

RTC4AO 2016

Green Flash

A Real Time Controller For E-ELT Julien BERNARD



Project #671662 funded by European Commission under program H2020-EU.1.2.2 coordinated in H2020-FETHPC-2014



Green Green Flash

- Public and private actors
 - Paris Observatory
 - University of Durham
 - Microgate
 - PLDA



Laboratoire d'Études Spatiales et d'Instrumentation en Astrophysique







- Part of Horizon 2020 : EU Research and Innovation programme
- 3 years project
- 3,8 million €
- Involve about 30 people
- Research axes
 - Real time HPC with accelerators and smart interconnects
 - Energy efficient platform based on FPGA
 - Real Time Controller (RTC) prototype for European Extremely Large Telescope Adaptive Optics (AO) system





Green Adaptive Optics

- Compensate in real-time the wavefront perturbations
- Using a wavefront sensor - WFS to measure them
- Using a deformable mirror – DM to reshape the wavefront
- Commands to the mirror must be computed in real-time (~ms rate)



GPU ROMA 2016 Vorente Durham MICROSOGATE Paris

Flash

Green RTC AO prototype for E-ELT

- European Extremely Large Telescope
 - Next telescope generation
 - 39 m diameter telescope
 - 100m dome, 2800 tones structure rotating at 360°, seismic safe (Chile)
 - first light foreseen in 2024
 - European project led by ESO funded by 15 European countries
 - 1.2 billion € project







	Simple Conjugate Adaptive Optics	Multiple Conjugate Adaptive Optics
Wave Front Sensor	1 (10k measures)	6 (60k measures)
Deformable mirror	1 (5.3k degree of freedom)	3 (16k degree of freedom)
Frequency	1kHz	500Hz







Green RTC concept for ELT AO



GPU ROMA 2016

Green Computation : MCAO

- MCAO process
 - 1) Calculate the pseudo-open loop measurement vector with the two last command vector and the interaction matrix

 $\vec{M}_{ol}[k] = \vec{M}[k] - D(a\vec{c}[k-2] + (1-a)\vec{c}[k-1])$

2) Calculate the raw tomographic vector

 $\vec{e}[k] = R \vec{M}_{o1}[k]$

3) Get the vector command by smoothing the raw vector with the last command

 $\vec{c}[k] = q\vec{e} + (1-q)\vec{c}[k-1]$

- Sizing :
 - SCAO : 1 matrix x 10 048 x 5316 = 53 MMAC per iteration and 53 GMACS @ 1kHz
 - MCAO : 2 matrix x 60288 x 15000 = 1.8GMAC per iteration and 1 TMACS @ 500 Hz

- Matrix Vector Multiply MVM take more than 90% of the time
- I/O bound computation
- Depends of the sustained bandwidth of the device

Bandwidth of GPU (theoretical, real without ECC, with ECC) in giga byte per second

OI (Operation Intensity) : number of MACS by quantity of non cached data byte to compute the MACS

I.E. GEMV processing in :

• f32
$$\frac{2n^2+2n}{4(n^2+3n)} \approx 0.50$$

• f64 $\frac{2n^2+2n}{8(n^2+3n)} \approx 0.25$

Number of GPU needed for the MCAO case computed in f32

ECC	K20C	K40	K80	P100
Off	12	9	6	5
On	14	10	6	5

Green Flash

Classic implementation

Persistent kernel implementation

Green Hardware

DGX-1 server

- Ubuntu 14.04
- Dual Intel Xeon E5-2698
 @ 2.2 GHz
- 8 Nvidia P100
 @ 1480:715 MHz
 - ECC disabled
 - CUDA 8.0

• 1-4 GPU for computing

GPU ROMA 2016

Green Result 1/2 : Scalability and jitter

Strong scalability

Constant case with 10,048 slopes x 15,000 commands

Histogram

Case with 10,048 slopes x 15,000 commands on 4 devices

Average : 0.45ms Jitter : 17µs

Result 2/2: Sync & Intercom time

Conclusion : Persistent kernel

Pros

GPU ROMA 2016

- Much lower jitter
- Simpler execution stream
- Cons
 - Beware of kernel mapping
 - No concurrency
 - Hard to debug
 - Prohibited features
 - No cuda library using kernel (i.e. cuBlas)
 - No cuda runtime synchronization
 - Useless features
 - stream, memcpy
 - No intercommunication through QPI

l'Observatoire Durham MICRO GATE PLDA

Thank you

Question?

Project #671662 funded by European Commission under program H2020-EU.1.2.2 coordinated in H2020-FETHPC-2014