Third AO4ELT Conference - Adaptive Optics for Extremely Large Telescopes
Florence, Italy. May 2013
ISBN: 978-88-908876-0-4
DOI: 10.12839/AO4ELT3.13189

REAL-TIME CONTROL SYSTEM VERIFICATION FOR ELT AO SYSTEMS

Alastair Basden[1a], Richard Myers[1], Tim Morris[1], Ali Bharmal[1], Urban Bitenc[1], Nigel Dipper[1], Andrew Reeves[1], Eric Gendron[2], Gerard Rousset[2], Zoltan Hubert[2], Fabrice Vidal[2], Olivier Martin[2], Damien Gratadour[2], and Fanny Chemla[2]

[1] Durham University, Department of Physics, South Road, Durham, DH1 3LE, UK
[2] Observatoire de Paris, LESIA, Paris, France

Abstract. ELT AO systems have demanding computational requirements for real-time control. These systems are required to be fully tested and robust before commissioning so that valuable on-sky time is not wasted. In this talk I will report recent work at Durham on our ELT AO real-time control system, algorithms that we use to improve robustness, and development of an end-to-end testing environment that will allow full testing of real-time control systems, including both Monte-Carlo simulation and hardware approaches. The talk will include experience gained with CANARY, how the robustness of this system has been improved, and our experience operating with four laser guide stars. Work carried out in this area on the DRAGON test-bench will also be described.

1 Introduction

Adaptive optics (AO) systems are required for Extremely Large Telescope (ELT) operation. A key part of an AO system is the real-time control system (RTCS), which is a highly complex component. For ELT systems, it is essential that the AO RTCS should be fully tested and robust before on-telescope commissioning so that valuable on-sky time is not wasted (at an estimated cost of about €100k per night on the European ELT (E-ELT)). However, these proposed systems are more complicated than existing RTCSs, due to the requirement for higher order correction, multiple laser guide stars (LGSs) and wide-field correction. A key question that requires addressing during the design of these systems is how to ensure that they are robust to external perturbations. Of equal importance, we need to be able to test the correctness of real-time implementations of algorithms before the RTCS reaches the telescope which is not trivial because real-time algorithm implementations can be significantly different from those used during algorithm development which are usually developed without computational efficiency in mind.

1.1 ELT real-time performance

It has been shown [1,6] that AO real-time control performance for the ELTs can be achieved using conventional commercial off-the-shelf (COTS) hardware. As an example, a single arm of the ELT Adaptive optics for GaLaxy Evolution (EAGLE) instrument (11 wavefront sensors (WFSs) all with 80x80 sub-apertures and 1 multi-object AO (MOAO)

---

[a] a.g.basden@durham.ac.uk

deformable mirror (DM)) has been demonstrated at 300 Hz using two PCs each containing four graphical processing units (GPUs), with less than 3.3 ms latency. In this case, the arms of EAGLE (of which there will be 20) are independent, so performance scaling is not required, rather the problem is simply duplicated 20 times. In this particular case, wavefront reconstruction was performed using GPU, which image calibration and slope calculation was performed in central processing unit (CPU).

Likewise, real-time control for the E-ELT multi-conjugate adaptive optics (MCAO) module multi-conjugate AO relay for the E-ELT (MAORY) has been shown to be achievable, with demonstration of about a sixth of the system achieved on a single PC with four GPUs (one 80x80 WFS and three DMs with a total of about 10,000 actuators), with the AO loop operating at more than 500 Hz with less than 2 ms latency. A scaling to the full MAORY system would involve six such nodes, and a latency increased by $\mathcal{O}(10)$-$\mathcal{O}(100)$ μs.

These performance estimates have been made using a full real-time system, the Durham AO real-time controller (DARC), and are not theoretical calculations. This system is also on-sky tested with CANARY [9], which is arguably the worlds most advanced AO system, using eight WFSs (four natural guide star (NGS), four LGS), and a large number of on-sky tested algorithms including brightest pixel selection [7], and with the flexibility to implement many others, including optical binning [5]. This system also saw the first use of GPUs on-sky in 2010.

2 Real-time verification

An end-to-end testing and verification environment is required for ELT AO systems, and should be able to operate without requiring the presence of physical components. This will allow system integration of the E-ELT AO systems without requiring, for example, access to the large deformable M4 mirror which is part of the telescope optical train, and without requiring access to the many WFS cameras which will also be individually expensive components.

This verification environment can be provided by interfacing DARC to a realistic numerical simulation code such as Durham AO simulation platform (DASP) [4,2,3]. There are several significant advantages in using numerical simulations over laboratory test benches including the ability to use many more turbulent atmospheric layers than is possible with bench systems, and the fact that the numerical simulation can be duplicated to many sites world-wide requiring only conventional PC hardware. A key part of a real-time simulation system is that the real-time control system does not need to know whether it is connected to a simulation or to real physical hardware, and therefore on-sky algorithms can be used without any change. Physical components can be modelled if not present, and used if they are present, allowing a mixing of present and non-present components to be realised, as demonstrated in Fig. 1.

2.1 Current uses for the real-time simulation

We have begun to make use of the real-time simulation capability in its initial form in several areas. These have included in DiCuRe/HWR testing [8], and in investigating correlation based wavefront sensing performance. When using DiCuRe/HWR on-sky, a

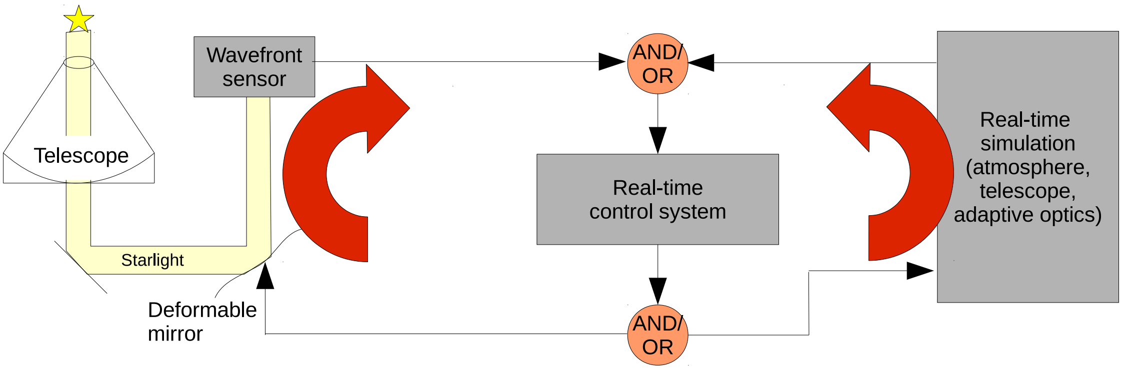

Fig. 1. A figure showing the concept of a real-time simulation code, used in harmony with real physical components, allowing testing and verification of the real-time control system to be realised.

global tilt across the DM was seen to develop with time, and led to performance that was below expectations. Extensive testing was performed using the real-time simulation and the problems were identified and removed. The real-time simulation capability was required in these instances because real-time algorithm implementations differ from the high level language implementations used for development.

Similarly, we were able to investigate real-time update of correlation reference images using this facility. Centre of gravity algorithms are not optimal for open-loop elongated LGS spots, and it is known that correlation techniques offer better performance. However, for astronomical AO systems, the correlation reference images require periodic update, which affects reference slopes. We have developed a technique allowing both slopes and correlation references to be updated simultaneously, including when the AO loop is closed, and used the real-time simulation to test this. The sequence of operations is to obtain a new reference (averaging many frames of calibrated WFS images), compute the update for the reference slopes, update in the RTCS, and repeat this operation over a long period of time, investigating AO performance after many updates. Updating while the loop is closed is essential because for long science images, correlation reference images must be updated during this exposure. We were also able to demonstrate this successfully on-sky with CANARY.

2.2 The DARC WFS simulation front-end

As part of the real-time simulation effort, we have developed a front-end module for DARC which can actively modify WFS images. This is a capability that is useful for algorithm testing and calibration. Real Shack-Hartmann sensor (SHS) spots, from a real WFS camera and a real DM can be used, for example in a laboratory. It may be necessary to investigate the AO performance as a function of signal level, or to investigate alternative spot point spread functions (PSFs), for example elongated spots. This can be achieved simply using the WFS simulation front-end, and can be used to test RTCS algorithms, robustness to noise etc. using a real AO loop.

Active spot modification is behind the core of this capability. A high light level (high signal to noise ratio (SNR)) SHS image is the starting point. Each sub-aperture is then convolved with an arbitrary shape, scaled to the appropriate brightness, and simulated shot and read-out noise are introduced. These resulting images are then treated by the RTCS as the raw images, and so this active spot modification is invisible to the rest of the RTCS. All this occurs in real-time, so the rest of the system is unaware that the spots they see are simulated. Fig. 2 demonstrates this capability in action.

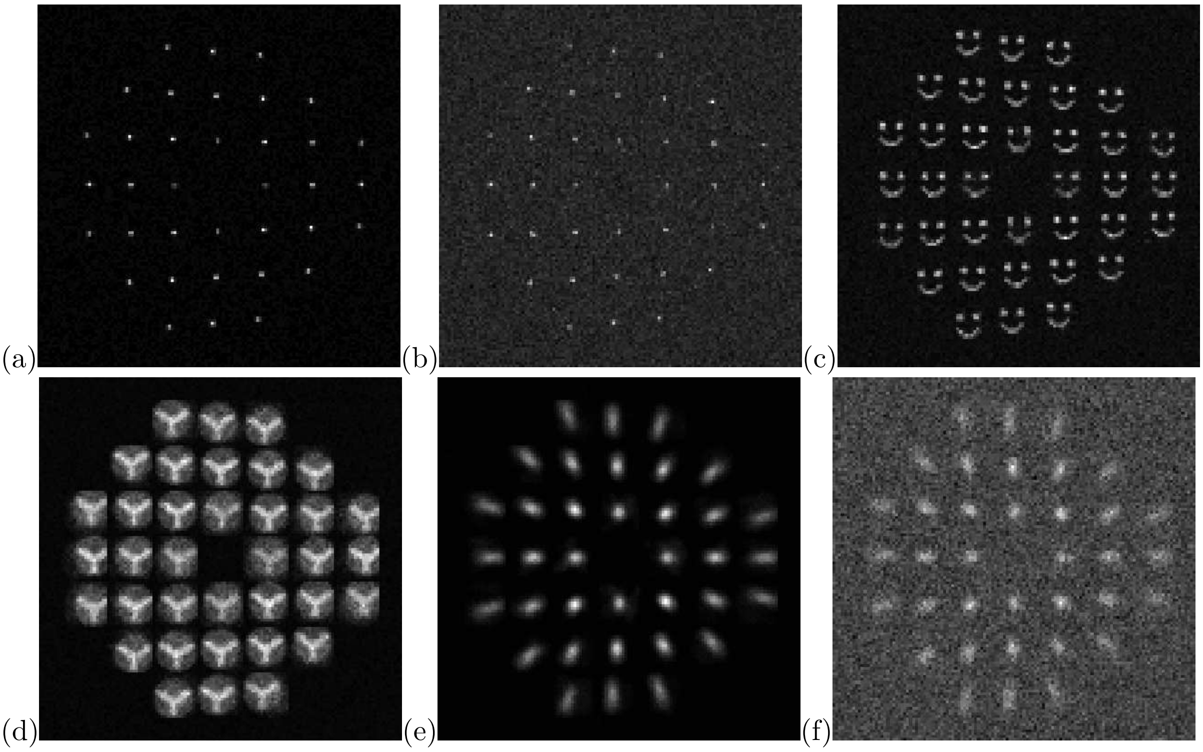

Fig. 2. A figure demonstrating the real-time simulation WFS front-end. (a) shows the SHS pattern from the camera. (b) shows the camera image after random readout noise is added in real-time. (c), (d) and (e) show the camera image after convolution with different PSFs, including a smiley face, the CANARY logo, and an elongated LGS pattern. (f) is identical to (e), but also with shot noise and readout noise added. These images are used within the AO loop, and are used to control the physical DM

It is important to realise that these modified SHS spot patterns are sensitive to changes to a DM (in closed loop systems), and thus can be used to investigate AO control with the loop engaged.

2.3 Hard real-time simulation: RTCS jitter measurement

A variation in the latency of an AO RTCS, on a frame by frame basis is known as the jitter. This must be measured to ensure a RTCS is meeting its performance specification, because large jitter can lead to critical performance degradation. Jitter can be estimated using the AO transfer function, but this requires a closed-loop AO system, and thus it is not possible to measure in this way for ELT systems until commissioning, which may be too late if there are serious jitter problems. Alternatively, it can be measured using a PC clock, however this technique is unreliable, excludes time for data acquisition and the sending of DM demands.

An alternative method for jitter measurement, which we propose here is to use a pixel simulator facility. We use a deterministic field programmable gate array (FPGA) board to generate a pixel stream sequence, which is sent to the RTCS using the same interface as the WFSs (cameraLink, ethernet, etc.). The RTCS then processes these as usual, invoking all the algorithms used on-sky, and computes DM demands. These are then sent back to the FPGA board, again using the same interface as would be used by the real DM. The RTCS therefore requires no changes to use this facility, and is unaware that it is not attached to a proper camera or DM. The FPGA is then used to measure the number of clock cycles between sending of the pixel data and receiving of the DM demands (the instantaneous latency), and this is repeated for millions of cycles. The whole set of instantaneous latency measurements can then be downloaded to PC for further processing, including jitter calculation and investigation of large single jitter events.

This facility has the additional benefit that it allows us to simulate camera and DM interfaces before the cameras and DMs themselves exist, providing information about potential interface issues and hardware throughput of data.

We are developing this facility using a Xilinx Kintex 7 FPGA development board. The pixel stream and DM demands will be implemented using 10G Ethernet in the first instance, to match the expected interface for the E-ELT. The generation of the pixel stream will be flexible, allowing multiple images to be cycled through, and the inter-packet delay and frame rate altered, allowing simulation of a true camera. The concept is shown in Fig. 3

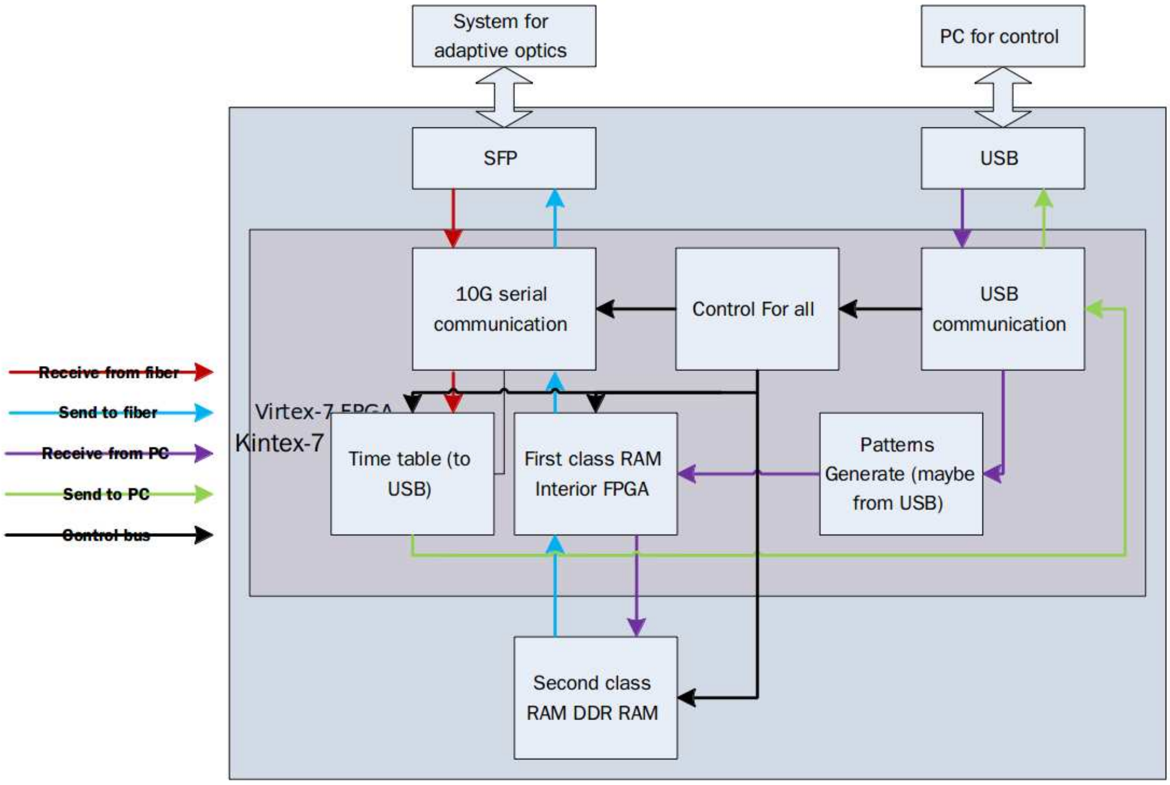

Fig. 3. A figure showing the FPGA pixel simulator concept.

2.4 GPU pipeline investigations

A key part of any ELT scale AO RTCS is likely to be the GPUs that perform the majority of the computational load. We have implemented a GPU pipeline including basic calibration, slope calculation and reconstruction using both OpenCL and NVIDIA CUDA with the goal of comparing different GPU models and programming techniques. Performance of an AMD Radeon HD7970, an NVIDIA GTX580 and an NVIDIA Tesla C2070 have been compared. In summary, we have found that jitter is much more significant for the AMD hardware (Fig.f̃ig:gpujitter, that CUDA and OpenCL performance is largely similar, and that the kernel launch overhead can become significant.

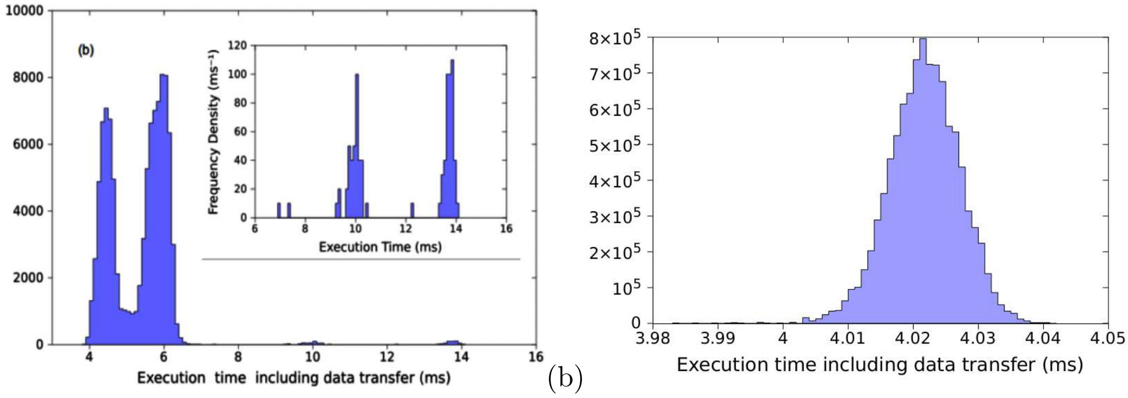

(a) (b)

Fig. 4. A figure comparing jitter for a real-time control system pipeline for (a) AMD and (b) NVIDIA GPUs.

Work is now progressing on a full GPU pipeline with more advanced algorithms.

3 On-sky experience using CANARY

The DARC RTCS used in CANARY is robust by design. However, external offload loops and interfaces required more attention, for example auto guiding and laser steering. The eight WFSs all requires synchronisation, a task made more difficult by unreliable hardware, which can result in occasional truncated image frames. To handle these, the DARC WFS front-end module is able to detect glitches, and handles them intelligently, maintaining WFS synchronisation.

CANARY operates many slow laser offload loops, including for laser beam combination, fast steering mirrors, and LGS asterism rotation. The control of these loops are based on signals from the RTCS, and can operate on any combination of reconstructed wavefront phase, DM demands, WFS slope measurements, and WFS flux. Intelligent flux monitoring is used to determine signal validity, and for example if a non-uniform flux pattern is detected, it probably implies that the laser spots are not centred on the SHS. Such detection and awareness is central to algorithms for AO system robustness.

The telescope autoguider is also operated from the RTCS, with options to guide from WFS slopes (any NGS), from reconstructed phase and DM demands. Again, flux in each sub-aperture is monitored to determine loss of star conditions and direction of travel. The selection of guiding mode is automatic depending on the AO system state, for example if the loop is open, there is no point guiding on DM demands, and so WFS slopes are chosen.

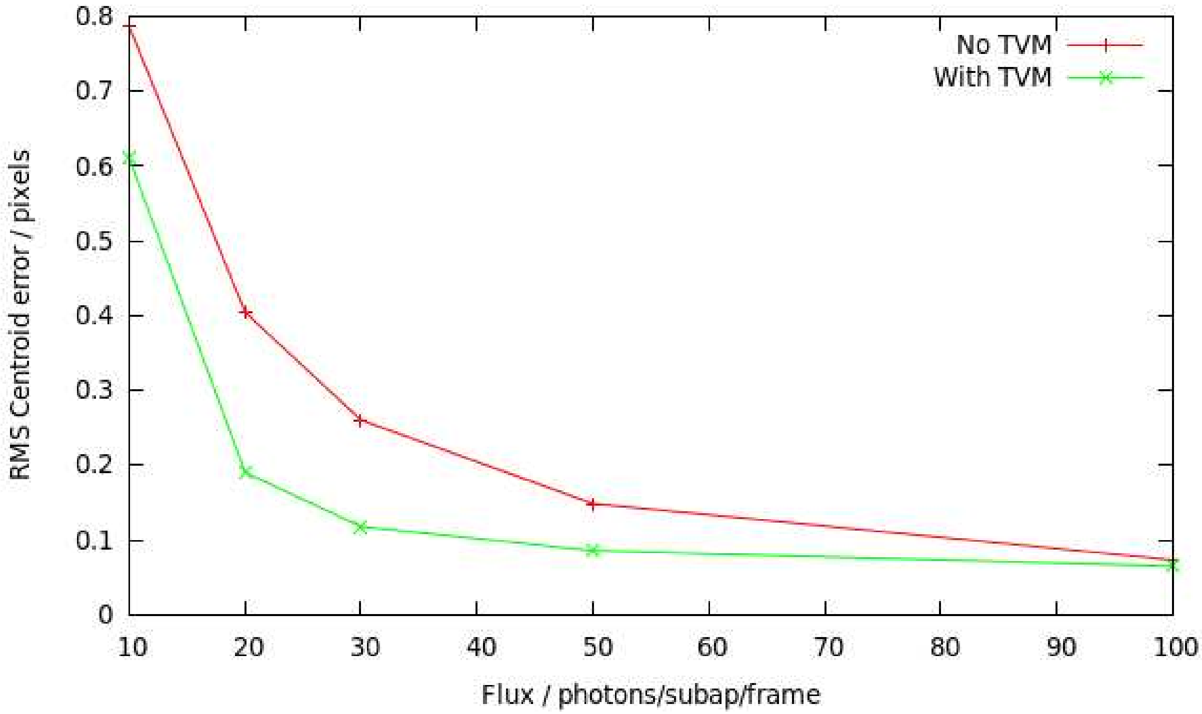

Fig. 5. A figure showing slope estimation accuracy as a function of WFS flux when the Gaussian noise reduction method (TVM) is used, and is not used.

However, when the loop is closed, the WFS slopes are minimised by the AO system, so autoguiding automatically switches to guide on DM demands.

3.1 CANARY algorithms

Many advanced and non-standard algorithms are used with CANARY to enable improved performance to be realised. Of key importance is a thresholding algorithm based on a set number of brightest pixels in each sub-aperture. This enables CANARY to operate will even when sky backgrounds are changing, for example at dawn, when conventional algorithms would fail or require recalibration every few minutes. Spot tracking is also used, allowing the sub-aperture sizes to be shrunk whilst the field of view is maintained. Individual spots, or arbitrary groups of spots can be tracked, and for CANARY we track individual spots for the NGS, and each LGS as a group.

Correlation and Matched filter as well as weighted centre of gravity slope calculation algorithms are available. We also have a Gaussian noise reduction algorithm based on total variation minimisation which can provide up to $\tilde{1}$ Mag gain in performance at low light levels, as demonstrated by Fig. 5.

Many wavefront reconstruction algorithms are also available, including Learn and Apply techniques, Linear Quadratic Gaussian (LQG) formulations, and Neural network based reconstructors for wide-field AO (all used on-sky), and for single conjugate AO (SCAO), we have used CuRe, DiCuRe/HWR, and conjugate gradient algorithms.

4 Conclusion

Conventional real-time control for the ELT is almost a solved problem, and certainly technologically demonstrable. However, implementation of essential more advanced algorithms could still be a challenge. AO system robustness can also be problematic, and so an advanced real-time simulation is required to provide a test harness for ELT RTCSs.

References

1. A. Basden, D. Geng, R. Myers, and E. Younger. Durham adaptive optics real-time controller. *Applied Optics*, 49:6354–6363, November 2010.
2. A. Basden, R. Myers, and T. Butterley. Considerations for EAGLE from Monte Carlo adaptive optics simulation. *Applied Optics*, 49:G1–G8, May 2010.
3. A. G. Basden, N. A. Bharmal, R. M. Myers, S. L. Morris, and T. J. Morris. Monte-Carlo simulation of ELT scale multi-object adaptive optics deformable mirror requirements and tolerances. *MNRAS*, 2013.
4. A. G. Basden, T. Butterley, R. M. Myers, and R. W. Wilson. Durham extremely large telescope adaptive optics simulation platform. *Applied Optics*, 46:1089–1098, March 2007.
5. A. G. Basden, D. Geng, D. Guzman, T. Morris, R. Myers, and S. Saunter. Shack-hartmann sensor improvement using optical binning. *Applied Optics*, 46:6136–6141, 2007.
6. A. G. Basden and R. M. Myers. The Durham adaptive optics real-time controller: capability and Extremely Large Telescope suitability. *MNRAS*, 424:1483–1494, August 2012.
7. A. G. Basden, R. M. Myers, and E. Gendron. Wavefront sensing with a brightest pixel selection algorithm. *MNRAS*, 419:1628–1636, January 2012.
8. N. A. Bharmal, U. Bitenc, A. G Basden, and R. M. Myers. A hierarchical wavefront reconstruction algorithm for gradient sensors. In *Adaptative Optics for Extremely Large Telescopes 3*, 2013.
9. E. Gendron, F. Vidal, M. Brangier, T. Morris, Z. Hubert, A. Basden, G. Rousset, R. Myers, F. Chemla, A. Longmore, T. Butterley, N. Dipper, C. Dunlop, D. Geng, D. Gratadour, D. Henry, P. Laporte, N. Looker, D. Perret, A. Sevin, G. Talbot, and E. Younger. MOAO first on-sky demonstration with CANARY. *A&A*, 529:L2, May 2011.