# Algorithmic Recipes for Extreme-Scale High Performance Computing

**Hatem Ltaief**

Senior Research Scientist
Extreme Computing Research Center, KAUST, Saudi Arabia

*RTC4AO Workshop in Paris, France*
October 25-26, 2018

# Outline


1. A Hostile Hardware Landscape
2. Linear Algebra Challenges in AO
3. The E-ELT
4. The Subaru Telescope
5. Summary and Future Work

# Students/Collaborators

- Extreme Computing Research Center @ KAUST
  **A. Charara, N. Doucet, D. Keyes and D. Sukkari**
- L'Observatoire de Paris, LESIA
  **R. Dembet, N. Doucet, E. Gendron, D. Gratadour, C. Morel, A. Sevin and F. Vidal**
- Subaru Telescope, National Astronomical Observatory of Japan
  **O. Guyon and J. Lozi**
- Innovative Computing Laboratory @ UTK
  **PLASMA/MAGMA/PaRSEC Teams**
- INRIA/INP Bordeaux, France
  **Runtime/HiePACS Teams**

# Vendors

- NVIDIA GPU Research Center
- Intel Parallel Computing Center
- Cray Center of Excellence

# Outline

1. A Hostile Hardware Landscape
2. Linear Algebra Challenges in AO
3. The E-ELT
4. The Subaru Telescope
5. Summary and Future Work

# High Performance Computing: The Top500 List

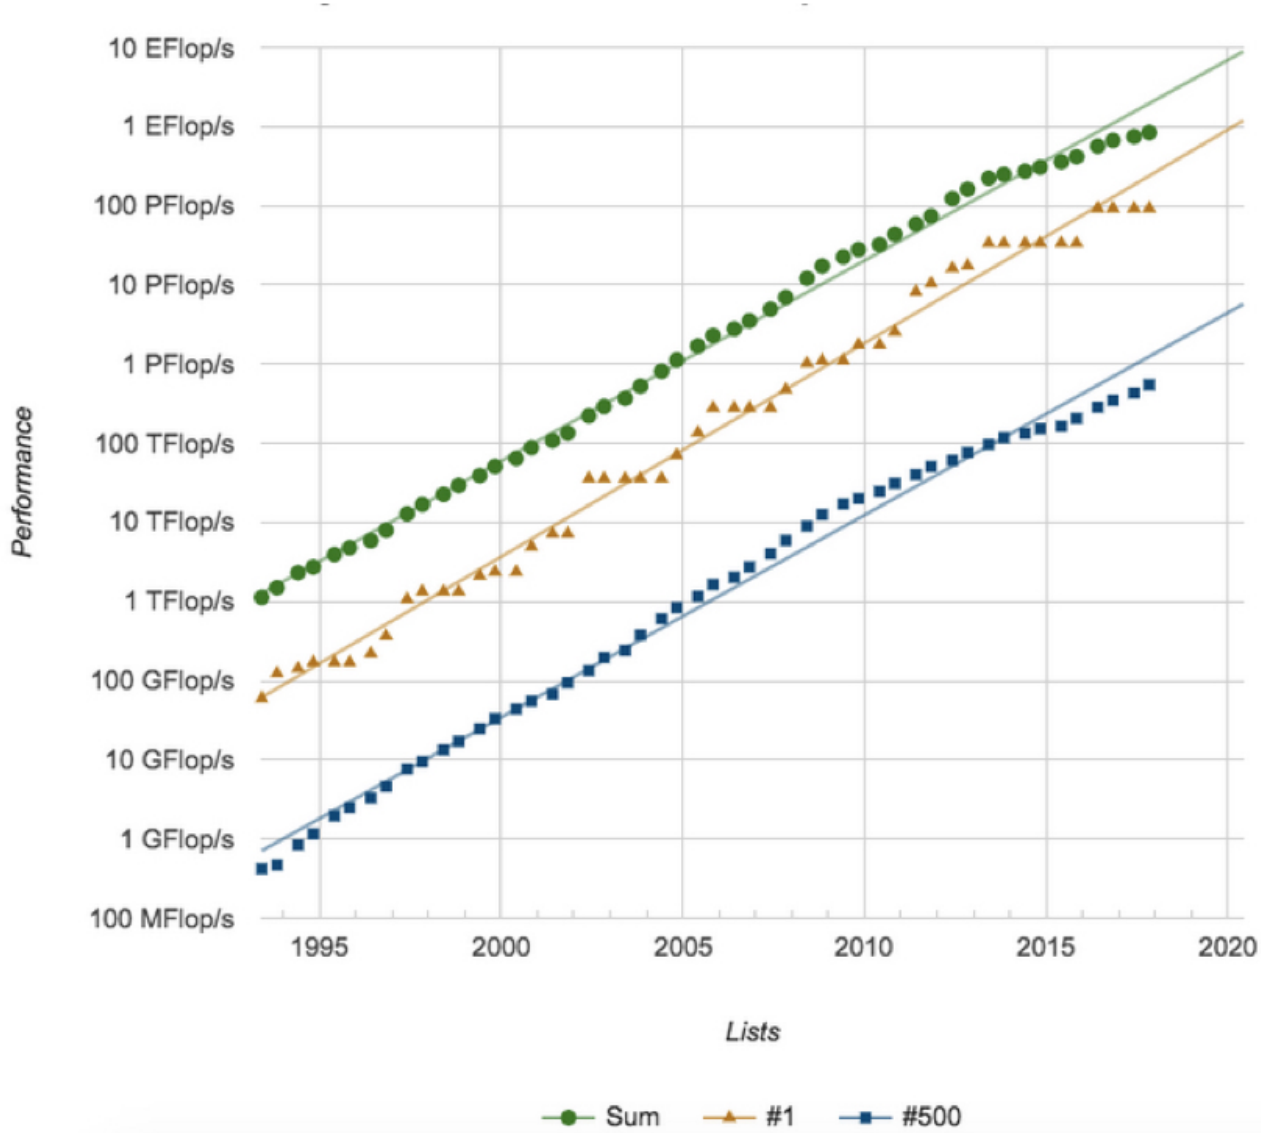

# High Performance Computing: The Top500 List

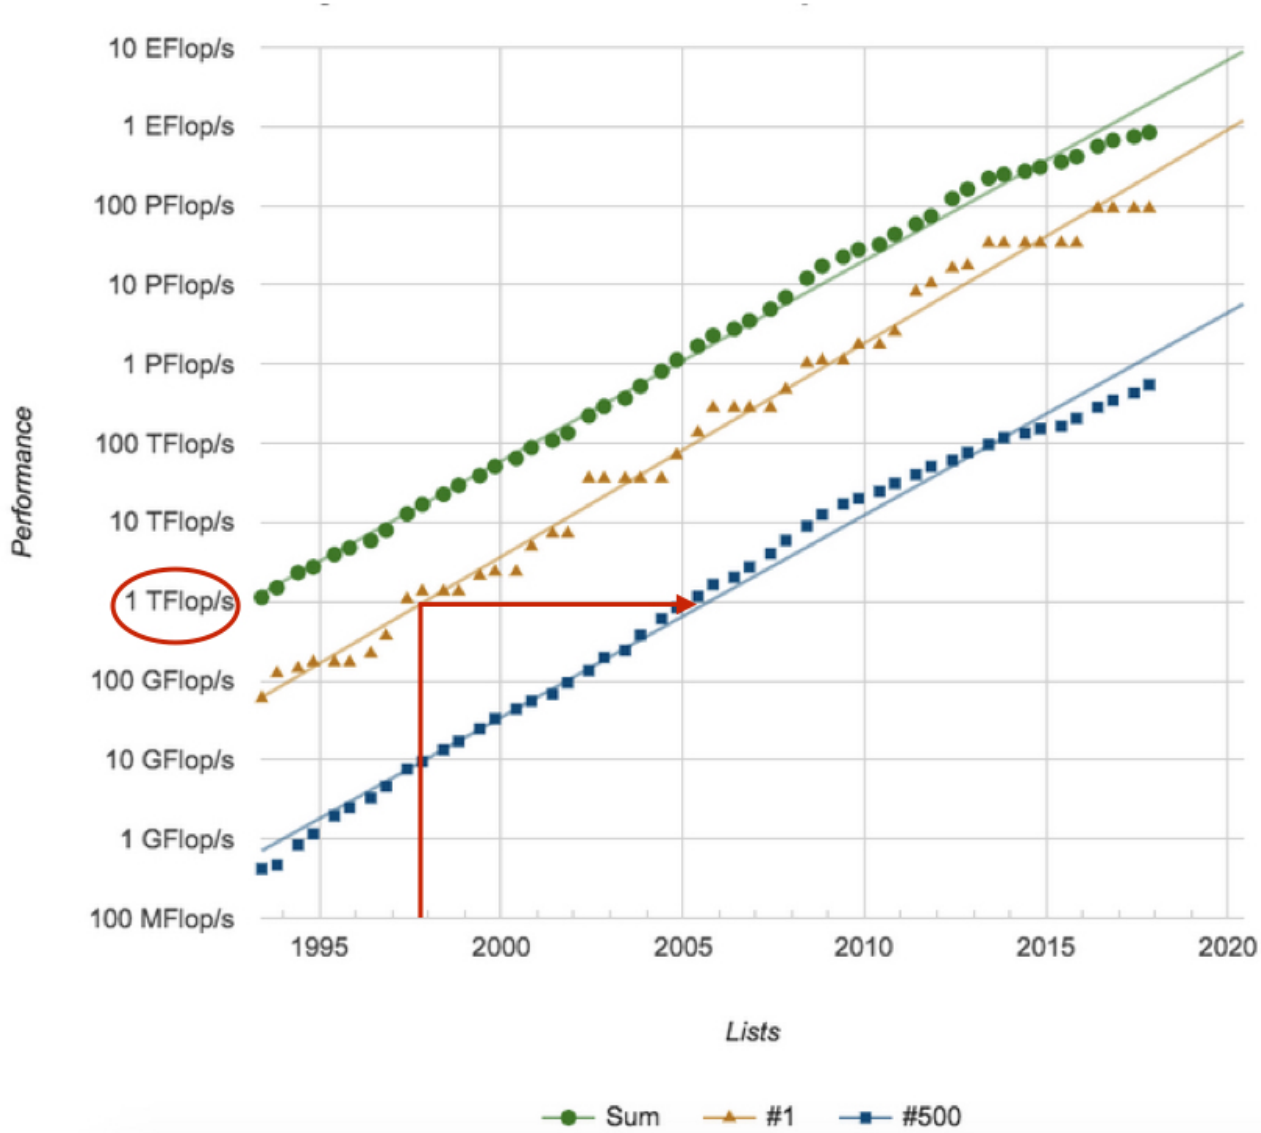

# High Performance Computing: The Top500 List

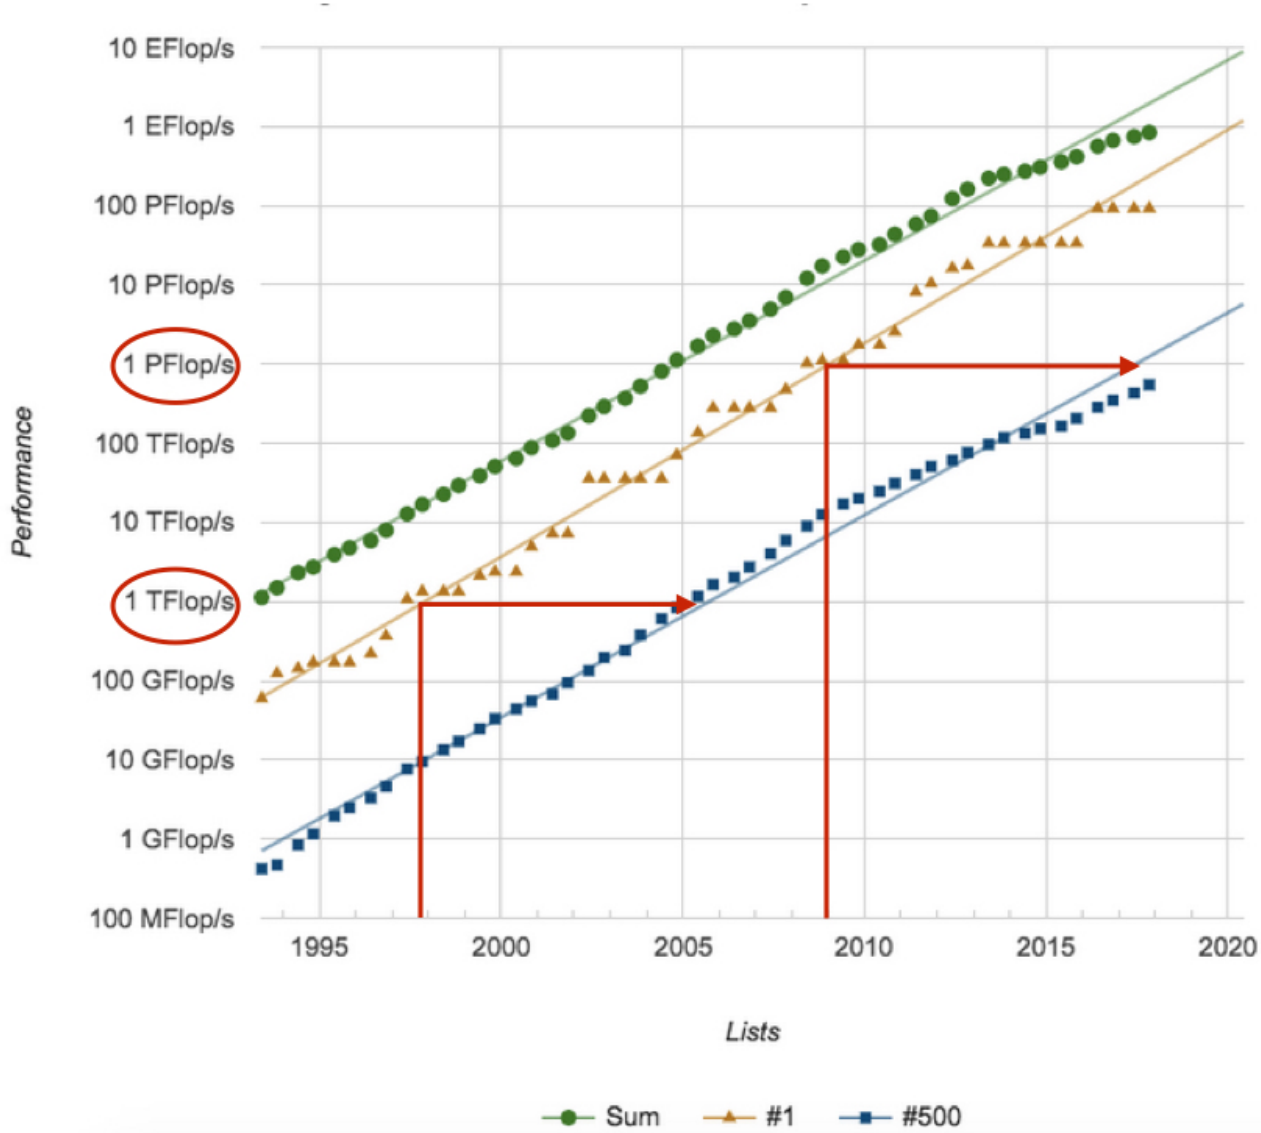

# High Performance Computing: The Top500 List

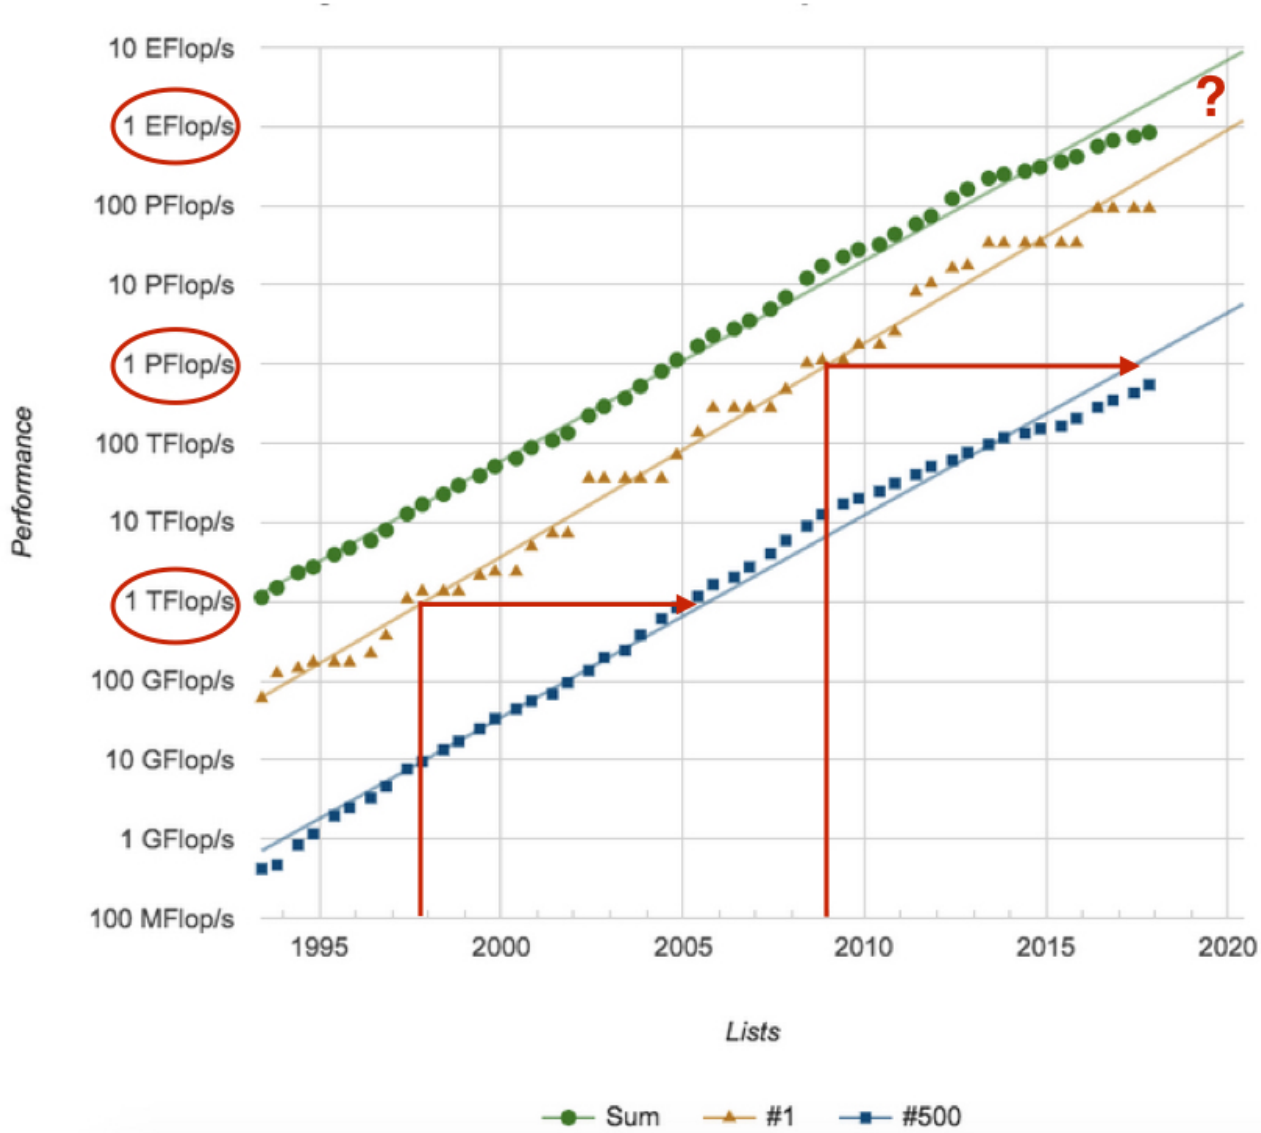

# High Performance Computing: The Top500 List

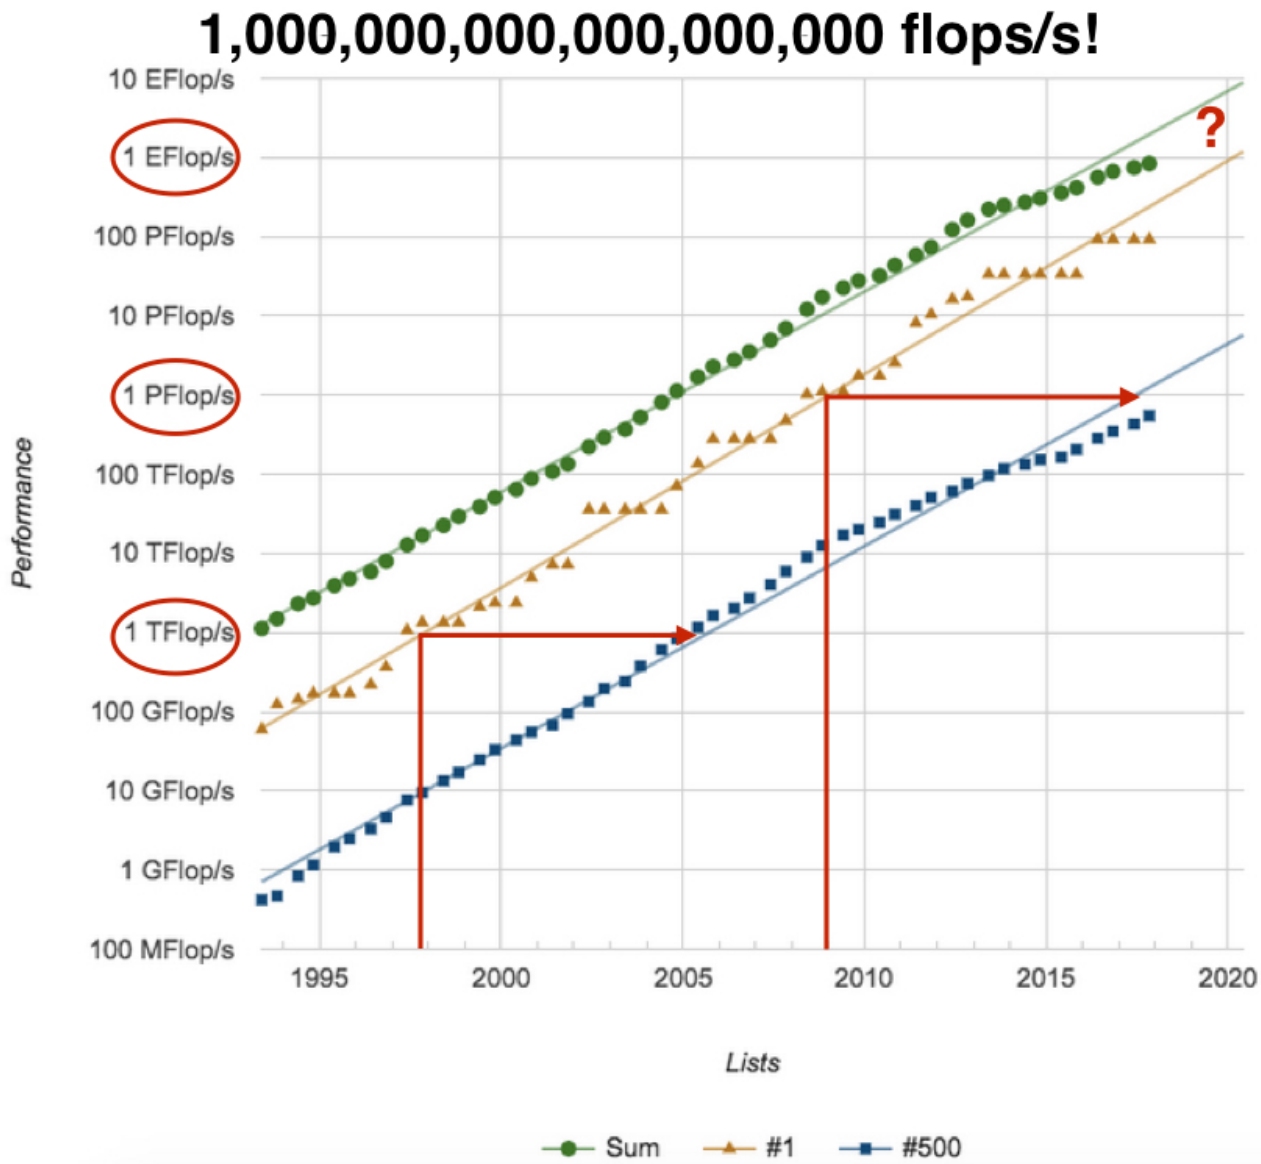

# High Performance Computing: The Top500 List

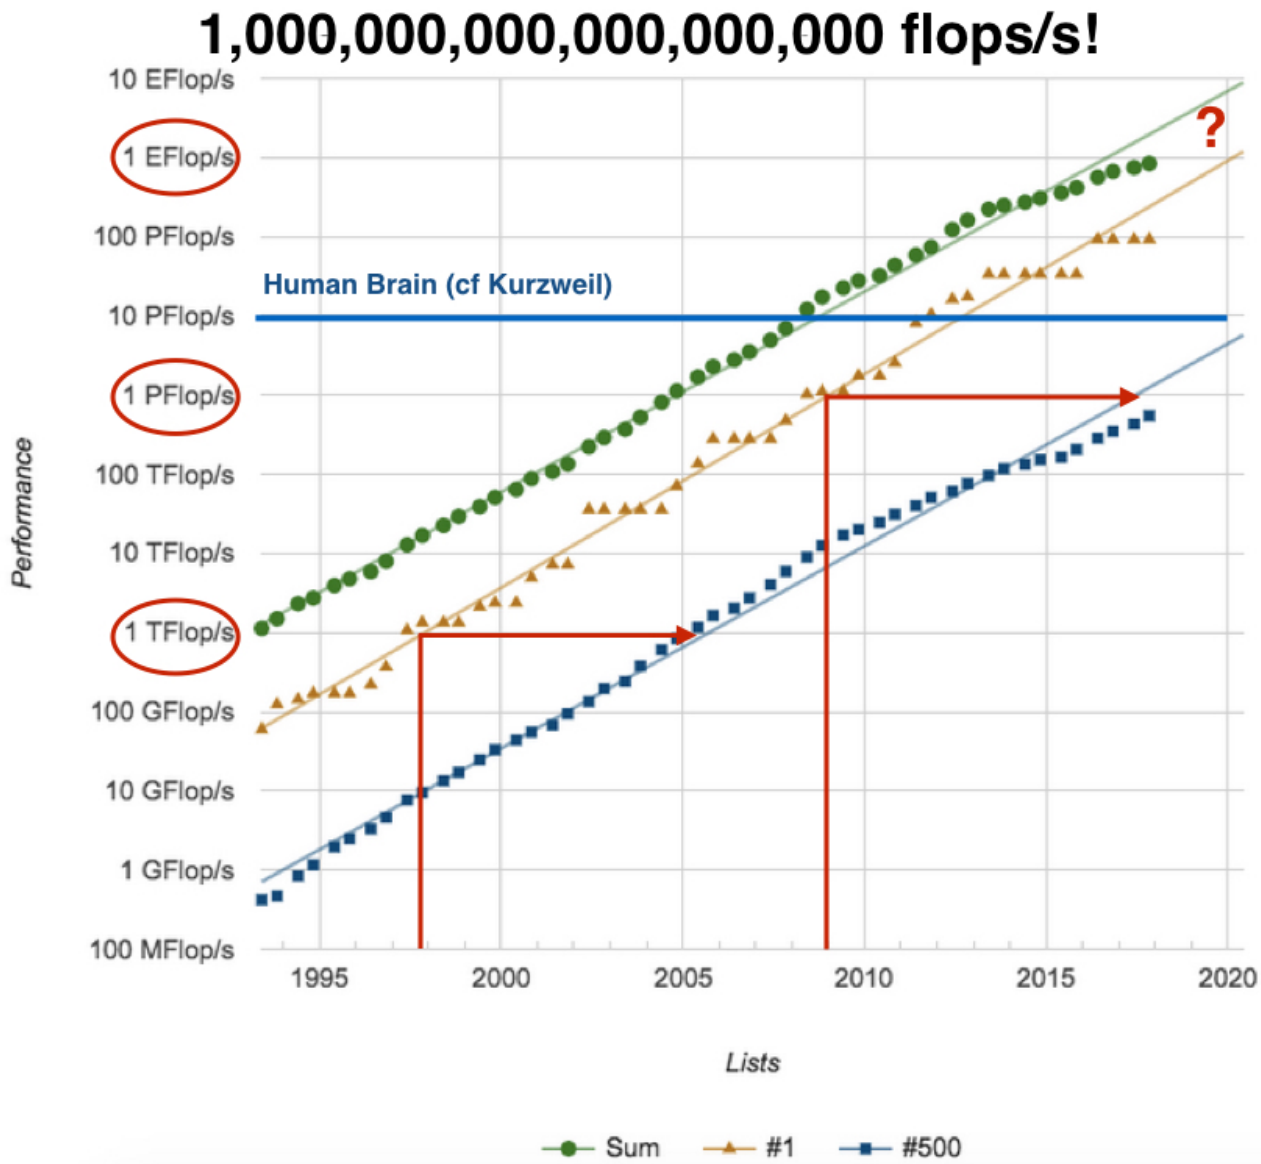

# It is getting *Moore and Moore* hot here!

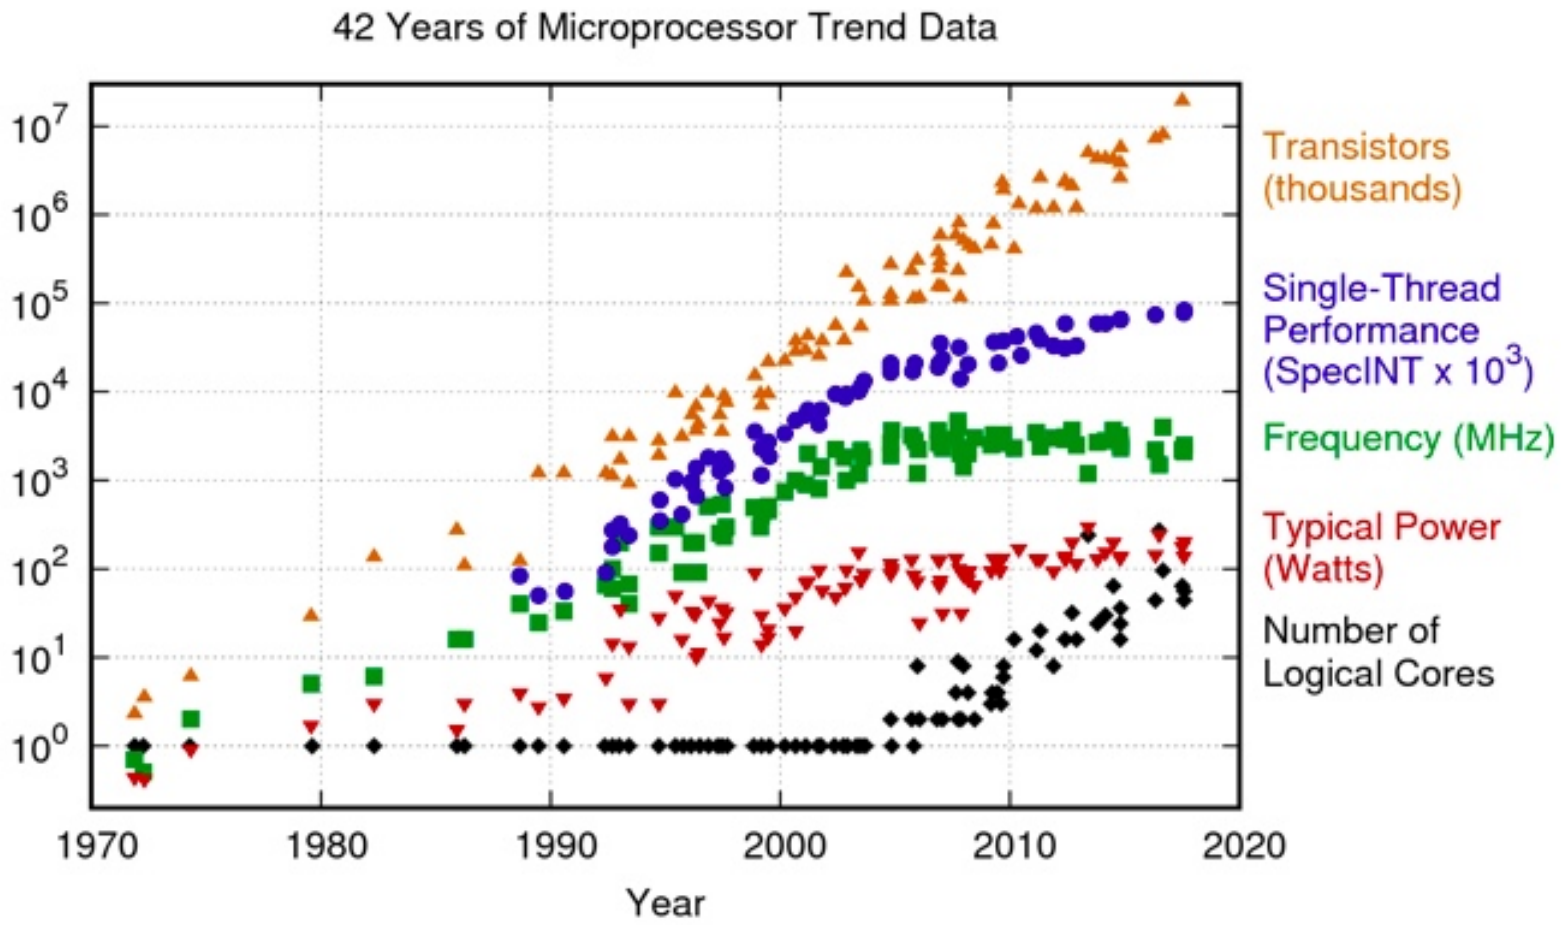


Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

https://www.karlrupp.net/2018/02/42-years-of-microprocessor-trend-data/

# Hardware Trends: Energy / Data Movement Matters!

| | 2011 | 2018 |
|---|---|---|
| **DP FLOP** | 100 pJ | 10 pJ |
| **DP DRAM Read** | 4800 pJ | 1920 pJ |
| **Local interconnect** | 7500 pJ | 2500 pJ |
| **Cross system** | 9000 pJ | 3500 pJ |

*John Shalf, LBNL*

# Algorithmic Recipes For Exascale Computing

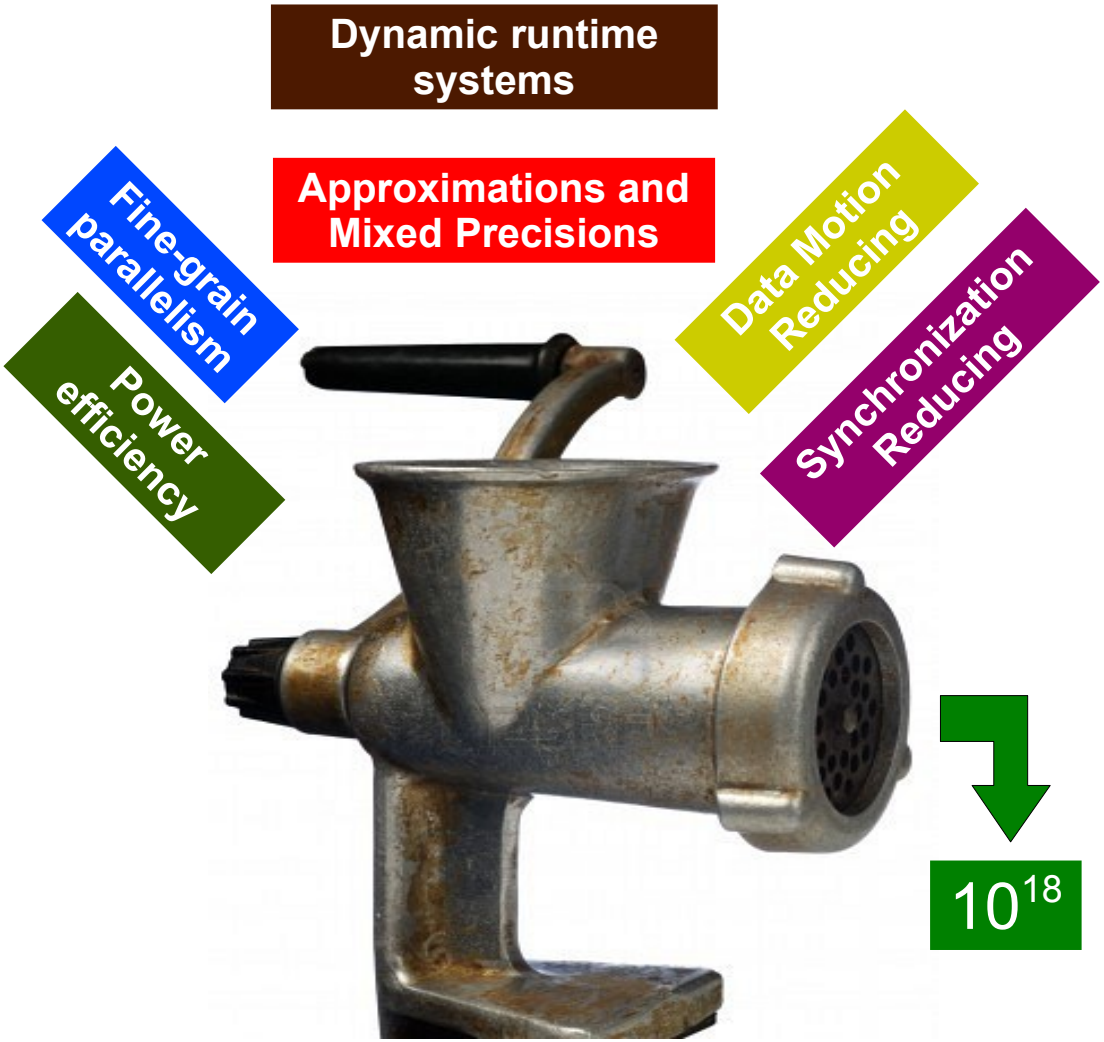

# Outline

1 A Hostile Hardware Landscape

2 Linear Algebra Challenges in AO

3 The E-ELT

4 The Subaru Telescope

5 Summary and Future Work

# 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

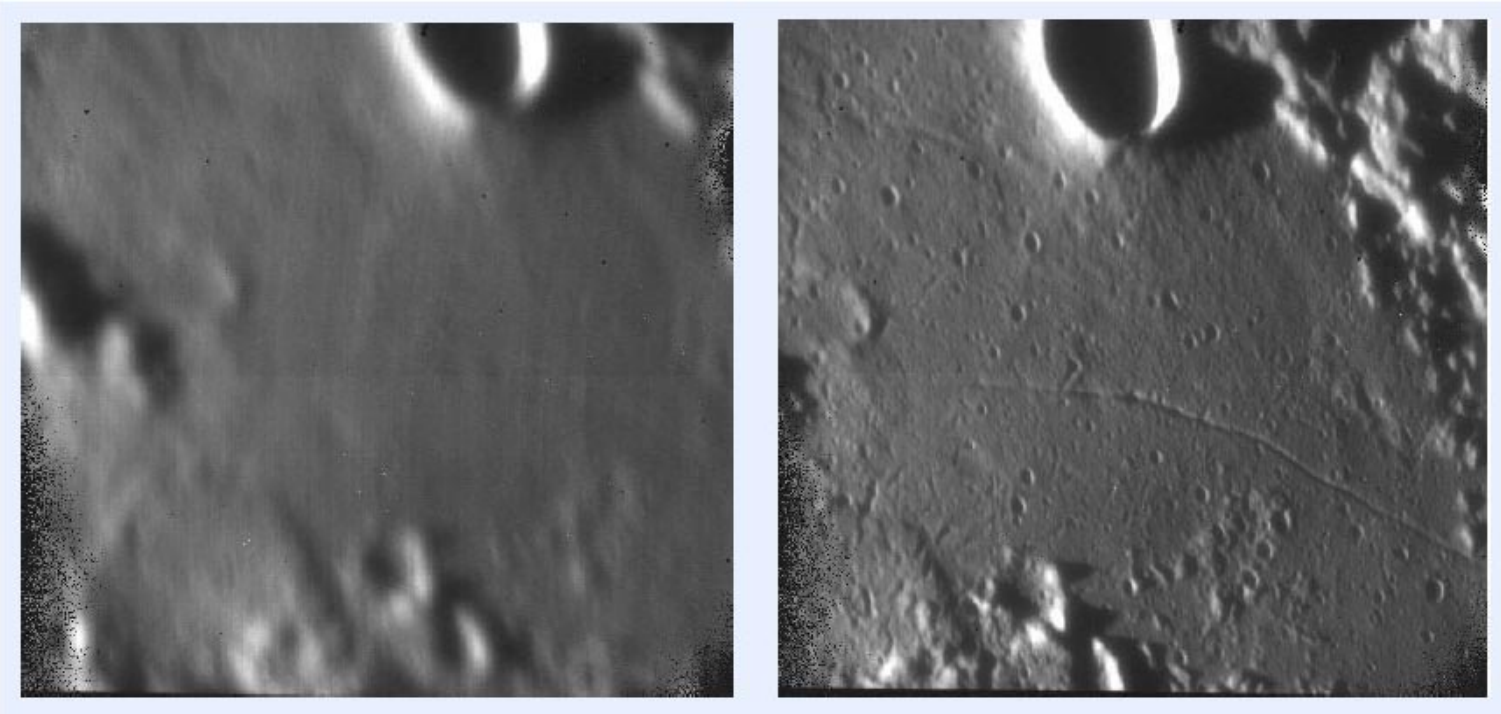

*The moon observed with a 8m telescope (left: no AO, right: with AO)*

# How AO works

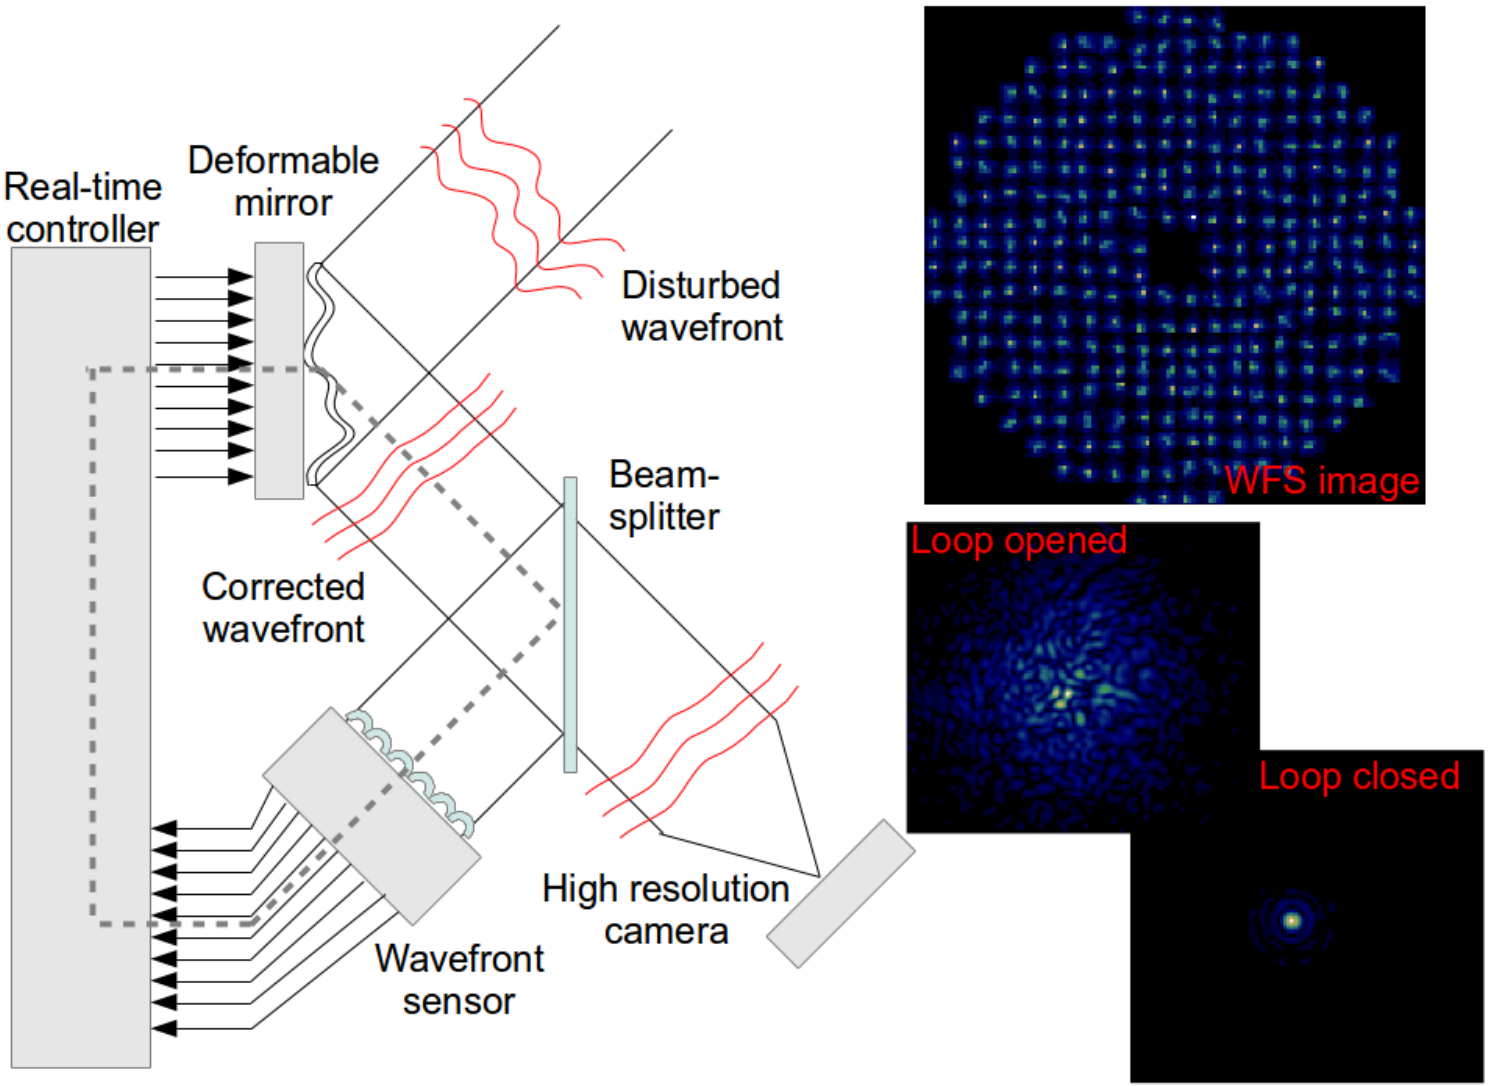
Real-time controller
Deformable mirror
Disturbed wavefront
Beam-splitter
Corrected wavefront
High resolution camera
Wavefront sensor
WFS image
Loop opened
Loop closed

# The World's Biggest Eye on The Sky

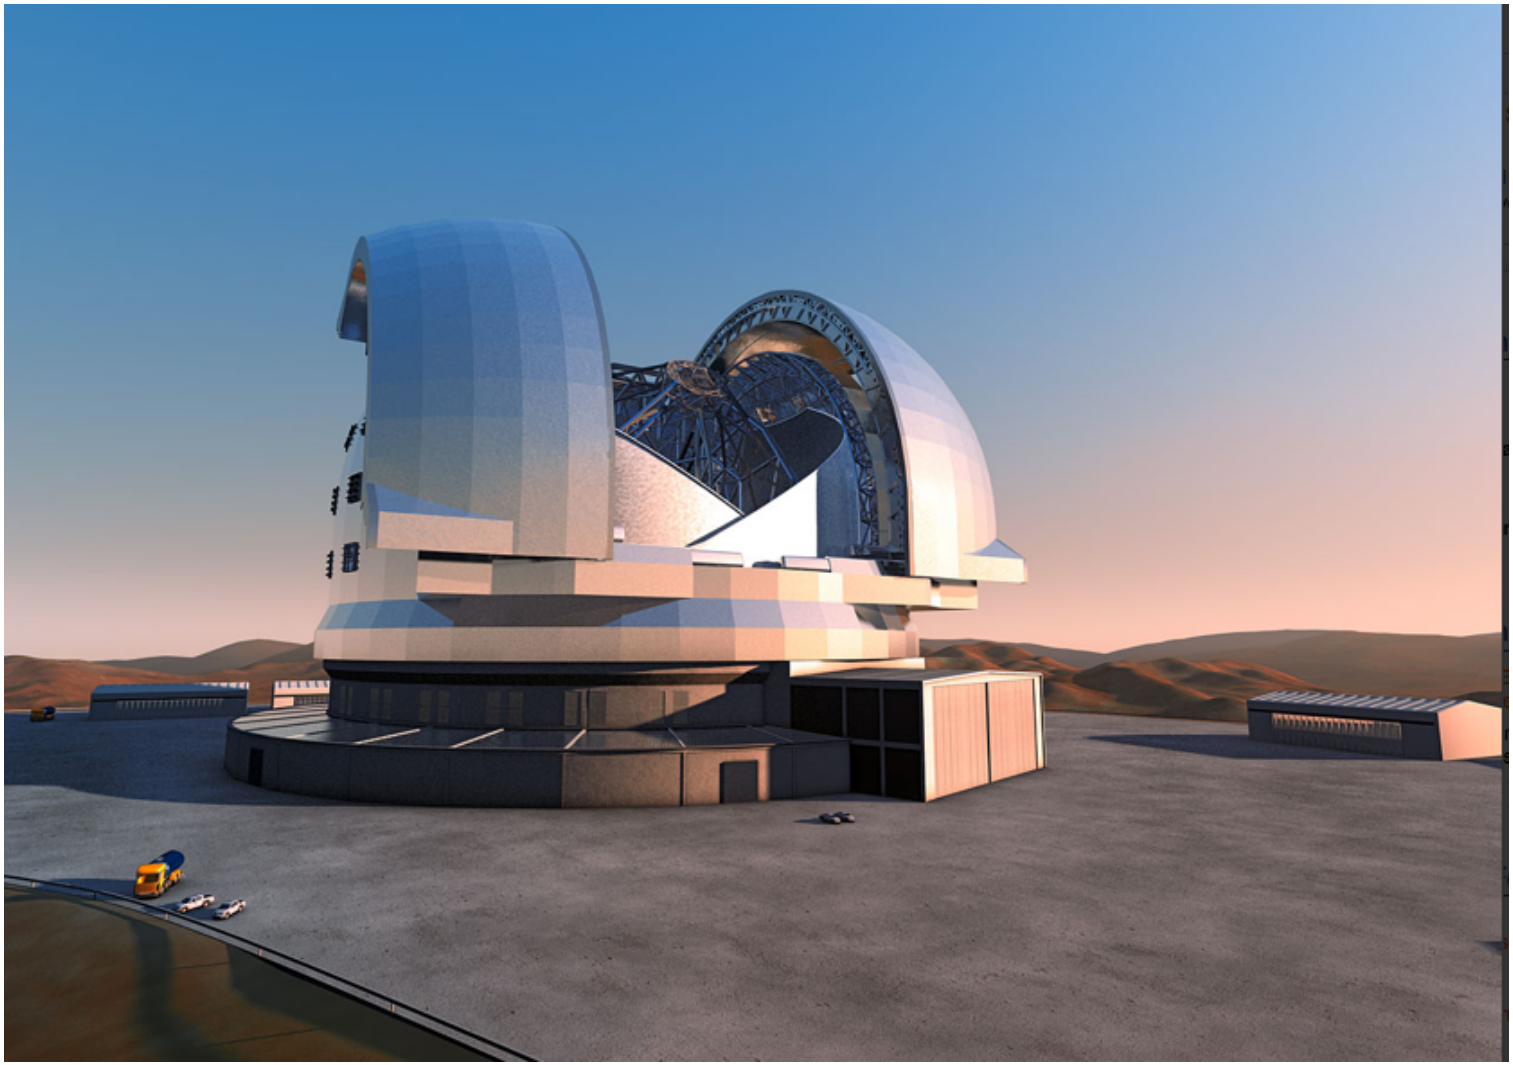

*Credits: ESO (http://www.eso.org/public/teles-instr/e-elt/)*

# And here comes the linear algebra...

- Multiple-Object AO approach (MOAO)
- 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)

# The Subaru Telescope

# And here comes *again* the linear algebra...

- Compute the pseudo inverse $A^+$

$$AA^+A = A\,, A \in \mathbb{R}^{m\times n}(m \geq n)$$

- The numerical challenge of the pseudo inverse are twofold:
  - Numerical: dealing with rectangular matrix which may engender numerical instabilities

$$A^\top A \;=\; V\Lambda V^\top$$

  - Computational: high algorithmic complexity, it should still be able to keep up with the overall throughput of the AO framework
- Using SVD: $A \;= U\Sigma V^\top$ then:

$$A^+ \;=\; V\Sigma^{-1}U^\top$$

- Only most significant singular values with their associated singular vectors are required ($\approx$**10%**)

# Outline

1 A Hostile Hardware Landscape

2 Linear Algebra Challenges in AO

3 The E-ELT

4 The Subaru Telescope

5 Summary and Future Work

# Blocked Algorithms: Fork-Join Paradigm

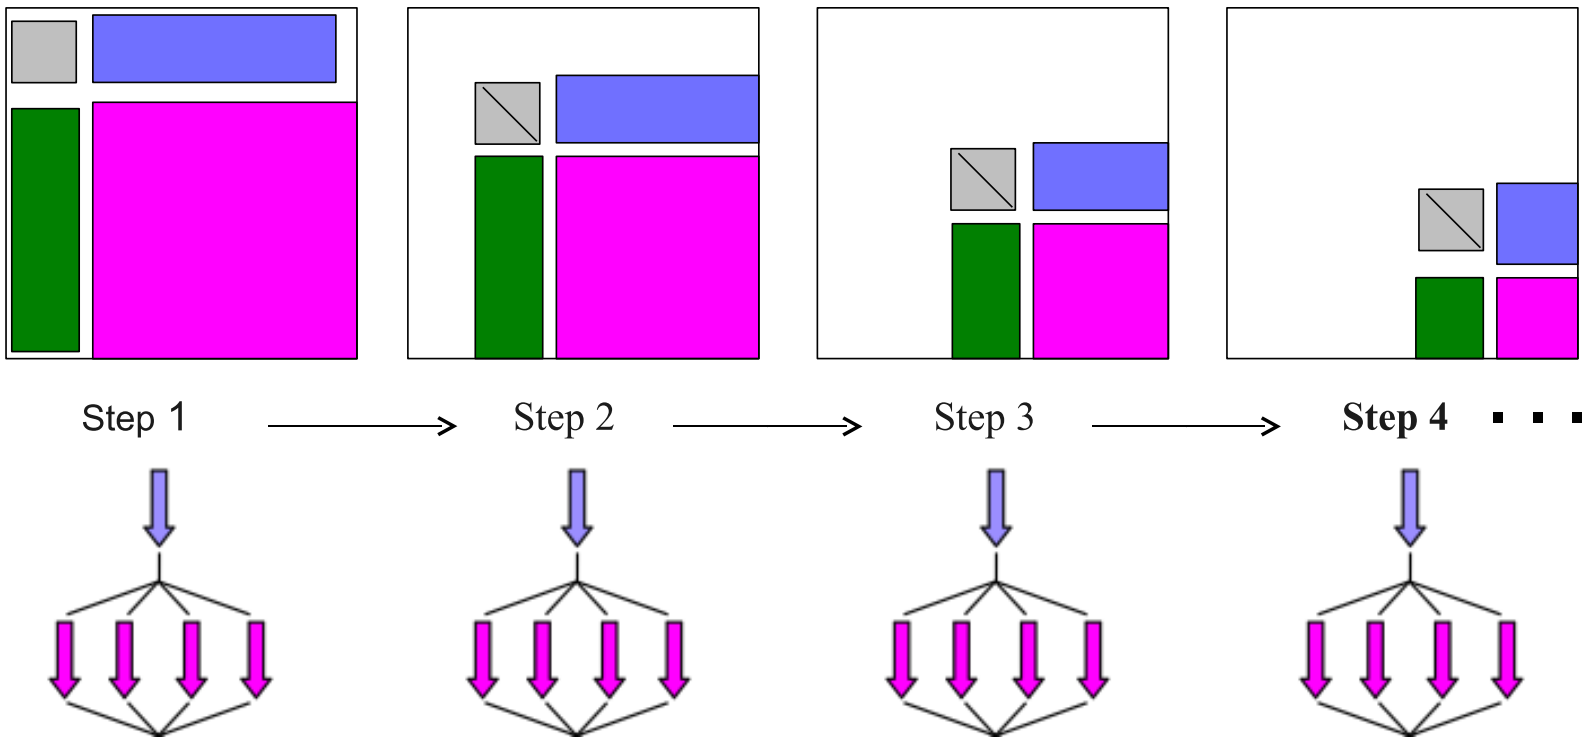

# Tile Algorithms: Asynchronous Execution

LAPACK: column-major format

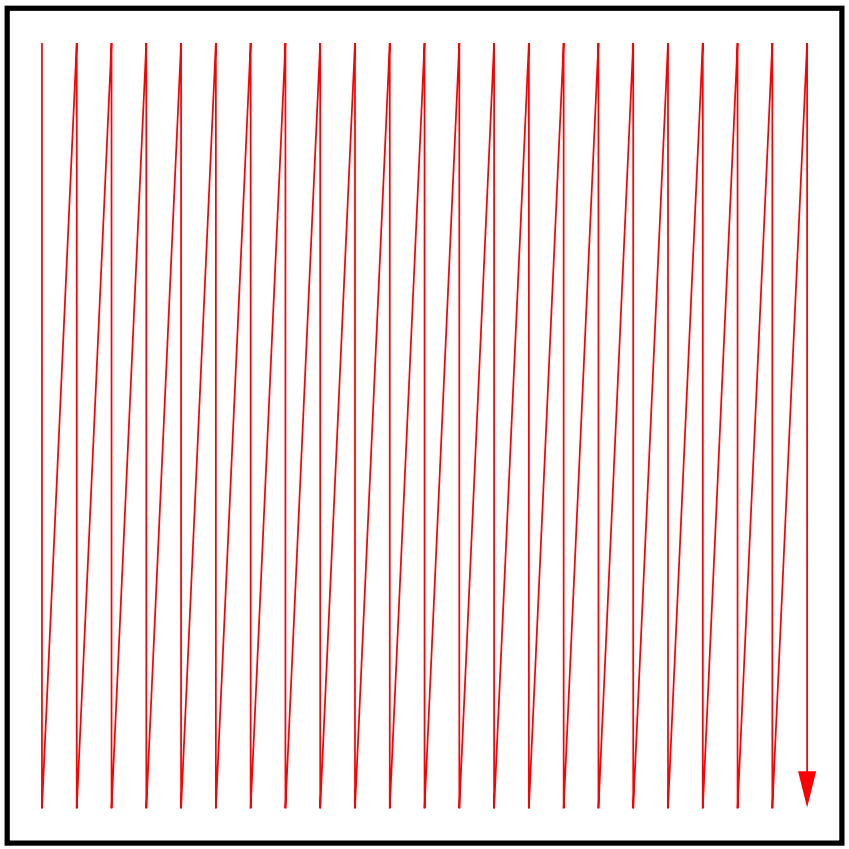

PLASMA/CHAMELEON: tile format

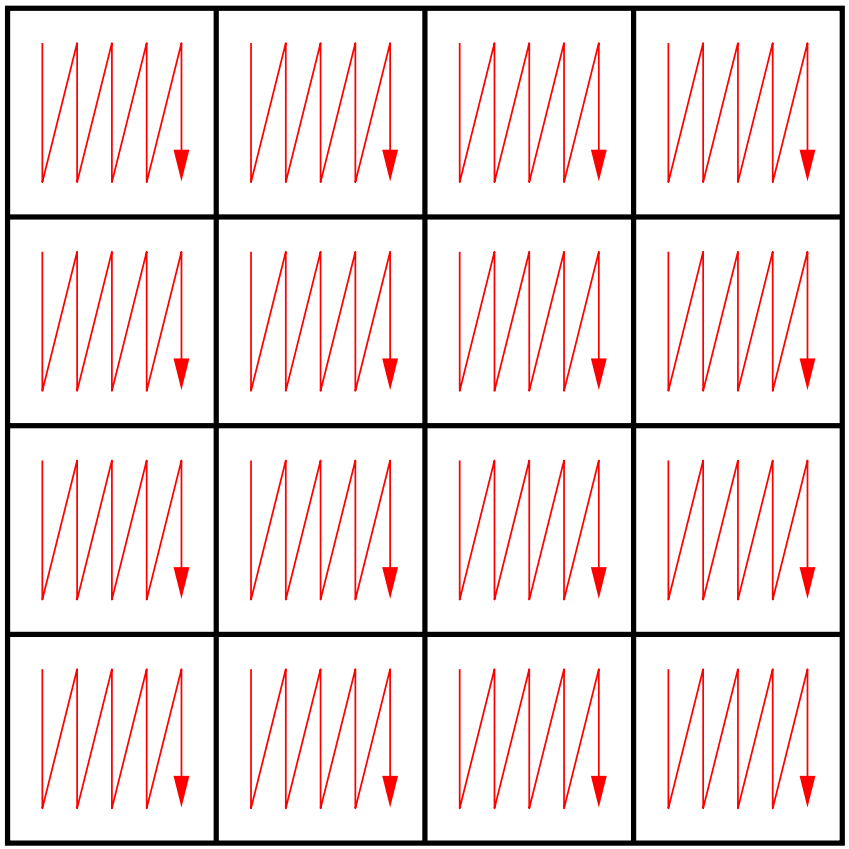

# Directed Acyclic Graph for MOAO

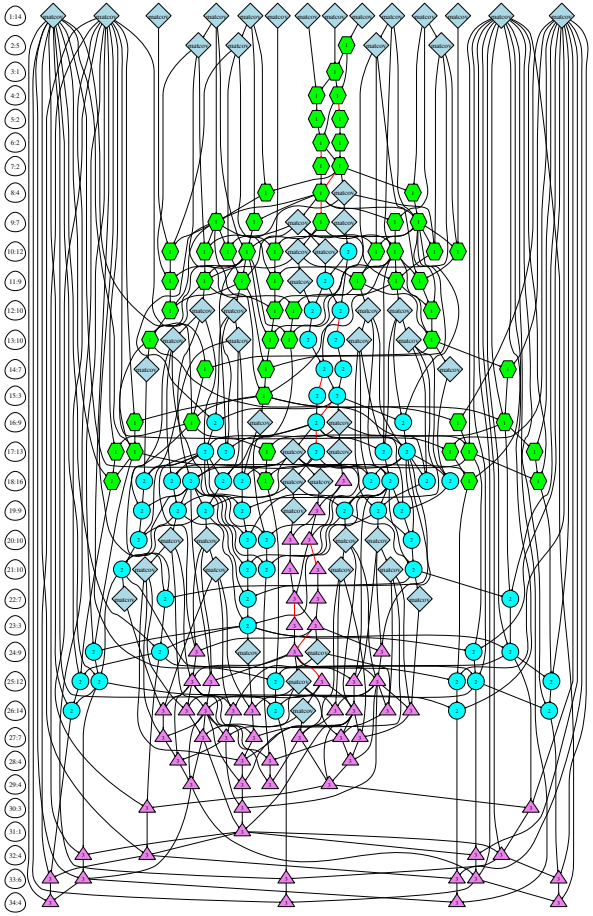

Figure: DAG representation of the overall MOAO framework execution for the calculation of three ToR (i.e., three iterations of the outer loop) and a single observing sequence (i.e., one iteration in the inner loop) on a two-by-two tile matrix.

# Directed Acyclic Graph for MOAO

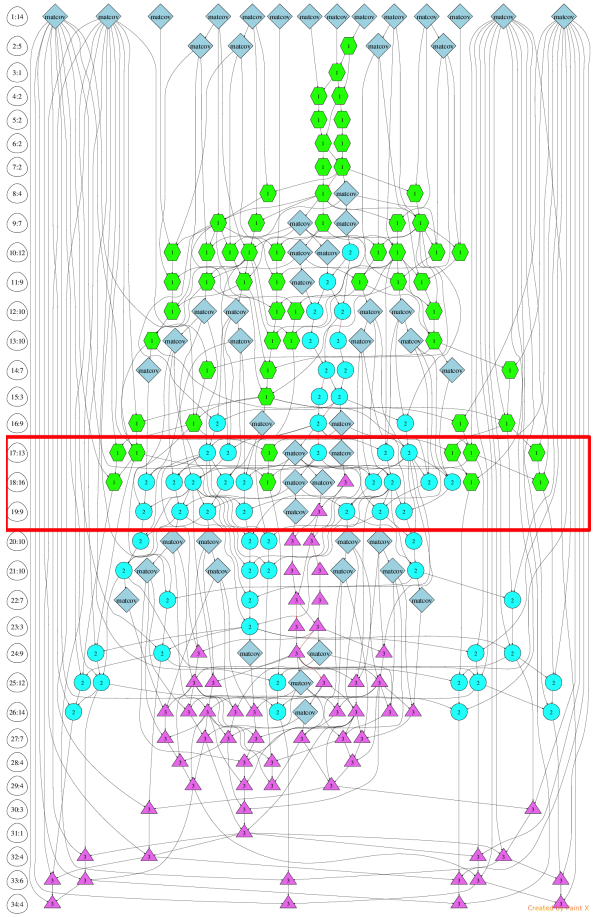

Figure: DAG representation of the overall MOAO framework execution for the calculation of three ToR (i.e., three iterations of the outer loop) and a single observing sequence (i.e., one iteration in the inner loop) on a two-by-two tile matrix.

# Zooming in...

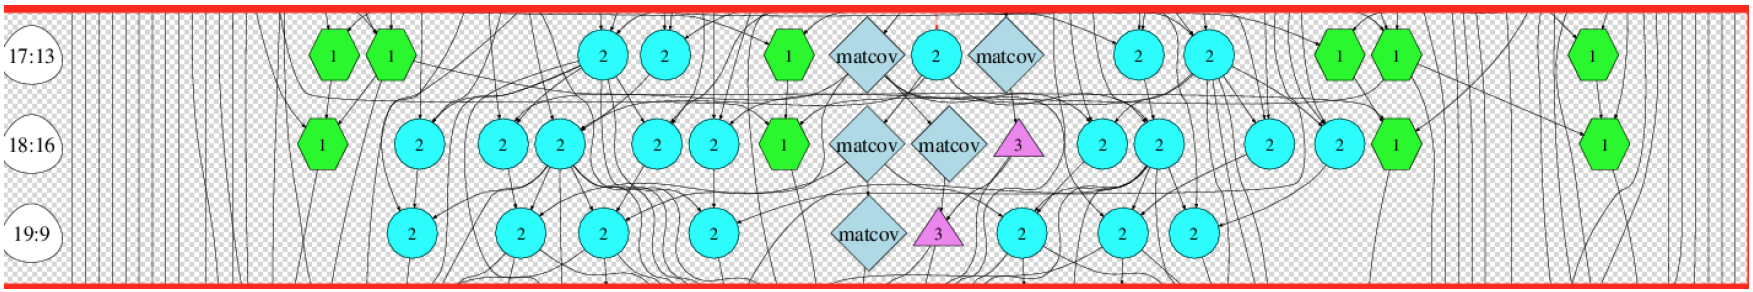

# E-ELT Application Apparatus

- Diameter telescope: 39m
- Number of measurements: up to 100K
- Number of measurements of the true sensor: 10240
- Number of actuators: 5120
- Performance of the ToR computation

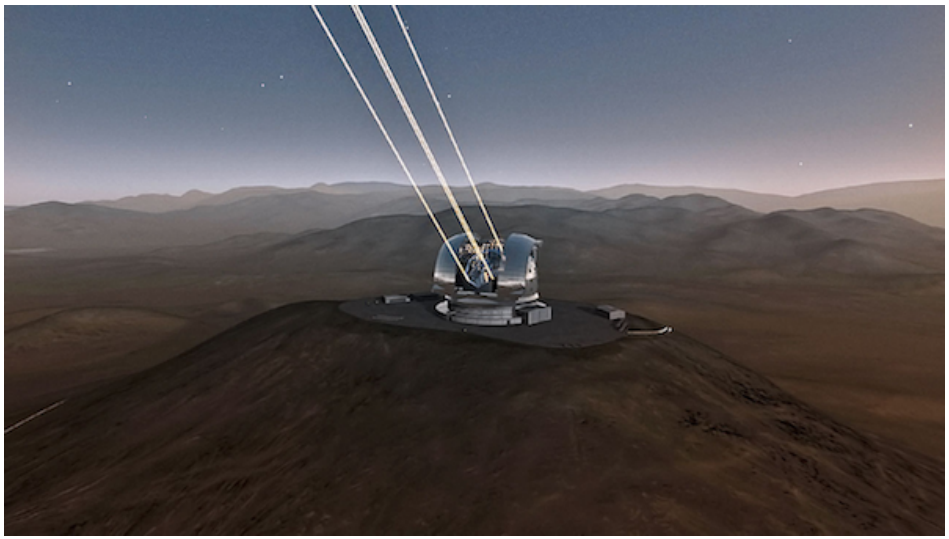

# Performance Evolution of the ToR Computations

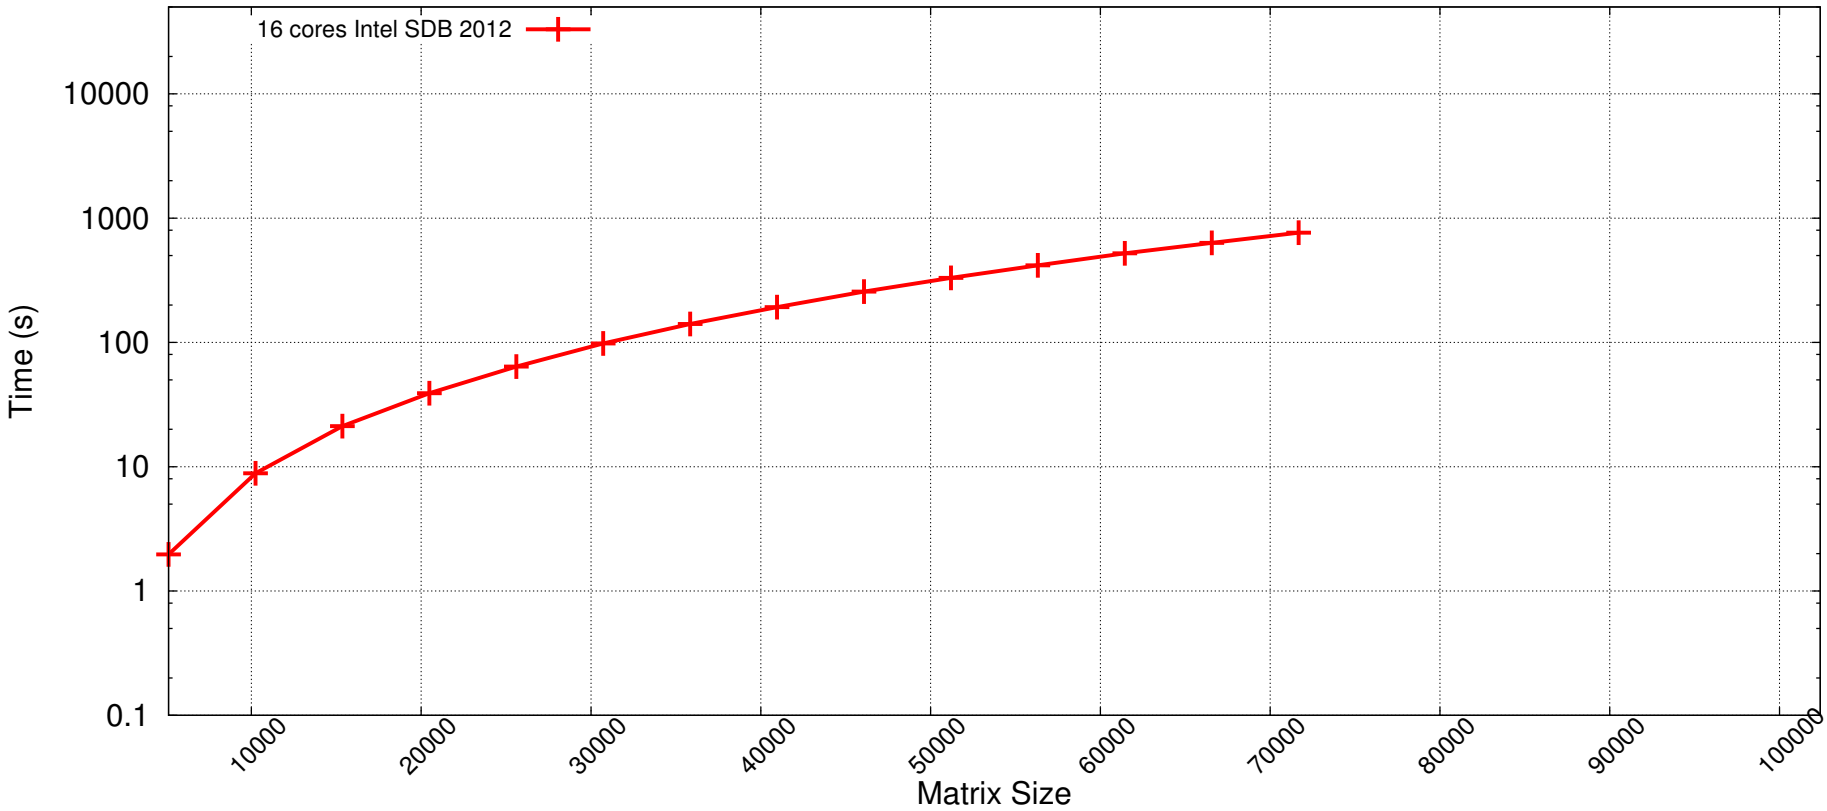

# Performance Evolution of the ToR Computations

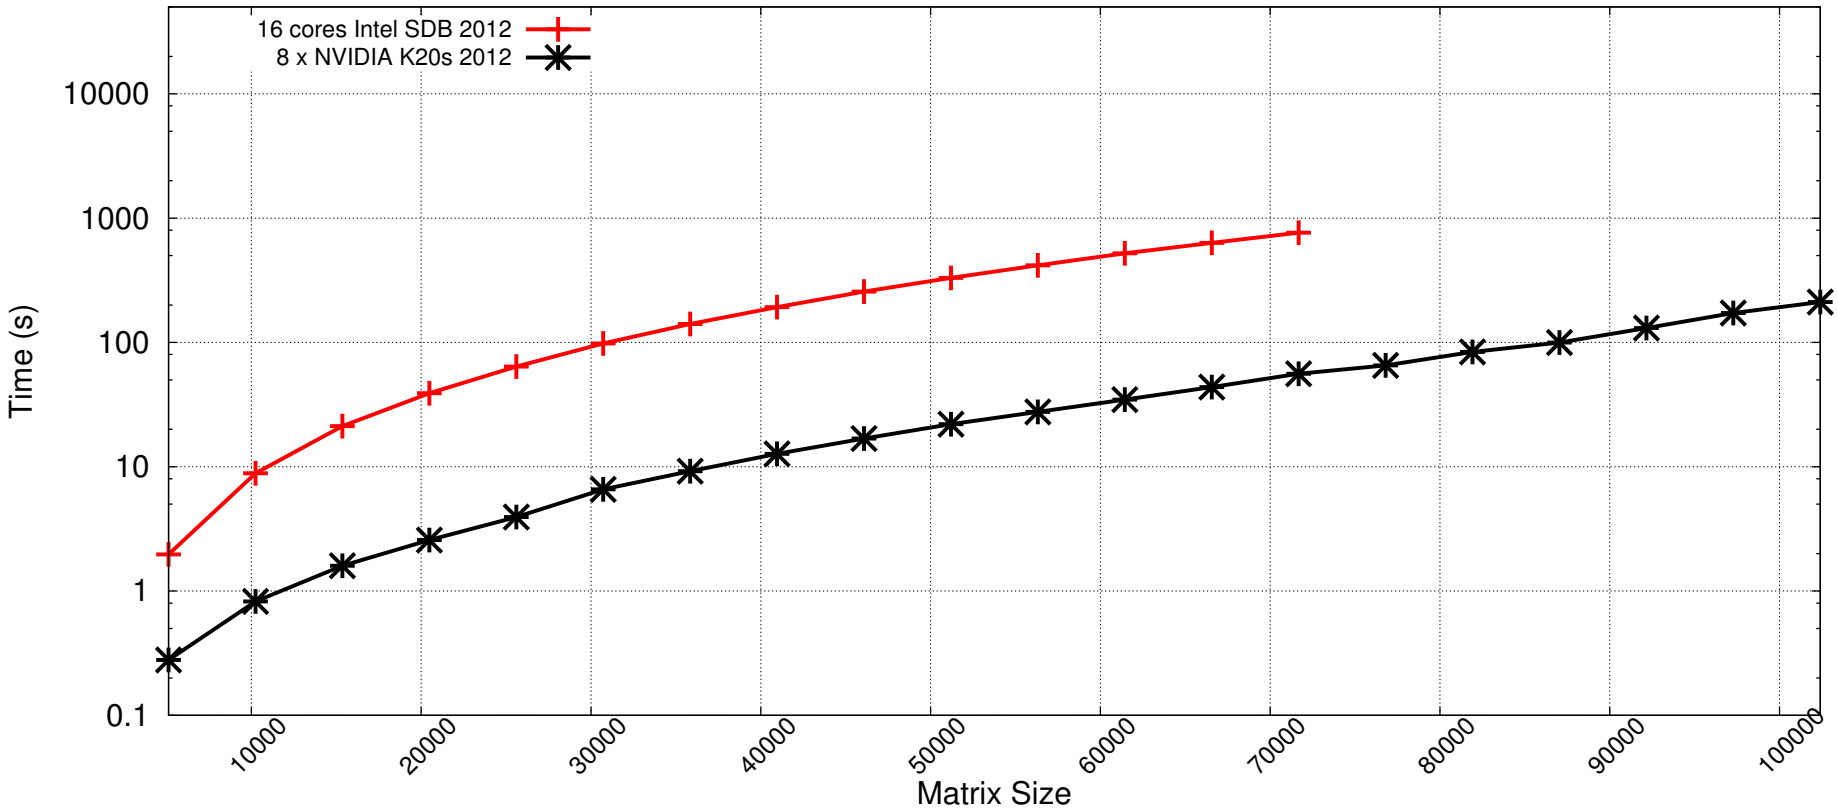

# Performance Evolution of the ToR Computations

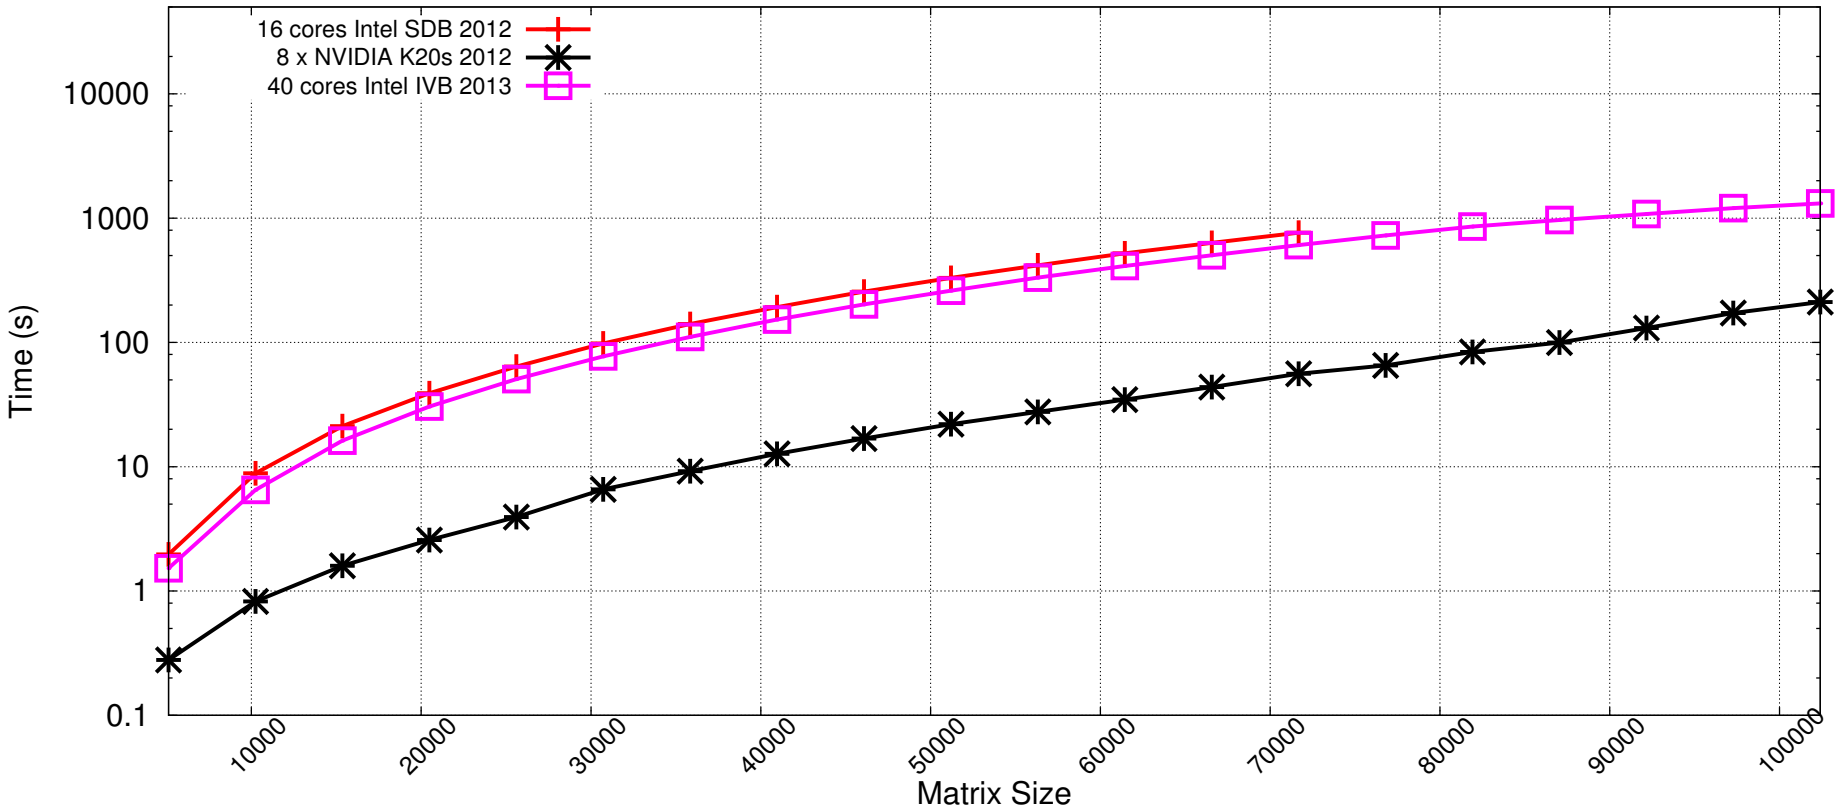

# Performance Evolution of the ToR Computations

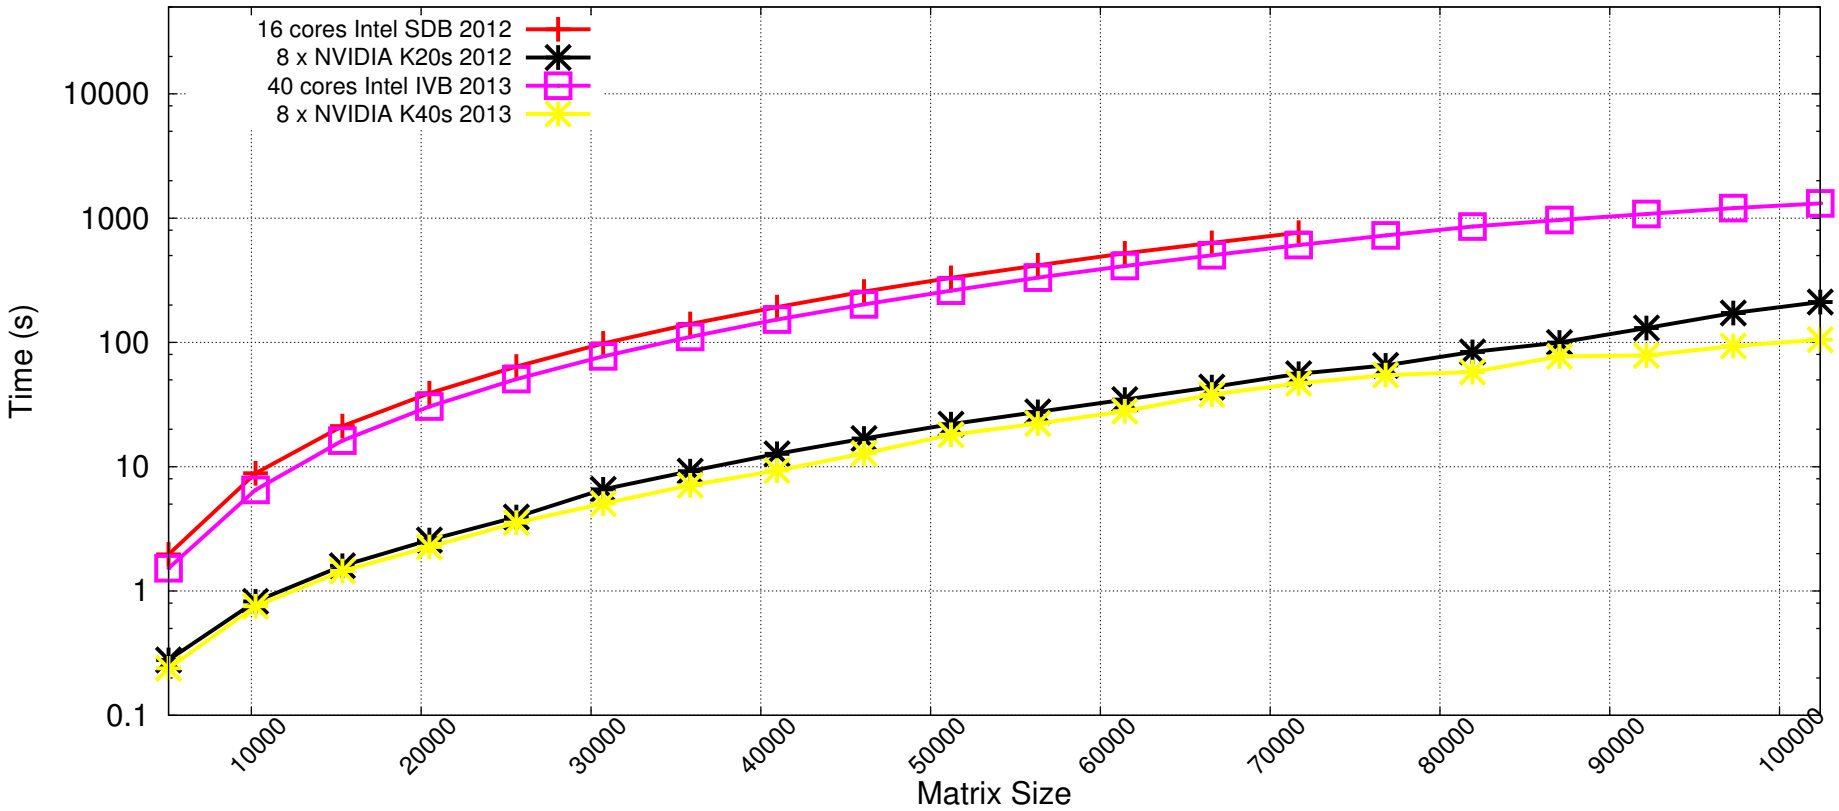

# Performance Evolution of the ToR Computations

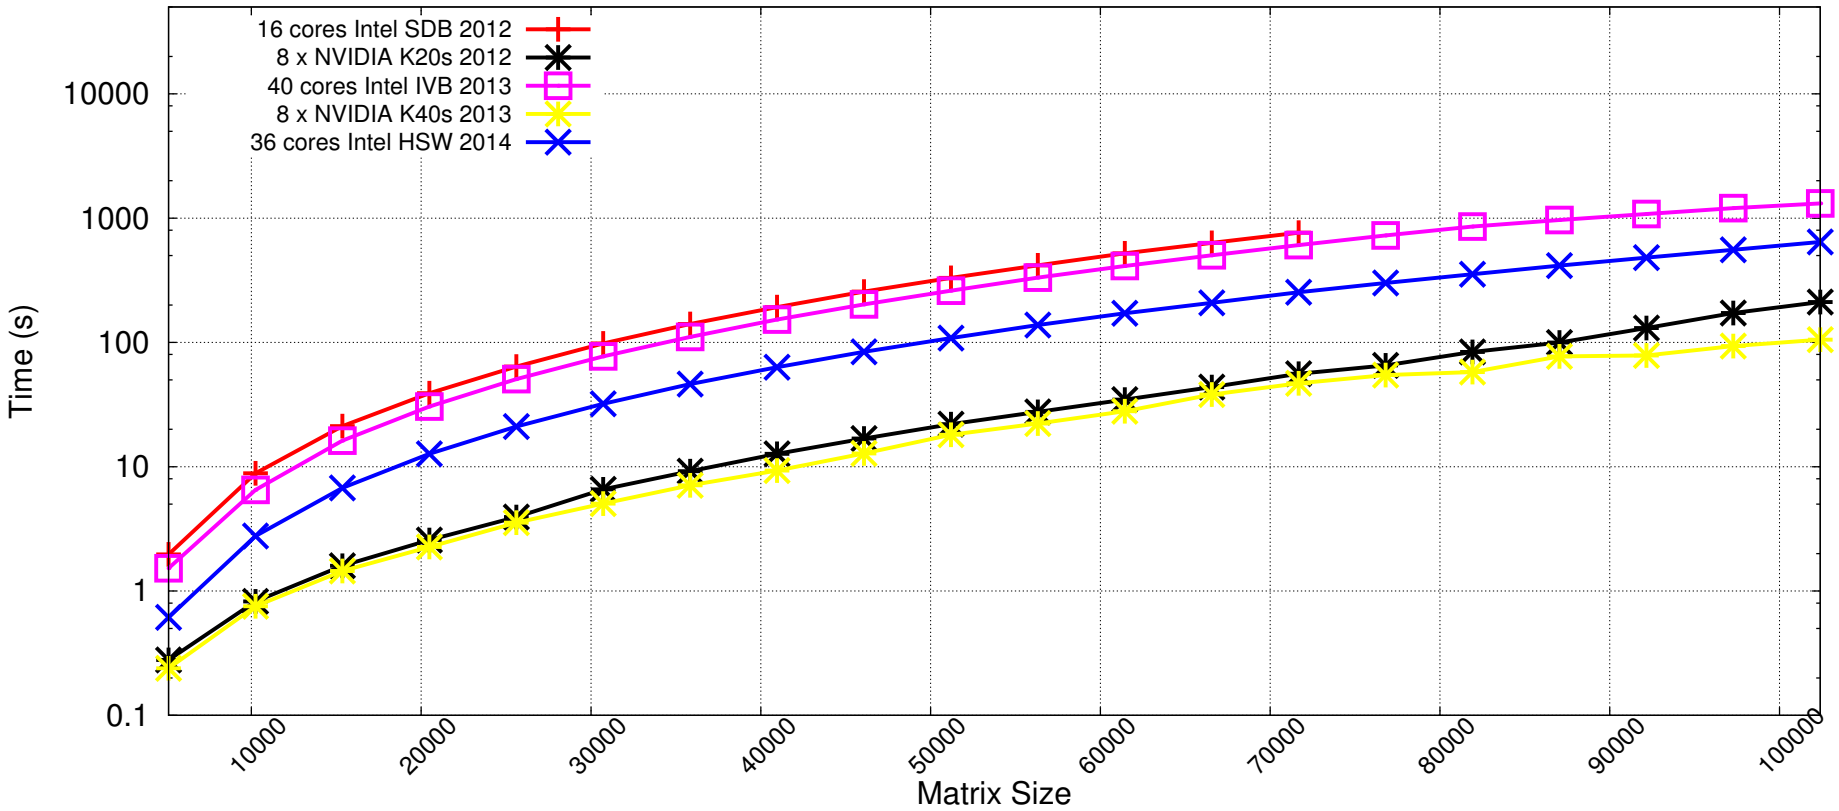

# Performance Evolution of the ToR Computations

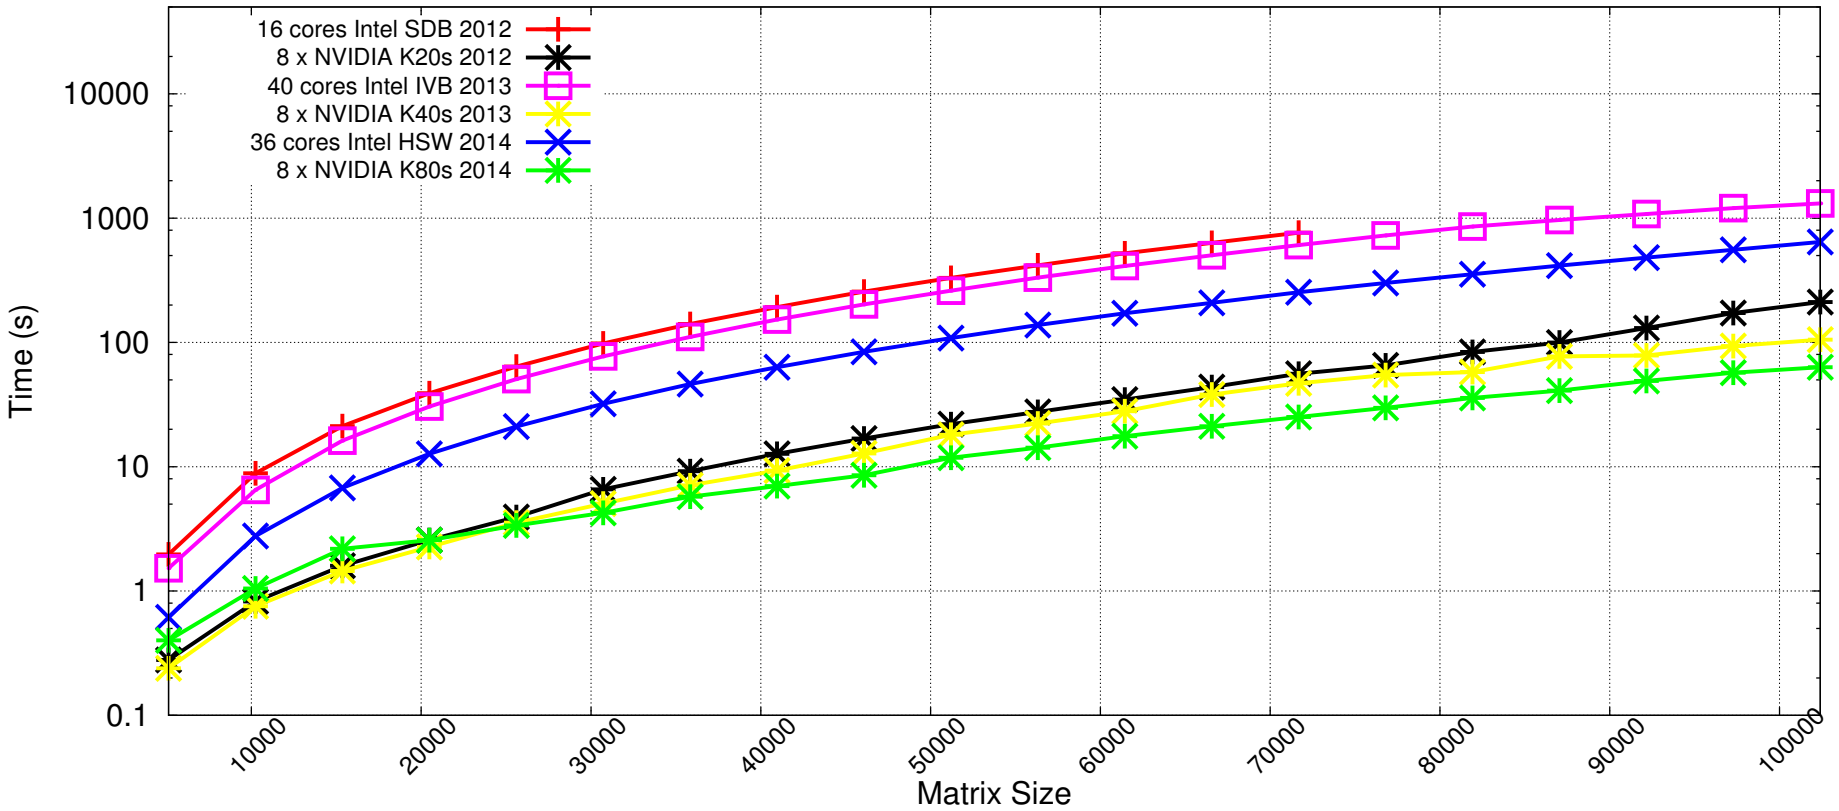

# Performance Evolution of the ToR Computations

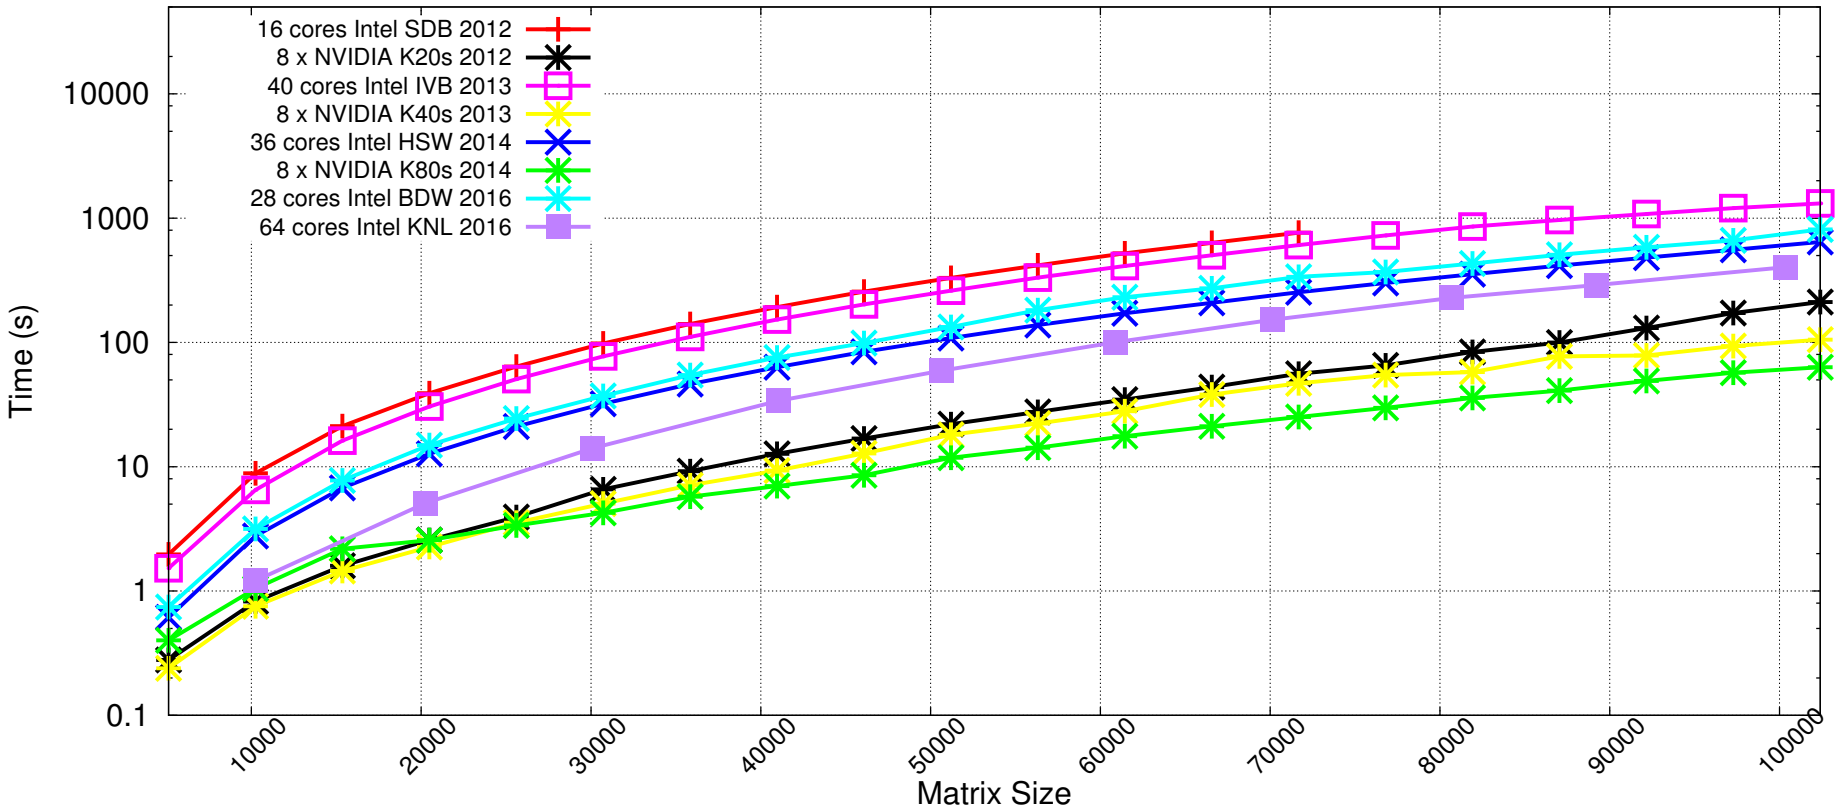

# Performance Evolution of the ToR Computations

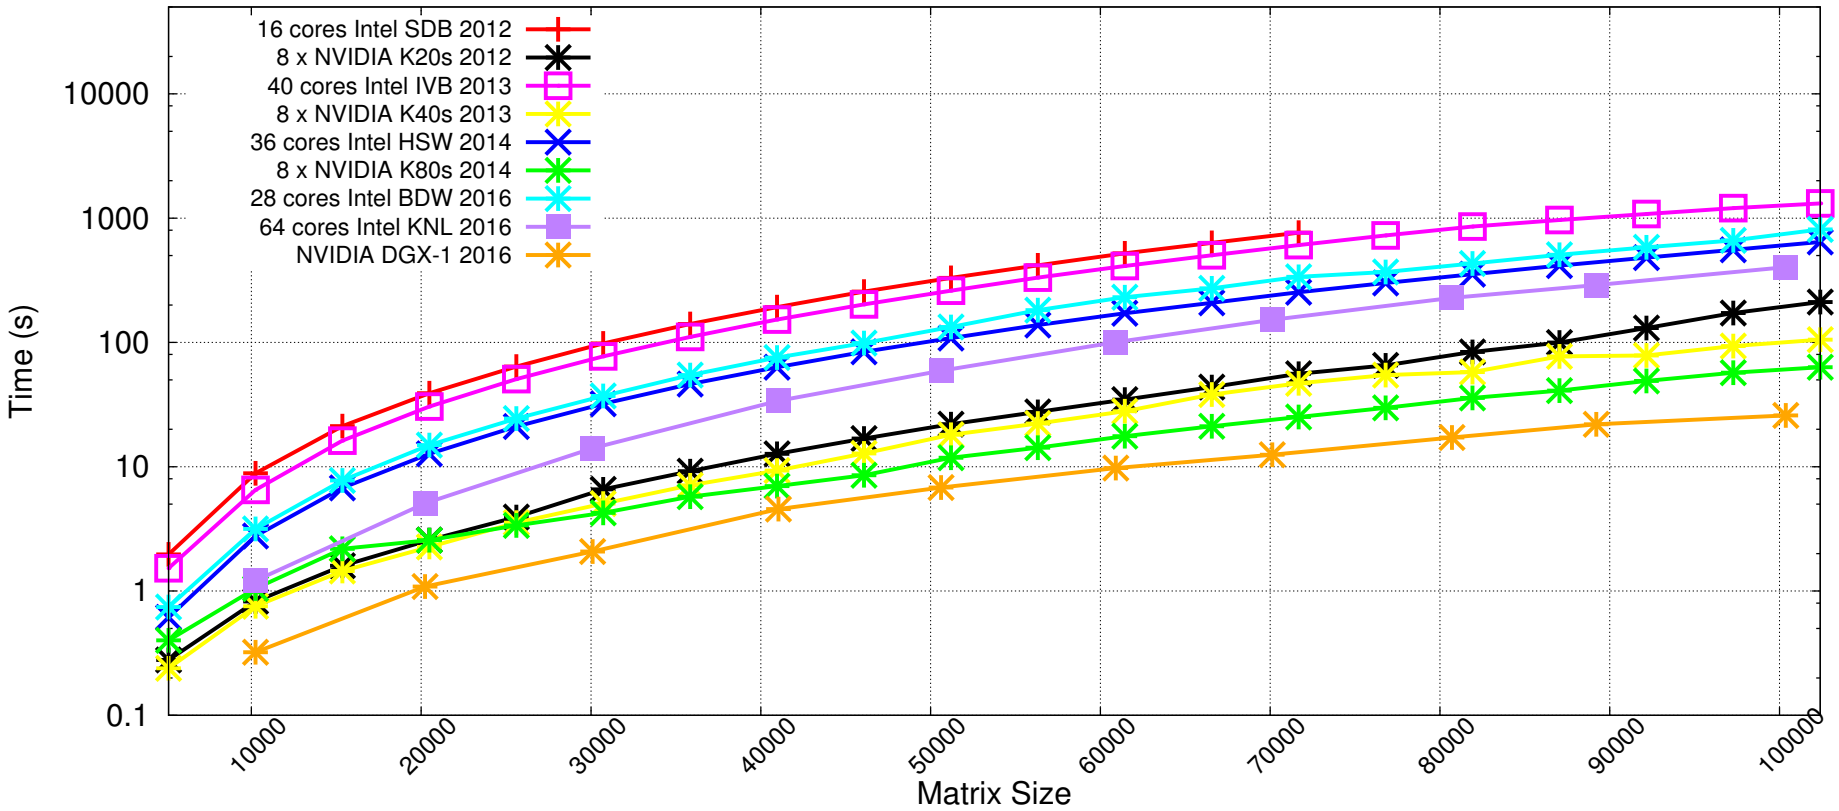

# MOAO Software Release – https://github.com/ecrc/moao

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

**MOAO**

King Abdullah University of Science and Technology | Extreme Computing Research Center | Observatoire de Paris LESIA

**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

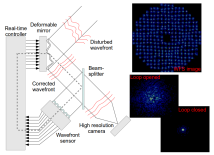


Single conjugate AO concept

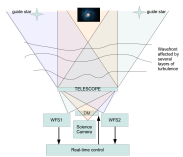


Open-Loop tomography concept

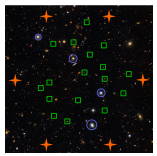

Observing the GOODS South cosmological field with MOAO

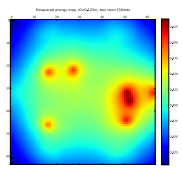

High res. map of the quality of turbulence compensation obtained with MOAO on a cosmological field

## THE PSEUDO-ANALYTICAL APPROACH

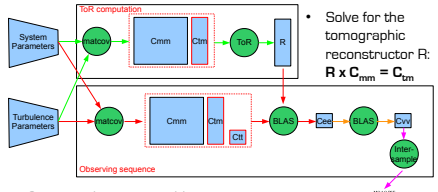


- Solve for the tomographic reconstructor R: $\mathbf{R \times C_{mm} = C_{tm}}$
- Compute the tomographic error: $\mathbf{C_{ee} = C_{tt} - C_{tm} R^T - R C_{tm}^T + R C_{mm} R^T}$
- Compute the equivalent phase map: $\mathbf{C_{vv} = D C_{ee} D^T}$
- Generate the point spread function image

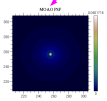

## MOAO 0.1.0

- 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)

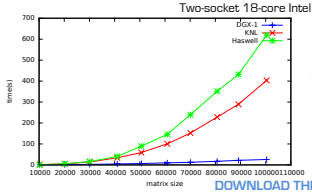


This is one tomographic reconstructor every 25 seconds!

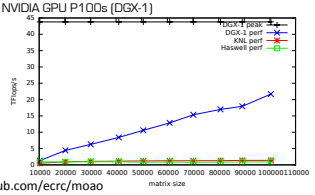


DOWNLOAD THE SOFTWARE AT http://github.com/ecrc/moao

A collaboration of | With support from | Sponsored by

# New DPOTRF @ KBLAS – https://github.com/ecrc/kblas

GPU-resident using recursive formulation

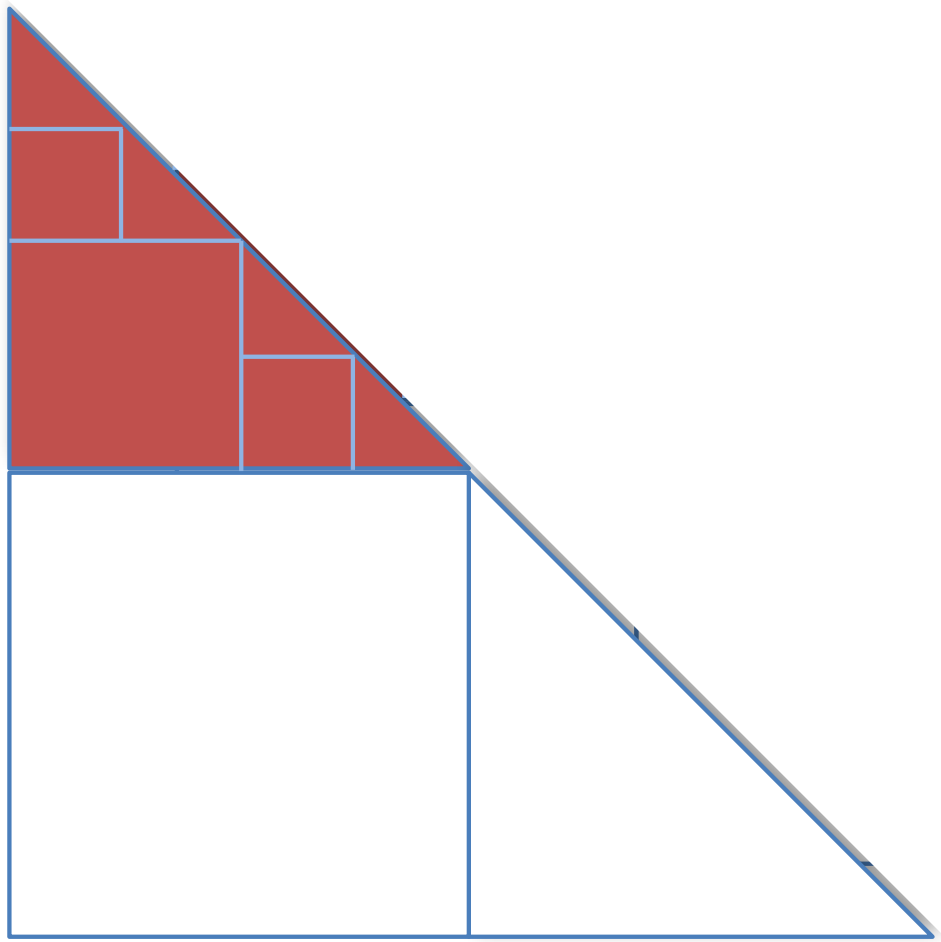

# ToR Performance using GPU-resident DPOTRF@KBLAS

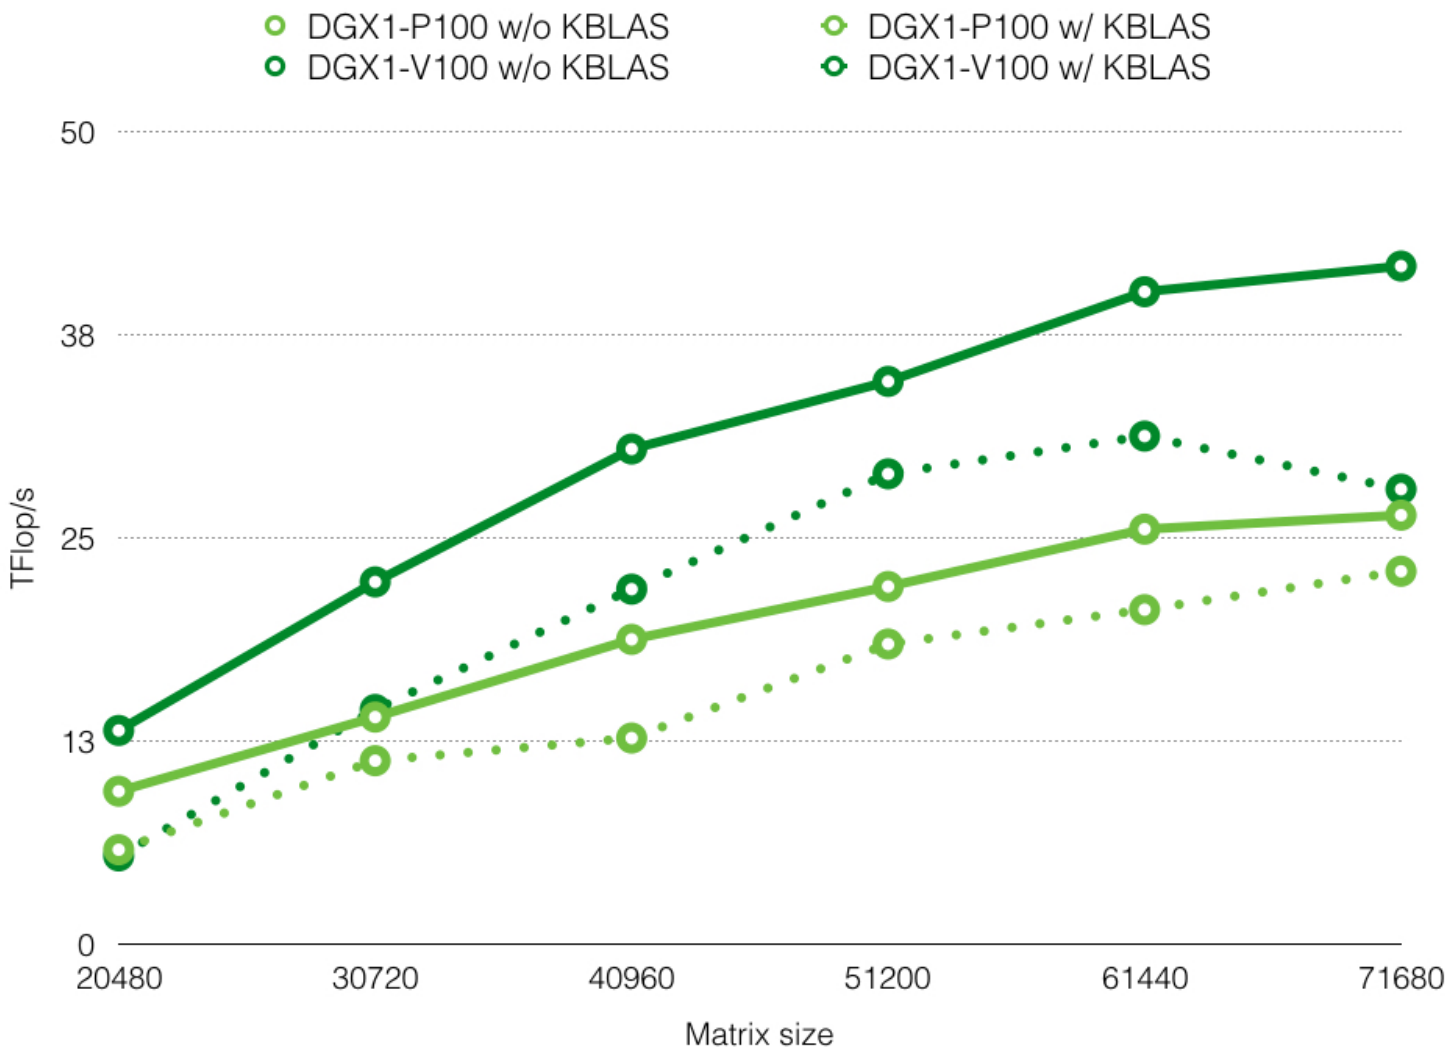

# ToR Performance Evolution

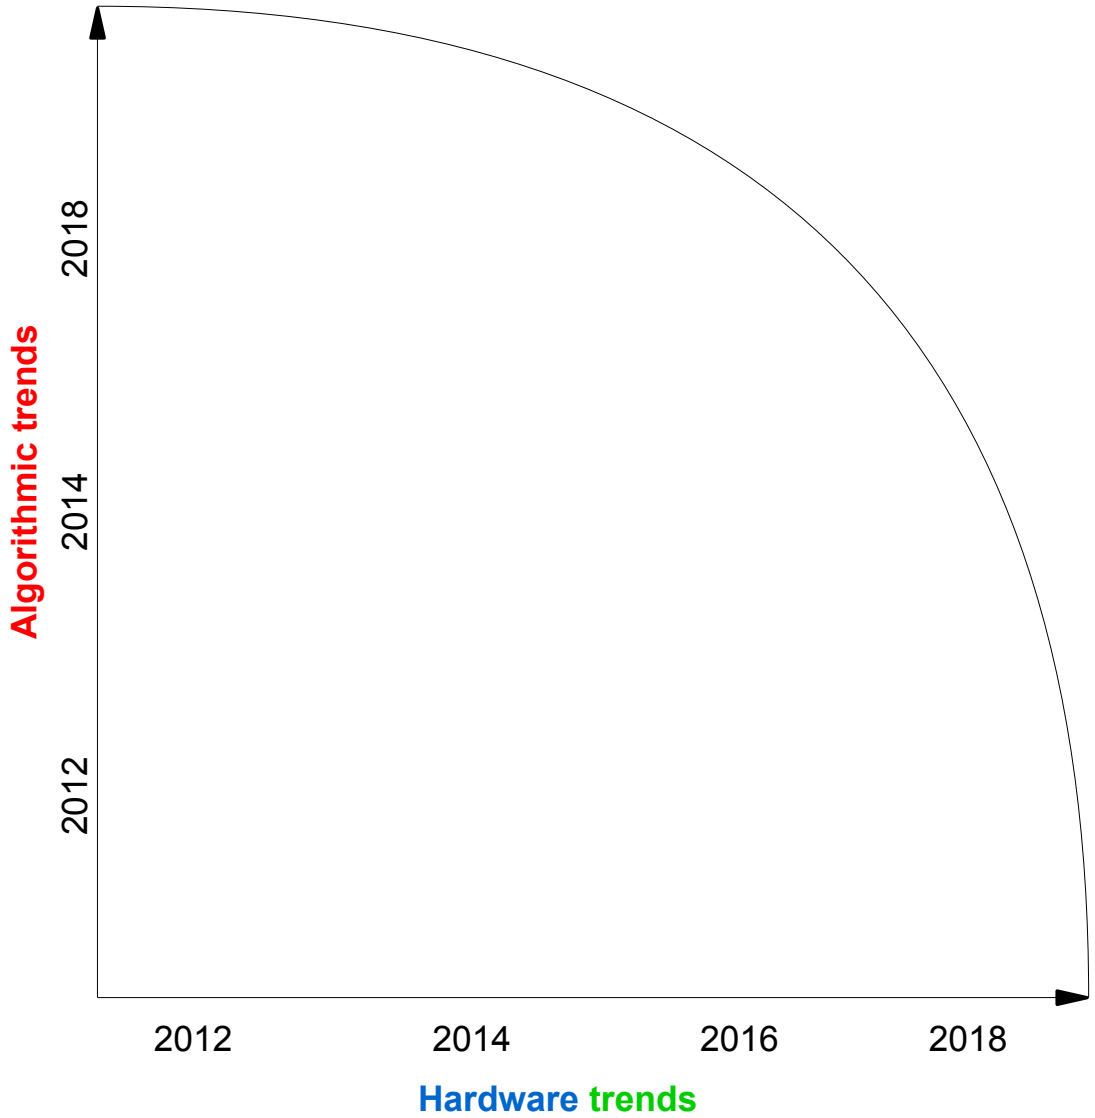

# ToR Performance Evolution

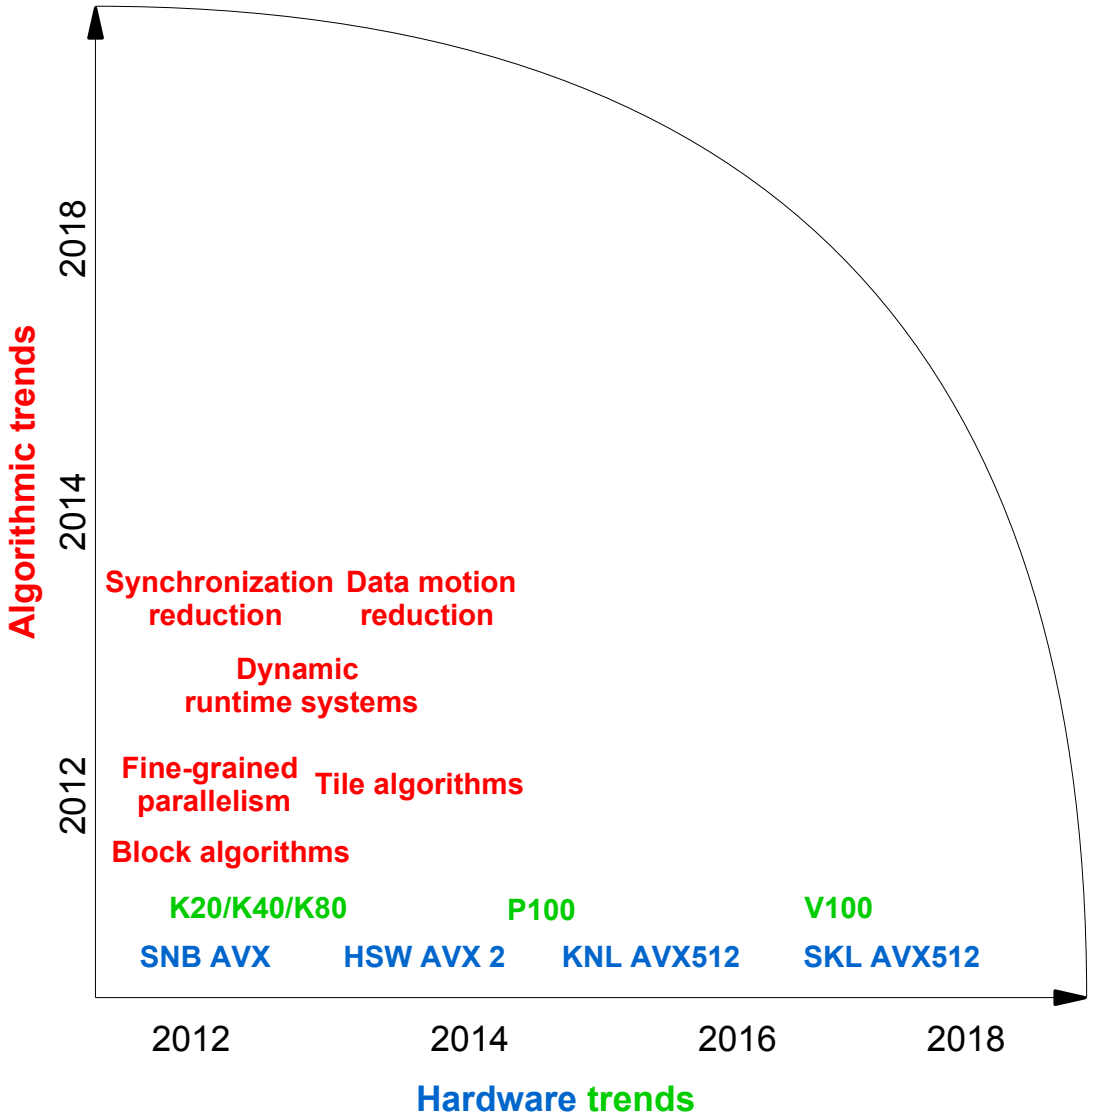

# ToR Performance Evolution

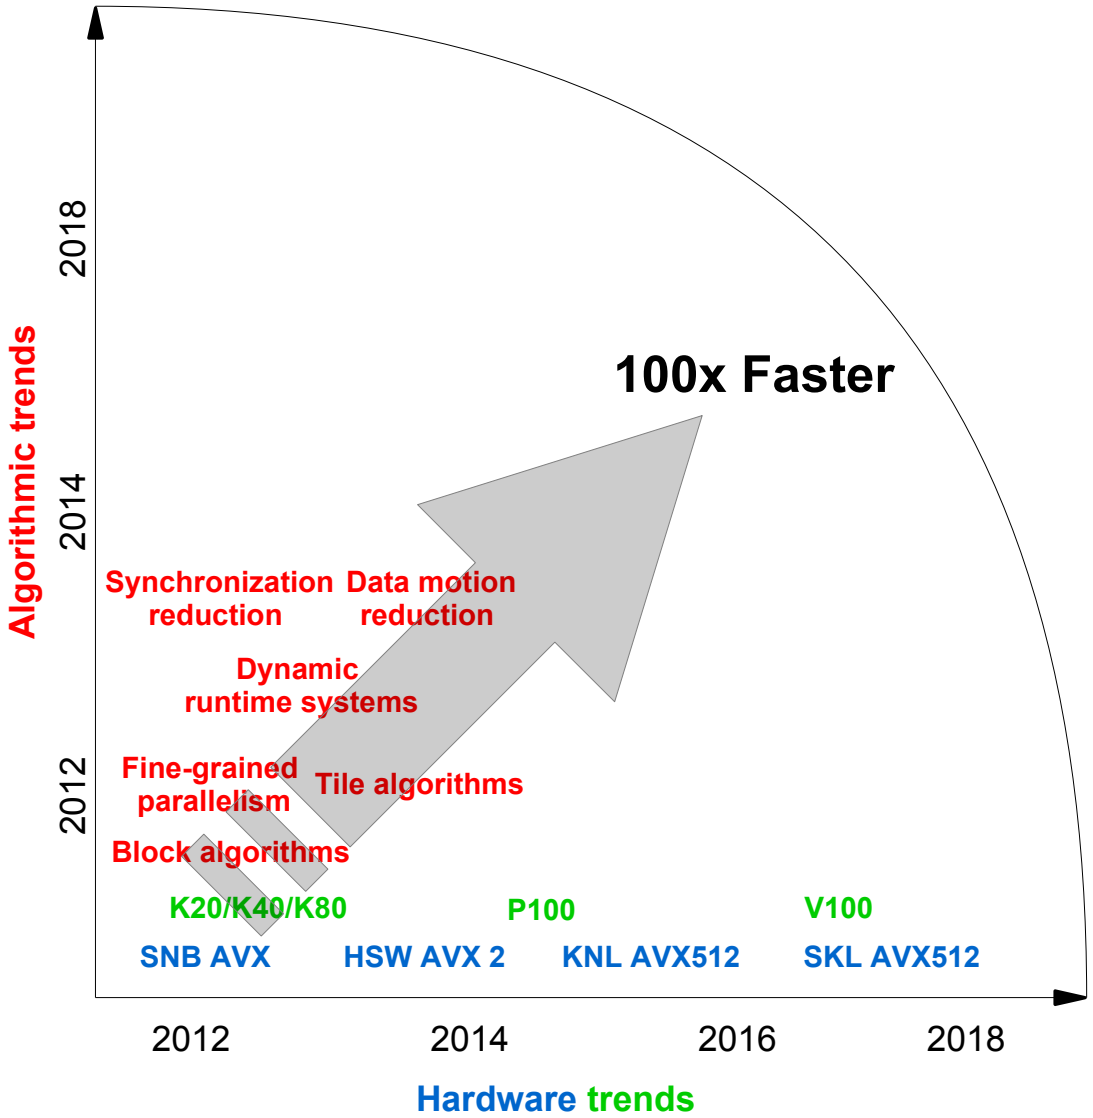

# ToR Performance Evolution

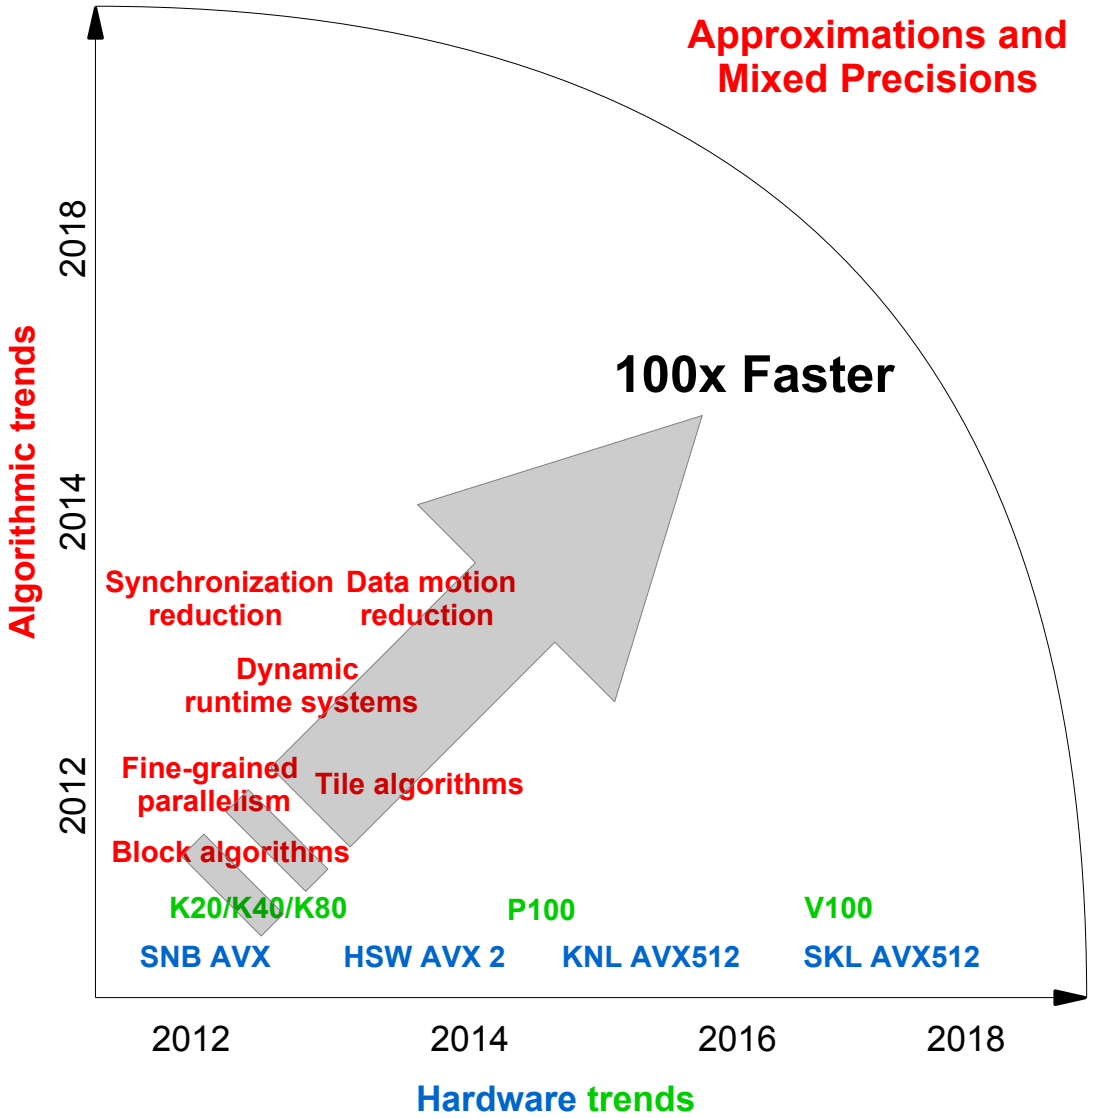

# Accuracy of Mixed Precisions ToR: Brute Force

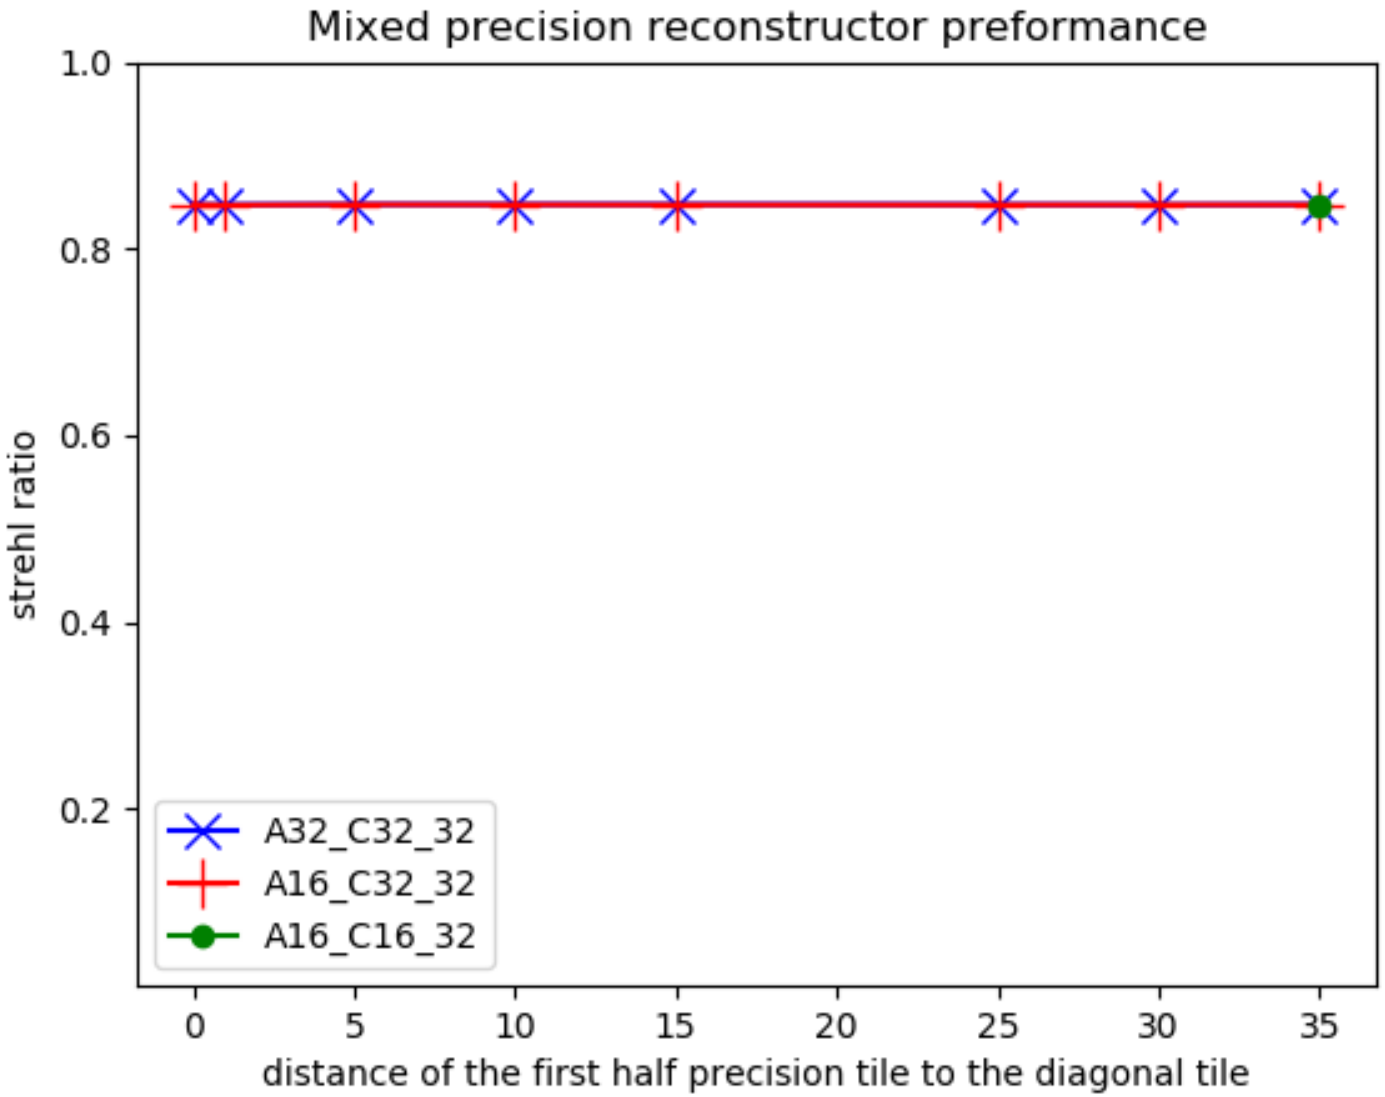

# Xray (SVD) of the Covariance Matrix

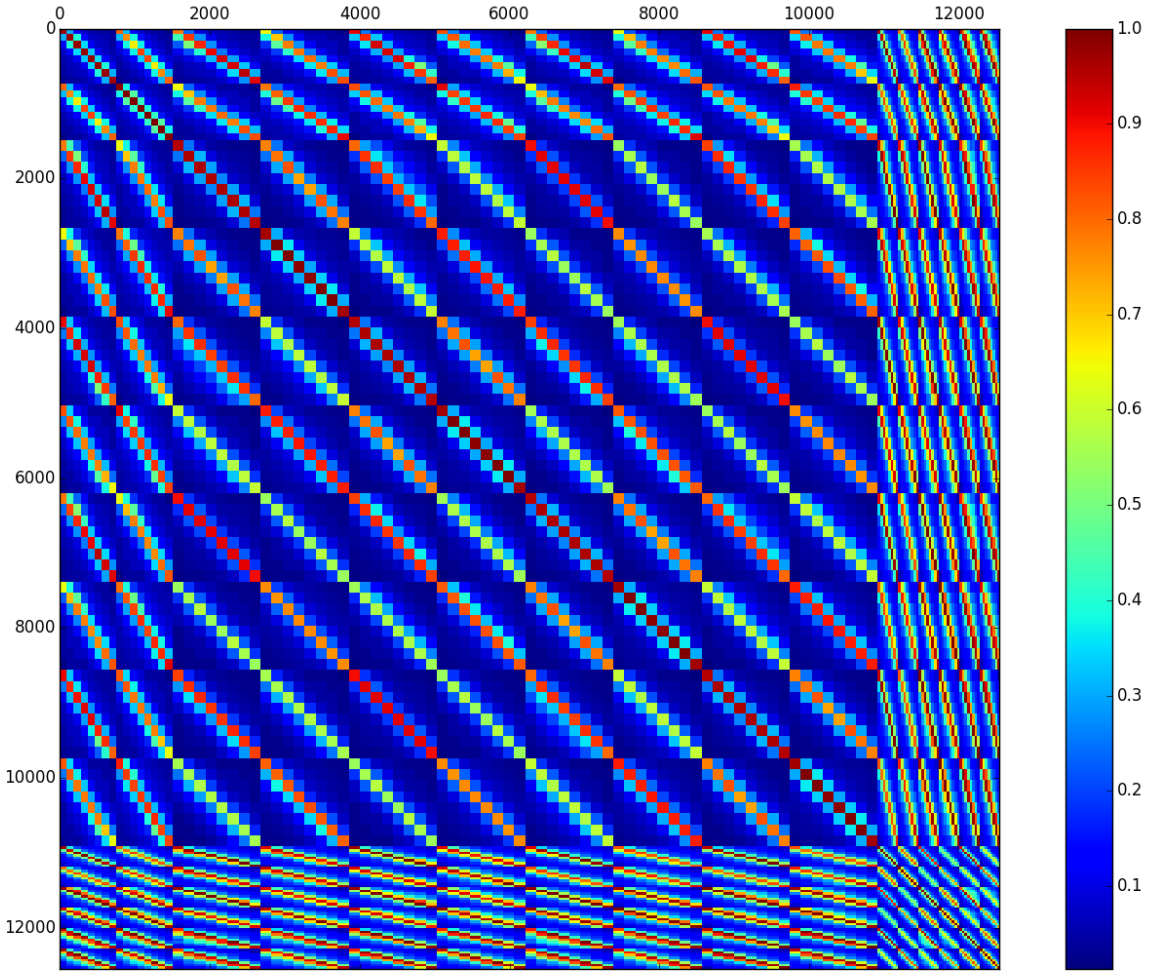

# Accuracy of Mixed Precisions ToR: Smart Heuristic

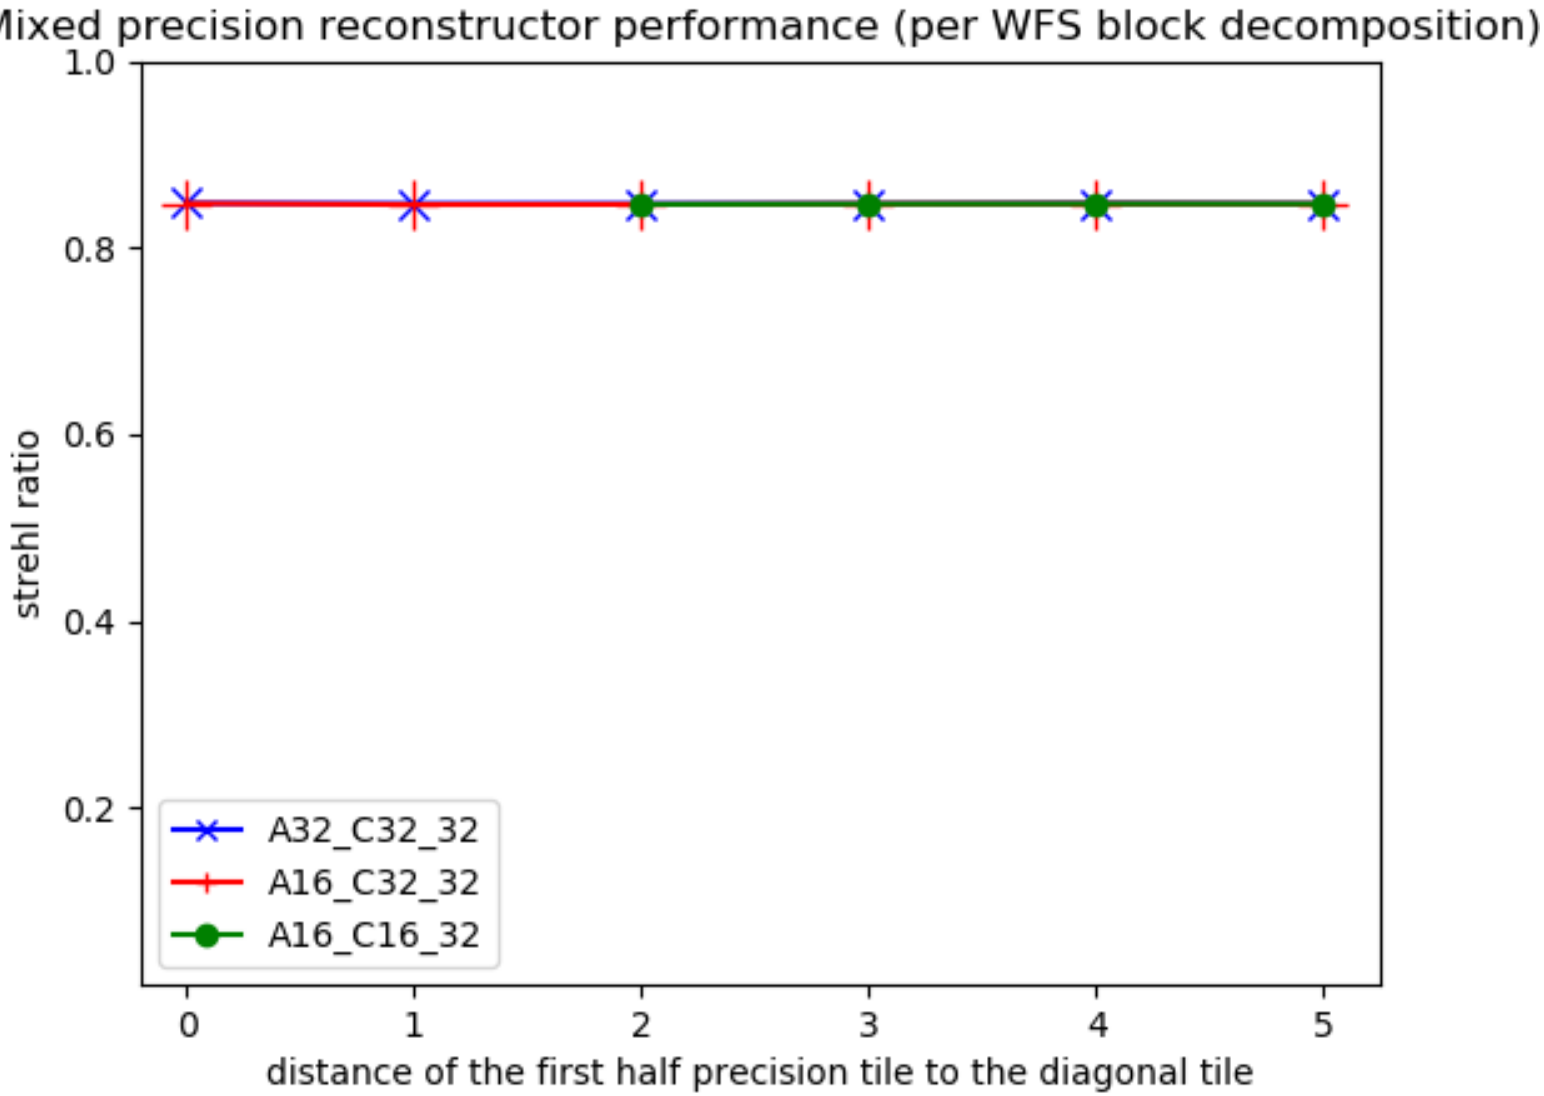

# Estimated Performance of Mixed Precisions ToR

- On NVIDIA V100 GPUs:
  - dgemm achieve about 6.4 Tflop/s on single V100
  - sgemm achieve about 14 Tflop/s on single V100
  - hgemm achieve about 27 Tflop/s on single V100
  - Tensor cores hgemm reach about 85 Tflop/s on single V100
  - Single precision ToR performance: **42** Tflop/s on 8 V100s
    - That is a ToR at ELT scale computed every 25 seconds
  - Speedup factor of **6** between sgemm and hgemm tensor cores
  - **6 x 42 TeraOps/s = 252 PetaOps/s on 8 V100s**
    - That is a ToR at ELT scale computed every 5 seconds
- Probably one of the first real applications amenable for tensor cores usage outside of the traditional AI workloads

# Outline

1 A Hostile Hardware Landscape

2 Linear Algebra Challenges in AO

3 The E-ELT

4 The Subaru Telescope

5 Summary and Future Work

# The Big Picture (Similar w/ SVD)

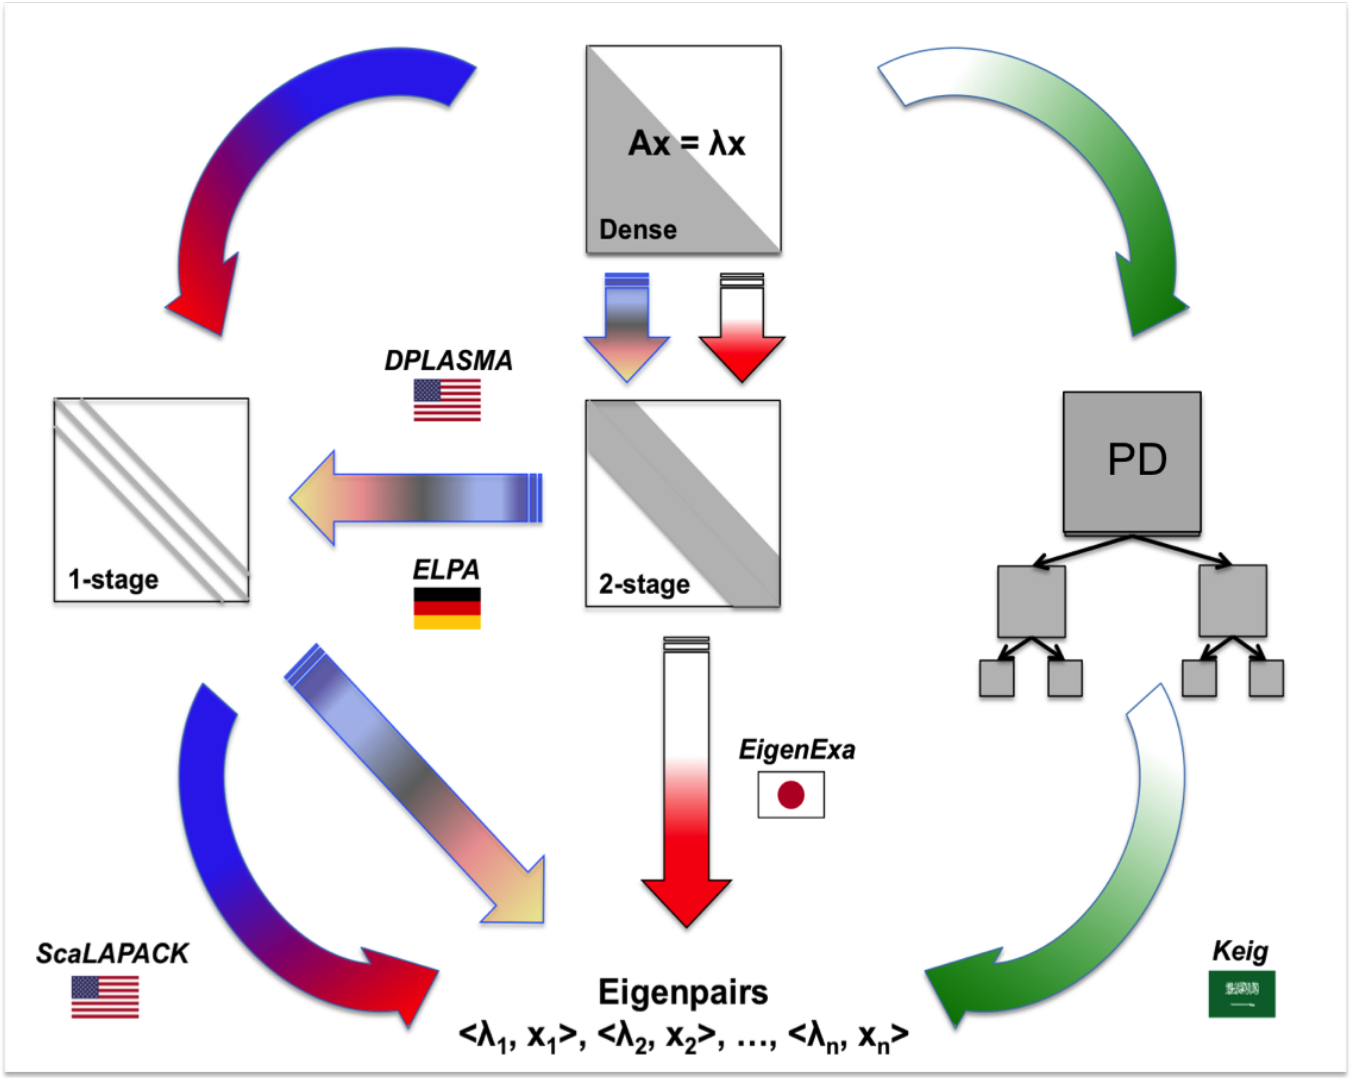


Cray LibSci 17.11.1

# Synthetic Ill-Conditioned matrices, K40

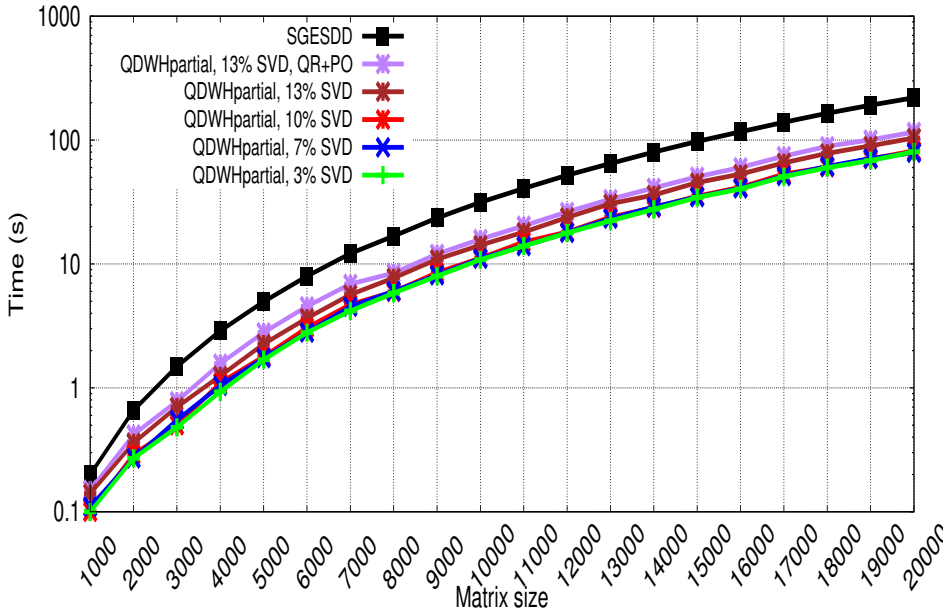


(a) In Seconds.

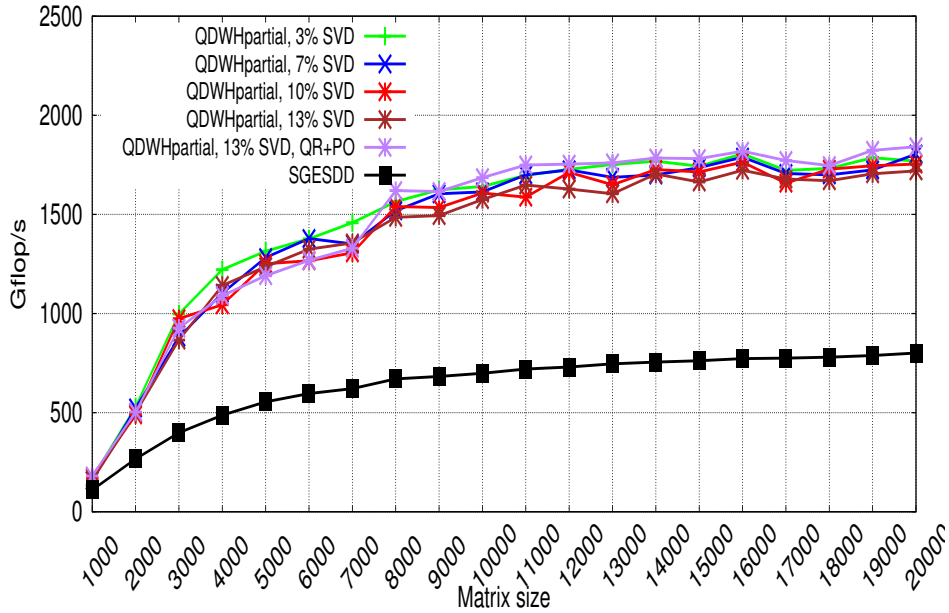


(b) In Gflops/s.

Up to **3X** speedup, 1.8Tflop/s, 45% of the theoretical peak performance

# Synthetic Ill-Conditioned matrices, P100

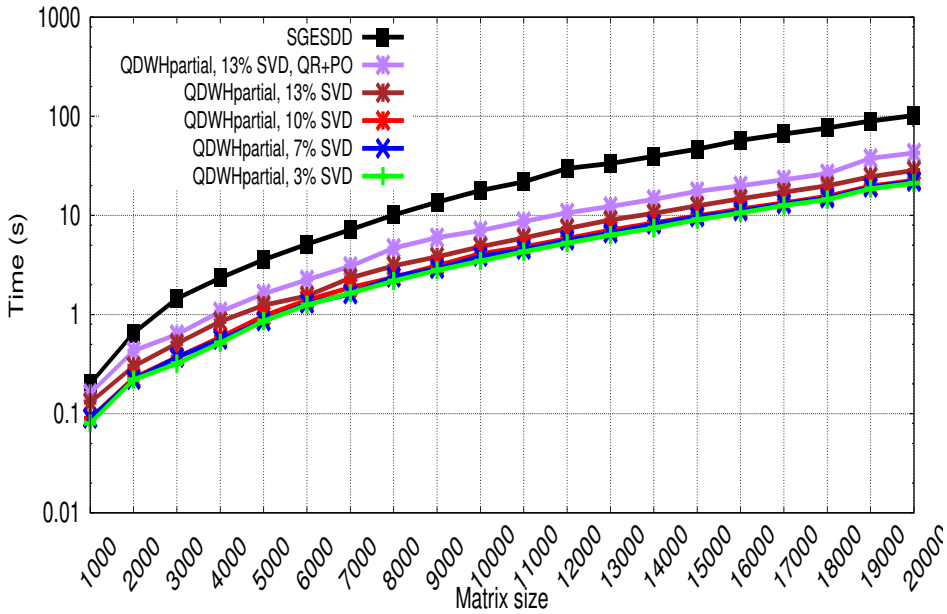


(c) In Seconds.

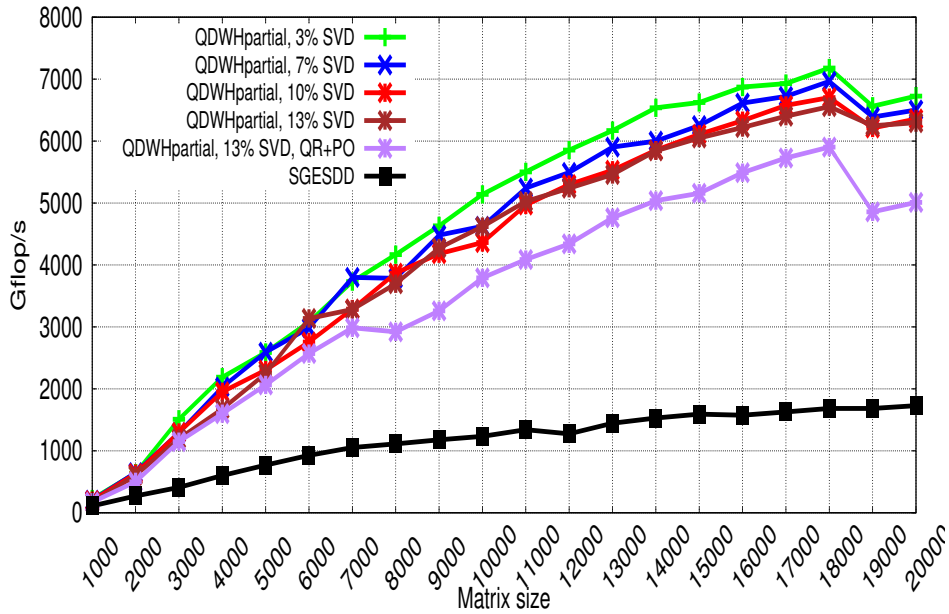


(d) In Gflops/s.

Up to **4X** speedup, 7Tflop/s, 75% of the theoretical peak performance

# Synthetic Ill-Conditioned matrices, V100

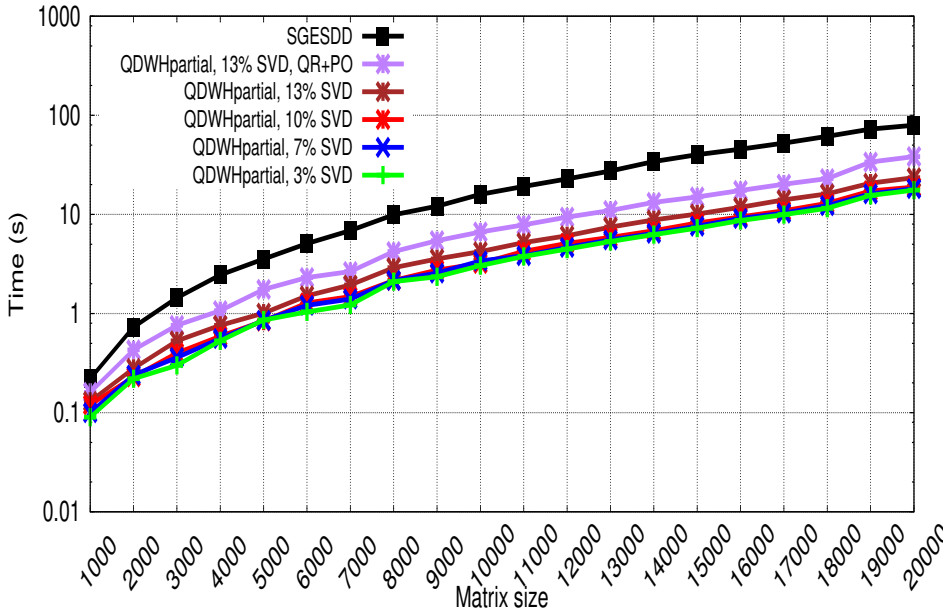


(e) In Seconds.

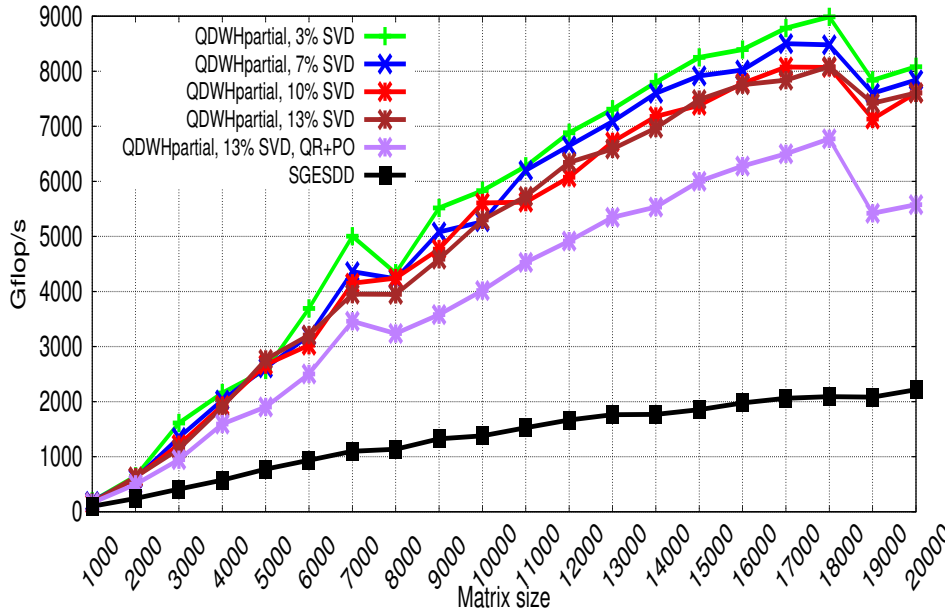


(f) In Gflops/s.

Up to **5X** speedup, 9Tflop/s, 65% of the theoretical peak performance

# Real Observational Datasets from Subaru, V100

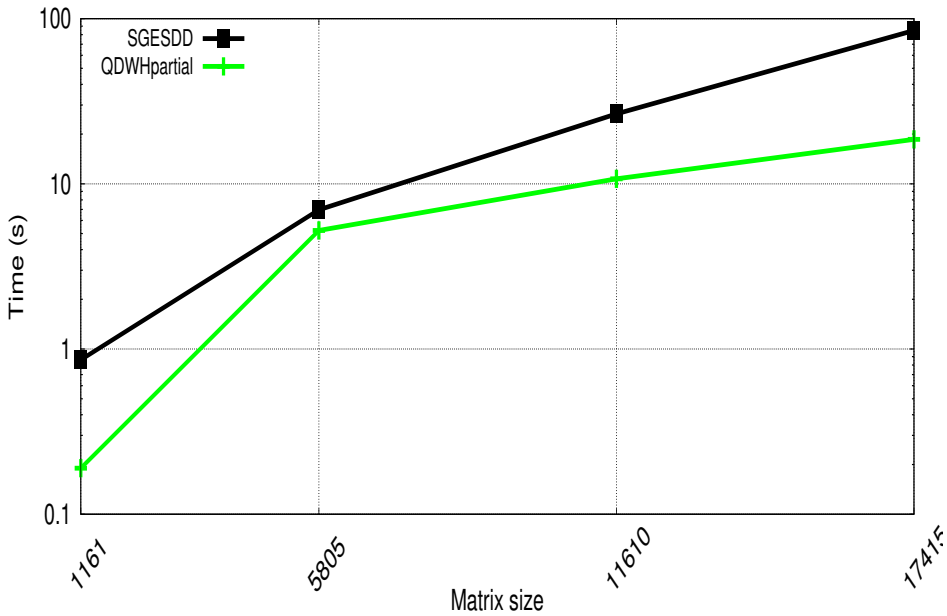


(g) In Seconds.

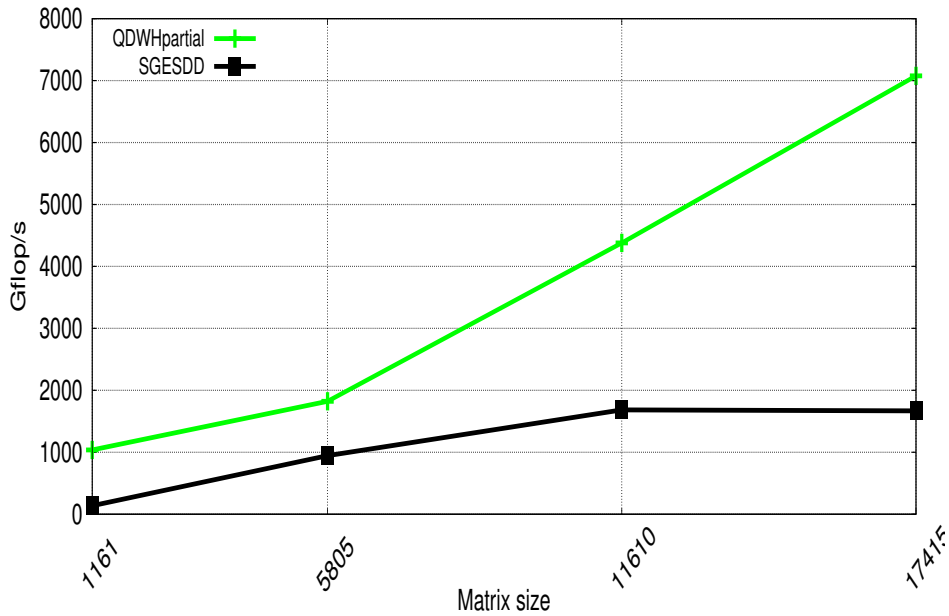


(h) In Gflops/s.

Up to **4X** speedup

# Porting to our Cray XC40 *Shaheen 2.0* Supercomputer

Compute Node x 6174 with dual socket 16-core Intel Haswell (200,000 cores total)

# Scaling up using 30,000 cores

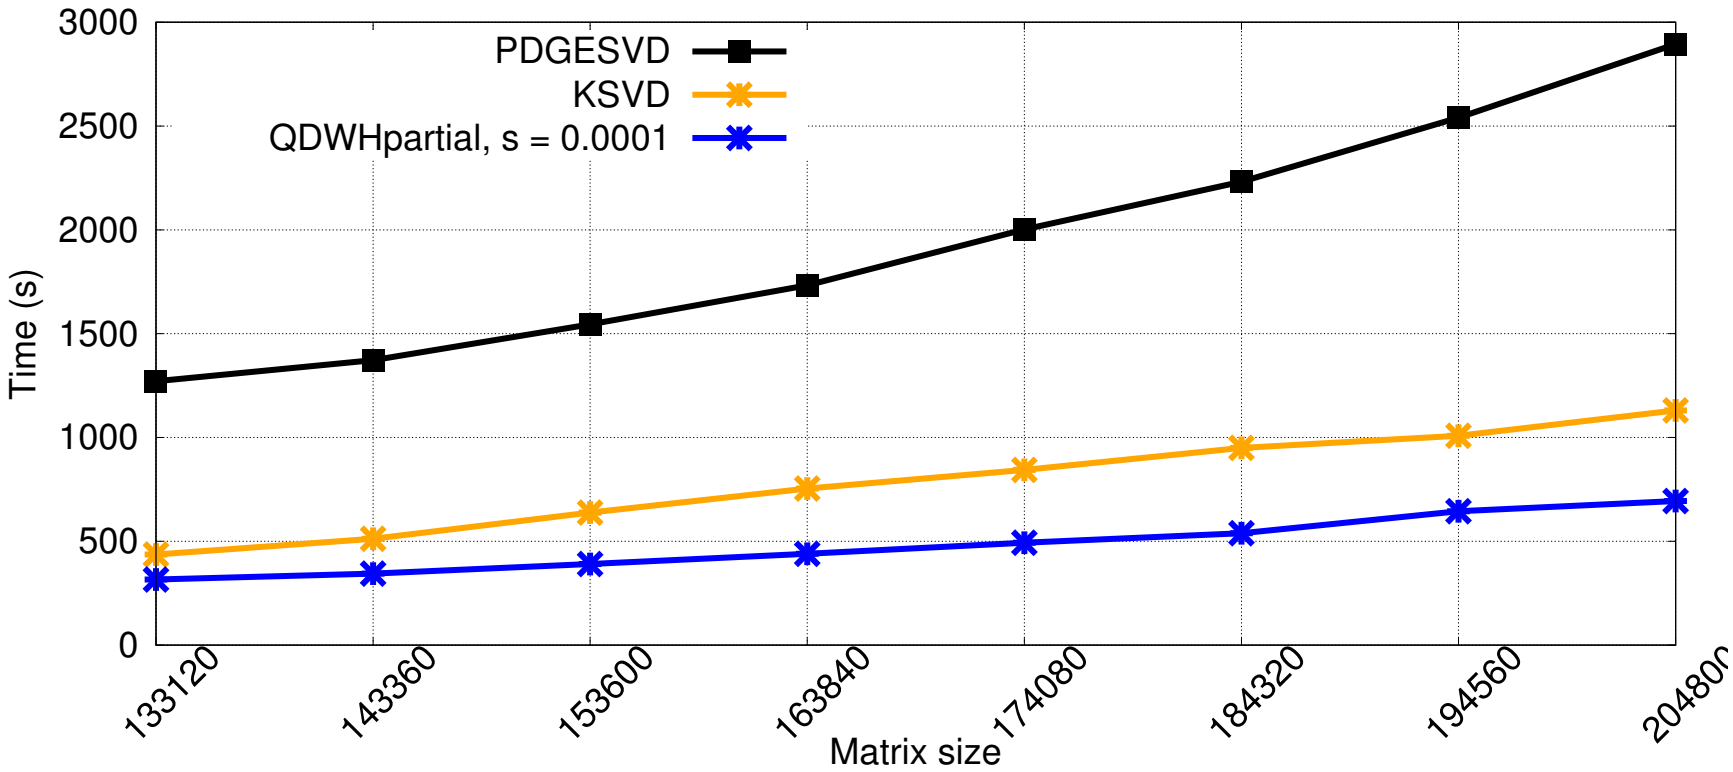

# Outline

1 A Hostile Hardware Landscape

2 Linear Algebra Challenges in AO

3 The E-ELT

4 The Subaru Telescope

5 Summary and Future Work

# Summary

- Adaptive optics require massively parallel hardware accelerators
- Efficient Task-based programming model
- Dense, tightly-connected GPU-based compute node (i.e., *DGX-1*) for real-time processing
- AO community effort for software standardization (w/ D. Gratadour and O. Guyon)
- Leverage the Toeplitz matrix structure (from cubic to quadratic algorithmic complexity)

# Exciting Time for Astronomy at KAUST/ECRC!

- Supporting two major worldwide ground-based astronomy efforts

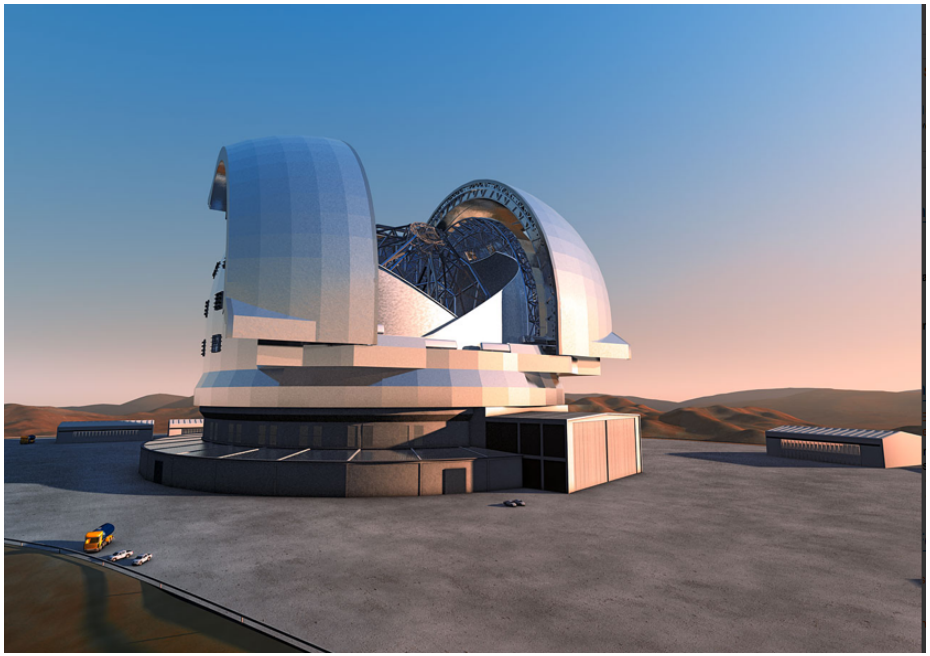

The E-ELT Telescope

The Subaru Telescope

# Bringing Astronomy Back Home ;-)

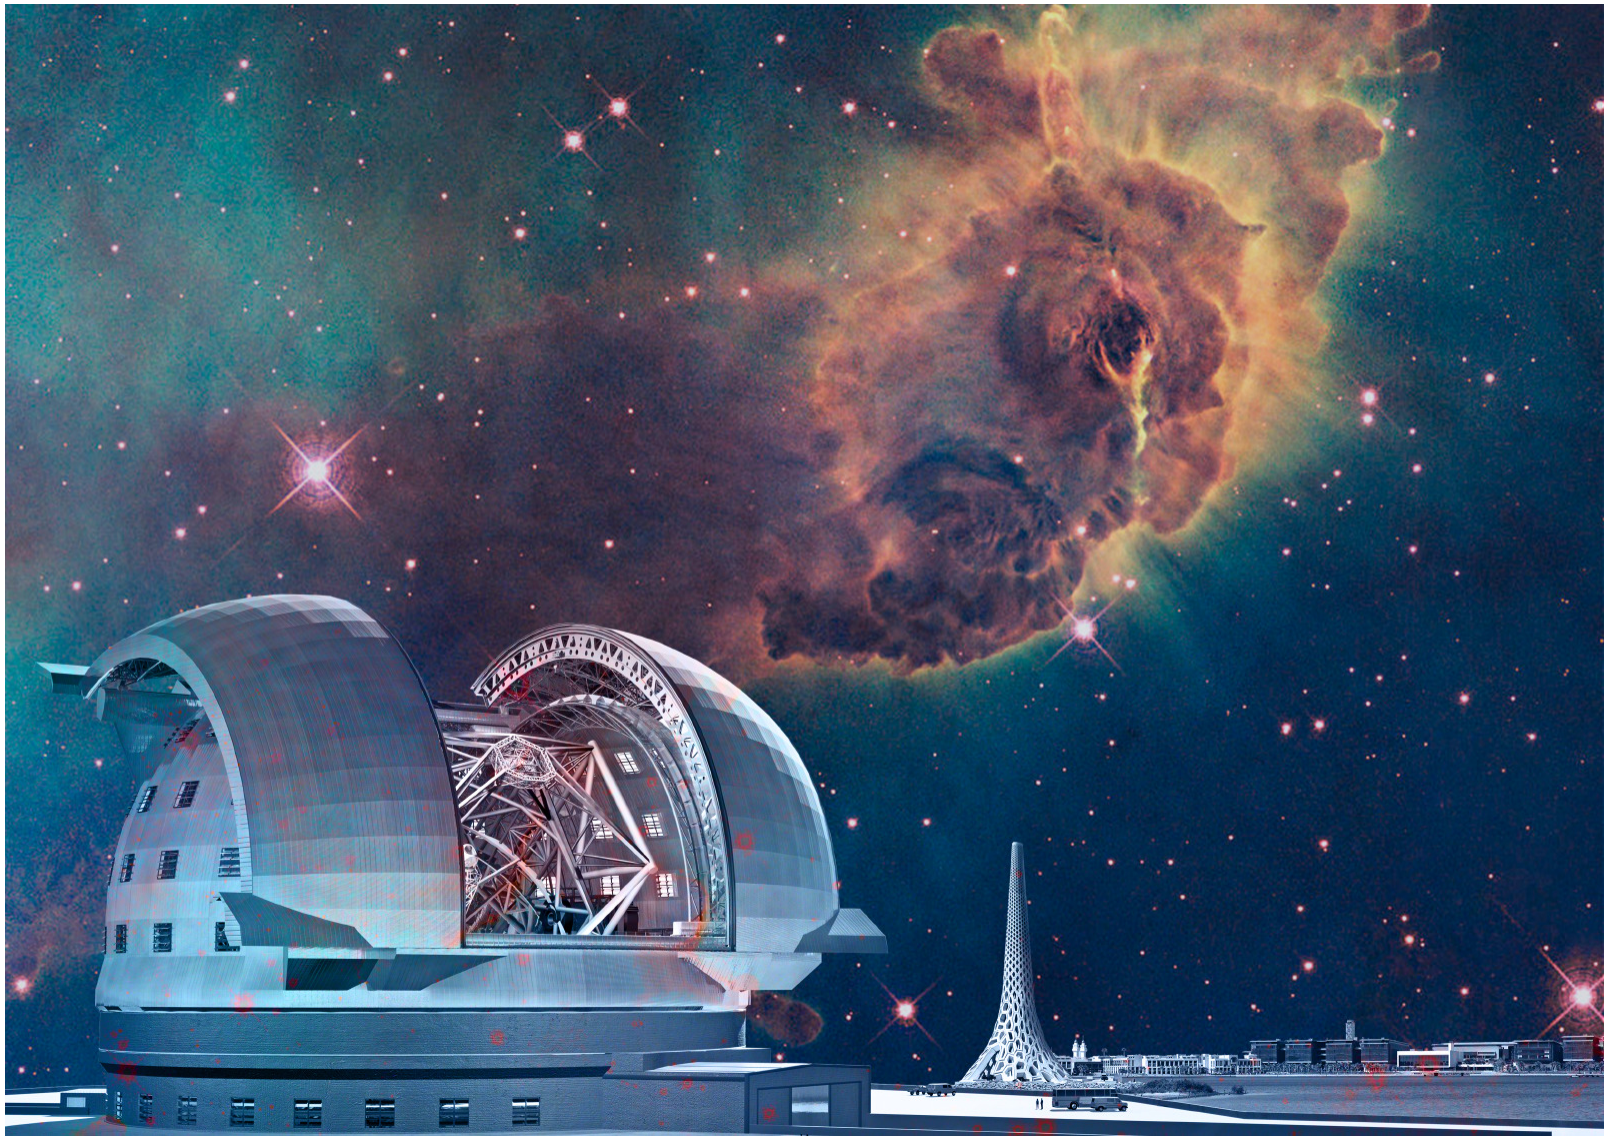

*Courtesy from CEMSE Communications, KAUST*

# The Hourglass Revisited

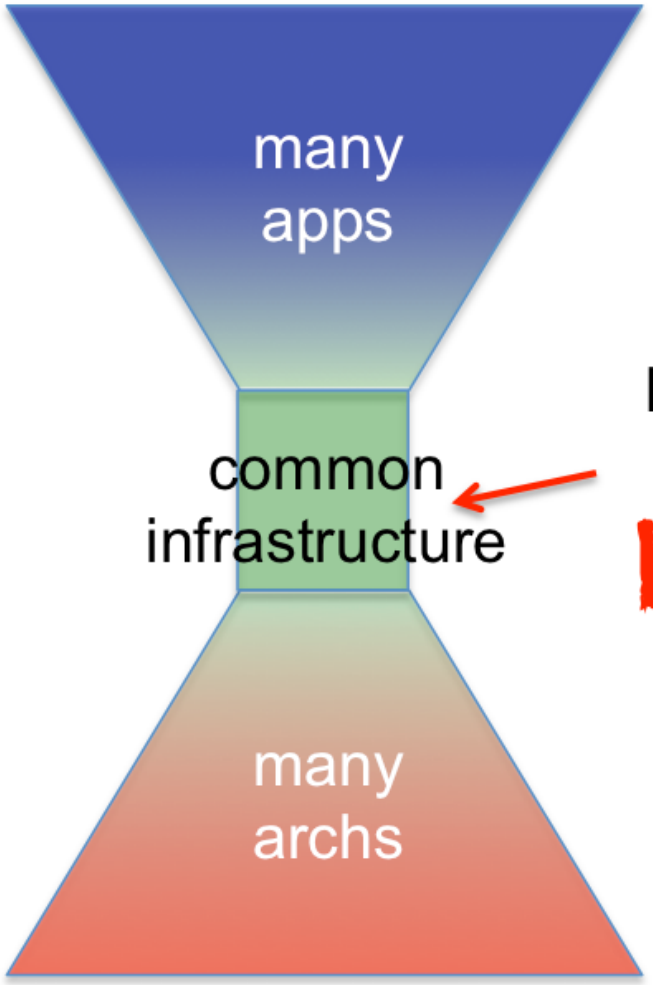


**ECRC is right**

here

@KAUST_ECRC

https://www.facebook.com/ecrckaust

# Questions?

Thank you ☺