l'Observatoire de Paris — LESIA
Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique

GPU accelerated simulations and real-time control of the E-ELT adaptive optics systems

Damien Gratadour, Arnaud Sevin, Julien Brulé
LESIA, Observatoire de Paris, France

PHASE

Introduction

Adaptive Optics (AO): instrumental technique for the **correction of dynamically evolving aberrations** in an optical system. It is used on astronomical telescopes to **compensate, in real-time, for the effect of atmospheric turbulence**, providing a significant **improvement in resolution** (up to a factor 10 for a 10m telescope), especially in the near-infrared (NIR).

European Extremely large Telescope (E-ELT): provide Europe with the biggest eye on the universe ever built. **39m diameter telescope** (800 hexagonal segments mounted on a 2800 tons structure). **AO is at the core of the telescope construction** plan: two first light instruments equipped with AO. First light foreseen in 2022. Designed to allow astrophysicists to make a leap forward in our knowledge of the Universe and the conditions of appearance of life.

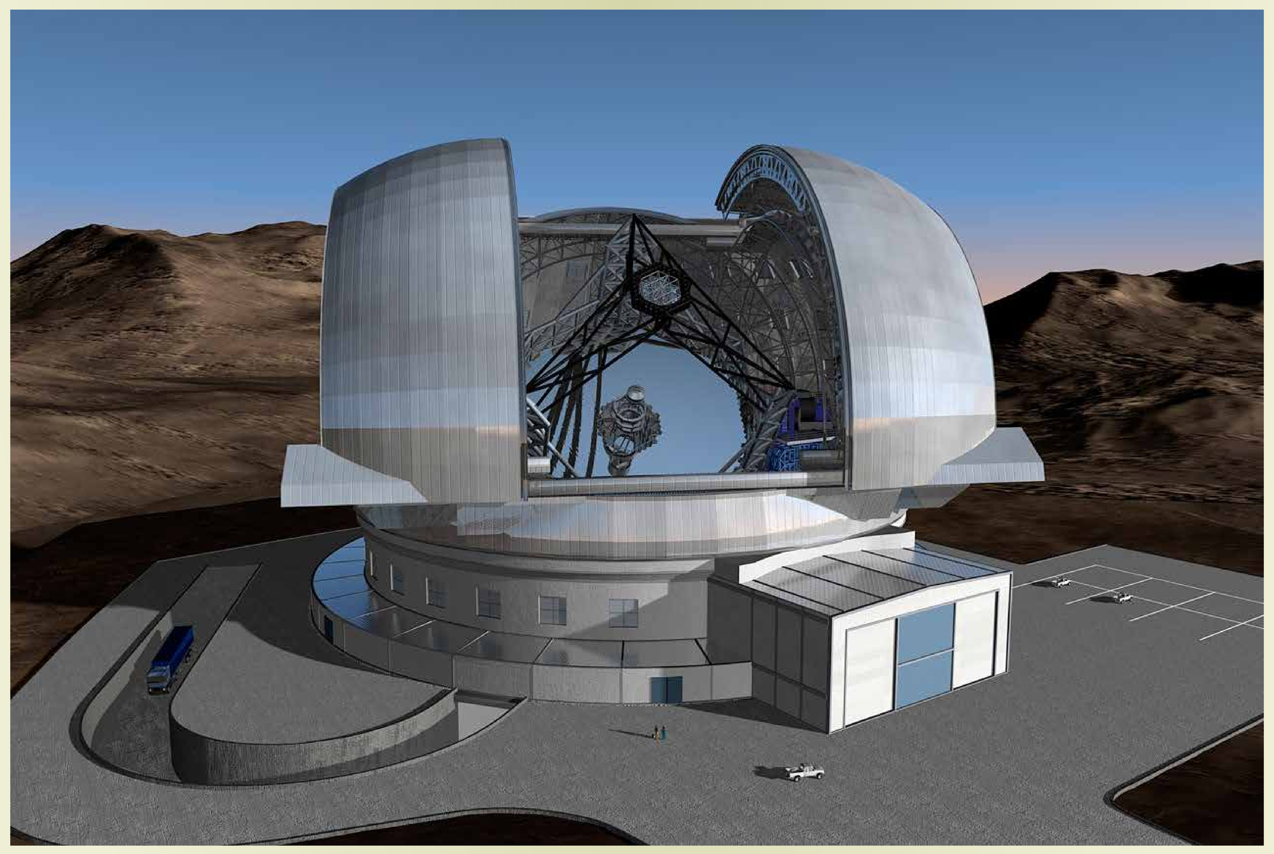

Several key technological developments for the E-ELT instrumentation have been identified during the initial design studies as top priorities. Among these, several are related to AO and **require realistic numerical simulations** to be studied at full scale. Additionally, the foreseen **AO systems will require an unprecedented amount of computing power** to be driven in real-time. GPUs are very attractive solutions to provide the required power.

YoGA: Yorick with GPU Acceleration

Complex systems simulations benefit from the use of an interpreted language providing a **comprehensive interface** to design and use the code.

YOGA — YOrick with Gpu Acceleration

Substantial efforts have been led at LESIA during the past 2 years to develop a simulation code able to target the E-ELT Size. YoGA is an **original binding** between **Yorick, an interpreted programming language** for scientific simulations or computations, and CUDA. YoGA_AO is an adaptive optics simulation code designed to run on NVIDIA GPUs and written as an extension of YoGA. It provides in **average a speedup of 10x as compared to available CPU simulation codes**.

https://dev-lesia.obspm.fr/projets/projects/yoga/wiki

It is highly configurable and allows to simulate a variety of systems from single conjugate natural guide star AO to multiple laser guide stars systems. It is available for free under BSD licence

End-to-end simulation of an AO system

The simulation of an AO system involves **multiple physics** from atmospheric turbulence models to tomographic reconstruction to control theory. Due to the **stochastic** nature of the turbulence, **Monte Carlo** simulations provide the most realistic results, the extremely large diameter of the E-ELT making each iteration a large scale problem.

To compensate for atmospheric turbulence in "**real-time**", an AO system must provide about **1000 corrections per second**. Hence the simulation of an observation represents at least few 10k iterations.

It includes a multiple layers turbulence model, a Shack-Hartmann wavefront sensor model for natural and laser guide stars, a simple real-time computer module and a realistic deformable mirror model.

Turbulence is generated on the fly on the GPU by the extrusion of Kolmogorov phase screen ribbons. The model mainly involves fast Fourier transforms (FFT) and random numbers generation. While the individual FFTs are small, the code benefits greatly from an optimal dispatch of the workload using native batched FFTs in CUDA.

Simulations speeds of few **1000 iterations/s** are reached for classical AO systems on 10m telescopes: this is faster than real-time control alone ! Additionally, several 100 iterations/s are reached for classical AO on the E-ELT: **quasi-real-time simulations of the whole system and its environment.**

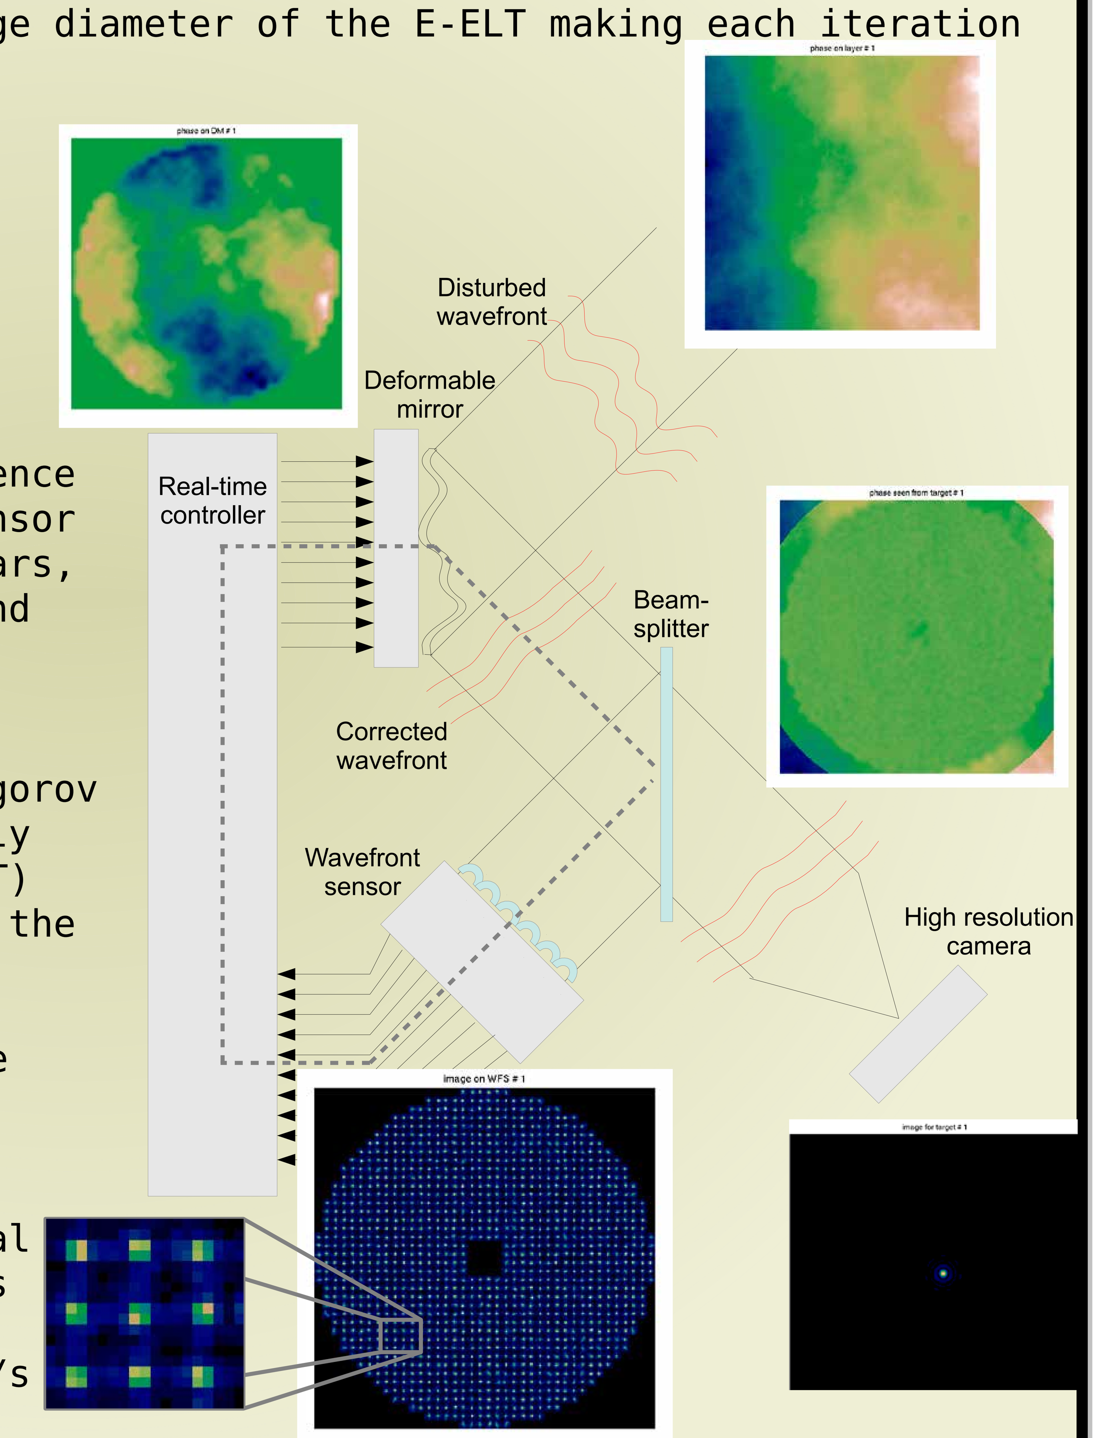

Simulations performance

The following tables contain profiles (in ms) for a full end-to-end simulation of a SCAO system (sub-apertures of 50cm diameters) for various telescope diameters and on 2 different GPUS (NVIDIA Tesla C2070 and NVIDIA Tesla M2090). Most of the computing time is spent on WFS modeling and DM shape computation. One striking result is that **an off-the-shelf high end GPU already provide the computing power** to drive such systems on sky (even at the E-ELT scale) as shown by the sixth and seventh columns. The second GPU has about 30% more processors and the performance for the most intensive computing tasks (WFS and DM) is about 30% better for this GPU showing the ability of the code to auto-tune and deliver better performance when possible.

Telescope diam.	Turbu generation	Raytracing turbu	Raytracing DM	WFS	COG	Control	DM shape computation	Raytracing target
4m	0.131	0.010	0.011	0.179	0.015	0.022	0.170	0.010
8m	0.238	0.028	0.028	0.591	0.024	0.068	0.477	0.029
20m	0.680	0.178	0.178	4.06	0.106	0.435	3.59	0.179
30m	1.173	0.394	0.396	8.40	0.324	1.164	8.67	0.391
40m	1.83	0.692	0.695	12.21	0.408	3.075	15.67	0.691

Tesla C2070

Telescope diam.	Turbu generation	Raytracing turbu	Raytracing DM	WFS	COG	Control	DM shape computation	Raytracing target
4m	0.107	0.008	0.008	0.138	0.013	0.019	0.137	0.008
8m	0.192	0.022	0.023	0.459	0.031	0.060	0.562	0.023
20m	0.550	0.135	0.136	3.07	0.079	0.363	3.22	0.137
30m	0.927	0.299	0.300	6.73	0.168	0.915	7.39	0.302
40m	1.44	0.526	0.525	11.9	0.320	2.263	13.62	0.527

Tesla M2090

Unifying simulations and real-time frameworks

Simulation models and the real-time core can be developed in a **unified framework** so as to provide optimal simulation performance and full flexibility for the integration of complex control schemes in real systems. While accelerator-based architectures provide the perfect tools for simulations, the main technological bottle-neck for real-time applications is the **latency in data transfers**. In the real-time framework, efficiently hiding latency requires a full control over the **data acquisition protocol**.

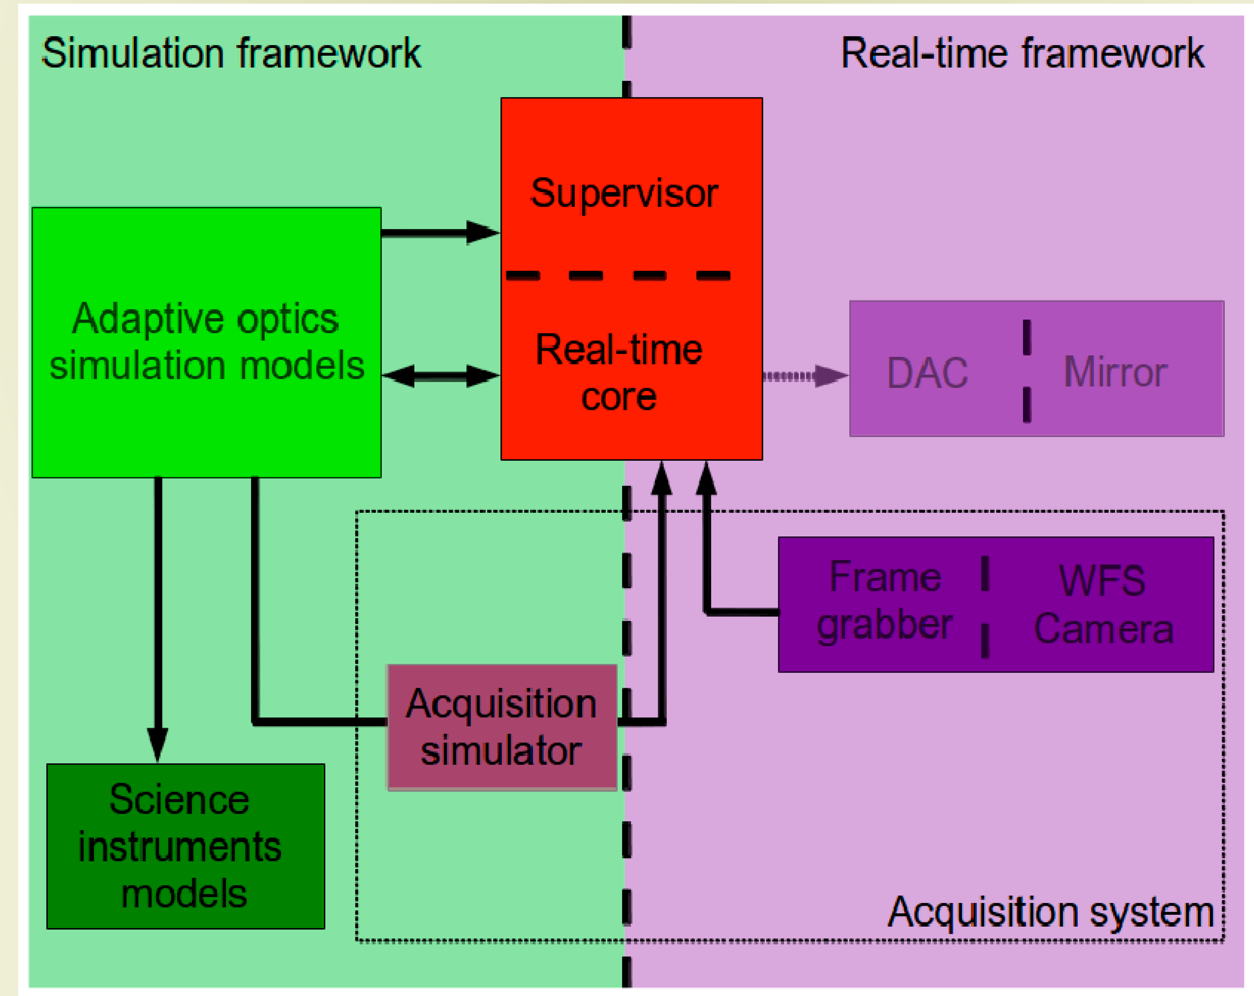

Toward a GPU-driven AO system

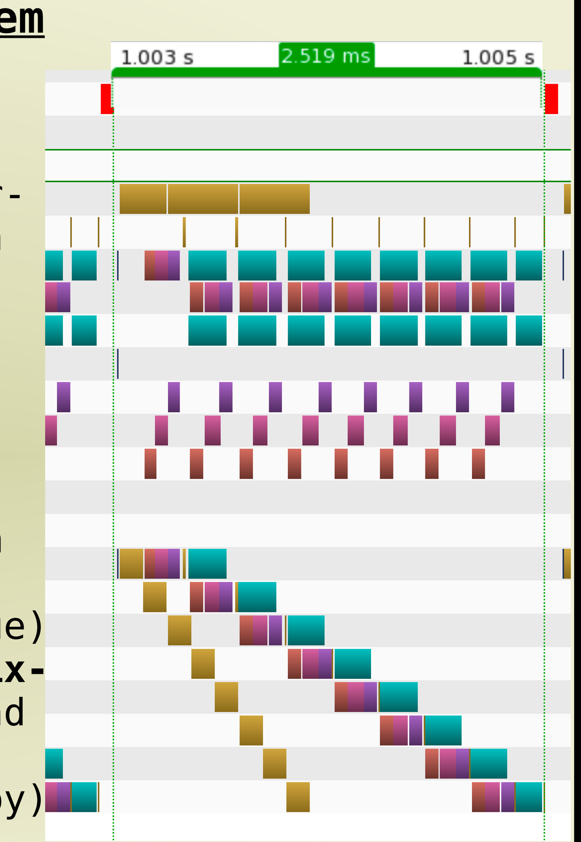

Profile of a RTC simulator demonstrating efficient hiding of memcopy latency when transferring the WFS image to the RTC in the case of a Shack-Hartmann.

In this code, we first transfer a frame asynchronously (dark yellow) from the host memory to the device, we then do several computations: **parallel reduction** on each sub-aperture (dark red), **centroiding** (violet and dark blue) on these sub-apertures and **matrix-vector multiply** using the command matrix (light blue). The overall execution time (including memcopy) is 2.52ms, compatible with real-time control at 400Hz. This profile was obtained on a Tesla C2070.

The COMPASS project

The main objective of the COMPASS project is to provide a **full scale end-to-end AO development platform** to the french high angular resolution community, able to address the E-ELT scale and including a real-time core that can be directly integrated on a real system. The end product, **a unified and optimized computing framework** (software and hardware), will address several major needs for research in AO including **a pathfinder for accelerator-based AO control**. The COMPASS project has been **awarded a 800k€ grant** by the French National Funding Agency and represents a **total investment of 2.5M€** from the 6 associated partners.