

## **UCLA**

### **Adaptive Optics for Extremely Large Telescopes 4 - Conference Proceedings**

#### **Title**

The use of CPU, GPU and FPGA in real-time control of adaptive optics systems

#### **Permalink**

<https://escholarship.org/uc/item/2vj6w3gm>

#### **Journal**

Adaptive Optics for Extremely Large Telescopes 4 - Conference Proceedings, 1(1)

#### **Authors**

Rodríguez Ramos, Luis Fernando  
Diaz Garcia, Jose Javier  
Fernández Valdivia, Juan Jose  
[et al.](#)

#### **Publication Date**

2015

#### **DOI**

10.20353/K3T4CP1131563

#### **Copyright Information**

Copyright 2015 by the author(s). All rights reserved unless otherwise indicated. Contact the author(s) for any necessary permissions. Learn more at <https://escholarship.org/terms>

Peer reviewed

# The use of CPU, GPU and FPGA in real-time control of adaptive optics systems

L.F. Rodríguez Ramos<sup>\*a</sup>, J.J. Díaz García<sup>a</sup>, J.J. Fernández Valdivia<sup>b</sup>,  
H.M. Chulani<sup>a</sup>, C. Colodro Conde<sup>c</sup>, J.M. Rodríguez Ramos<sup>b</sup>

<sup>a</sup>Institute of Astrophysics of the Canary Islands (IAC). Vía Láctea s/n 38200 La Laguna (SPAIN);

<sup>b</sup>University of La Laguna, Tenerife (SPAIN)

<sup>c</sup>Universidad Politécnica de Cartagena, Murcia (SPAIN)

## ABSTRACT

The use of deformable mirrors for the compensation of the effect of the atmospheric turbulence in Adaptive Optics (AO) systems requires a very fast computation of the adequate commands with the information provided by wavefront sensors. Providing accurate results with minimum latency is vital to achieve optimum performance. Conventional CPUs, GPUs and also FPGAs have been successfully used in real-time control for AO at IAC for a number of projects (EDIFISE, AOLI, AOonFPGA, GTCAO). Based in this experience, a comparative description is made in this paper, pointing out advantages and drawbacks of each solution as seen by each of the projects.

**Keywords:** CPU, GPU, FPGA, Real-time control, latency, closed-loop control, turbulence compensation, adaptive optics

## 1. INTRODUCTION

Adaptive Optics control is undoubtedly one of the most interesting control problems being addressed nowadays. The huge amount of actuators required for the compensation of the atmospheric turbulence in Extremely Large Telescopes (ELT), distributed in several physical mirrors when wide field correction is required, approaching to the range of  $10^4$ , the cycle time and low latency involved in an effective compensation of the turbulence (a few  $10^{-4}$  s), and the number of sensing inputs, normally wavefront sensors with a number of elements matching the number of actuators ( $10^4$ ), define an overwhelming control problem which will require a correspondingly high computing capability to cope with.

Our group identified this control problem long time ago<sup>1,2</sup> and have explored the use of commercially available specialized hardware in a number of AO projects, each of them having their own scientific objectives and specific requirements, but having in common the need of a low-latency control loop capable of providing optimum atmospheric turbulence compensation. Three different commercial products have been used as a basis for the implementation of the control system, or combinations of the three: CPU, GPU and FPGA. Having used all three of them, we are now in a situation where we can cross-examine their pros and cons, and provide a judgement based on real experience which may help other groups in the decision making process of future control systems.

This introduction will be followed by a brief summary of the main characteristics of the three technologies being compared. Then the three scientific projects will be separately described and finally the technologies will be compared with respect to a number of relevant concepts, pointing out their pros and cons.

## 2. TECHNOLOGIES FOR CONTROL COMPUTATION

Three basic technologies are mature and commercially available for high-end control applications among many other less common solutions ranging from analog computing to proprietary imaging processing units. They are widely known after their acronyms, the classical CPU (Central Processing Unit), the GPU (Graphics Processing Unit), and the FPGA (Field Programmable Gate Array).

\*lrr@iac.es; phone +34 922 605 200; fax +34 922 605 210; www.iac.es

## 2.1 CPU

Central Processing Units, CPUs, are by far the most widespread processing unit existing nowadays. Present in each PC computer, laptop or mobile phone, the CPU is probably the paradigm of the success of the microelectronics development and one of the pillars of the internet era. The continuous improvement of its processing power, outlined in Moore's Law doubling roughly every 18 months, and their very easy commercial availability at extremely low prices, have made CPU option the easiest way to cope with the processing needs of adaptive optics systems.

CPUs traditionally carries out the instructions of a sequential computer program, previously written, performing low level operations (arithmetic, logical...) on the data made available to it through input/output operations, with the help of intermediate registers, memory, and other peripherals. CPUs are normally offered to the user as a core part of a much higher system, a computer, which relies in an underlying program, the operating system (OS), to undertake a huge number of tasks required for obtaining results with a reasonable low amount of effort from the user side. Today's CPUs are really a combination of a number of processing units, "cores", each of them capable of performing simultaneously their tasks separately and in a parallel way.

Programming of the control algorithms to be executed by CPUs should not strictly require the use of any operating system, but it is almost unthinkable to undertake manually the many features which are managed effortlessly by the OS. Even the task of splitting the amount of operations between the available cores could be left to the OS, if so desired. On the contrary, the presence of an operating system often generates unexpected side effects in latency, jitter and determinism of the control behavior which will be discussed later.

## 2.2 GPU

Graphics Processing Units, GPUs, are microelectronic circuits developed for very fast processing of relatively simple tasks to be performed in a frame buffer intended for immediate display. They were designed for computer graphics and image processing, with an extremely parallel structure which allows to outperform CPUs by orders of magnitude in algorithms where a big amount of processing over the same data is required. Intended to accelerate geometric calculations such as the rotation and translation of vertices into different coordinate systems, adaptive optics control systems may benefit from the fact that these operations involve many matrix and vector operations, as it is required in many steps of the algorithms for turbulence compensation.

GPU circuits may be included in PC motherboards depending on the features provided, but in the case of adaptive optics control system they will normally be placed in dedicated add-on boards, installed using either PCIe or AGP interface ports. Data exchange between CPU and GPU must be made using those interfaces, and may become the limiting factor in many cases.

GPU programming is normally done using extensions of the C programming language such as CUDA or OpenCL. CUDA is devoted to NVIDIA GPUs whilst OpenCL can be used with a number of architectures, including also CPU and DSP.

## 2.3 FPGA

Field Programmable Gate Arrays, FPGAs, are also microelectronic circuits designed to be configured by a designer, many times, long time after it was manufactured, hence the name "field programmable". Instead of having a fixed core capable of performing sequentially a number of operations (out of a certain set) over some data, the processing itself may be coded using the silicon real state existing at the chip, providing virtually any combination of operations, and more important, which can be executed in a real parallel way because they are physically done at different locations within the integrated circuit. The coding might also be changed, sometimes even during its execution, in order to adapt to new situations or to cope with different scopes. Sequential processors can also be implemented using the FPGA logic made available to the designer, and hence they become a particular case of the more general FPGA processing capability.

FPGA programming is done mostly in HDL languages (VHDL, Verilog...) which imposes some burden to the programmer due to its low-level nature. There are also higher level cross-compilers commercially available which may be used up to a point to ease the programming task, with interesting results in some general cases. The normal design approach is to use a top-down paradigm, splitting the processing into modules which are defined with very clear input/output interfaces and performed tasks, and which are replicated and interconnected as needed to provide the required overall processing.

The FPGA superb flexibility comes, as always, with a price: the look-up tables normally used as core processing blocks can only be made to work roughly at a clock speed one order of magnitude lower than fixed processors.

### 3. ADAPTIVE OPTICS PROJECTS CONSIDERED

The approach used in this paper is strongly oriented to the use of real experience when evaluating pros and cons of the technologies under comparison. Three different AO scientific projects will be used to this end because they are all being developed with the participation of our group at IAC, with different consortia. They have rather different requirements sets and scientific objectives, and they are not in the same developmental status, but they are advanced enough to allow for a reasonable comparison, and thus to point out the advantages and drawbacks of the hardware solution chosen in each case. A short description of each of them follows, with links to more detailed information.

#### 3.1 EDiFiSE

EDiFiSE (Equalized and Diffraction-limited Field Spectrograph Experiment) is a prototype featuring an Adaptive Optics subsystem, with both high and low order corrections, and a spectrograph with equalized integral field unit (EIFU)<sup>3</sup>. It is intended to be a test platform to verify the technology and its viability in bigger projects for large telescopes, where available information regarding the real-time statistics of the atmosphere is used for optimizing the configuration of the system. Both subsystems, AO and EIFU should then be designed concurrently to fulfill this purpose. The main scientific case aiming the technique is related to compact objects with high intensity contrast. The resolved detection of the spatial components will be used for improving both spatial and spectral resolution.

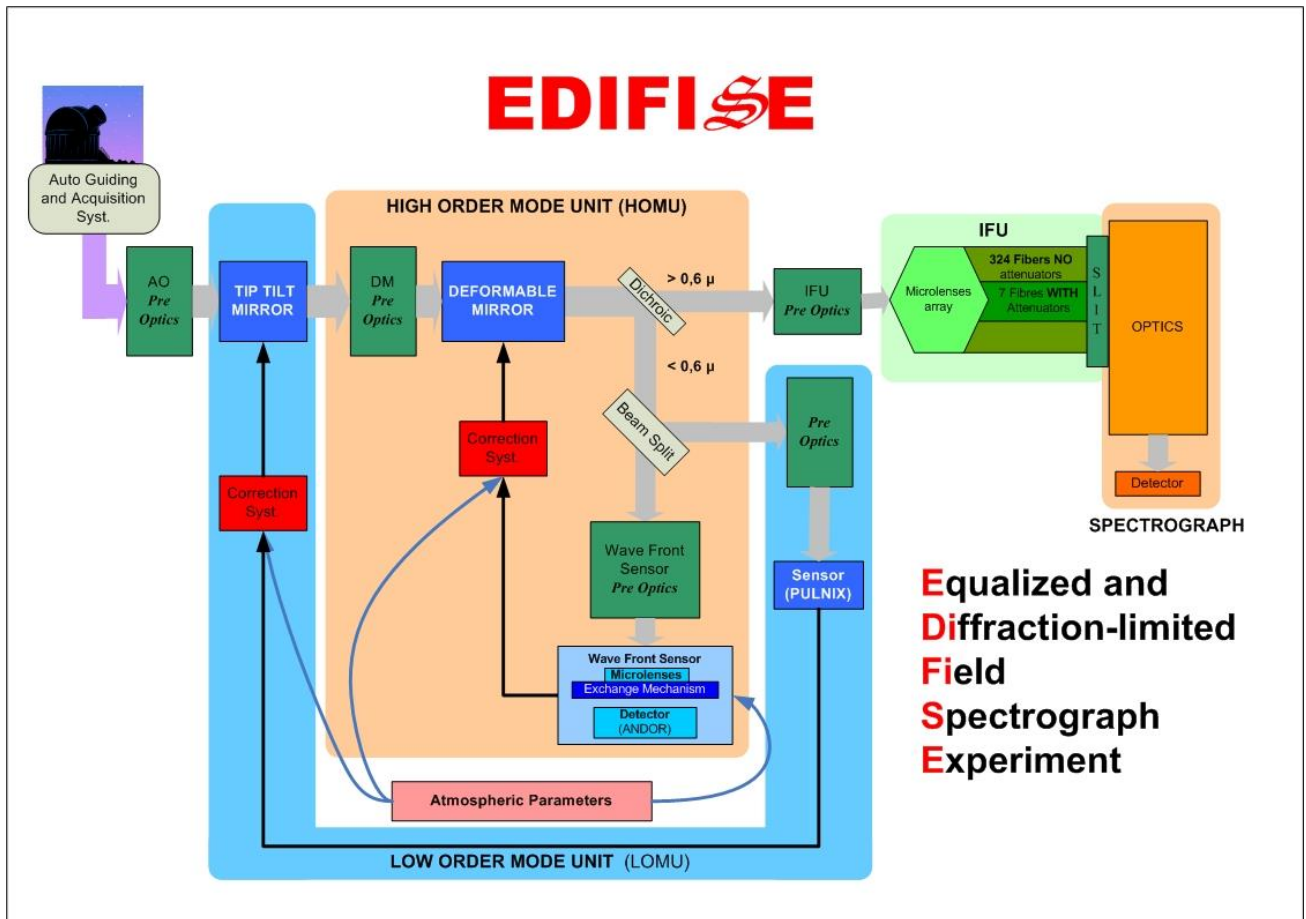


Figure 1. Overall diagram for EDiFiSE project. The three main subsystems are depicted sequentially, and the capability of adaptation to existing turbulence parameters is pointed out.

A two-loop arrangement has been selected for the correction subsystem: An outer loop devoted to low-order correction (tip-tilt) plus an inner loop dealing with high-order corrections. The two loops are conceptually and physically independent in order to guarantee the stability of the control system, and they are designed to work at very different speeds. This arrangement provides also a maximum of quality in the measurement of the low-order aberration by devoting a specific sensor to this only end, instead of deriving this information from the high order wavefront sensor. This higher quality measurement is done at the expense of the use of about only 5% of the light flux available for atmospheric turbulence correction, which is considered rather cost effective from the systems engineering point of view.

The sensors selected for both low and high-order correction are electron-multiplied CCDs, specifically the model IXON 128x128 pixels manufactured by ANDOR, which provides virtually no read-out noise and are thus adequate for low level fluxes. For the tip-tilt measurement, a combination of both windowing and binning capabilities provides the required readout speed and thus the servo loop working frequency.

Low-order correction is implemented by a commercial S-330.2SL platform from Physik Instrumente, and high order correction is based on the deformable mirror DM 97-15 from ALPAO, featuring 11 actuators at the diameter with 1.5 mm spacing. Real time control electronics are also independent for low and high-order correction and have been implemented directly on FPGA under VHDL programming, as a stream processor instead of a conventional Von-Neumann arrangement, using a data driven scheme without a "program" to be executed by a processor. An ancillary PC computer with a LABVIEW program serves as user interface and for system setup and telemetry data storage.

The relevant characteristics from the adaptive optics control point of view can be summarized as follows. The ANDOR camera output is split from the proprietary interface board in order to make it available to the FPGA processor. An input module compensate the Shack-Hartmann wavefront sensor image from bias and flat, and up to 16x16 centroids are computed, with thresholding and sub-pixel resolution. Errors are obtained by comparing actual centroid positions with a reference and then a matrix-vector multiplication, with a pre-loaded matrix, is computed to obtain a zonal or modal decomposition of the measured wavefront. After filtering modes with pre-loaded coefficients, they are converted to actuations by a second matrix-vector multiplication and the result is sent to the deformable mirror controller using UDP/IP Ethernet link.

## 3.2 AOLI

The Adaptive Optics Lucky Imager (AOLI)<sup>4</sup> instrument is designed to combine the techniques of AO and Lucky Imaging into a single instrument. It consists of a geometric curvature WFS and a low order adaptive optics wavefront corrector using a deformable mirror in conjunction with a wide-field, array detector Lucky Imaging camera. A calibration subsystem is also present in order to provide a reference flat wavefront and also to simulate atmospheric-like wavefront distortions for day-time system tuning.

The AOLI setup is depicted in figure 2. The light is collimated and passed through an atmospheric dispersion corrector before it strikes a deformable mirror (DM). We selected the ALPAO DM241 unit as it has excellent stability and provides an unusually long stroke. The light is then reflected and, after a fold mirror, is reimaged on to a pickoff mirror. The pickoff mirror deflects light towards the science camera, where a 2x2 array of EMCCDs provide a wide field selectable from 0.5 to 2 arc minutes on sky. The EMCCDs are back illuminated (thinned) with very high quantum efficiency (peak >95%) from E2V Technologies. Custom electronics developed in Cambridge give up to 30 MHz pixel rate, and 25 (full) frames per sec.

The light from the reference star goes directly on to the wavefront sensor (WFS). The telescope pupil is reimaged down to approximately 2 mm diameter from its original 4.2 m (WHT telescope, Canary Islands, Spain). The wavefront sensor arrangement uses two images at both sides of the pupil, obtained by beam splitting, in order to measure the shape of the incoming wavefront. A smaller EMCCD is used for the WFS in order to obtain speeds of 100 frames per second, adequate for the very faint stars planned to be used for guiding.

From the control point of view, the two images coming out from the geometric wavefront sensor, once digitized, arrives at the processing PC computer through an ADLINK frame grabber board. A shared memory mechanism allows the simultaneous storage and transfer to the K40 NVIDIA GPU board installed in the PC. The GPU undertakes the processing required for extracting the Zernike expansion which better describe the wavefront, and delivers it to the CPU, which is then in charge of filtering the modes (also mode converting if so desired) and computing the optimal actuations for the

deformable mirror. These actuations are directly sent to the ALPAO mirror driver using the proprietary interface supplied by the manufacturer.

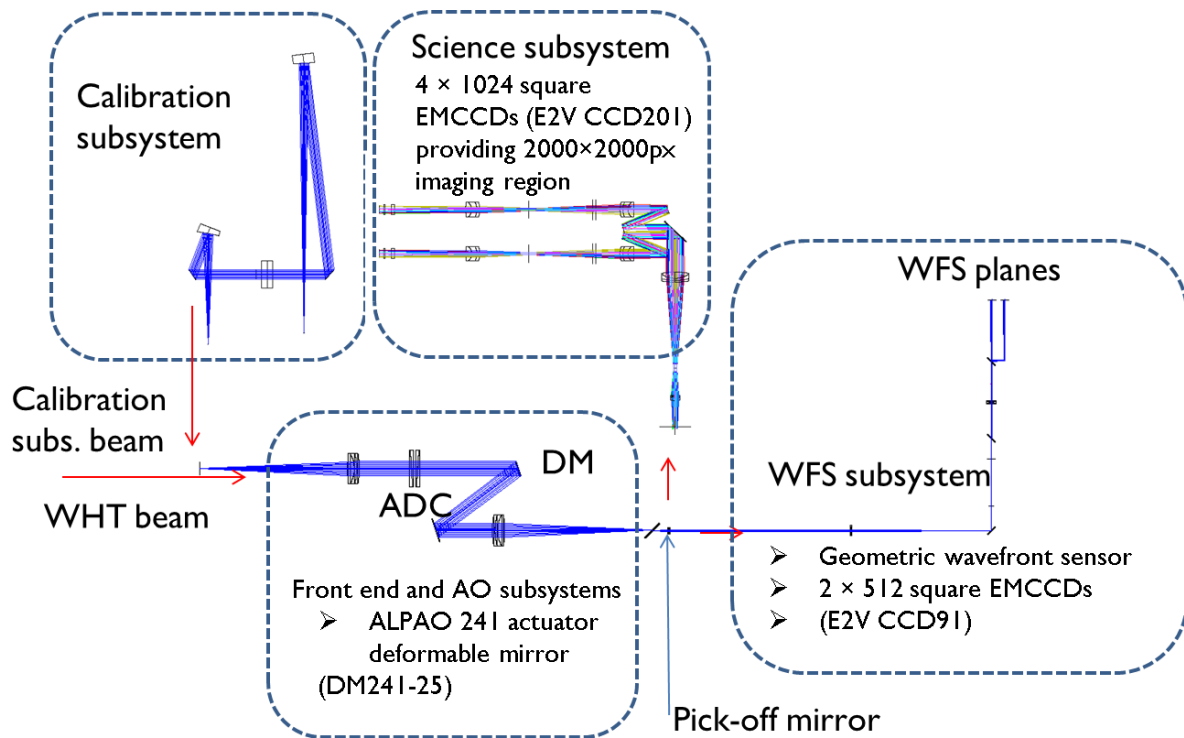


Figure 2. Overall diagram for AOLI project. Four main subsystems are depicted: Calibration, science, wavefront correction and wavefront sensor.

### 3.3 GTCAO

GTCAO is the first step in the Adaptive Optics capability of the 10m GTC telescope. The system provides a single deformable mirror conjugated to the telescope pupil. The wavefront sensor will be a Shack-Hartmann sensor, operating in the visible. The instrument is designed to provide a corrected beam that will achieve a Strehl Ratio (SR) of 0.65 in K-band with bright guide stars. The size of the transmitted (and de-rotated) field of view is 1.5 arcmin diameter. The optical layout of the GTCAO system includes an ADC in order not to degrade its performance with increasing zenith angle. The ADC is designed to work up to  $60^\circ$  zenith angle. The GTC Adaptive Optics Instrument (FRIDA) will operate together with the GTCAO system, sharing the Nasmyth platform.

Figure 3 depicts the optical arrangement of GTCAO. The beam entering the Nasmyth platform is de-rotated by a K-mirror and collimated by an off-axis parabola (OAP), which forms an image of the telescope pupil, with the help of a folding mirror, at the deformable mirror (DM). The wavefront corrector is a custom made CILAS SAM373, with 21 actuators in diameter spaced 7 mm horizontally and 6.96 mm vertically. A second OAP recovers the original F# for the output beam, which is split in wavelength by a dichroic, delivering the infrared radiation to the science instrument, after correcting for atmospheric dispersion.

The WFS has two operation modes: a high order mode (HOWFS) with a  $20 \times 20$  lenslet array and a low order mode (LOWFS) for tip-tilt and defocus sensing with a  $2 \times 2$  lenslet array, both with a Natural Guide Star. It also features the capability of ADC and pupil position adjustment, and is equipped with three collimator-camera stages for optical relaying and scale adjustment (not shown in the diagram). The detector used is an OCAM2 camera manufactured in France by FirstLight, based on the E2V CCD220 EMCCD sensor, running up to 1500 frames per second.

From the real-time control point of view, the relevant parameters are the Camera-link nature of the output interface of the WFS camera, which is connected to a frame grabber board plugged on the PC computer. The wavefront sensor image is

recomposed and processed using multicore CPU software (DARC<sup>5</sup> under evaluation), and the computed actuations are sent to the deformable mirror using sFPDP fiber optic interface.

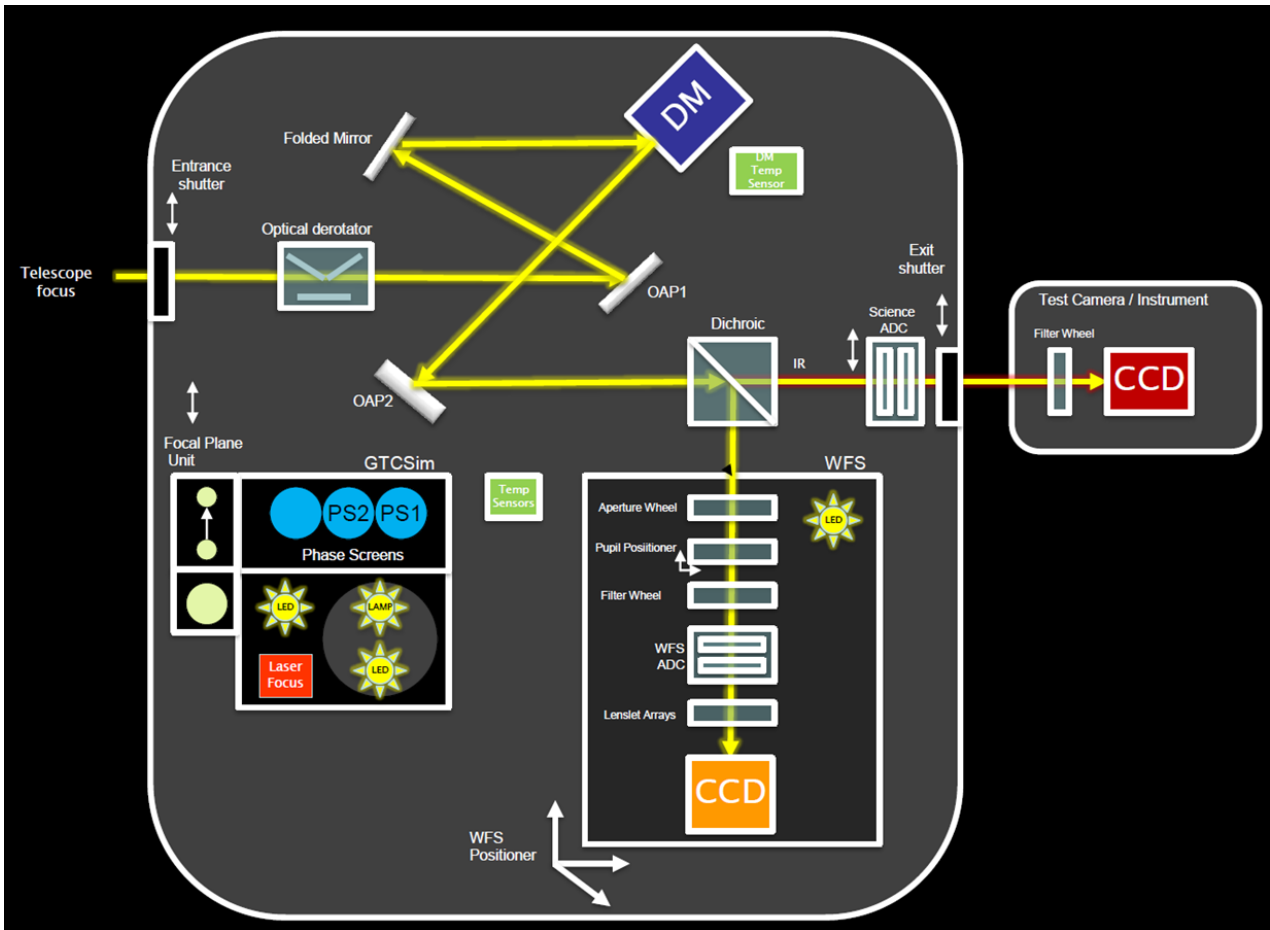


Figure 3. Overall diagram for GTCAO project. Main components are drawn, paying special attention to the calibration simulator GTCSim and the Wavefront sensor.

#### 4. TECHNOLOGY COMPARISON

A number of parameters have been identified as relevant for the development, and all other lifetime stages, of a real-time control system for an adaptive optics instrument. They are not the only existing ones, for sure, but undoubtedly they play an important role in the decision making process regarding how the required computations are going to be made and which architecture should be used to fulfill requirements in a cost effective way. In this chapter the parameters will be defined and an assessment of up to what point each of the technologies, used in the three projects, behave with respect to these parameters will be made.

The parameters have been listed in alphabetical order, pointing out that no particular analysis of their relative relevance has been made, which is intentionally left to the reader. In each case we have summarized the assessment with a graphical character showing advantage, disadvantage or equality, to our understanding.

##### 4.1 Development cost

We have named as “development cost” as the amount of engineering effort required for the development of the real-time control system of each project. These estimations have been obtained partly from personal time reports but also with some

corrections for taking into account the use of subcontracting in some cases, and also some extrapolation needed for the cases where the work is not 100% completed at the time of writing.

Development cost		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
Estimated 3 engineer-year, including the subcontract of part of the FPGA development.	Estimated 2 engineer-year, including both GPU program and CPU control	Estimated 1 engineer-year, making full re-use of available software
👎	≈	👍

#### 4.2 Flexibility

We understand as “flexibility” as the capability of the system for accepting changes, specifically in code. These changes may be originated by improvements in algorithms, system architectures, components, or much simpler things like debugging when integration or verification.

Flexibility		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
Any change in the system requires the synthesis of the new code, which can take dozens of minutes and may present design problems.	Software change only requires compilation and building.	Software change only requires compilation and building.
👎	👍	👍

#### 4.3 Hardware cost

The cost of the hardware required for the execution of the algorithms is not normally an issue when comparing with the rest of the very specialized components of an adaptive optics system. However it could regain relevance in the future when adaptive optics systems might eventually be commercially available in reasonable quantities.

Hardware cost		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
A couple of general purpose development boards: ≈3 K€	Desktop PC, highly equipped, plus K40 GPU board: ≈9 K€	Rack-mounted PC, highly equipped in both memory and disk: ≈5 K€
👍	👎	≈



#### 4.4 Jitter

We have named as “jitter” the variation in the latency (see below). The way in which the different software packages are managed by the operating systems, and specifically their distribution between different processors and cores, normally generates a certain variation in the latency, which may be relevant when long time delays are found, because they will for sure generate a sporadic loss of performance, and, more important, may lead to instabilities in the closed loop behavior.

Jitter		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
The FPGA logic is completely deterministic and its state can be known down to every clock cycle. Jitter is thus negligible.	GPU processing is virtually jitter-free, but the operating system (Linux) is running all time. Jitter peaks of several iterations are common.	The use of a real-time Kernel provides a reasonably low jitter behaviour. Jitter peaks of several iterations can be observed.
👍	👎	≈

#### 4.5 Latency

Latency is a key figure for the overall correction capability of a closed loop control system. The CCD nature of the camera detector and the physical behavior of the deformable mirror put some constraints to the minimum latency to be obtained, but the amount of time devoted to the calculations should be kept to a minimum, comparatively with those. In this framework, calculation latency can be considered as the amount of time elapsed since the arrival of the last pixel from the camera to the availability of the actuation command.

Latency		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
Latency in the FPGA processing is really low, due to the intrinsically parallel nature of the logic. Less than 10% of loop time (~150 μs)	Key figure for latency is the “bottleneck” of the delivery of the images to the GPU (~2ms). However, it can be accepted due to the relative low speed of this loop (100 Hz)	The processing can be broken down among many threads and cores, minimising latency. It has been estimated to ~250 μs
👍	👎	≈

#### 4.6 Power consumption

The amount of power consumed by the computing hardware may not be negligible in comparison to other components of the adaptive optics system, and is gaining relevance nowadays where global warming has become an issue along the whole world. It has also an extra relevance in adaptive optics due to the fact that this power shall be specifically handled, and by no means may be allowed to produce further turbulence in the telescope area.

Power consumption		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
Power consumption of FPGA itself and related logic, can be estimated in the really low figure of 4 watts.	The power consumption directly related to the CPU processor can be estimated in 100 watts. GPU board is specified at 235 W.	As already cited in the AOLI case, the CPU related power consumption is in the order of 100 W, depending on the number of cores being used.
👍	👎	👎

#### 4.7 Reusability

The concept of reusability of the software codes is very important nowadays because it can save a lot of engineering effort by using already developed and tested, ready to use, existing code instead of developing and verifying new pieces.

Reusability		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
Design in the VHDL field is normally based on reusable IP modules.	CUDA programs are specific to NVIDIA GPU boards, but there are many of them available.	CPU real-time control is making full use of available software.
👍	≈	👍

#### 4.8 Skill level

The skill level of the engineers undertaking the development could eventually become an issue, depending on the amount of well-trained people existing available to be hired. The nature of the skills and the required background of the engineers is not the same for different technologies, and comments are summarized in the table below:

Skill level		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
High level of VHDL design knowledge is required, not easily found. A background in electronics is preferable.	C and C++ under Linux programming experience required, plus CUDA.	C, C++ and python programming experience required, under Linux.
👎	≈	👍

#### 4.9 Volume

The volume associated to processors is not normally an issue due to its extremely low dimensions, especially when comparing with other components of an earth-based adaptive optics system. However, if eventually such a system will be made to fly in a satellite, or any moving vehicle, this parameter could become very important. We are not considering strictly the size of the integrated circuit, but instead the real component which allows it to function properly.

Volume		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
A 1U 19" rack is enough for allocating both the low-order and the high-order processing FPGAs	A desktop PC computer, with extra powerful power supply, has been selected for allocating the GPU board and frame grabbers.	A 4U, 19" rack mounted PC, short depth, has been used for the allocation of all processors and frame grabbers.
👍	👎	👎

#### 4.10 Weight

Weight is again a parameter of no concern in an earth-based system. However, as in the case of volume, it might become important if the adaptive optics system is intended to be installed in any moving vehicle, or even if it is going to be attached to any moving part of a telescope.

Weight		
EDiFiSE (FPGA)	AOLI (GPU+CPU)	GTCAO (CPU)
A fairly simple PCB board is enough for the computation and servicing of the real-time loop. A Xilinx general purpose board weighting a few hundred grams is used.	The rack mounted which allocates the K40 GPU board weights 13 Kg. The GPU board is specified at 826 gr.	The use of a highly featured PC computer, rack mounted, as previously said for the AOLI case, is in the order of 13 Kg, two orders of magnitude greater than the FPGA case.
👍	👎	👎

### CONCLUSIONS

A number of relevant parameters have been used, to cross-evaluate the advantages and drawbacks of the three different technologies, readily available to implement all calculations required in an adaptive optics control system. The assessment has been made using the experience obtained after developing at IAC three adaptive optics instruments, each of them having different scientific requirements and approaches to the computation architecture.

As it could have been expected, each technology has its pros and cons and, to our understanding, no overall technology winner can be set as the outcome of this assessment. Instead, each specific problem in the future should make their own study and examine its requirements and restrictions, in order to select the best option.

## REFERENCES

- [1] Marichal-Hernández, J. G.; Rodríguez-Ramos, L. F.; Rosa, F. & Rodríguez-Ramos, J. M. (2005), 'Atmospheric wavefront phase recovery by use of specialized hardware: graphical processing units and field-programmable gate arrays', *Applied optics* 44(35), 7587--7594.
- [2] L.F. Rodríguez Ramos et al., "FPGA adaptive optics system test bench", *Proc. SPIE* 5903, 120-128 (2005).
- [3] García-Lorenzo, B, et al. "EDiFiSE: equalized and diffraction-limited field spectrograph experiment", *Proc. SPIE* 7014, *Ground-based and Airborne Instrumentation for Astronomy II*, 70144B (July 11, 2008); doi:10.1117/12.787848;
- [4] Craig Mackay et al. " AOLI: Adaptive Optics Lucky Imager: diffraction limited imaging in the visible on large ground-based telescopes ", *Proc. SPIE* 8446, *Ground-based and Airborne Instrumentation for Astronomy IV*, 844621 (September 24, 2012); doi:10.1117/12.925618.
- [5] Basden, A. G., & Myers, R. M. (2012). The Durham adaptive optics real-time controller: capability and Extremely Large Telescope suitability. *Monthly Notices of the Royal Astronomical Society*, 424(2), 1483-1494.