Real-Time Massively Distributed Multi-Object Adaptive Optics Simulations for the European Extremely Large Telescope

Hatem Ltaief
Senior Research Scientist
Extreme Computing Research Center, KAUST

Intel eXtreme Performance Users Group Conference
Thuwal, Jeddah, KSA
April 22-25, 2018

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

KAUST

Bringing Astronomy Back Home ;-)

Courtesy from CEMSE Communications, KAUST

Outline

Students/Collaborators

- Extreme Computing Research Center @ KAUST
  A. Charara and D. Keyes
- L'Observatoire de Paris, LESIA
  R. Dembet, N. Doucet, E. Gendron, D. Gratadour, C. Morel, A. Sevin and F. Vidal
- KAUST Supercomputer Lab
  B. Hadri and S. Feki
- Innovative Computing Laboratory @ UTK
  PLASMA/MAGMA/PaRSEC Teams
- INRIA/INP Bordeaux, France
  Runtime/HiePACS Teams

Support/Funding

- Funded partially by the French National Center for Scientific Research (CNRS, 2016)
- Funded partially by the European Commission (Horizon 2020 program, FET-HPC grant# 671662)
- Intel Parallel Computing Center
- KAUST/Cray Center of Excellence: **Aniello Esposito and Adrian Tate**
- For core-hours:
  - **Shaheen 2.0 from KAUST Supercomputer Lab**
  - **Cray KNL internal system**

Outline

The World's Biggest Eye on The Sky

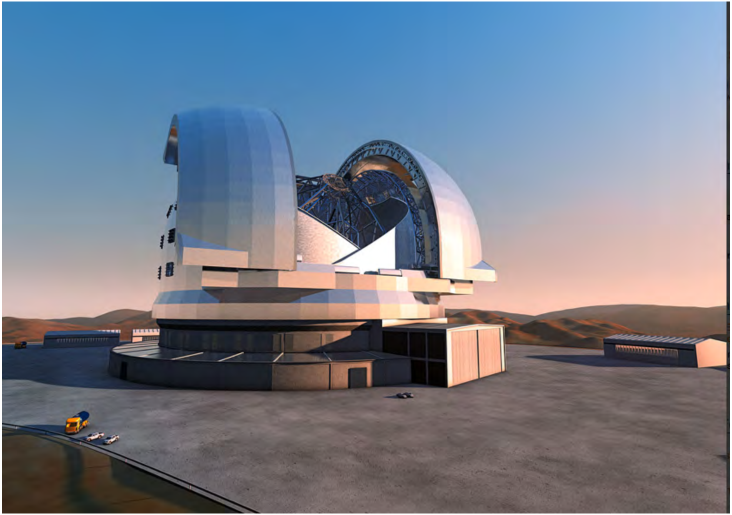

Credits: ESO (http://www.eso.org/public/teles-instr/e-elt/)

The World's Biggest Eye on The Sky

Credits: ESO (http://www.eso.org/public/teles-instr/e-elt/)

- The largest optical/near-infrared telescope in the world.
- It will weigh about 2700 tons with a main mirror diameter of 39m.
- Location: Chile, South America.

The Top 10 (present and future) Radio and Optical Ground-based Telescopes

Rank	Name	Location	Diameter	Cost	Year
10	Large Synoptic Survey Telescope (LSST)	Chile	8.4m	450 million	2014
9	South African Large Telescope (SALT)	South Africa	9.2m	36 million	2005
8	Keck	USA	10m	100 million	1996
7	Gran Telescopio Canarias (GTC)	Spain	10.4m	130 million	2009
6	Aricebo Observatory	Puerto Rico	305m	9.3 million	1963
5	Atacama Large Millimeter Array (ALMA)	Chile	12m	1.4 billion	2013
4	Giant Magellan Telescope (GMT)	Chile	24.5m	2.2 billion	2024
3	Thirty Meter Telescope (TMT)	USA	30m	1.4 billion	2030
2	Square Kilometer Array (SKA)	Australia	*90m*	2 billion	2020
1	**European Extremely Large Telescope (E-ELT)**	**Chile**	**39m**	**1.3 billion**	**2024**

Consortium: multiple nation initiatives

Src: http://www.space.com/14075-10-biggest-telescopes-earth-comparison.html

Outline

MOSAIC

- Represents the most challenging embedded instrument in the E-ELT.
- Observes the most distant galaxies in parallel to understand their evolution with a high multiplex
- Exploits a Field of View (FoV) of 7 to 10 arcminutes (1/4 of the full moon) with a resolution of few tens of milli-arcsec (1/20,000 of the full moon)
- Uses multiple guide stars (tomographic measurement of the turbulence) and multiple deformable mirrors (direction specific compensation).

The Effect of Atmospheric Turbulence

The sun observed with a compact camera

- Disturbs the trajectory of light rays (wavefront perturbations)
- Reduces astronomical images quality

Adaptive optics (AO)

- AO: technique used compensate in real-time the wavefront perturbations providing a significant improvement in resolution

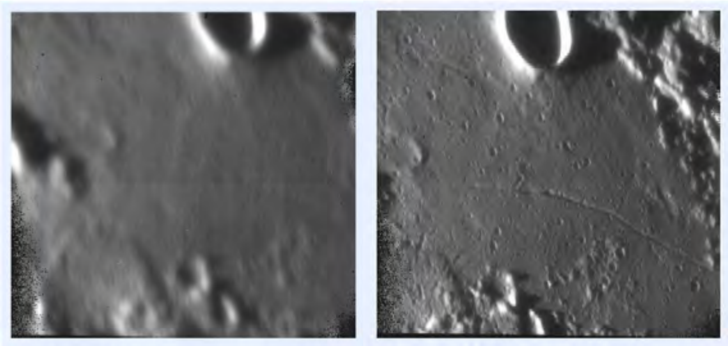

The moon observed with a 8m telescope (left: no AO, right: with AO)

How AO works

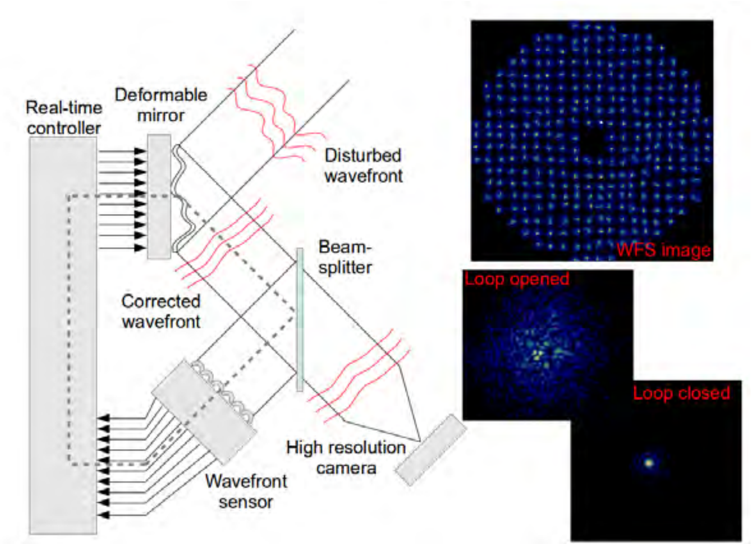
Real-time controller
Deformable mirror
Disturbed wavefront
Beam-splitter
Corrected wavefront
High resolution camera
Wavefront sensor
WFS image
Loop opened
Loop closed

Multi-object Adaptive Optics (MOAO)

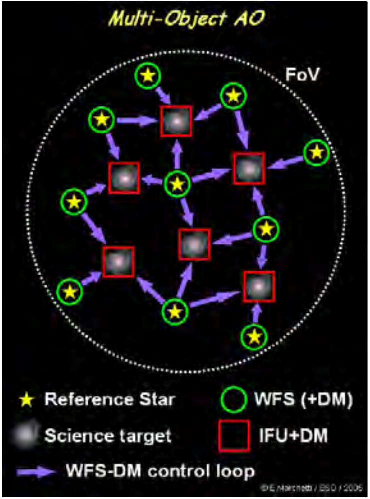

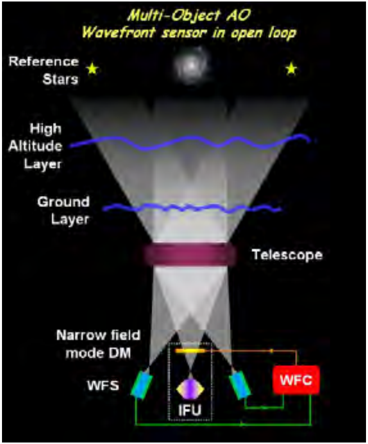

Credits: ESO

Multi-object Adaptive Optics (MOAO)

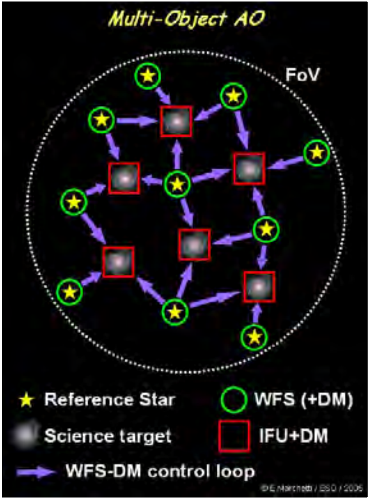

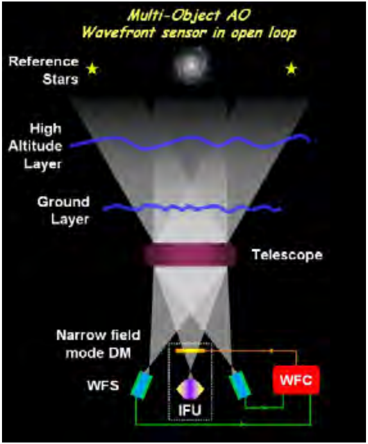

Credits: ESO

Good news: Extremely compute intensive at full telescope scale

Outline

One of The Key Operations: Tomographic reconstructor

- Compute the tomographic reconstructor matrix using covariance matrix between direction specific *truth* sensor and other sensors and the inverse of measurements covariance matrix
- $R' = C_{tm} \cdot C_{mm}^{-1}$
- Factorize and Solve for R' with C_{mm}, a 100k x 100k matrix, is extremely compute intensive
- At the core of system operations (soft real-time, should be achieved in seconds to update the real-time box)
- Also a critical component for the numerical simulation of the system behavior (observation forecast for today's design studies)

AO simulation pipeline: observations forecast

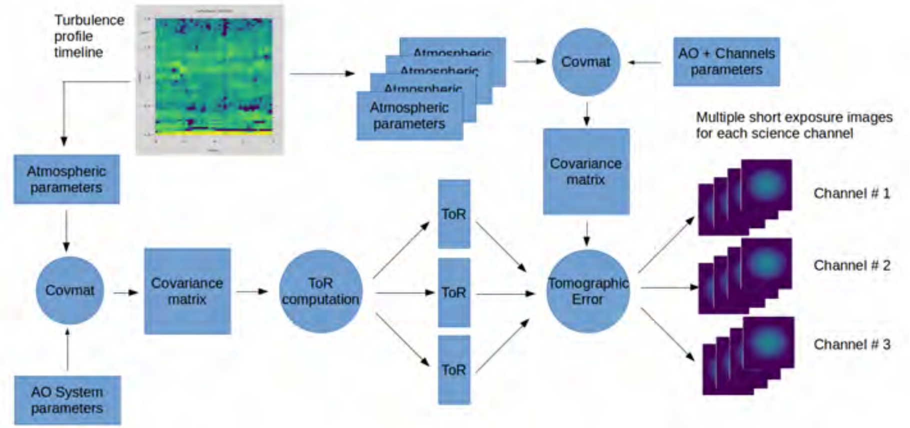

- From system parameters and turbulence evolution timeline to image quality at the output of the system

Global Workflow Chart

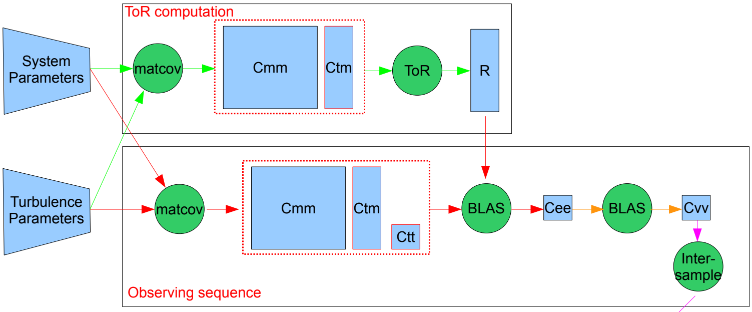

$$R = C_{tm}.C_{mm}^{-1}$$

$$C_{ee} = C_{tt} - C_{tm}R^t - RC_{tm}^t + RC_{mm}R^t$$

$$C_{vv} = D^{\dagger}C_{ee}D^{\dagger t}$$

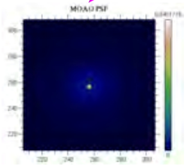

Global Workflow Chart

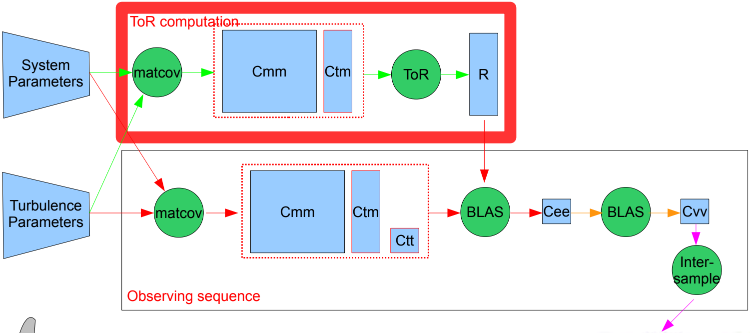

$$R = C_{tm}.C_{mm}^{-1}$$

$$C_{ee} = C_{tt} - C_{tm}R^t - RC_{tm}^t + RC_{mm}R^t$$

$$C_{vv} = D^{\dagger}C_{ee}D^{\dagger t}$$

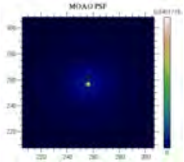

Global Workflow Chart

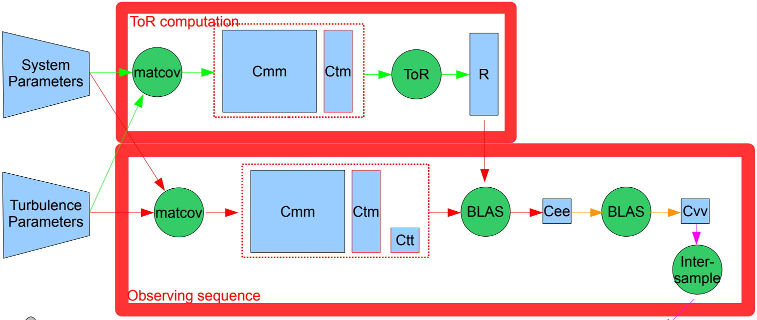

$$R = C_{tm}.C_{mm}^{-1}$$

$$C_{ee} = C_{tt} - C_{tm}R^t - RC_{tm}^t + RC_{mm}R^t$$

$$C_{vv} = D^{\dagger}C_{ee}D^{\dagger t}$$

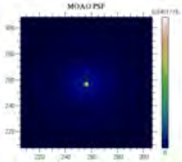

Observing sequence (close to real-time)

- ToR computation (outer loop)
- A serie of Level 3 BLAS operations (inner loop): GEMM, SYR2K, SYMM, etc.

Outline

Environment Settings - *Shaheen 2.0*

Compute Node x 6174

Environment Settings - *Shaheen 2.0*

Hardware and software description:

- CPU system (22nm)
  - Cray XC40 $\approx 5 Pflops/s$ **LINPACK**
  - Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.3GHz (32 Haswell cores)
  - 128 GB of DDR4 main memory
  - Aries Interconnect
- Software
  - GCC compiler suite
  - Intel MKL 2018.1.163

E-ELT Application Apparatus

- Diameter telescope: 39m
- Number of measurements of the true sensor: 10240
- Number of actuators: 5336
- Number of science channels (i.e., #galaxies): 16, 32 and 64
- Performance of the ToR computation

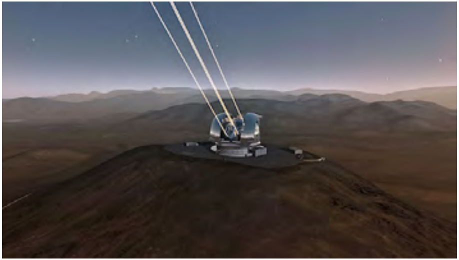

ScaLAPACK 2D Block-Cyclic Data Distribution

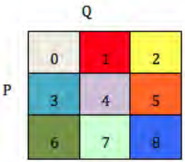

Processes Grid

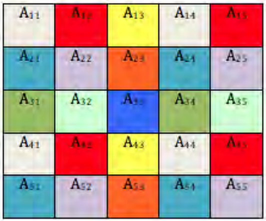

Logical View (Matrix)

Local View (CPUs)

AO simulation pipeline: observations forecast

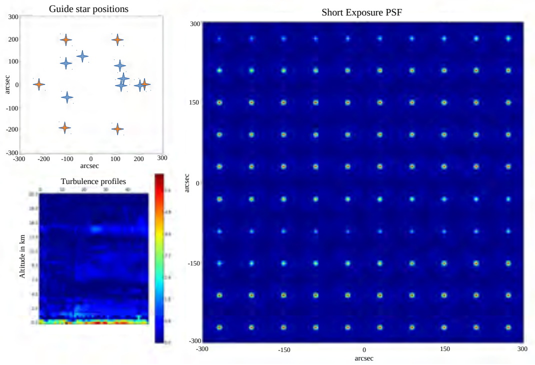

Figure: Overview of the simulation input parameters and output short and long exposure images for a single science channel and during a full night.

High Resolution Galaxy Maps

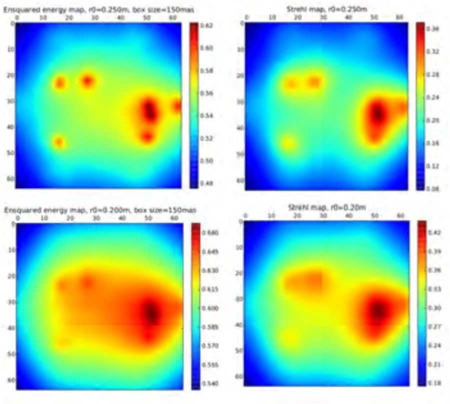

- Goal is to produce image quality maps over the instrument's full FoV depending on turbulence conditions evolution over the night

H. Ltaief et. al, *Real-Time Massively Distributed Multi-Object Adaptive Optics Simulations for the European Extremely Large Telescope*, IPDPS'18, Vancouver, BC

AO simulation pipeline: observations forecast

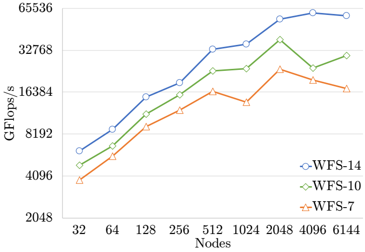

(a) Strong scalability.

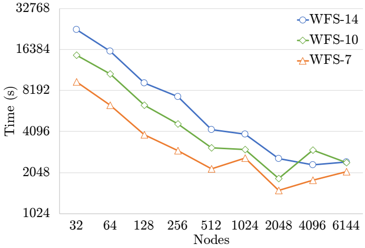

(b) Time to solution.

Figure: Strong scaling and time to solution of the large scale MOAO simulation up to 6144 nodes (196,608 cores) on *Shaheen-2* system.

H. Ltaief et. al, *Real-Time Massively Distributed Multi-Object Adaptive Optics Simulations for the European Extremely Large Telescope*, IPDPS'18, Vancouver, BC Canada

Performance Bottlenecks

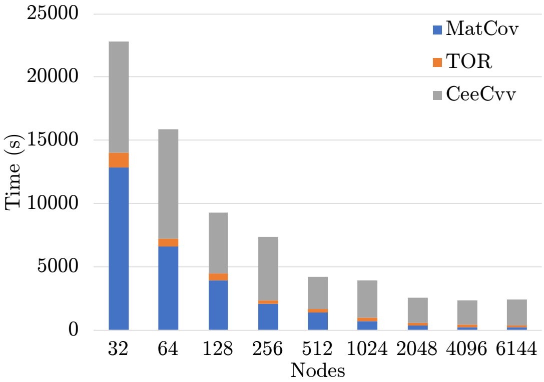

Figure: Time breakdown of WFS-14.

H. Ltaief et. al, *Real-Time Massively Distributed Multi-Object Adaptive Optics Simulations for the European Extremely Large Telescope*, IPDPS'18, Vancouver, BC Canada

Outline

Global Workflow Chart

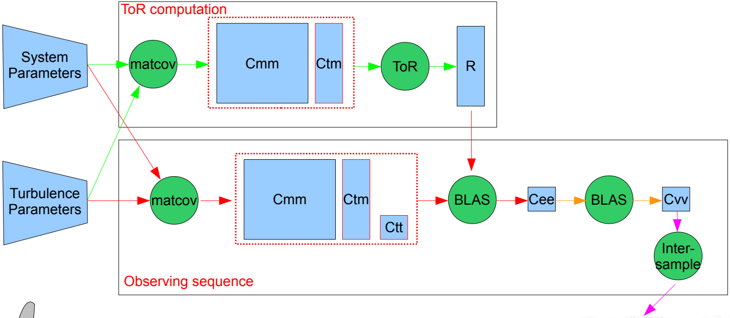

$$R = C_{tm}.C_{mm}^{-1}$$

$$C_{ee} = C_{tt} - C_{tm}R^t - RC_{tm}^t + RC_{mm}R^t$$

$$C_{vv} = D^{\dagger}C_{ee}D^{\dagger t}$$

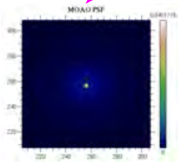

Blocked Algorithms: Fork-Join Paradigm

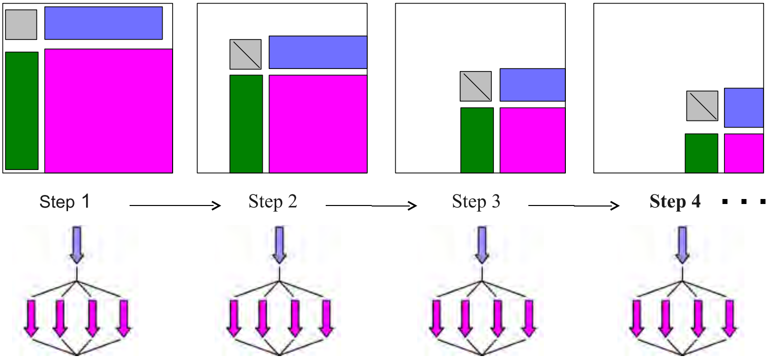

ScaLAPACK: Blocked Algorithms

Principles:

- Panel-Update sequence
- Transformations are blocked/accumulated within the Panel
  Level-2 BLAS
- Transformations applied at once on the trailing submatrix
  Level-3 BLAS
- Parallelism hidden inside the BLAS
- Fork-join model
- A broken model!

Tile Data Layout Format

LAPACK: column-major format

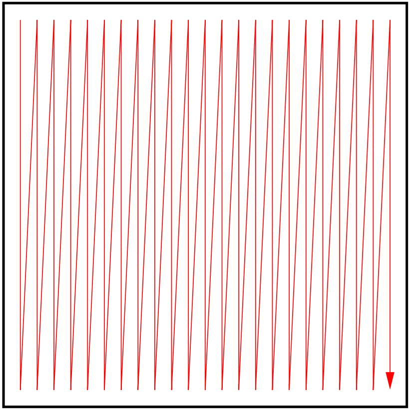

PLASMA/CHAMELEON: tile format

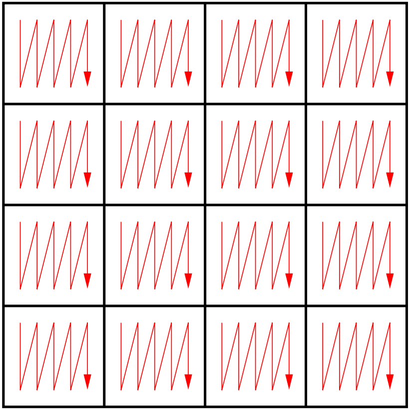

PLASMA/CHAMELEON: Tile Algorithms

PLASMA $\Longrightarrow$ http://icl.cs.utk.edu/plasma/
CHAMELEON $\Longrightarrow$ https://gitlab.inria.fr/solverstack/chameleon.git

- Break the bulk synchronous programming model
- Parallelism is brought to the fore
- May require the redesign of linear algebra algorithms
- Tile data layout translation
- Remove unnecessary synchronization points between Panel-Update sequences
- DAG execution where nodes represent tasks and edges define dependencies between them
- Default dynamic runtime system environment **StarPU** (but could use Quark, PaRSEC, OmpSs, OpenMP etc.)

StarPU Runtime System 101

- Provides:
  ⟹ Task scheduling
  ⟹ Memory management
  ⟹ Out-of-core
- Supports:
  ⟹ SMP/Multicore Processors (x86, PPC, . . .)
  ⟹ NVIDIA GPUs (e.g., multi-GPU)
  ⟹ Hybrid architectures
  ⟹ Shared and Distributed-memory

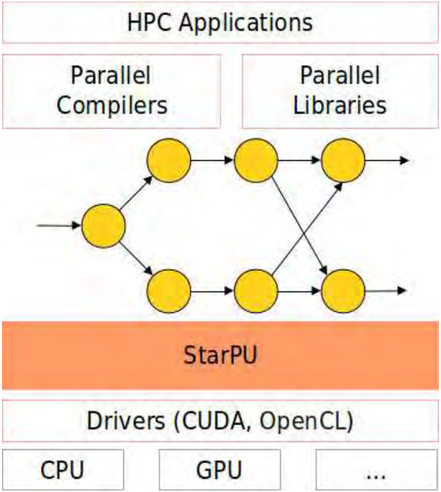

StarPU Runtime System: User Productivity!

```
starpu_Insert_Task(&cl_dpotrf,
    VALUE,   &uplo,         sizeof(char),
    VALUE,   &n,            sizeof(int),
    INOUT,   Ahandle(k, k),
    VALUE,   &lda,          sizeof(int),
    OUTPUT,  &info,         sizeof(int)
    CALLBACK, profiling?cl_dpotrf_callback:NULL, NULL,
    0);
```

Global Workflow Chart

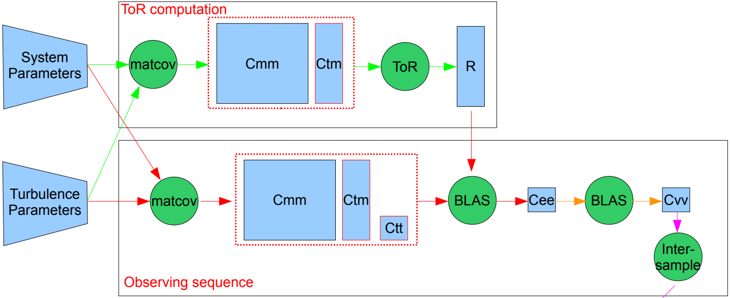

$$R = C_{tm}.C_{mm}^{-1}$$

$$C_{ee} = C_{tt} - C_{tm}R^t - RC_{tm}^t + RC_{mm}R^t$$

$$C_{vv} = D^{\dagger}C_{ee}D^{\dagger t}$$

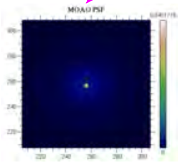

Directed Acyclic Graph for MOAO

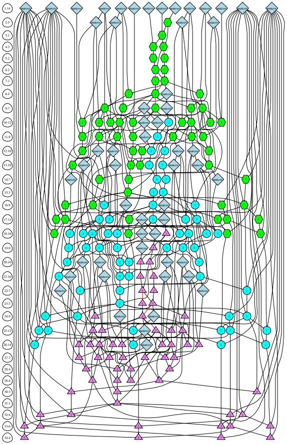

Figure: DAG representation of the overall MOAO framework execution for the calculation of three ToR (i.e., three iterations of the outer loop) and a single observing sequence (i.e., one iteration in the inner loop) on a two-by-two tile

Directed Acyclic Graph for MOAO

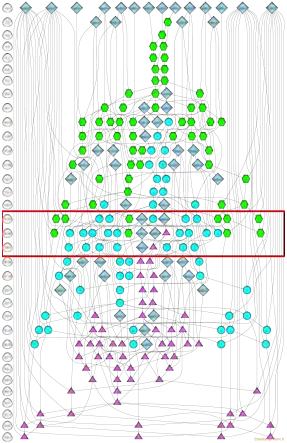

Figure: DAG representation of the overall MOAO framework execution for the calculation of three ToR (i.e., three iterations of the outer loop) and a single observing sequence (i.e., one iteration in the inner loop) on a two-by-two tile

Zooming in...

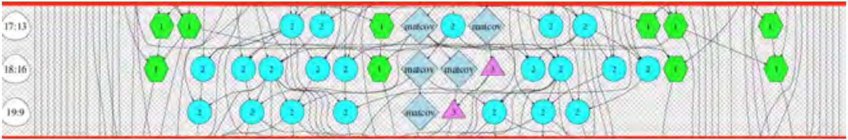

MOAO Software Release – https://github.com/ecrc/moao

A HIGH PEFORMANCE MULTI-OBJECT ADAPTIVE OPTICS FRAMEWORK FOR GROUND-BASED ASTRONOMY

MOAO

Extreme Computing Research Center

The Multi-Object Adaptive Optics (MOAO) framework provides a comprehensive testbed for high performance computational astronomy. In particular, the European Extremely Large Telescope (E-ELT) is one of today's most challenging projects in ground-based astronomy and will make use of a MOAO instrument based on turbulence tomography. The MOAO framework uses a novel compute-intensive pseudo-analytical approach to achieve close to real-time data processing on manycore architectures. The scientific goal of the MOAO simulation package is to dimension future E-ELT instruments and to assess the qualitative performance of tomographic reconstruction of the atmospheric turbulence on real datasets.

THE MULTI-OBJECT ADAPTIVE OPTICS TECHNIQUE

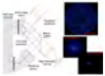

Single conjugate AO concept

Open-Loop tomography concept

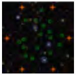

Observing the GOODS South cosmological field with MOAO

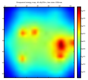

High res. map of the quality of turbulence compensation obtained with MOAO on a cosmological field

THE PSEUDO-ANALYTICAL APPROACH

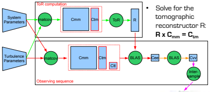

- Solve for the tomographic reconstructor R: $R \times C_{mm} = C_{tm}$

- Compute the tomographic error: $C_{ee} = C_{tt} - C_{tm} R^T - R C_{tm}^T + R C_{mm} R^T$
- Compute the equivalent phase map: $C_{vv} = D C_{ee} D^T$
- Generate the point spread function image

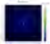

MOAO 0.1.0

- Tomographic Reconstructor Computation
- Dimensioning of Future Instruments
- Real Datasets
- Single and Double Precisions
- Shared-Memory Systems
- Task-based Programming Models
- Dynamic Runtime Systems
- Hardware Accelerators

CURRENT RESEARCH

- Distributed-Memory Systems
- Hierarchical Matrix Compression
- Machine Learning for Atmospheric Turbulence
- High Resolution Galaxy Map Generation
- Extend to other ground-based telescope projects

PERFORMANCE RESULTS TOMOGRAPHIC RECONSTRUCTOR COMPUTATION – DOUBLE PRECISION

Two-socket 18-core Intel HSW – 64-core Intel KNL – 8 NVIDIA GPU P100s (DGX-1)

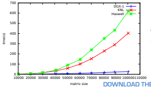

This is one tomographic reconstructor every 25 seconds!

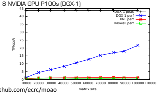

DOWNLOAD THE SOFTWARE AT http://github.com/ecrc/moao

A collaboration of INRIA Centre de recherche BORDEAUX – SUD-OUEST ICL UT INNOVATIVE

With support from intel NVIDIA CRAY

Sponsored by OSR

Outline

Summary

- Adaptive optics require compute-intensive hardware
- Pipelining computational stages
- Efficient Task-based programming model
- MOAO effort for standardization (w/ O. Guyon from Subaru telescope)
- Machine learning for the control matrix (e.g., maximum likelihood)

Exciting Time for Astronomy at KAUST/ECRC!

- Supporting two major worldwide ground-based astronomy efforts

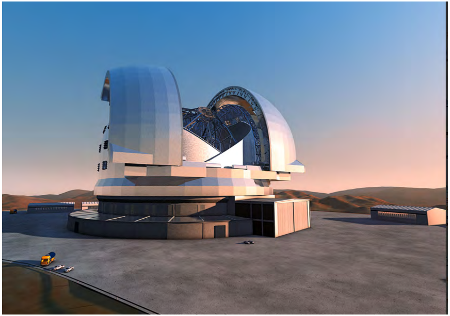

Future Work

- we need more (*smarter*) DLA (echoing A. Gara)
- We are not cheating by doing approximation! (echoing A. Heinecke)

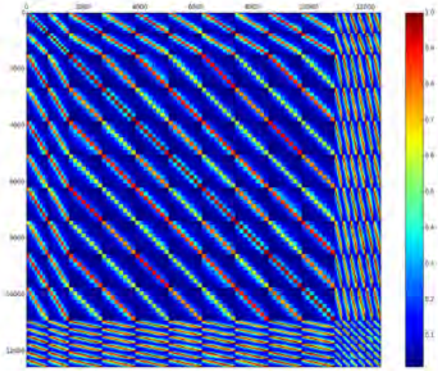

The Hourglass Revisited

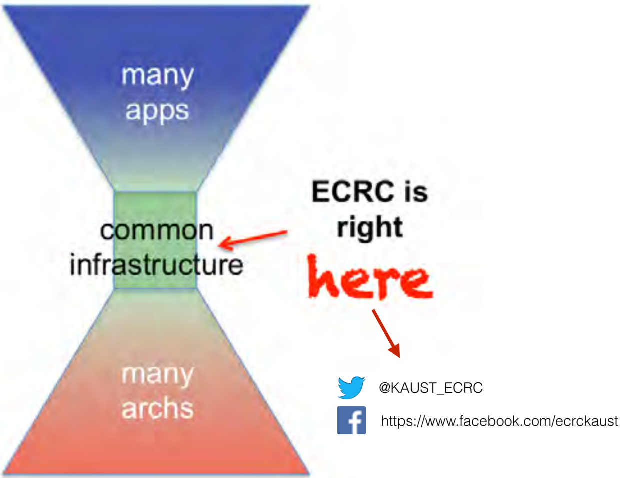

Questions?