and time-resolved astronomical studies. In this paper, we present the scientific motivation behind *pt5m*, as well as the specifications and unique features of the facility. We also present an example of the science we have performed with *pt5m*, where we measure the radius of the transiting exoplanet WASP-33b. We find a planetary radius of $1.603 \pm 0.014 R_J$.

Key words: instrumentation: miscellaneous – instrumentation: photometers – telescopes – stars: individual: WASP-33 – planetary systems.

1 INTRODUCTION

Robotic observing facilities requiring no human input during night time are not new or uncommon,[1] and are becoming ever more popular (e.g. Zerbi et al. 2001; French et al. 2004; Tsapras et al. 2009; Gillon et al. 2011; Brown et al. 2013; Gorbovskoy et al. 2013; Andersen et al. 2014; Schmitt et al. 2014; Strolger et al. 2014). The largest example is the 2 m Liverpool Telescope (Steele et al. 2004), providing a fully autonomous, large aperture telescope at a high-quality site, with a suite of instrumentation available. Such facilities can sharply reduce operation costs as they require much less staff interaction. They can also be programmed to make complex scheduling decisions in order to observe the most appropriate targets, given the observing conditions. In addition, using automated transient alert-listening software, they can respond to unexpected events much faster than human-controlled telescopes.

pt5m[2] is a 0.5 m robotic telescope located on the roof of the William Herschel Telescope (WHT) on La Palma (Fig. 1) and is hosted by the Isaac Newton Group of Telescopes (ING). It functions as an astrophysics research facility, an atmospheric turbulence profiler, and a teaching resource. *pt5m* has already been used to study pulsating stars, transiting exoplanets, Solar system objects, interacting binaries (Littlefair et al. 2013; Kupfer et al. 2014; Campbell et al. 2015), and transients (Gandhi et al. 2014; Hardy et al. 2015), but no detailed description of *pt5m* itself has appeared in the refereed astronomical literature to date. In this paper, we therefore detail the hardware and software architectures (Sections 3 and 4), the telescope performance (Section 5), and an example photometric study of the transiting exoplanet WASP-33b (Section 6).

2 BRIEF HISTORY OF *pt5m*

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

In order to fully utilize the telescope when not being used for SLODAR observations, we decided to add a science imager in 2011. We recognized that even a small telescope located at such an excellent site could provide a powerful tool for time-domain astronomy. The science CCD and integrated filter wheel (see Section 3.2) were originally mounted on a linear slide to move the science camera in and out of the light path of the SLODAR optics. By 2012 we had implemented the robotic mode (Section 4) and started using *pt5m* on a routine basis for time-resolved and transient astronomy. In 2014, the original 0.5 m Orion Optics Modified Dall–Kirkham telescope was replaced with a new Ritchey–Chretien telescope (Section 3.3). This refurbishment was made to improve throughput and image quality, which were not a high priority for the original SLODAR studies. At the same time the slide mechanism was removed and replaced with a flip mirror that allows light to be diverted to either one of two optical ports. One of these ports hosts the science camera
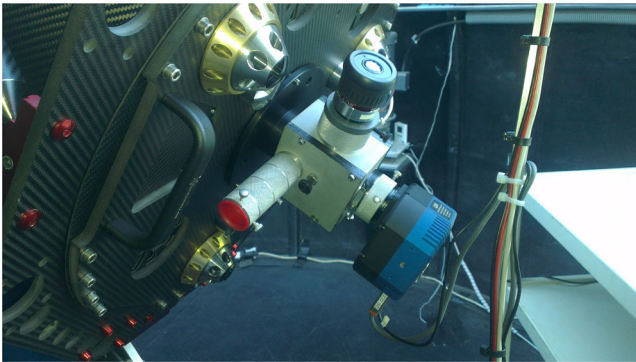

★ E-mail: vik.dhillon@sheffield.ac.uk
[1] For a global list of robotic telescopes, see http://www.astro.physik.uni-goettingen.de/~hessman/MONET/links.html
[2] The name is an abbreviated form of 'point-five-metre'.
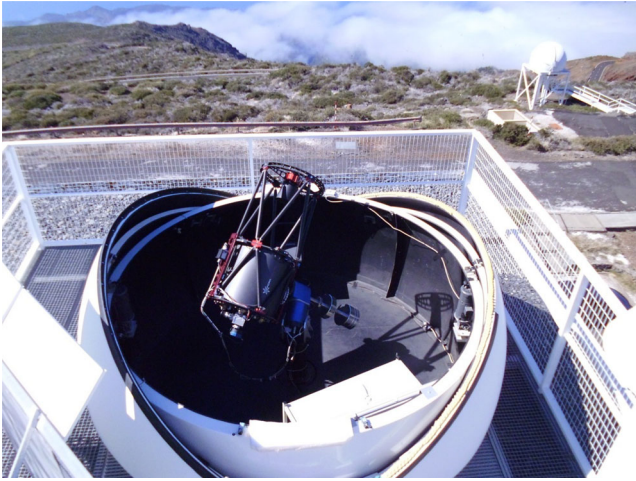
Figure 1. Photograph of *pt5m* on the roof of the WHT taken in April 2015. The WHT dome is out of view to the left.

Figure 2. Photograph of the Cassegrain focus of *pt5m*, showing the flip mirror mechanism (grey), and the CCD camera (blue) with integral filter wheel (black). In this photo, the SLODAR port (pointing upwards) temporarily holds an eyepiece.

and the other hosts the atmospheric turbulence profiler (see Fig. 2 and Section 3.3).

3 HARDWARE

The following subsections describe the hardware components of *pt5m*. Fig. 3 shows the components schematically, whilst Table 1 summarizes the facility specifications.

3.1 Dome and mount

pt5m is enclosed in a 12-foot Astrohaven clamshell dome. The use of a clamshell design is vital for rapid follow-up work, as there is no delay whilst the dome rotates to catch up with the telescope.

The telescope mount, a direct-drive Astelco Systems NTM-500, is of German-Equatorial design. This mount design requires a 'pier-flip' when tracking a target across the observer's meridian from east to west. The pier flip introduces a discontinuity in the *x*-axis of time series observations of up to several minutes, as the mount must wait for the target to pass through the observer's meridian before reacquiring. The pier flip also rotates the field of view through 180°, placing the target and comparison stars on different parts of the chip, with different levels of vignetting. This results in a small step in the differential photometry which, in principle, can be corrected for
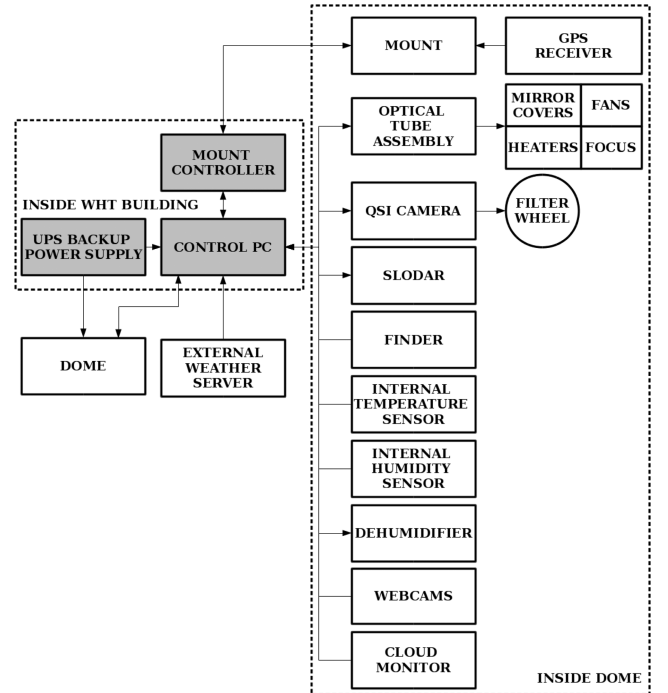
Figure 3. Schematic showing the *pt5m* hardware components and their connections.

with flat-fields. The only major negative effect of the pier flip is when it occurs in the middle of an interesting event, e.g. during an eclipse.

In its implementation in *pt5m*, the mount has an RMS pointing accuracy of $\sigma = 0.4$ arcmin around the mean pointing position, and an average tracking drift of 12 arcsec h^{-1}. These values are larger than expected from the manufacturer's specification, most probably due to flexure in the telescope and roof-top foundations, and small errors in the polar alignment.

Although it is stable during normal use, both the telescope and mount suffer from vibrations in high winds (8–10 m s^{-1}), as the clamshell dome leaves them exposed. They also vibrate when the main WHT dome motors are engaged, as the facility's roof-top foundations are not fully isolated from the building. We are in the process of erecting substantial wind baffling around the dome, particularly on the south-west side where the wind is forced between the WHT dome and the rest of the building.

An advantage of this particular mount is its fast slewing speed of 20° s^{-1}.[3] This allows new targets to be acquired very quickly, which is essential for rapid follow-up work, as well as reducing overheads between observations.

3.2 Science camera

The 'Quantum Science Imaging' QSI-532ws camera uses a Kodak KAF-3200ME CCD detector. The science camera is equipped with an integral five-position filter wheel, which at present holds the following Astrodon filters: Johnson *B*, Johnson *V*, Cousins *R*, Cousins *I*, and Hα 5 nm. The camera employs a thermoelectric cooling system, and is maintained at -20°C when in use. The CCD specifications are listed in Table 1. Of particular note is the low dead time of the CCD (5 s in 2 × 2 binned mode), which reduces

---

[3] In fact, with a smaller payload, the mount can slew at up to 100° s^{-1}.

Table 1. Specifications of *pt5m*. Location information is as measured by GPS (WGS84).

Parameter	Value
Telescope	
Longitude	28°45′38.7″N
Latitude	17°52′53.2″W
Altitude	2383 ± 3 m
Primary diameter	0.5 m
Primary focal length	1500 mm
Primary conic constant	−1.071
Secondary diameter	0.148 m
Secondary focal length	−609.89 mm
Secondary conic constant	−4.182
Dome diameter	3.66 m
Focal length	5000 mm
Platescale	41.25 arcsec mm^{-1}
Telescope unvignetted FoV	9 arcmin
Slew speed	20° s^{-1}
Pointing accuracy (RMS)	0.4 arcmin
Tracking accuracy	12 arcsec h^{-1}
Finder field of view	7°.2 × 5°.4
Instrument	
Detector format	2184 × 1472 pixels
Pixel size	6.8 μm
Pixel scale	0.28 arcsec pixel^{-1}
CCD field of view	10.2 arcmin × 6.9 arcmin
Full well depth	83 000 e^{-}
Readout noise (binned 1 × 1)	10.4 e^{-}
Readout noise (binned 2 × 2)	14.4 e^{-}
Readout noise (binned 3 × 3)	18.5 e^{-}
Gain	1.3 e^{-}ADU^{-1}
Flat-field noise (grain)	0.15 ± 0.05 per cent
Dark current	8 ± 4 e^{-} pixel^{-1} h^{-1}
Dead time (binned 1 × 1)	8 s
Dead time (binned 2 × 2)	5 s
B-band zero-point	21.91 mag
V-band zero-point	22.06 mag
R-band zero-point	22.16 mag
I-band zero-point	21.05 mag

overheads and improves the cadence of time series observations. The CCD camera uses a shutter, and takes timing information directly from the control PC (see Section 3.5). The timestamps are moderated by the Network Time Protocol which has a typical accuracy of <100 ms. As well as a science camera, *pt5m* also has an ASTROVID StellaCam finder camera mounted on the telescope.

3.3 Telescope

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

A tertiary mirror is housed in a *Van Slyke Instruments Flipper* (Fig. 2). It uses a mirror on a hinged support that flips up and down. This allows the light beam to either pass straight through to the filter wheel and CCD or to be folded by 90° into the port dedicated to atmospheric turbulence profiling.

3.4 Weather system

In order to keep *pt5m* safe from weather damage, a reliable weather station, with backup systems available in case of failure, is essential. Since *pt5m* is hosted by the ING, we are able to benefit from access to their weather stations.[4] These include the main weather system for the whole observatory located on a mast close to the Jacobus Kapteyn Telescope (JKT), as well as a local mast located on the WHT. In addition, *pt5m* employs its own internal sensors, including a humidity and temperature sensor located in the dome. The Conditions Monitor (see Section 4.4) monitors the weather information and decides when the dome should be open or closed. The dome is forcibly closed if this weather data becomes unavailable, or returns exactly the same values for a period of more than 20 min (we assume this only happens if the ING weather server has crashed and is not supplying reliable data).

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

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

3.5 Electronics

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

pt5m is also equipped with three physical emergency stop buttons. The first stops the telescope and mount, engaging the brakes immediately to stop any potential damage if the telescope is about to crash into the pier, for example. This button is mounted on the telescope mount. The second button is mounted on the interior dome wall, near to the access hatch. This button shuts off power to the dome motors to prevent accidents when climbing through the hatch. The third button is located inside the WHT elevator control

---

[4] http://catserver.ing.iac.es/weather/index.php
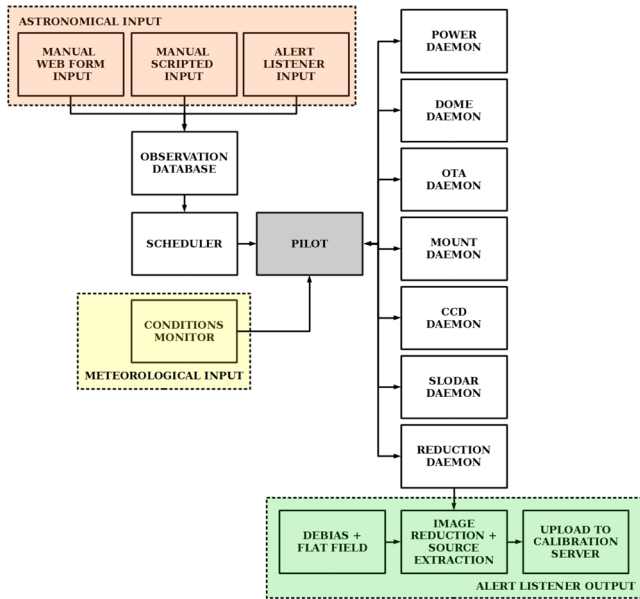
Figure 4. Schematic showing the software components of *pt5m* and their interaction.

room, and is used to shut down the telescope and close the dome in the event of software malfunction or severed connectivity to the external network.

4 SOFTWARE

All of the software used to control *pt5m* is written in PYTHON and runs on the LINUX control PC (Section 3.5). *pt5m* can be used in either 'manual' or 'robotic' mode. In manual mode, a human user controls the facility in real time by sending a series of instructions via a command-line terminal that has been opened on the control PC. This mode is used for testing and teaching purposes, and can be used to control *pt5m* locally or remotely over the Internet. In robotic mode, *pt5m* operates autonomously, opening and closing the dome, monitoring the weather, taking and processing calibration data, and selecting and observing science targets.

A schematic showing the interactions between the *pt5m* software components is shown in Fig. 4.

4.1 Hardware control: daemons

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

The *pt5m* hardware daemons are as follows:

(i) Power: communicates with the remote-controlled power switches (Section 3.5), used to power the individual hardware components on and off. This daemon runs at all times.

(ii) Dome: controls opening and closing of the dome via an Arduino interface. It constantly checks the current weather conditions via the Conditions Monitor (Section 4.4). It will automatically close

the dome (or refuse to open it) if any of the Conditions Monitor flags are set to 1, if the external power fails or if the network connection to the outside world fails. This daemon runs at all times.

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

(iv) Mount: communicates with the telescope mount controller via an internal network. This daemon controls the pointing and tracking of the telescope. It also constantly monitors a wide range of status information that is available from the mount and makes this available to other software processes.

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

4.2 System control: pilot

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

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

(1) Sun at 10°: power on CCD and mount.
(2) Sun at 7°: take bias frames (dome still closed).
(3) Bias frames complete; open dome and then open mirror covers.
(4) Sun at 5°: point telescope to a blank field in the sky and take flats.
(5) Focus the telescope using a bright star.
(6) Process flats and biases into master frames if enough individual frames were taken in good conditions, i.e. no cloud presence detected during flat-field observations.
(7) Optionally, observe flux standards and collect dark frames.
(8) Park the telescope and wait for the end of twilight.

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

The pilot is considered to be in the night-time state while the Sun is below $-10°$ altitude. During this time, its behaviour is governed by the 'scheduler' (see Section 4.3). The pilot periodically queries the scheduler and, if the scheduler suggests a new target should be observed, the pilot slews the telescope and adds the specified exposures to the CCD queue. Once the pilot verifies that the

telescope is tracking on the target, the CCD exposures begin. The system then remains in this state, taking exposures, until one of the following occurs:

(i) all exposures in the CCD queue have been taken and the observation is complete,

(ii) the scheduler returns a new, higher priority target to observe,

(iii) conditions force the dome to close, or

(iv) the end of the night is reached.

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

In the event of bad weather, the dome is closed immediately by the dome daemon (see Section 4.1) without the intervention of the pilot. This is an essential precaution to protect the instrumentation from wind and water damage as the weather conditions can change quickly. The pilot then responds as follows: if the telescope is currently observing a valid target, it will pause the queue of CCD exposures and leave the telescope tracking. If the conditions improve after a short time, then the pilot simply reopens the dome and resumes the CCD exposures. If/when the current observation becomes invalid (e.g. the target sets below the horizon), the pilot parks the telescope and clears the CCD queue.

The morning-twilight shutdown procedure is as follows:

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

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

(3) Power off mount and CCD.

4.3 The scheduler

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

Each new target (also referred to as a 'job' or 'pointing') is described by an ASCII text file that defines the position of the target, the CCD exposures required, the priority (an integer from 0 to 5, where 0 is the highest priority), and various other constraints on the required observations, as shown in Fig. 5. Only one target per job is allowed. Each job is given a unique ID number, which is also the name of the ASCII file. The jobs define the 'queue', and all of the awaiting job files are stored in a single directory on the control PC. When a job is complete, the file is moved from the queue directory to a 'completed' directory. Alternatively, if the job is aborted, deleted, expired, or interrupted, it is moved to a correspondingly named directory.

Table 2. Flags used by the scheduler. Each flag can only take the value 0 (legal) or 1 (illegal). The artificial horizon refers to a map of minimum observable altitude for any given azimuth, which ensures that *pt5m* never tries to observe targets which are vignetted by the WHT dome or building. See Fig. 5 for an explanation of the terms StartUTC, StopUTC, Sun Alt, Min.Alt., and Min.Time.

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

Each time it is polled, the scheduler runs through every job in the queue and calculates a set of flags for each job, with a result of 0 indicating the flag is legal and 1 indicating it is illegal. An explanation of each of the flags is listed in Table 2. The job is legal only if all the flags are 0. The scheduler then updates the priority of each job in the queue as follows:

(i) If the job is defined as a Target of Opportunity (ToO = 1 in Fig. 5), its priority is given by the user-defined integer priority (priority in Fig. 5) plus airmass/100, where the airmass of the target is calculated at a point in time mid-way between the current time and the minimum end time of the job (Min.Time in Fig. 5). Hence, a priority 2 job with a current airmass of 1.234 becomes 2.01234.

(ii) If the job is not a ToO (ToO = 0 in Fig. 5), its priority is given by the user-defined integer priority plus airmass/10. Hence, a priority 2 job with a current airmass of 1.234 becomes priority 2.1234.

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

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

The scheduler takes the job with the highest priority (equal to the lowest numerical value) in the queue and compares it with the

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

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

updated priority of the job currently being observed. On the basis of this comparison, the scheduler decides what to do next as follows:

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

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

(3) If the current observation is no longer the highest priority in the queue, and the highest priority job *is* a ToO, interrupt the current observation by aborting the current exposure and immediately slewing to the new target. This rule guarantees that time-critical observations can be obtained. Note that if the current observation is also a ToO, it will not be interrupted unless the interrupting ToO lies in a higher priority band: for example, if the current observation is a ToO of priority 3.01234, it would not be overridden by a ToO job of priority 3.01000, but would be overridden by a ToO job of priority 2.01234.

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

4.4 Conditions Monitor

The Conditions Monitor is a PYTHON script that reads the weather information provided by up to three online sources from the ING, and is updated every 60 s. These sources record the meteorological conditions at three different locations at the Observatorio del Roque de los Muchachos: the weather masts at the JKT, Isaac Newton Telescope and WHT. The various parameters that are read from each mast are listed in Table 3, along with the limits defining when the dome should be closed. In order to prevent repeated opening and closing of the dome when conditions are close to any of the limits, we define different limits for when the dome should be reopened following a closure due to bad weather. These reopening limits are also listed in Table 3.

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

Occasionally, there is a problem with the JKT mast, and one or more of the meteorological parameters reported by the system

Table 3. Weather parameters considered by the Conditions Monitor, with different limits applicable depending on whether the dome is open or closed. The parameters are more lenient when the dome is already open compared to when the dome is closed, as the latter is likely to be because of recent bad weather, and some additional contingency is needed. This also helps prevent the dome from repeatedly opening and closing when conditions oscillate around the limits. The final two parameters (Max Sun and Max Moon) obviously do not affect the opening and closing of the dome, but dictate when observing can commence.

Parameter	Limit (open)	Limit (closed)
Wind speed (m s^{-1})	10.0	9.0
Wind speed (km h^{-1})	36.0	32.4
Relative humidity (per cent)	70	63
Internal humidity (per cent)	70	63
Dry/wet sensor	DRY	DRY
Temperature (Celsius)	1.0	2.0
Age of weather data (s)	700	700
Max time for parameters to be unchanged (s)	1200	1200
Max time without contact to Durham server (s)	300	30
Max Sun altitude (°)	−10	−10
Max Moon altitude for dark time (°)	−5	−5

is either unavailable, corrupted, or not updated. In this case, the Conditions Monitor automatically switches to taking its input from one of the other masts. If this mast is also in error, the Conditions Monitor will shut the dome. The Conditions Monitor also continuously checks that the internet link between *pt5m* and the UK is working. If the link is broken, the Conditions Monitor will shut the dome. Once again, this is adopting a conservative approach: with no internet link, there is no possibility of us logging in remotely to shut the dome in the case of a problem. The system also notes the timestamps of all weather data packets, and treats the data as unreliable if it is more than 700 s old, or if the values do not change for more than 20 minutes.

The Conditions Monitor works by setting flags. All of the parameters listed in Table 3 have an associated flag, which is set to 0 when the parameter is legal (i.e. it is safe to open the dome) and 1 when the parameter is illegal (i.e. it is not safe to open the dome). Only when all of the flags are 0 can the dome be opened, and if any of the flags are subsequently set to 1, the dome will close. When observing locally, and when sure that it is safe to do so, it is possible to override this system and allow the dome to remain open even if one or more of the flags are set to 1.

We have analysed the fraction of time that the dome remains open during night time hours, given the current weather parameters. Between 2013 January 20 and 2015 September 17 (971 nights), 21 per cent of night time experienced high humidity, 17 per cent experienced high wind speeds, and 10 per cent experienced low temperatures. These fractions will overlap with each other, as often high humidity is coupled with high winds and/or low temperatures. Overall, 32 per cent of night time was lost because of adverse weather, while 5 per cent was lost due to incomplete or corrupted meteorological information. Were we to be more liberal with our weather limits, e.g. setting the humidity limit to 80 per cent, the wind speed limit to 11 m s^{-1}, and the temperature limit to 0°C for both open and closed dome scenarios, we could expect the fractions of lost night time to drop to 14 per cent, 11 per cent, and 6 per cent for humidity, wind speed, and temperature, respectively. Despite this, we choose to maintain our current weather limits, as the risk of equipment damage outweighs the marginal gains made in observing time.

4.5 Observation database

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

Submission of jobs to the SQL data base is performed either via a web form[5] which interfaces to the data base through PHP scripts, or by custom-written PYTHON scripts which can be run anywhere and interface to the same PHP scripts. The PHP scripts send the job details to the SQL data base, which writes the ASCII files (Fig. 5) to be picked up by the scheduler. Access to the data base requires a username and password, which we administer directly to collaborators.

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

4.6 Alert Listener

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

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

The data reduction component of the Alert Listener is a software daemon (Fig. 4) that automatically processes every science frame as soon as it has been read out. The following tasks are performed:

(i) Debias and flat-field using the master calibration frames (Section 4.2). If no master frames are available from that night, the daemon searches backwards through previous nights until a master frame is found, up to a maximum lookback time of 30 nights.

(ii) Using the SEXTRACTOR software, extract a catalogue of all stellar sources in the image.

(iii) Match the extracted catalogue to published *Vizier* catalogues (Ochsenbein, Bauer & Marcout 2000) and conduct an astrometric calibration, allowing for rotation, shift, and stretch.

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

The third component of the Alert Listener automatically publishes the transient data. Each extracted catalogue associated with a transient alert is automatically uploaded to the Gaia Science Alerts Calibration Server[6] which calculates and publishes a calibrated magnitude for the transient. We use this server because we have been providing photometric follow-up for the Gaia Science Alerts project (see e.g. Campbell et al. 2015), and it is a useful repository for transient data.

4.7 Autoguiding

Since *pt5m* does not have a separate autoguider, we have implemented autoguiding using the science frames. For each new image, the catalogue of sources extracted by the data reduction daemon (Section 4.6) is compared with that of the first science frame in the sequence. A minimum of 11 stars are needed for an accurate comparison; any fewer than this and the autoguiding process is skipped for that image. The median offset between the catalogues is then used to adjust the telescope position. The offset is applied during the readout of the following image. Guiding can be enabled or disabled at any time and is automatically re-established following a pier flip. The downside to our autoguiding implementation is that when particularly long exposures are needed, the guiding is unable to correct for tracking errors on shorter time-scales than the exposure time. Equally, when very short cadence observations are required (<10 s including exposure time and read out), the data reduction process can sometimes cause minor delays. However, observations of such short cadence are very rare. Section 5.4 describes the performance of the telescope tracking with and without autoguiding enabled.

5 PERFORMANCE

All of the performance data given in this section refer to the final configuration of the telescope described in Section 2.

5.1 Detector

The CCD readout noise, gain, and flat-field noise were measured from bias frames and dome flats using the photon transfer curve technique (Janesick 2001) and are shown in Table 1. The gain is unchanged when binning, but the readout noise depends on the binning factor. The flat-field noise is the same in all filters, to within the measurement uncertainties, suggesting the dominant source of this noise is variations in the physical size of the pixels, rather than some wavelength-dependent response.

The dark current was measured at night with the dome, mirror, and camera shutters closed. Exposure times of 20 min were used, and the dark current was measured as the median count level after the median bias level (measured only minutes earlier) had been subtracted. The detector shows negligible dark current (8 ± 4 e$^-$pixel^{-1} h^{-1}), but this tends to be concentrated in hot pixels and hence it can be important to use median-combined dark frames when reducing data.

By observing a star field with a range of exposure times in quick succession, we were able to analyse the linearity of the detector. We found that the detector performs linearly to within 0.5 per cent at light levels of up to 45 000 counts pixel^{-1}, but departs from linearity by up to 2 per cent at higher light levels.
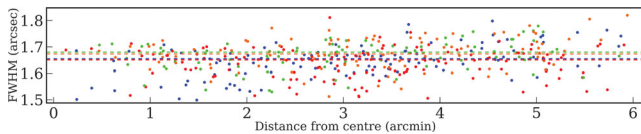
Figure 6. FWHM in the *R* band as a function of radial distance from the centre of the CCD, measured from an image of the open cluster IC 4996. The standard deviation of the FWHM is 0.06 arcsec in this image. The four colours represent four different quadrants of the CCD: blue is top right, red is bottom right, green is top left, and orange is bottom left. The horizontal lines represent the median FWHM in each quadrant.

5.2 Optics

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

The image quality appears to be largely independent of wavelength; the difference in FWHM from filter to filter is less than 10 per cent (approximately 2σ). The image quality degrades by approximately 0.05 arcsec from the centre to the corners of the CCD, as shown in Fig. 6, in accordance with the expectation from ray tracing of the telescope optics.

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

5.3 Zero-points

The photometric zero-points in the four broad-band filters were measured using the standard stars SA101-315 and SA102-620. Each star was observed twice on two separate photometric nights, at different airmasses. We corrected for extinction using the coefficients given in the ING Observers Manual, 1995.[7] The above-atmosphere zero-points representing the magnitude of a star giving one photoelectron per second in each filter were then calculated. The final zero-points, representing the mean of all four standard star measurements in each filter, are shown in Table 1.

Fig. 7 shows the 5σ limiting magnitudes that *pt5m* can achieve as a function of exposure time and moon brightness, calculated using the measured zero-points.

5.4 Pointing, tracking, and guiding

The pointing accuracy of *pt5m* is 0.4 arcmin (RMS), with marginally worse pointing at low elevations and high declinations. The tracking accuracy of *pt5m* was measured from observations following the same targets for 4 h. The worst measured tracking drifts were up to 30 arcsec h^{-1} for certain areas of the sky (at low elevations and at high declinations), but on average tracking drifts of 12 arcsec h^{-1}

[7] http://www.ing.iac.es/astronomy/observing/manuals/ps/general/obs
_guide.pdf
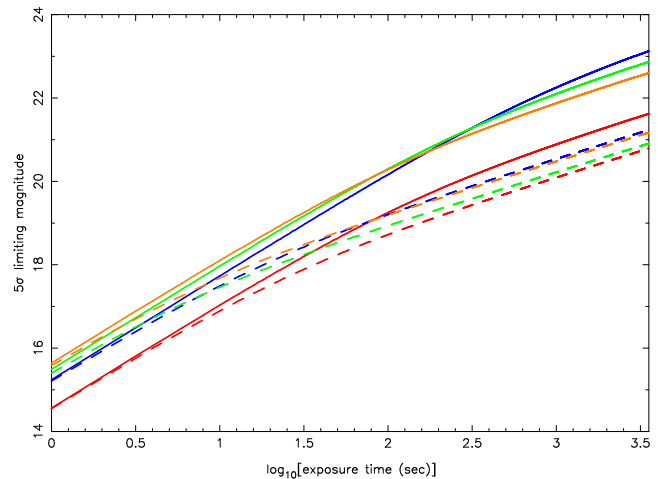
Figure 7. 5σ limiting magnitudes of *pt5m* as a function of exposure time. The blue, green, orange, and red curves show the results for the *B*, *V*, *R*, and *I* filters, respectively. Solid lines show the results for dark time and dashed lines for bright time. The calculations assume seeing of 1.5 arcsec, airmass of 1.0 and 2 × 2 binning.

were seen. This drift is larger than expected according to the mount specifications, most probably due to small errors in the mount's polar alignment, and flexure in the telescope and roof-top platform. The drift limits exposure times to a maximum of 5 min. With the introduction of autoguiding (Section 4.7), positional fluctuations of only 1.3 arcsec (RMS) from the original position are seen.

5.5 Response times

It takes less than 2 s from the announcement of a gamma-ray burst as a VOEvent for the Alert Listener (Section 4.6) to submit a job for follow-up observation to the *pt5m* queue. The pilot checks the scheduler for the highest priority job every 15 s. We estimate a maximum slew time of 30 s, assuming that the telescope has to conduct a pier flip and move to the opposite side of the sky. Adding these together we can expect to be tracking on target less than a minute after an alert is announced. Unfortunately, we are yet to receive an observable night time trigger during safe weather conditions.

5.6 Residual images

In October 2013, we serendipitously observed an apparent transient event whilst monitoring the active blazar 4C+38.41 (Gasparrini 2013). The bright (*B* ∼ 16) transient was observed in the same field, but was located several arcminutes away and was not thought to be associated with the blazar. The transient appeared to fade quickly (0.2 mag min^{-1}) over the course of around 15 min, and moved slowly (0.5 arcsec min^{-1}) with respect to background stars. We could not identify the transient as a known comet, asteroid, or minor planet, and the coordinates did not coincide with the orbit of any known man-made satellite. The object appeared to have a stellar point spread function. Bizarrely, the transient reappeared two weeks later, 18 arcsec away from the original position, and this time appearing as a double source.

After exhausting all possible astrophysical origins, we later discovered that this fast-fading transient signal was in fact the residual image of a bright star used for focusing just minutes earlier. Photoelectrons appear to have been trapped in impurity sites in the
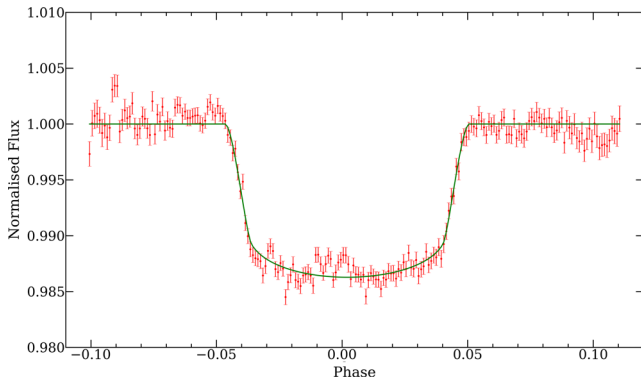
Figure 8. Transit light curve of WASP-33b, comprised of 15 individual transits. The observations have been phase folded, binned, and combined using inverse-variance weighting. The solid curve shows our best-fitting model. See the text for details.

substrate of the chip. After the CCD was read out and cleared, these trapped electrons gradually leaked out during the next few exposures, creating a spurious signal in the resulting images.

This residual image phenomenon is not uncommon in infrared detectors, but is rare in optical detectors (Rest et al. 2002). Our Kodak KAF-3200ME CCD is known to exhibit this behaviour.[8] Clearly, care must be taken when analysing images taken immediately after exposures containing bright stars.

6 EXAMPLE SCIENCE

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

We observed a total of 15 transits of WASP-33b in 2011 and 2012. We also collected several hours of out-of-transit data in order to study the stellar pulsations. The observations were phase folded using our own ephemeris, binned and combined using inverse-variance weighting. The combined transit light curve is shown in Fig. 8. We model the planetary transit using the geometrical analysis of Sackett (1999), finding a planetary radius of 1.603 ± 0.014 R_J and an orbital inclination of $87°04 \pm 0°26$. Uncertainties were estimated using the 'prayer-bead' technique of residual permutation (Désert et al. 2009).

Our measured inclination matches well with previous studies (Collier Cameron et al. 2010; Herrero et al. 2011; Smith et al. 2011; Kovács et al. 2013; von Essen et al. 2014), but the planetary radius is found to be up to 3σ larger. This highlights the fact that WASP-33b is substantially bloated in comparison to other exoplanets. Our measurement of the radius is plotted on a mass–radius diagram and compared to theoretical models of exoplanetary structure in Fig. 9. The measured radius is almost 20 per cent larger than expected from planetary models (Fortney, Marley & Barnes 2007), although it should be noted that these models are not tailored to such a close-in exoplanet as WASP-33b.

The origin of an enlarged planetary radius could be attributed to a number of factors, including tidal heating (Bodenheimer, Lin &
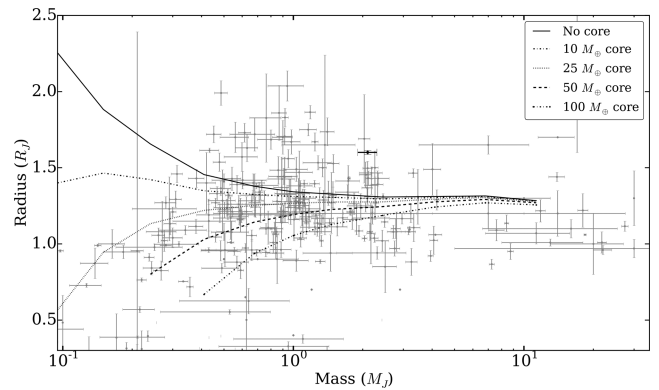
Figure 9. Mass–radius diagram showing the known exoplanets with measured or estimated masses and radii, as of July 2015 (Schneider et al. 2011), plotted as grey points. WASP-33b is shown in black, with the planetary mass measured by Lehmann et al. (2015). Also shown are planetary models of a H/He planet of age 300 Myr and a scaled orbital separation of 0.02 au, with five different core masses (Fortney et al. 2007).

Mardling 2001), Ohmic heating (Batygin & Stevenson 2010), and kinetic heating (Guillot & Showman 2002). The circular orbit of WASP-33b (Smith et al. 2011; Deming et al. 2012) rules out tidal heating, but Ohmic and kinetic heating may play a role. However, we propose an additional and simpler reason for its enlarged radius, the young age of the system. The age of WASP-33 could be as young as 10 Myr (Moya et al. 2011), making it one of the youngest known exoplanet hosts. The size of an exoplanet is expected to decrease with time after formation, as the planet cools and contracts (Fortney et al. 2007). In addition, standard planetary models usually use arbitrary starting conditions, with their main scientific interest being the state of the planet some gigayears after formation, by which time their starting conditions have little effect (Fortney & Nettelmann 2010). In this sense, it is not surprising that WASP-33b is found to be bloated with respect to these models, as it is much younger than the planets the models are designed to simulate.

Our phase-folded light curve (Fig. 8) shows a mid-transit bump, which is present in several of the individual transit light curves. The star spot explanation for this phenomenon proposed by Kovács et al. (2013) is unlikely, as the host is an A-type star. Instead, we propose that the bump may be due to a pulsation in the host star resonating with the orbital period, and therefore is still visible after phase folding. However, von Essen et al. (2014) find no conclusive evidence for such star–planet interactions, and we found no pulsation periods in resonance with the orbital period in power spectra of the out-of-transit light curves. The nature of the bump remains unclear.

7 CONCLUSIONS

We have described the design and performance of *pt5m* – a 0.5 m robotic telescope on the roof of the WHT on La Palma. The facility is primarily used for research, and we present an example study of the transiting exoplanet WASP-33b performed with *pt5m*. The facility also functions as a valuable teaching resource, and is used regularly for undergraduate projects at both the University of Sheffield and Durham University. Interested readers are encouraged to visit the *pt5m* website[9] for further information.

[8] http://canburytech.net/QSI532/RBI.html

[9] https://sites.google.com/site/point5metre/

REFERENCES

Andersen M. F. et al., 2014, Rev. Mex. Astron. Astrofis. Ser. Conf., 45, 83

Batygin K., Stevenson D. J., 2010, ApJ, 714, L238

Bodenheimer P., Lin D. N. C., Mardling R. A., 2001, ApJ, 548, 466

Brown T. M. et al., 2013, PASP, 125, 1031

Campbell H. C. et al., 2015, MNRAS, 452, 1060

Catala L. et al., 2013, MNRAS, 436, 590

Collier Cameron A. et al., 2010, MNRAS, 407, 507

Deming D. et al., 2012, ApJ, 754, 106

Désert J.-M., Lecavelier des Etangs A., Hébrard G., Sing D. K., Ehrenreich D., Ferlet R., Vidal-Madjar A., 2009, ApJ, 699, 478

Fortney J. J., Nettelmann N., 2010, Space Sci. Rev., 152, 423

Fortney J. J., Marley M. S., Barnes J. W., 2007, ApJ, 659, 1661

French J. et al., 2004, in Fenimore E., Galassi M., eds, AIP Conf. Proc. Ser. Vol. 727, Gamma-Ray Bursts: 30 Years of Discovery. The Sun. Am. Inst. Phys., New York, p. 741

Gandhi P., Dhillon V. S., Tomsick J. A., Butterley T., Littlefair S. M., Wilson R. W., Kennea J. A., 2014, Astron. Telegram, 6447, 1

Gasparrini D., 2013, Astron. Telegram, 5232, 1

Gillon M., Jehin E., Magain P., Chantry V., Hutsemékers D., Manfroid J., Queloz D., Udry S., 2011, EPJ Web Conf., 11, 6002

Gorbovskoy E. S. et al., 2013, Astron. Rep., 57, 233

Guillot T., Showman A. P., 2002, A&A, 385, 156

Hardy L., Littlefair S., Dhillon V., Butterley T., Wilson R., 2015, Astron. Telegram, 7681, 1

Herrero E., Morales J. C., Ribas I., Naves R., 2011, A&A, 526, L10

Janesick J. R., 2001, Scientific Charge Coupled Devices. SPIE Press, Washington

Kovács G. et al., 2013, A&A, 553, A44

Kupfer T. et al., 2014, in van Grootel V., Green E., Fontaine G., Charpinet S., eds, ASP Conf. Ser. Vol. 481, 6th Meeting on Hot Subdwarf Stars and Related Objects. Astron. Soc. Pac., San Francisco, p. 293

Lehmann H., Guenther E., Sebastian D., Döllinger M., Hartmann M., Mkrtichian D. E., 2015, A&A, 578, L4

Littlefair S. P. et al., 2013, MNRAS, 431, 2820

Morris T. et al., 2009, in Assoufid L., Quarles G., Testorf M., eds, Frontiers in Optics 2009/Laser Science XXV/Fall 2009 OSA Optics & Photonics Technical Digest, CANARY: An On-Sky Laser Guide Star Multiple Object AO Demonstrator. Optical Society of America, Washington, DC, p. AOTuA5

Moya A., Bouy H., Marchis F., Vicente B., Barrado D., 2011, A&A, 535, A110

Ochsenbein F., Bauer P., Marcout J., 2000, A&AS, 143, 221

Rest A., Mndermann L., Widenhorn R., Bodegom E., McGlinn T. C., 2002, Rev. Sci. Instrum., 73, 2028

Sackett P. D., 1999, in Mariotti J.-M., Alloin D., eds, NATO ASI Series C, Vol. 532, Planets Outside the Solar System: Theory and Observations. Kluwer, Dordrecht, p. 189

Schmitt J. H. M. M. et al., 2014, Astron. Nachr., 335, 787

Schneider J., Dedieu C., Le Sidaner P., Savalle R., Zolotukhin I., 2011, A&A, 532, A79

Seaman R., Allan A., Williams R., 2008, in Argyle R. W., Bunclark P. S., Lewis J. R., eds, ASP Conf. Ser. Vol. 394, Astronomical Data Analysis Software and Systems XVII. Astron. Soc. Pac., San Francisco, p. 213

Smith A. M. S., Anderson D. R., Skillen I., Collier Cameron A., Smalley B., 2011, MNRAS, 416, 2096

Staley T. D. et al., 2013, MNRAS, 428, 3114

Steele I. A. et al., 2004, in Oschmann J. M., Jr, ed., Proc. SPIE Conf. Ser. Vol. 5489, Ground-based Telescopes. SPIE, Bellingham, p. 679

Strolger L. G. et al., 2014, AJ, 147, 49

Tsapras Y. et al., 2009, Astron. Nachr., 330, 4

von Essen C. et al., 2014, A&A, 561, A48

Wilson R. W., 2002, MNRAS, 337, 103

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

Zerbi R. M. et al., 2001, Astrono. Nachr., 322, 275

This paper has been typeset from a TEX/LATEX file prepared by the author.