A Real-time Control Computer for the E-ELT

Document: GF-PDR-01

Introduction to Green FLASH

Version 1.0

18th January 2016

Change Record

Version	Date	Author(s)	Remarks
0.1	11 Jan 2016	D. Gratadour	Initial version
1.0	11 Jan 2016	D. Gratadour	Final version

Applicable Documents (AD)

These are the Green Flash PDR documents

No.	Title	Reference	Issue	Date
AD01	Introduction	GF-PDR-01		
AD02	Management plan and WP definition	GF-PDR-02		
AD03	Requirements Specification	GF-PDR-03		
AD04	System Architecture	GF-PDR-04		
AD05	Distributed GPUs for real-time HPC	GF-PDR-05		
AD06	FPGA Solution for hard real-time	GF-PDR-06		
AD07	Interconnect Strategy	GF-PDR-07		
AD08	Interface Control Document	GF-PDR-08		
AD09	Supervisor Strategy	GF-PDR-09		

Reference Documents (RD)

These are documents external to the Green Flash project

No.	Title	Reference	Issue	Date

Acronyms and abbreviations

Table 1 Acronyms and Abbreviations

AD	Applicable Document
AO	Adaptive Optics
CANARY	Durham/LESIA on-sky AO demonstrator
CPU	Central Processing Unit
CUDA	NVIDIA GPU based software development language
DARC	Durham AO Real-time Controller
DDS	Data Distribution Service
DM	Deformable Mirror
DRAGON	Durham laboratory-based AO demonstrator bench
ELT	Extremely Large Telescope
E-ELT	European ELT
ESO	European Souther Observatory
FPGA	Field Programmable Gate Array
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HLS	High Level Synthesis
HPC	High Performance Computing
MIC	Many Integrated Core
MVM	Matrix-Vector Multiplication
NIC	Network Interface Controller
PCIe	Peripheral Component Interconnect express
RD	Reference Document
RTC	Real-Time Control
RTL	Register Transfer Level
SIMD	Single Instruction Multiple Data
SPARTA	ESO VLT AO Real-time Control System
SHERE	VLT Planet finder instrument
UDP	User Datagram Protocol
UK ATC	United Kingdom Astronomical Technology Centre
VLT	Very Large Telescope
WFS	Wave-Front Sensor
WP	Work Package

Table of Contents

1 Scope

This document provides an overview of Green FLASH context, goals and expected impact.

2 Context

High Performance Computing (HPC) has become a critical component to the availability of large scientific instruments such as the European Extremely Large Telescope (E-ELT)[1], a 39m diameter telescope project to see first light by the horizon 2022. For this new category of extreme scale scientific equipment, HPC is not only used to optimize the design of the facility and its instrumentation and save money in construction but also to operate its complex sub-systems and produce high-level scientific data for hundreds of research teams.

This major European project can be used as an example of the requirements of HPC computing in science and technology in the coming decade. In this joint academic and industrial proposal, meeting the computational requirements of the E-ELT is used to demonstrate what can be achieved in real-time HPC using existing technologies and to forge a path towards Exascale computing.

2.1 Challenges for E-ELT AO RTC

At the core of telescope operations is the adaptive optics (AO) module, used to compensate in real-time the effect of atmospheric turbulence on the wavefront to maintain the best image quality out of the telescope during the observations. It requires the control, at the millisecond rate, of deformable optics with thousands of degrees of freedom. Designing, optimizing, building and eventually operating and exploiting this critical sub-system require access to HPC resources from early simulations for system design to the real-time control of the optics for routine operations. Indeed, the latter, at the E-ELT scale, requires the processing of several 100Gb/s of data streaming from several sensors, through batched-centroiding and matrix-vector multiplies (MVM), to provide several millions of commands per second to the deformable optics. This is the role of the real-time controller (RTC), one of the core subsystems of the AO module. To ensure adequate turbulence compensation, hence the stability of the system, the RTC computation time must be deterministic at the level of tens of microseconds, and the latency in data transfers between the sensors, the computing cores and the deformable optics should be minimal.

To build this critical component of the telescope operations, the astronomical community is facing technical challenges, emerging from the combination of high data transfer bandwidth, low latency and high throughput requirements, similar to the identified critical barriers in the ETP4HPC road map and, as debated more generally in the HPC community, on the road to Exascale. Indeed, meeting the specifications of these real-time multi-sensors data-intensive computing facilities, to be installed on a remote site and to be operated continuously over several years, involves addressing the four strategic topics in Exascale systems design:

- optimize the balance between processors performance, memory capacity and data access between sensors, compute nodes and storage
- design and implement efficient computing schemes addressing the concurrency and locality challenges
- follow an energy efficient approach in the system design
- build a modular, scalable and resilient system

1 http://www.eso.org/public/teles-instr/e-elt/

2.2 Concept

For the real-time control of AO systems a standard platform, SPARTA, has been implemented by the European Southern Observatory (ESO) in collaboration with Durham University, one of the partners in the project, to fulfill the needs of the Very Large Telescope (VLT) instrumentation. This successful project, partly funded by the OPTICON/JRA-1 program, has delivered a standard platform used to develop the real-time control system of 10 VLT instruments in operation or in development.

However, as stated by the SPARTA architects: "*the E-ELT with its instruments poses new challenges in terms of cost and computational complexity. Simply scaling the current SPARTA implementation to the size of E-ELT AO system would be unnecessary expensive and in some cases not even feasible*"[2]. Because the E-ELT AO instruments RTC have to be based on the For standard platform for maintainability issues, one of our primary goal is to provide the building blocks of the next generation of SPARTA platform. Incidentally, our philosophy and part of our specifications are greatly inspired by the initial guidelines for the SPARTA concept developed more than 10 years ago.

The output of the green Flash project is not an evolution of SPARTA in itself. We propose to develop new core components and an architecture consistent with a possible evolution of this platform. We believe that our work goes beyond this sole application and we try to follow a generic approach using or providing mainstream solutions as much as possible.

2.3 Proposed approach

At the core of telescope operations is the adaptive optics (AO) module, used to cTo meet the specifications imposed by the E-ELT AO module's RTC dimensioning, while mitigating potential technology risks, our approach is twofold:

- develop a tailored solution based on the FPGA technology, through hardware developments and the deployment of a comprehensive ecosystem (FPGA development environment and IP blocks, drivers, middleware and software)
- assess a solution based on distributed accelerators using smart interconnects, with the real-time constraint, and compare (in terms of performance, energy efficiency and resilience) to the tailored FPGA solution

The main goal of Green Flash is to design and build a prototype for an AO RTC targeting the E-ELT first-light AO instrumentation. We will propose technical solutions, assess the enabling technologies through prototyping and assemble a full scale demonstrator to be validated with a simulator and tested on sky. With this research and development program we aim at feeding the E-ELT AO systems preliminary design studies, led by the selected first-light instruments consortia, with technological validations supporting the designs of their RTC modules.

Our strategy is based on a strong interaction between academic and industrial partners. Components specifications and system requirements are derived from our main application. Industrial partners lead the development of enabling technologies aiming at innovative tailored solutions with potential wide application range. The academic partners provide the missing links in the ecosystem, with possible wider collaborations with specialized research teams, through generic tools supporting their own application but in mainstream open source projects. This increases both the value and market opportunities of the developed products. A prototype harboring all the features is used to assess the performance. It also provides the demonstration of concept for a resilient modular solution to equip a large scale European scientific facility, while containing the development cost by providing

2 E. Fedrigo and R. Donaldson, Proceedings of the SPIE, Volume 7736, id. 77364O (2010), DOI: 10.1117/12.857109

opportunities for return on investment.

3 Project objectives

3.1 Real-time HPC using accelerators and smart interconnects

Meeting the first light AO module's RTC specifications for the E-ELT means building a HPC facility able to process up to 250 Gb/s[3] of streaming data from sensors with a throughput of up 1.5 TMAC/s[4] (Tera Multiply ACcumulate per seconds). With the current mainstream GPGPU technology, assuming a sustained peak performance of 75 GFLOP/s per GPU for the cuBLAS single precision *gemv* algorithm[5], this figure means distributing the computation over more than 45 GPUs. Additionally, the response of the RTC must be deterministic with a maximum jitter of the order of few tens of microseconds. Reaching this deterministic performance with a distributed architecture challenges both the job scheduling and the interconnect strategies.

- The primary objective of this axis of development is to prototype a cluster, able to reach a sustained performance of 1.5 TMAC/s of computing power and processing 250 Gb/s of streaming data with a maximum jitter of 100 µs over 1s of operation and using commodity accelerators such as GPUs as compute engines. (obj. 1.1)
- One objective of this axis of development is to provide a COTS NIC solution, based on the FPGA technology, with smart features, compatible with existing high performance switch solutions, based on standard serial protocols (TCP/UDP through 10G Ethernet and 40G Infiniband), and supported in mainstream non proprietary middleware. (obj. 1.2)
- Another objective of this axis of development is to complement the ecosystem of an existing integrated FPGA development environment (QuickPlay from PLDA), by providing data handling and computational blocks tailored to our application, and support for several FPGA options and board designs. The end user API will be made interoperable with mainstream non proprietary middleware. (obj. 1.3)
- The last objective of this axis of development is to assess the performance of Cholesky factorization on the prototype cluster, as well as the overall control matrix computation algorithm and the control matrix upload to the controller with minimal introduced latency and jitter in the control process. (obj. 1.4)

3.2 Energy efficient platform based on FPGA for HPC

An alternative to the commodity accelerator solution for real-time control is to build tailored processor boards, based on high cell density FPGAs for instance, designed for the application and relying on custom data flow models. The great advantage of such approach is to provide full control over hardware and critical software (device driver) to ensure real-time and deterministic performance and to guarantee the support and availability of the solution over the system lifetime. However, this is usually done at the cost of an expensive and rather long development cycle (as compared to custom arithmetic designs) to produce specialized components with very few options for return on investment.

- The primary objective of this axis of development is to prototype a main board, hosting a high

3 6 cameras with 1600 x 1600 16bit pixels at 1000 frames per second

4 90000 x 15000 MVM at 1kHz

5 http://arxiv.org/abs/1410.1726

cell density FPGA with an integrated ARM-based HPS, including a PCIe Gen3 rootport as an internal interface, and using 10G Ethernet and 40G Infiniband as network interface and back plane. (obj. 2.1)

- Another objective of this axis of development is to provide support for this prototype microserver in an integrated FPGA development environment (QuickPlay), deployable in the FPGA BSP and allowing to build a custom design on the FPGA, handling complex data flows through the internal (PCIe), the network interfaces and computing blocks, as well as the driver and end user API to use these features. Blocks developed in the first research axis on smart interconnects will be reused. (obj. 2.2)
- The last objective of this axis of development is to build a small scale cluster by interconnecting several prototype microservers through a standard network protocol (10G Ethernet and 40G Infiniband). The performance achieved for the MVM algorithm and more generally the AO control application on single and multiple boards configurations will be assessed in terms of performance, determinism and energy efficiency as compared to an accelerator based solution. (obj 2.3)

3.3 AO RTC prototyping and performance assessment

The main goal of the prototyping and validation phase of the Green Flash project is to assess the enabling technologies through prototyping and assemble a full scale demonstrator for an E-ELT first-light AO RTC. As described above, such facility, designed to drive in real-time a large scientific experiment, is composed of a real-time core, processing streaming data from sensors and controlling deformable optics, and a supervisor module, optimizing the control loop by providing updated versions of the control matrix at a regular rate depending on the observing conditions evolution.

- The primary objective of this axis of research is to assemble a full functionality prototype for an AO RTC, scalable to the dimensioning of the E-ELT first light instrumentation, including a real-time core and a supervisor module. Architecture choices will be made on the basis of the individual prototyping stages output. (obj. 3.1)
- Another objective of this axis of development is to implement a real-time simulator, designed to emulate the AO system I/O with various levels of accuracy, using existing AO sub-systems models. (obj 3.2)
- The last objective of this axis of development is to fully characterize the AO RTC prototype performance with several configurations (single and multi-conjugate AO) and propose a strategy for its integration on sky. (obj. 3.3)

4 Building blocks of Green FLASH

4.1 COMPASS an efficient AO computing platform based on GPUs

The main objective of the COMPASS project (COMputing Platform for Adaptive opticS Systems[6]) is to provide a full scale end-to-end AO numerical development platform, able to address the E-ELT scale and designed as a free, open source numerical tool with a long term maintenance plan. This project, led by Observatoire de Paris is federating the efforts of 6 laboratories in France and is funded by ANR grant ANR-12-MONU-0022 of the French Ministry of Research.

6 D. Gratadour et aL, Proceedings of the SPIE, Volume 9148, id. 91486O 8 pp. (2014), DOI:10.1117/12.2056358

The development of this platform is based on a full integration of software with hardware and relies on an optimized implementation on heterogeneous hardware using GPUs as accelerators. The end product, a unified and optimized computing framework (software and hardware), addresses several major needs for research in AO, as it provides:

- a software development platform to validate several key components for the E-ELT AO such as Laser Guide Stars (LGS) with extreme elongation or fast high precision tomographic reconstruction
- an efficient computing environment to run large scale AO systems simulations in almost-real-time in order to test the design of the E-ELT instrumentation and prepare the scientific return of the foreseen instruments
- a link to end-to-end simulations including instruments models so as to perform technical trade-off of the full chain during instrument design studies

One key contribution of the COMPASS project is the development of a common framework for high performance simulations and real-time control of AO systems based on GPUs. The latter application is based on the development of a custom solution implementing Direct Memory Access (DMA) to the GPU memory from a third party device (frame grabber) bypassing the operating system running on the node. Thanks to the GPUdirect RDMA API, available from the latest CUDA distributions, NVIDIA is providing the tools to implement tailored solution to access the GPU memory efficiently and transparently so as to enable a given real-time application.

In the context of the COMPASS project, a demonstrator has been set up, codenamed PRANA[7], to validate the achievable bandwidth of DMA of pixel data from an external camera to GPU memory and emulate the RT box in an AO RTC. The custom low latency 10 Gb Ethernet frame-grabber developed for this purpose is able to handle any GigEVision device. Beyond the first demonstration of DMA from a 10 GbE GigEVision camera to a GPU, the measured latency in this experiment shows that significant reduction of latency and jitter can be achieved when exploiting the DMA enabling features of the NVIDIA driver. The underlying real-time core, processing the data stream, is based on the use of control algorithms implementations developed in the framework of the COMPASS project. This unified approach allows for the seamless integration of complex AO control schemes that can be fully debugged, validated and characterized using simulations, maximizing the robustness and updatability of the real-time core while minimizing the development cost.

Another key building block for the Green Flash project is the work on an efficient implementation of the core algorithms of the supervisor module on GPU, led in the context of the COMPASS project and in collaboration with the Matrices Over Runtime System at Exascale[8] (MORSE) project. MORSE is a high-performance numerical library for solving large dense (and sparse) linear algebra problems on multicore systems with hardware accelerators. For the AO loop supervision application, we developed an efficient implementation the control matrix computation for MOAO on multicore system with multi-GPUs. Based on the Cholesky matrix inversion, it combines tile algorithms from MORSE and relies on the StarPU dynamic scheduler to efficiently schedule and pipeline successive computational stages. This new implementation outperforms asymptotically previous state-of-the-art control matrix computation implementations up to 13-fold speedup and runs at an unprecedented scale. This work has been presented at the IEEE/ACM Super Computing 14 conference[9].

---

7 A. Sevin et al., Proceedings of the SPIE, Volume 9148, id. 91482G 9 pp. (2014), DOI: 10.1117/12.2055770

8http://icl.cs.utk.edu/morse/

9 A. Charara et al. Proceedings of the International Conference for High Performance Computing '14, DOI: 10.1109/SC.2014.27

4.2 DARC : on-sky demonstration of AO RTC with commodity components

The Center for Advanced Instrumentation (CfAI) at Durham University has a long-standing research program in adaptive optics and in real-time control systems. This work has been funded by the UK Science and Technology Facilities Council and by the EU funded Opticon program. RTC systems have been designed for both laboratory and telescope facility instruments over more than two decades utilizing several generations of different computer hardware. In recent years, this work has contributed to the development of the ESO SPARTA system[10], where Durham were responsible for the FPGA technology used, and to a novel mainly CPU based RTC system, the Durham AO Real-time Controller, DARC[11]. Whilst SPARTA is targeted mainly at medium to high order AO systems such as the SPHERE VLT planet finder, DARC was initially developed as a flexible RTC for the wide-field CANARY experiment[12], a collaboration between Durham and Paris Observatory. This experiment is designed to demonstrate the multiple new techniques in AO that have emerged in the last few years. In order to show that these techniques are practical and can be used on the E-ELT, this experiment is installed at the 4.2 metre William Herschel Telescope (WHT) in La Palma, Canary Islands. Canary, supported by DARC, has already made the first observations demonstrating both MOAO and LTAO and will evolve into a higher order system over the period of this proposal.

DARC has now been demonstrated successfully at ELT-scales. It has been designed specifically to take advantage of the parallel processing made possible by multi-core and many-core processors. It implements a parallel data transfer and data processing policy that allows an efficient and optimum regime for the full data pipeline in multiple threads. As well as the data pipeline, the DARC RTC provides all of the facilities required by an AO system including data telemetry and configuration, calibration and control.

In addition to its role in Canary, DARC is also used as the RTC for the Durham laboratory-based AO bench known as DRAGON. In this environment, as well as serving low, medium and high order systems on the bench, it can be used with simulated data, to test RTC hardware and software for high order systems such as those required for the E-ELT. A program of investigation of HPC accelerator hardware is already in progress. The FPGA based camera interconnect developed for SPARTA has been used in Canary, mainly in the role of allowing direct access to camera pixel streams (which is not provided by commercial frame-grabber hardware) rather than as a pixel pre-processor. A 10G-Ethernet wavefront sensor has also been integrated with DARC, and real-time performance tested at ELT-scales with real data, again with pixel stream access. GPU cards have been used on-sky with Canary (being the first demonstration of GPUs on-sky worldwide) and are being used to accelerate the DRAGON RTC. The data pipeline has been successfully ported to GPU allowing frame rates as high as 500Hz for a single wave-front sensor camera at an E-ELT scale (80x80 sub-apertures) on a single GPU. As an alternative to the use of GPUs, a program of investigation of many integrated core (MIC) processors has also been started using the Xeon Phi device[13]. In each case, as well as testing the offloading of specific algorithms to the accelerator, a full implementation of a data pipeline is being tested.

---

10 E. Fedrigo et al., Proceedings of the SPIE, Volume 6272, id. 627210 (2006), DOI: 10.1117/12.671919

11 A. Basden et al., Proceedings of the SPIE, Volume 7736, id. 77364N (2010), DOI: 10.1117/12.856415

12 R. Myers et al., Proceedings of the SPIE, Volume 7015, article id. 70150E, 9 pp. (2008), DOI: 10.1117/12.789544

13 D. Barr et al., Proc. SPIE 9148, Adaptive Optics Systems IV, 91484B (4 August 2014); doi: 10.1117/12.2055092

DARC, both in its CPU based and hardware accelerated forms provides an excellent starting point for investigating these technologies at the E-ELT scale within this proposed program.

4.3 QuickPlay a comprehensive FPGA development environment

This tool has been developed by PLDA, with the ultimate objective of drastically speeding-up the development and validation of FPGA projects. It enables direct use of FPGA's by users without specific hardware skill, hence making easier the penetration of the FPGA technology in new domains.

The overall process of implementing a design using QuickPlay is straightforward. It consists of:

1. Developing a C/C++ functional dataflow model of the hardware engine
2. Verifying the functional model with standard C/C++ debug tools
3. Specifying the target FPGA boards and interfaces (PCIe, Ethernet, DDR, QDR, etc)
4. Compiling the HW engine

This process is comprehensive from the design of a functional model to the effective validation of the FPGA design on hardware platforms supported by QuickPlay. In order for this simple process to work seamlessly, the generated hardware engine must be guaranteed to function identically to the original software model. Another way of stating this is that the functional model must be deterministic so that, no matter whether executed in software or in any possible hardware implementation, execution will give the same results, albeit at different speeds. Most parallel systems suffer from non-deterministic execution. Multi-threaded software execution, for example, depends on the CPU, on the OS and on non-related processes running on the same host. Multiple runs of the same multi-threaded program can produce different behaviors.

Such non-determinism in hardware would be a nightmare, as it would require debugging the hardware engine itself, at the electrical waveform level. QuickPlay thus uses an intuitive dataflow model that mathematically guarantees deterministic execution, regardless of the execution engine. Such model consists of concurrent functions, called "kernels", communicating with streaming channels, which correlates well with how applications would be sketched on a whiteboard. In order to guarantee deterministic behavior, these kernels must communicate with each other in a way that prevents data hazards, such as race conditions. This is achieved with streaming channels that are:

- Blocking (read and write),
- Lossless, and
- Point-to-point.

Numerous technical problems have been solved to deliver a first workable version of QuickPlay, which has been used by two early Customers in spring 2014. This first prototype did clearly demonstrate that the objective of drastically speeding up the application development time was reachable, and a first production version of QuickPlay has been released in November 2014. QuickPlay is by design an open project and numerous improvements of the tool are targeted by PLDA. QuickPlay does perfectly match the objectives of the Green Flash project, as long as FPGA usage is concerned, and will boost up the productivity of the whole project.

4.4 Custom hardware developments at Microgate

Microgate has developed in the past a custom hardware platform dedicated to specific control tasks for adaptive optics. More specifically, the custom hardware has been initially designed to cope with the specific goal of controlling large adaptive mirrors based on contactless voice-coil technology. These adaptive mirrors rely on force actuators to control the shape of a thin mirror shell that levitates on the magnetic field generated by voice coil motors, which moving part is attached to the rear, non-optical side of the shell. A co-located, contactless capacitive sensor provides the position sensing for

each actuator. This peculiar configuration has several advantages with respect to competing technologies, but it requires a quite sophisticated control system. The demanded computational throughput for a typical correcting mirror of an 8-10m telescope, with about 1000 actuators, is in the range of few hundreds of Gflops, while the communication interface requires a bandwidth in the range of few Gbit/s with very demanding latency and jitter requirements, in the range of few microseconds. Moreover, signal integrity issues, noise pickup of the sensing devices and harness complexity impose to place the control electronics very close to the controlled device, thus demanding very tight size, low power dissipation and housing temperature constraints. All these requirements led Microgate to develop a very specific electronics architecture, encompassing at an initial stage most of the goals of the present design: high computational throughput, strict real-timeliness, large communication bandwidth with extremely low latency and jitter, low power consumption and scalability. Thanks to its flexibility and modular design, the same electronics has been also adapted and deployed for other adaptive optics real time tasks, like wavefront sensing and real time reconstruction. The described architecture has been deployed on the adaptive optic systems of the several large telescopes, including MMT, LBT, Magellan, Keck, VLT.

From the technological standpoint, the custom electronics described above has been initially based on a mixed DSP/FPGA design, in which the FPGA performed as glue logic and handler/router of the low level DMA communication within the large cluster, besides some basic computational task, while most of the computations were executed by DSPs. More recently, specifically in the frame of the development of the adaptive mirrors for the ELTs (E-ELT and GMT) with several thousands of controlled channels, the availability of more powerful and flexible FPGAs allowed to move here all computational tasks, abandoning the DSP technology. This provided an additional benefit in terms of power consumption and performance. To further reduce dissipation and complexity of the electronics embarked on the adaptive mirror, only the co-located control and the power/analog electronics required to interface to the actuators has remained on-board. We moved the computational-heavy global control tasks, in the range of several hundreds of Gflops, to an external control unit, which prototyping is one of the goals of this proposed activity.

4.5 Starting point

The main driver of the prototyping phase of the green Flash project is the combined need for

- new core components in the design of SPARTA
- an optimized strategy to handle the complex data flows and their interactions in the system
- a long term maintainable solution, based on evolving standards

Our goal is to design and prototype new solutions addressing these issues, assemble a demonstrator to validate the achievable performance and propose an evolution of SPARTA harboring these solutions. While we believe that they could have a wider application range, these solutions will thus be specified to be integrable in the future SPARTA platform. However, our strategy enables sufficient modularity to ensure ease of portability to other applications and reuse as much as possible of SPARTA software components to minimize the development effort.

Such strategy could be facilitated by several properties of the SPARTA platform:

- the RT box and co-processing cluster are well decoupled so that they may be reused / updated separately
- the co-processing cluster employs standard interfaces and middleware that can be entirely reused
- most of the co-processing cluster SW can be reused. Compute-intensive parts use standards

library (MKL) and are well encapsulated

- the HW dependencies in the RT box SW are strongly encapsulated which limits the impact of porting to other architectures
- the SPARTA light project, targeting small scale AO systems, provides a single node, CPU based solution, running Linux, and entirely consistent with the SPARTA specifications

Our starting point, depicted in the figure below, is a simplified view of the AO RTC architecture, following the same functional blocks breakdown as SPARTA (ensuring compatibility) but including a unified interconnection concept between these blocks through switches and removing all the complexity of a facility instrument in terms of operability and integration in a larger infrastructure.

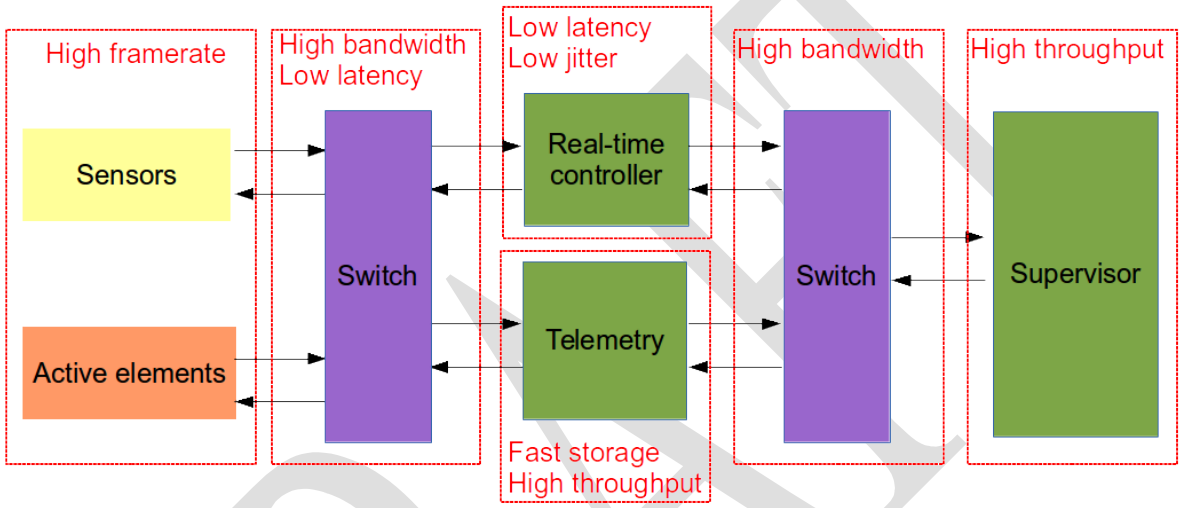

4.6 Ambition

Our main ambition is to demonstrate the potential of new core components for the control of AO modules in real-time at the scale of extremely large telescopes and targeting the E-ELT first-light instrumentation dimensioning and functions.

Current state-of-the-art solutions already in operation on telescopes, mostly involve application specific hardware and core software / firmware as it is the case for instance for the SPARTA platform already integrated on several instruments on the ESO VLT or the real-time controller for the Gemini Multi-conjugate AO system (GeMS) on the Gemini telescope[14] both based on the use of DSP boards. While these solutions are robust in terms of AO performance and resilient, they rely on a costly development cycle to produce non portable end products with a low expandability level. Up scaling these solutions to the ELT dimensions would be not practical (and even not feasible for some subsystems) and very costly.

Several concepts have been proposed in the community to provide new architectures for AO RTC at the ELT scale and go beyond state-of-the-art. Among these we find:

- a solution based on large boards containing several hugh cell density FPGAs[15] and relying on new control algorithms, less mature in terms AO performance reliability than MVM, but efficiently mapped to this hardware
- solutions based on accelerators such as GPUs[1617] or Intel Xeon Phi[18], where the goal is to accumulate enough computing power to cope for data transfer latency
- solutions based on distributed Intel Xeon CPUs[19] where the main strategy is to rely on standard libraries and stack enough nodes to meet the requirements

With Green Flash, we propose to go beyond these approaches and develop custom solutions, addressing the critical aspects of the AO loop to provide an optimized design relying on energy efficient core components. The end products are the building blocks of a new standard platform expandable to fit the need of the various AO flavors to be installed on the E-ELT and possibility other ELTs. We believe they could also apply to a wider application spectrum going beyond the usual application specific design of AO module's RTC.

These ambitious objectives are supported by a strong academic – industrial partnership, following a co-design strategy driven by performance optimization, expandability, energy efficiency and cost containment.

One of the innovative aspect of the proposed work is to focus on a unified interconnect strategy across

---

14 F. Rigaut et al., Monthly Notices of the Royal Astronomical Society, Volume 437, Issue 3, p.2361-2375 (2014)

15 H. Zhang et al., Proc. SPIE 8447, Adaptive Optics Systems III, 84472E (13 September 2012); doi: 10.1117/12.925479

16 D. Gratadour et al., Proceedings of the Third AO4ELT Conference. Firenze, Italy, May 26-31, 2013 sciencesconf.org:ao4elt3:13354

17 L. Wang et al., roceedings of the Third AO4ELT Conference. Firenze, Italy, May 26-31, 2013 sciencesconf.org:ao4elt3:13172

18 D. Barr et al., Proc. SPIE 9148, Adaptive Optics Systems IV, 91484B (4 August 2014); doi: 10.1117/12.2055092

19 M. Smith at al., Proc. SPIE 9148, Adaptive Optics Systems IV, 91484K (21 July 2014); doi: 10.1117/12.2057345

the system, and propose a solution able to handle data streams from sensors and from intra and inter-nodes communications each transmitted through dedicated protocols. To meet this objective, we will explore the concept of smart interconnect in the context of the AO application, but with a methodology that could apply to other applications. We believe that this process, including the final validation of performance, provides a proof of concept that could be reproduced by other teams for their application, adding value to the developed solution and a potentially wider impact than in the AO community.

Another aspect of the work on smart interconnect is to study the possibility of including the smart features at various levels in the programming strategy either for optimized communication schemes or for adding compute blocks. Here again our ambition in to propose, through the AO application, a proof of concept, relying on collaborations with developers of dynamic run time systems and programming models to demonstrate the potential of this new concept.

Among the potential innovations proposed in Green Flash, the concept of energy efficient FPGA based microserver is an ambitious attempt to address the question of energy efficiency in a HPC facility. In this approach, we propose to centralize the functions of host CPU, high bandwidth NIC and accelerator on a single energy efficient component and explore the possibility of clustering several boards interconnected through standard high bandwidth protocols. The design of each board can be tailored for dedicated data intensive work flow using a mainstream, user oriented development environment maximizing the modularity of the facility with a minimum time-to-application. Moreover, we propose to develop microserver solutions expandable with accelerators through PCIe expansion ports. The potential wider application range of these new products will be explored by the industrial partners relying on the demonstration of performance led in the Green Flash project.

Another ambition is to promote the use of a new product: QuickPlay from PLDA, that will reach the market during the course of the project. It will be done not only through an extensive use of the existing features as a demonstration of concept, but also by adding value to the product through the implementation of a number of extensions and support for new boards driven by our main application needs. QuickPlay provide the essential tool to dramatically speed-up the development and validation process of FPGA projects, hence making easier the penetration of the FPGA technology in domains where its intrinsic benefits have not been exploited yet such as HPC. QuickPlay delivers on the promise that FPGA based computing should be available to a broader audience. It does so by removing the one and only barrier to FPGA adoption: its complex and specialized programming model.

5 Partners

5.1 LESIA at Observatoire de Paris

Observatoire de Paris (http://www.obspm.fr/obsparis.en.shtml) is a national research centre in astronomy and astrophysics. The Observatory employs approximately 1000 people (750 on permanent positions), comprising one third of researchers and two thirds of engineers, technicians and administrative personnel.

Observatoire de Paris is the largest astronomy centre in France and one of the most important ones in the world. It represents alone one third of astronomy in France. It depends on the Ministry of Higher Education and Research, belongs to the category of 'Great Establishments', and has the status of an independent university. The Observatory is structured in 5 laboratories (GEPI, LESIA, LUTH, LERMA and SYRTE), one scientific unit (Nançay radio astronomy station), and one institute (IMCCE) covering all fields of astronomy and astrophysics.

LESIA, Laboratoire d'Etudes Spatiales et d'Instrumentations en Astrophysique is one of five science laboratories of Observatoire de Paris. It is also a CNRS Laboratory (UMR 8109). LESIA is associated with the University Pierre et Marie Curie - Paris 6 (UFR de Physique 925 & Institut de la Vision) and with the University Paris Diderot - Paris 7 (UFR de Physique & UFR Sciences de la Terre, de l'Environnement et des Planètes - STEP). The laboratory counts around 250 staff members, with 140 permanent members including 65 engineers and technicians. More than 30 PhD theses are currently undertaken. LESIA's primary role is: design and implementation of scientific instrumentation in space and on the ground, analysis and interpretation of scientific observations made by use of the built instruments and development of innovative advanced techniques to fulfill the new requirements of its scientific programs.

LESIA is at the forefront of scientific advances in all areas where it is involved, including instrumentation, numerical simulation and theoretical work. Its major achievements in innovative instrumentation are adaptive optics (COME-ON, ADONIS, PUEO, NAOS), four quadrant phase mask coronagraph and optics fiber recombination in multi-telescope interferometry (FLUOR, OHANA). It is currently involved in a number of major projects in astronomy like: SPHERE and GRAVITY for the VLT (ESO) and MIRI for the JWST (NASA, ESA). It has participated to four Phases A studies for the instrumentation of the E-ELT (EAGLE, EPICS, MICADO, ATLAS) with contributions in different topics such as AO system design and prototyping, LGS wave-front sensing, high contrast imaging and speckle suppression.

Currently, LESIA is a member of MICADO consortium and is responsible for its first light SCAO module. For this project, the team at LESIA has work closely with the MAORY consortium to share a number of developments such as the wavefront analysis on natural star and the RTC. LESIA is also a member of the new MOSAIC consortium for the development of multi-object spectrograph equipped with MOAO. LESIA is responsible for the development of the AO subsystem. This MOAO system needed a sky demonstrator to prove its feasibility: the CANARY project developed in collaboration with the University of Durham. This sky demonstration with CANARY will continue for several years.

During the past two years, LESIA has been leading the COMPASS project, a numerical development platform for AO, supported by the ANR through grant ANR-12-MONU-0022, which is the result of a national collaborative multi-disciplinary (astrophysics, instrumentation tion and HPC) effort led by 5 laboratories in France. Additionally, thanks to this project, special links have been forged with international teams around the optimization of AO dedicated work-flows on heterogeneous architectures. The main objective of the COMPASS project is to provide a full scale end-to-end AO development platform, able to address the E-ELT scale and designed as a free, open source numerical tool with a long term maintenance plan. The development of this platform is based on a full integration of software with hardware and relies on an optimized implementation on heterogeneous hardware using GPUs as accelerators. One key contribution of the COMPASS project is the development of a common framework for high performance simulations and real-time control of AO systems based on GPUs. The latter application is based on the development of a custom solution implementing Direct Memory Access (DMA) to the GPU memory from a third party device (frame grabber) bypassing the operating system running on the node. In the context of the COMPASS project, a demonstrator has been set up, codenamed PRANA, to validate the achievable bandwidth of DMA of pixel data from an external camera to GPU memory and emulate the RT box in an AO RTC. The custom low latency 10 Gb Ethernet frame-grabber developed for this purpose is able to handle any GigEVision device.

5.2 CfAI at University of Durham

Durham University is a legally independent and autonomous institution established by Act of Parliament in 1832 and Royal Charter in 1837. It is also an educational charity. The University receives about a third of its income directly from the UK Government, via HEFCE (Higher Education Funding Council for England). Durham University is the third oldest university in England and one of the leading universities of the United Kingdom. The university exists both for the education of undergraduate and postgraduate students and as a centre of excellence for research in a very wide range of fields.

Durham has one of the leading physics and astronomy departments in the UK. Each year we admit about 170 students to study for degrees in Physics, Physics & Astronomy and Theoretical Physics. Our research ranges from fundamental topics such as elementary particle physics and cosmology to applied areas such as material physics and biophysics.

The Centre for Advanced Instrumentation (CfAI) is one of the major research groups in the Physics Department with approximately 50 staff and research students. It is distributed across two physical locations, one of which is in the Department of Physics Rochester Building on the University science site in Durham, and the other in the NetPark Research Institute about ten miles away near Sedgefield. CfAI develops state-of-the-art instruments for application across a wide range of disciplines including astronomical instrumentation, biophysics, remote sensing and fusion diagnostics.

The CfAI has been using state of the art high performance computing (HPC) for many years both for simulations and for the real-time handling of data in astronomical telescopes. We are also a part of the wider HPC community within the university led by the Institute for Advanced Research Computing (iARC). Our staff includes engineers who specialize in software engineering and design and in the application of state of the art HPC hardware such as FPGAs and GPU accelerators. The scale of HPC applications, particularly in astronomy, has led to strong relationships between the university and HPC providers such as IBM and NVidia. Durham University is an NVidia CUDA (GPU) Research Centre.

The CfAI has a long record of effective management of substantial research funding for large projects both from UK research councils and from the EU. We have delivered large astronomical instruments both to ESO and to other observatories worldwide. Our staff structure provides for the roles of project scientists, project managers, system engineers and engineering and science specialists. The project will be led at Durham by Professor Richard Myers who is currently the head of the research section and a world leader in the field of adaptive optics. The day-to-day progress of the project will be overseen by Dr. Nigel Dipper, the head of software at CfAI who leads the HPC real-time control research program at Durham.

5.3 Microgate

Microgate is a SME based in Bolzano, Italy. The company operates in four fields: *Professional Timing*, *Training & Sport*, *Medical Rehab* and *Engineering*. The company expertise covers the whole development process of its products, including hardware, firmware and software design, as well as the production of medium quantities. All processes inside the company are monitored by means of certified quality systems (ISO9001 and ISO13485 certifications).

The *Engineering* department is mainly focused on large project for astronomy, and more specifically on adaptive optics systems for large telescopes. In this specific field, we have developed over the past 20 years the large, contactless deformable mirror technology, together with other Italian partners (ADS International, INAF-Osservatorio Astrofisico di Arcetri and the Aerospace Engineering Department of Politecnico di Milano). This technology has been already deployed realizing the

adaptive secondary mirror of several large telescopes, including the Multiple Mirror Telescope (MMT - Arizona), the Large Binocular Telescope (LBT – Arizona), the Magellan Telescope (Chile) and the ESO Very Large Telescope (VLT – Chile). Microgate is currently engaged in the design of the adaptive mirrors for the next generation of Extremely Large Telescopes (ELTs), in particular the ESO European Extremely Large Telescope (E-ELT) and the Giant Magellan Telescope (GMT).

These mirrors are 'massively controlled' structures which shape is determined by a large number – up to several thousands – of actuators. The real time control of such mirrors requires extreme computational throughput in a hard real time environment; moreover, physical and layout constraints impose a minimization of power consumption. In this frame, Microgate has developed a very specific know how in developing fully customized solutions based on large clusters of mixed DSP-FPGA boards and, more recently, on FPGAs only. This modular specific hardware, which embeds also the analog front-end for the digital control of the actuated system, has been employed also to perform other fundamental tasks in the adaptive optics loop, like wavefront sensing and real-time wavefront reconstruction.

Besides the hard real-time applications, the company expertise covers also middleware and less stringent real time control. A good example in this field is the control system and metrology for the European antennas of the ALMA (Atacama Large Millimetric Array) project, developed by Microgate in the past years. Also the user interfaces are developed in house: this is crucial aspect, not only for the commercial products of the *Professional Timing*, *Training & Sport*, *Medical Rehab*, but also for the engineering interfaces of the adaptive optics system.

The company know-how comprehends also dynamic simulation of complex multi-physics multiple-input multiple output systems by means of standard tools (Matlab, Comsol) and specifically developed simulation codes.

Concerning production, Microgate covers the full process from the components procurement, inspection, assembly of electronics boards (mainly performed though external well qualified suppliers with which we collaborate since long time), final integration of the products and testing, including also burn-in and environmental tests. Specifically automated test tools are mainly developed and maintained in-house, tracing all processes though the centralized quality assurance system.

The company facilities comprehend also a clean room (ISO 7), used mainly for the integration of the delicate adaptive optics instruments.

5.4 PLDA

PLDA aims to design and sell tools and building blocks for designers of FPGA's or ASIC components. It consists in:

- **Intellectual Property** (IP) reusable blocs targeting high performance standard communication protocols (currently PCIe and TCP/UDP)
- **Hardware boards** for prototyping and development, enabling the implementation of those functions in systems used by Customers in production,
    - Those boards are based on up to date versions of FPGA's provided by suppliers leading the FPGA market
    - Key features are
        - Data acquisition, transmission and restitution, with defined format and at determined characteristics of bandwidth and latency

    - Data pre- or post-processing
  - Field of applications include data captured from various analog or digital sensors and restituted similarly to analog or digital actioners, after some processing, or also data being used in large applications from various market domains (medicine, financial, scientific…)
- A fully integrated **FPGA application development system**, drastically enhancing the speed of development of FPGA based projects, assisting penetration of the FPGA's as a technology of choice in the field of hardware acceleration and High Performance Computing.

High speed interconnection standard protocols, FPGA's and elaborated FPGA development environments, along with production hardware platform to implement them, are key contributors to tangible progress in the field of High Performance Computing systems. Large scale Adaptive Optics mechanisms and related control systems provide a perfectly matching field of research and of experiment for up-to-date technologies developed by PLDA.

5.5 Consortium as a whole

The LESIA team at OdP and the CfAI team at UoD have close links with all of the instrument consortia for the E-ELT and are involved in all of the instruments at some level. This will facilitate the capture of the exact requirements of the RTC system for AO on all of the instruments. In particular, a close collaboration exists between the MAORY consortium, Durham University and Paris Observatory allowing for the straightforward capture of the core requirements for the MCAO prototype.Many new techniques in AO are required for the efficient operation of all of these instruments. These techniques need to be demonstrated to be practical. To this end an AO experiment (known as CANARY) has been built by Durham University and Paris Observatory and used on the 4 meter William Herschel telescope in La Palma, Canary Islands. This platform has enabled the testing of both new AO techniques such as LTAO and MCAO and the associated RTC systems at a reduced scale. Laboratory based AO benches at both UoD and OdP support this work. The RTC systems used for these AO test benches provide the starting point for this project.

A critical part of the requirements capture and prototypes definition will be the interfacing of the RTC with the telescope control system. This information will come from ESO. Both Durham University and Paris Observatory have a close relationship with ESO. Durham has provided elements of existing RTC systems for existing ESO telescopes (the VLT). This interface is currently being designed at ESO and will be available in a preliminary form during the requirements capture phase.

PLDA aims at designing and selling tools and building blocks for designers of FPGA's or ASIC components such as Intellectual Property (IP), hardware boards and FPGA development environment. Some of the other partners (OdP, Mic) have been using PLDA solutions in the past. QuickPlay, a new product design by PLDA, should reach the market during the course of the Green Flash project and is by design an open project and numerous improvements of the tool are targeted by PLDA. Moreover, it does perfectly match the objectives of Grenn Flash, and will boost up the productivity of the whole project. High speed interconnection standard protocols, FPGA and elaborated FPGA development environments, along with production hardware platform to implement them, are key contributors to tangible progress in the field of High Performance Computing systems. Large scale Adaptive Optics

mechanisms and related control systems provide a perfectly matching field of research and of experiment for up-to-date technologies developed by PLDA.

Microgate has developed in the past a custom hardware platform dedicated to specific control tasks for adaptive optics and have deployed systems on large telescopes including MMT, LBT, Magellan, Keck, VLT. More specifically, the custom hardware has been initially designed to cope with the specific goal of controlling large adaptive mirrors based on contactless voice-coil technology. More recently, in the frame of the development of the adaptive mirrors for the ELTs (E-ELT and GMT) with several thousands of controlled channels, the availability of more powerful and flexible FPGAs provided an innovative solution for this application in terms of power consumption and performance. The same could apply to the AO RTC module. In this frame, Microgate has developed a very significant experience in designing custom electronics with similar goals to the specific ones of this project. This represents a solid background to assure the technical success of the proposed development.