THE COMEDY CLUB OF HIGH-PERFORMANCE COMPUTING: LOW-RANK MATRIX APPROXIMATION TAKES THE STAGE!

Hatem Ltaief
Principal Research Scientist
Extreme Computing Research Center

NHR PerfLab Online Seminar
May 16th 2023

Acknowledgments

- King Abdullah University of Science and Technology
  - Y. Hong, M. Ravasi, and D. Keyes
- Paris Observatory and Research School Of Astronomy & Astrophysics, Australian National University
  - J. Craney and D. Gratadour
- KAUST Supercomputing Lab
- Vendors
  - NVIDIA: S. Jones and I. Said
  - AMD: D. Cownie, M. Gontier, M. Jaghoub, and G. Markomanolis
  - Graphcore: P. Gepner and C. Goreczny
  - Cerebras: M. Jacquelin and L. Wilson

HPC Comedy Club

Intel #97

AMD #27

IBM/NVIDIA CPU/GPU #4

Fujitsu ARM A64FX #2

HPC Comedy Club

Intel #97

NVIDIA

intel

AMD

FUJITSU

NEC

AMD #27

IBM/NVIDIA CPU/GPU #4

Fujitsu ARM A64FX #2

HPC-AI Comedy Club

Intel #97

IBM/NVIDIA CPU/GPU #4

NVIDIA

intel

AMD

FUJITSU

NEC

GRAPHCORE

cerebras

AMD #27

Fujitsu ARM A64FX #2

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

SC22
Dallas, TX | hpc accelerates.

SHAHEEN III KEY FACTS

Shaheen III supercomputer with **25 HPE Cray EX supercomputer cabinets**

Expected to deliver over **100 Pflops/s**

20x faster than Shaheen II

4,608 CPU compute nodes, **AMD EPYC™ processors, "Genoa",** amounting to 884,736 cores in the entire system

2,800 NVIDIA Grace Hopper Superchips, tightly coupled CPU/GPU accelerators

Cray Slingshot interconnect

Cray ClusterStor E1000 with **additional 50 PB** of storage capacity

Operational by end of 2023

Accelerating research and developments in **energy, environment, food, water and healthcare**

2/3rds of KAUST faculty use computational modeling and simulation: **"to outcompute is to outcompete"**

Hardware Landscape

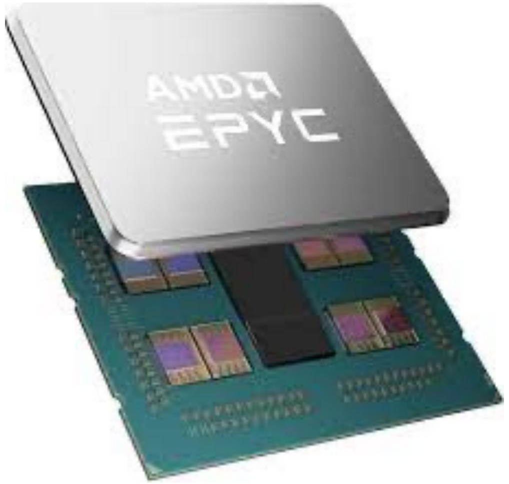

AMD Epyc Genoa

High cache capacity
High memory bandwidth
x86 programming env
Memory-bound workloads

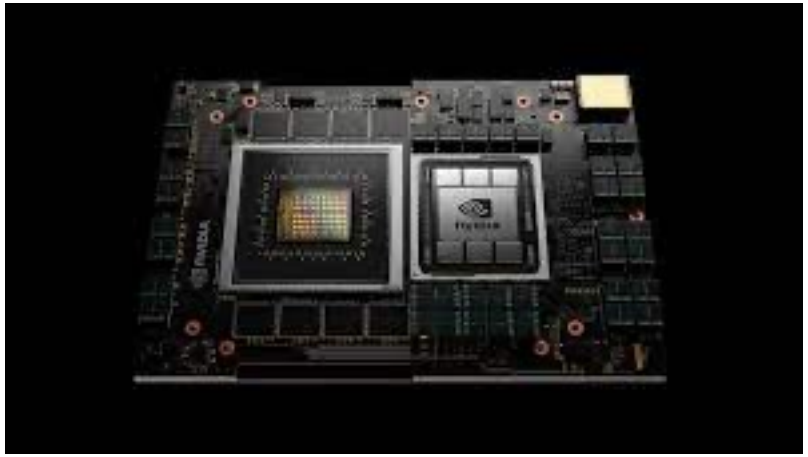

NVIDIA Grace Hopper

High speed CPU-GPU interconnect
Memory coherency
Support for mixed precisions
Compute-bound workloads

Hardware Landscape

AMD Epyc Genoa

High cache capacity
High memory bandwidth
x86 programming env
Memory-bound workloads

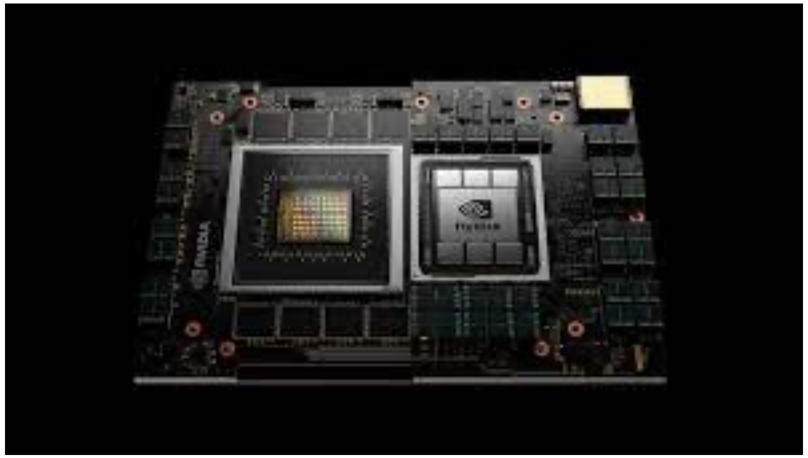

NVIDIA Grace Hopper

High speed CPU-GPU interconnect
Memory coherency
Support for mixed precisions
Compute-bound workloads

Cerebras CS-2
Graphcore IPU

AI-focused chip
Flat memory hierarchy
High SRAM bandwidth
Inference

Hardware Landscape

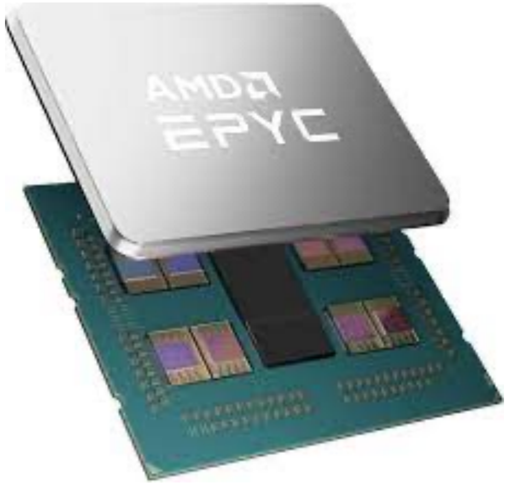

AMD Epyc Genoa

High cache capacity
High memory bandwidth
x86 programming env
Memory-bound workloads

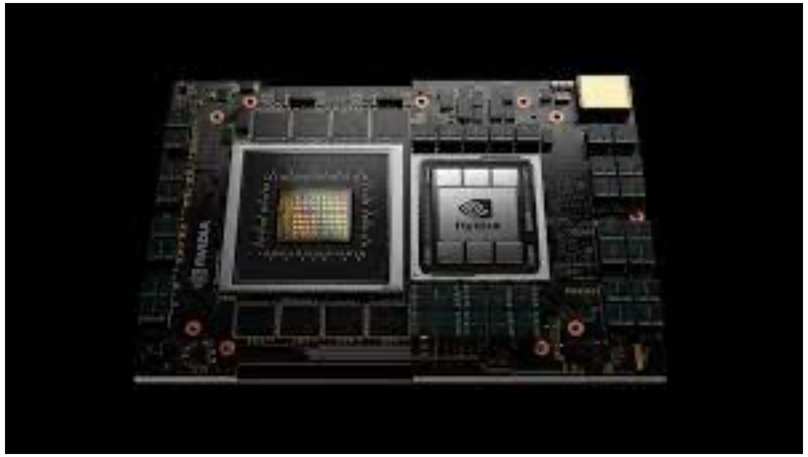

NVIDIA Grace Hopper

High speed CPU-GPU interconnect
Memory coherency
Support for mixed precisions
Compute-bound workloads

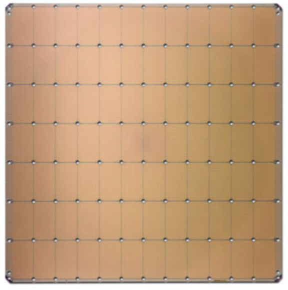

Cerebras CS-2
Graphcore IPU

AI-focused chip
Flat memory hierarchy
High SRAM bandwidth
Inference

How do we reconcile this hostile environment with HPC scientific applications?

Revisiting the Hourglass

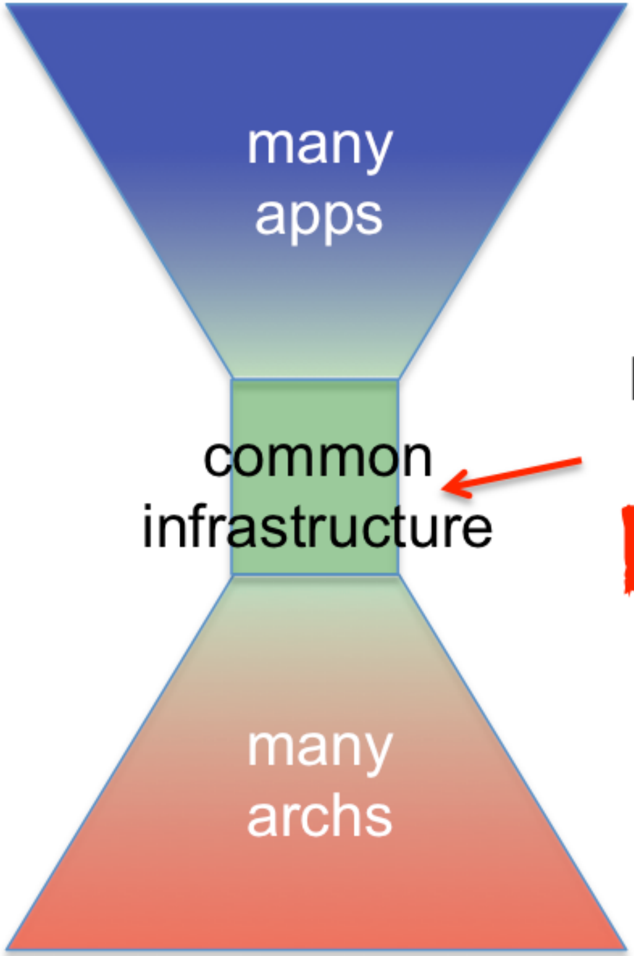

$$A\,x = b$$

ECRC is right

here

@KAUST_ECRC

https://www.facebook.com/ecrckaust

Reshaping Linear Algebra for Massively Parallel Architectures

- Enhance user-productivity using layers of abstraction
- Expose parallelism using fine-grained computations
- Achieve scalability using asynchronous executions
- Exploit data sparsity using low-rank approximations
- Maintain code portability using standard basic blocks

Are you willing to redesign your algorithm?

One possible productive solution: **Matricization**

HPC Scientific Applications

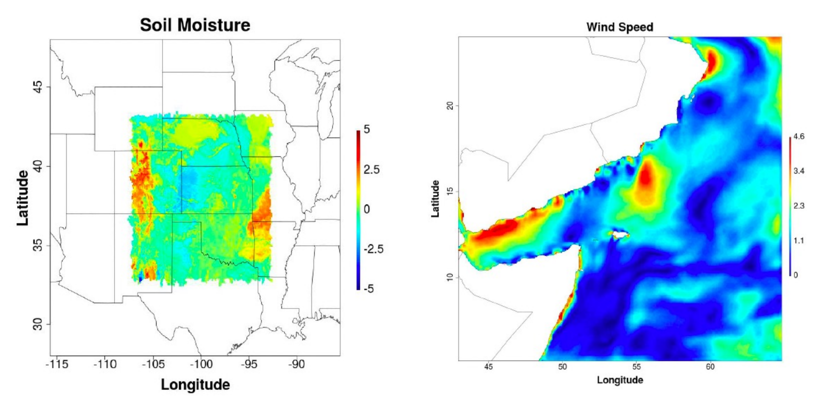

Fig. 4: **Left:** Soil moisture residuals at the topsoil of the Mississippi River basin. **Right**: Wind speed (m/s) in the Arabian Sea.

3D Geospatial Statistics

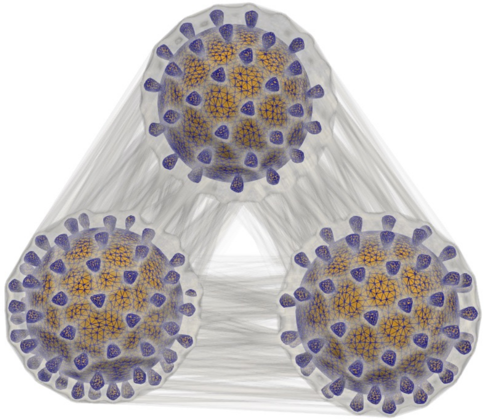

3D Mesh Deformations

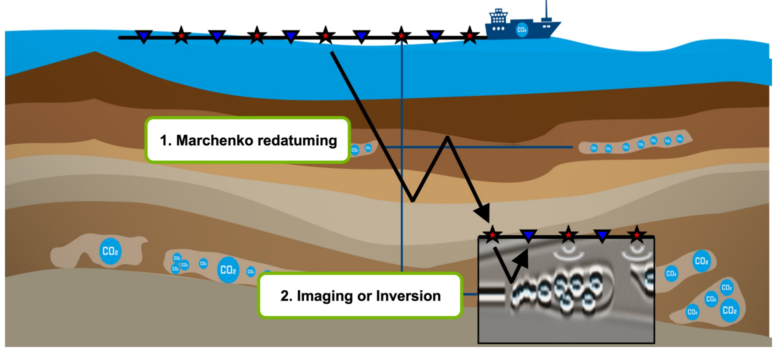

Seismic Imaging

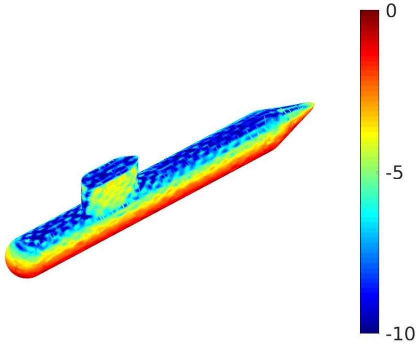

3D Computational Electromagnetics

Wireless Communications

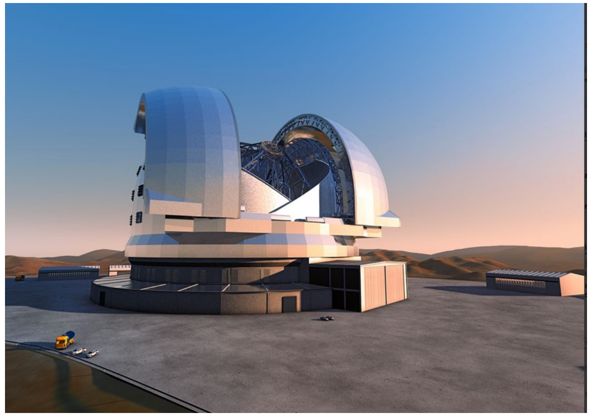

Computational Astronomy

HPC Scientific Applications

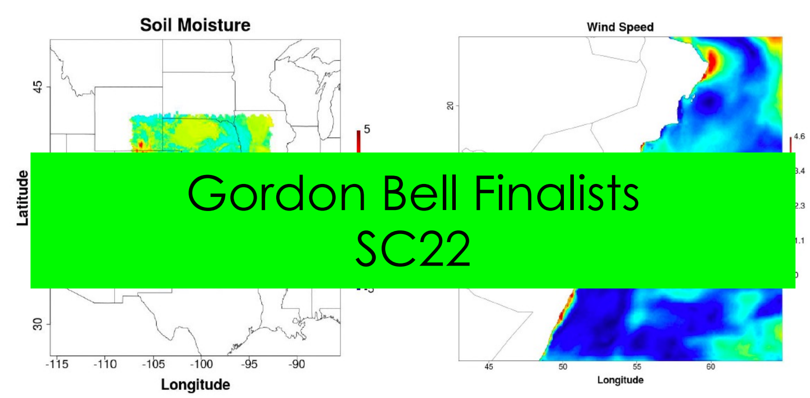

Gordon Bell Finalists
SC22

Fig. 4: **Left:** Soil moisture residuals at the topsoil of the Mississippi River basin. **Right**: Wind speed (m/s) in the Arabian Sea.

3D Geospatial Statistics

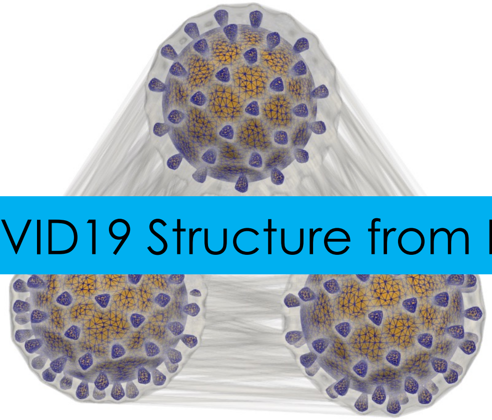

COVID19 Structure from PDB

3D Mesh Deformations

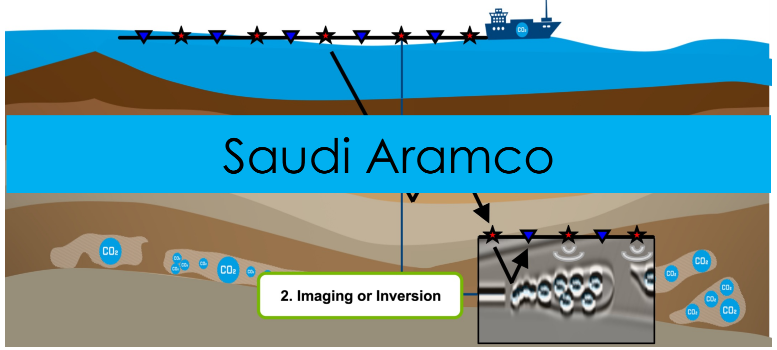

Saudi Aramco

Seismic Imaging

GAUSS Award

3D Computational Electromagnetics

Huawei

Wireless Communications

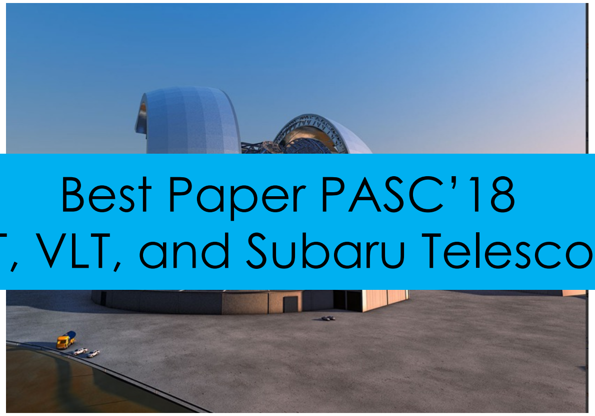

Best Paper PASC'18
E-ELT, VLT, and Subaru Telescopes

Computational Astronomy

HPC Scientific Applications

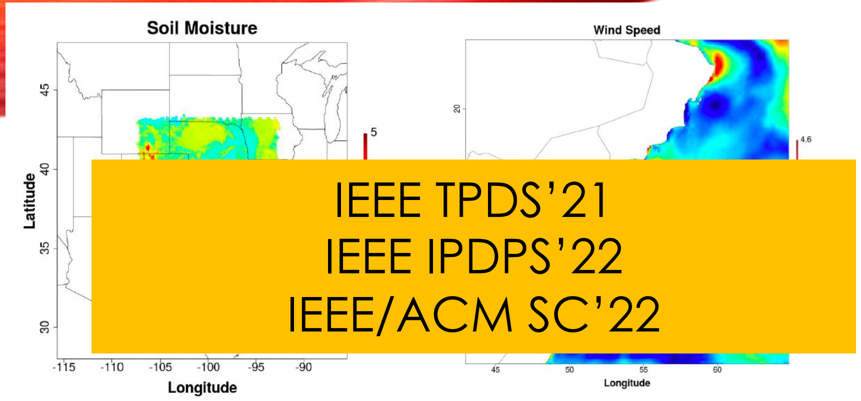

IEEE TPDS'21
IEEE IPDPS'22
IEEE/ACM SC'22

Fig. 4: **Left:** Soil moisture residuals at the topsoil of the Mississippi River basin. **Right**: Wind speed (m/s) in the Arabian Sea.

3D Geospatial Statistics

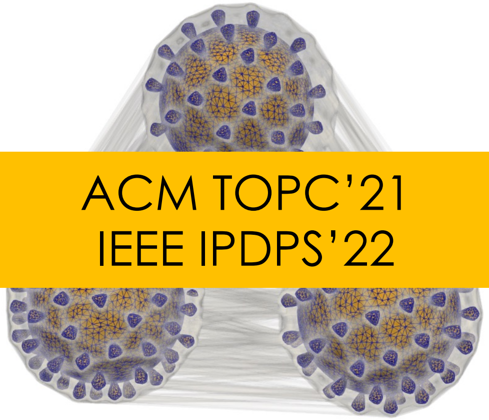

ACM TOPC'21
IEEE IPDPS'22

3D Mesh Deformations

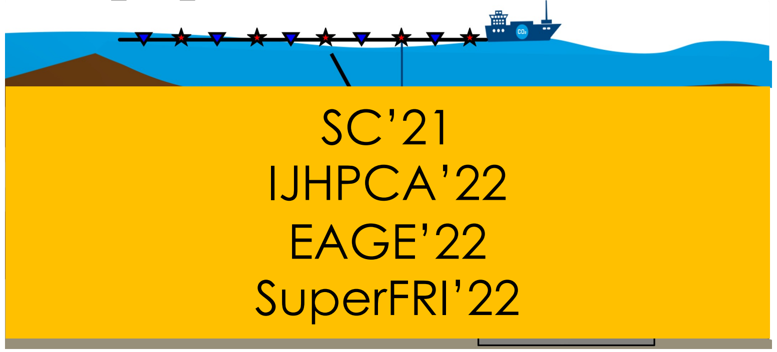

SC'21
IJHPCA'22
EAGE'22
SuperFRI'22

Seismic Imaging

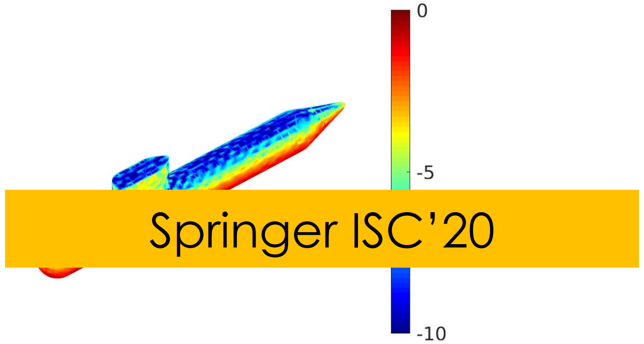

Springer ISC'20

3D Computational Electromagnetics

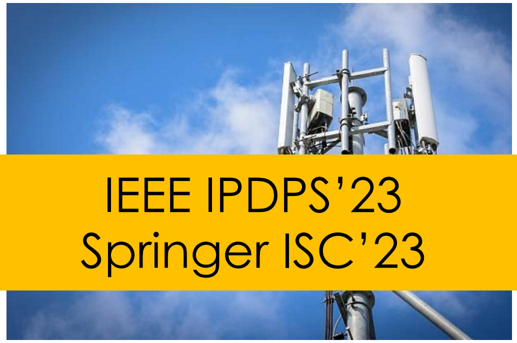

IEEE IPDPS'23
Springer ISC'23

Wireless Communications

IEEE IPDPS'18
ACM PASC'18
IEEE/ACM SC'21

Computational Astronomy

HPC Scientific Applications

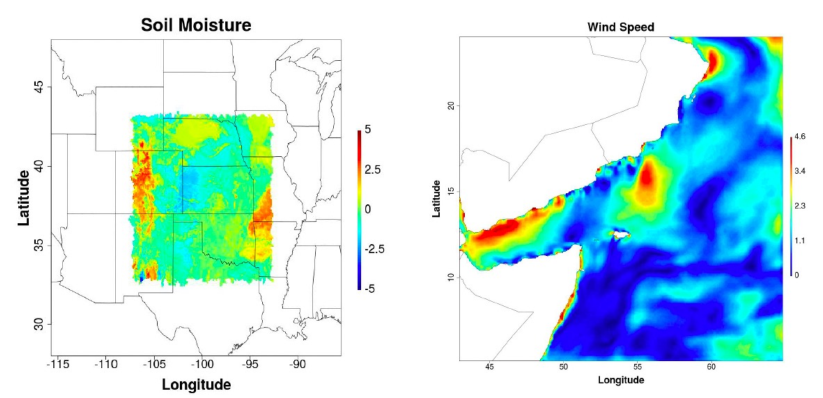

Fig. 4: **Left:** Soil moisture residuals at the topsoil of the Mississippi River basin. **Right**: Wind speed (m/s) in the Arabian Sea.

3D Geospatial Statistics

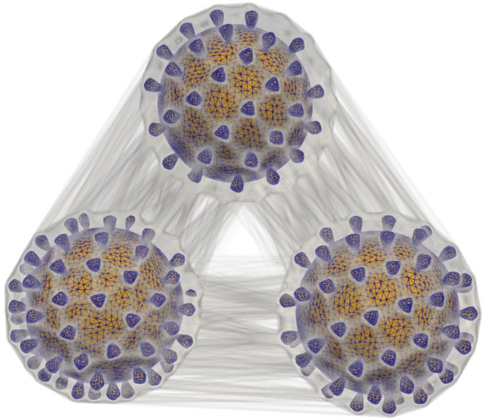

3D Mesh Deformations

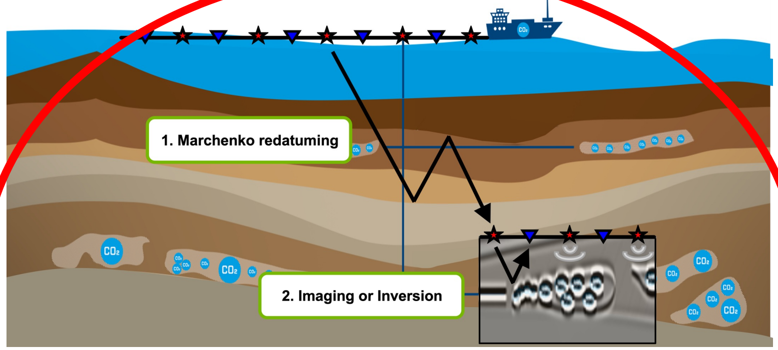

Seismic Imaging

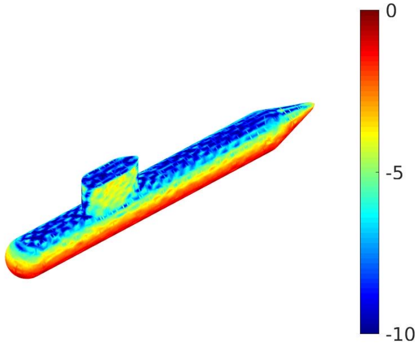

3D Computational Electromagnetics

Wireless Communications

Computational Astronomy

HPC Scientific Applications

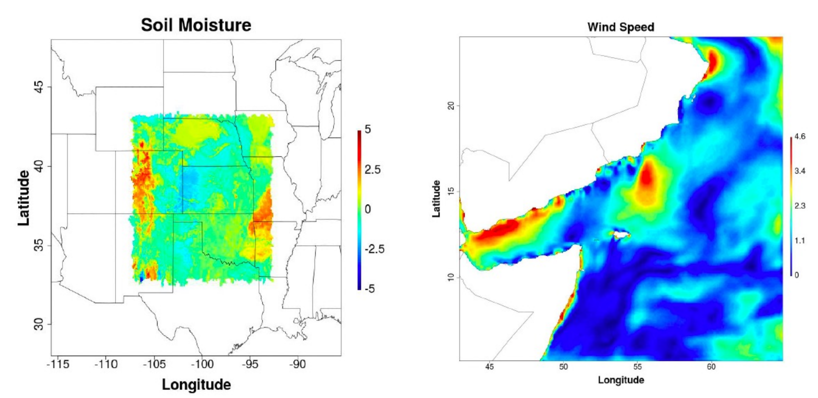

Fig. 4: **Left:** Soil moisture residuals at the topsoil of the Mississippi River basin. **Right**: Wind speed (m/s) in the Arabian Sea.

3D Geospatial Statistics

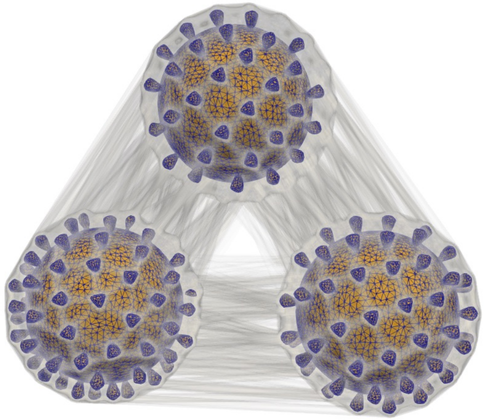

3D Mesh Deformations

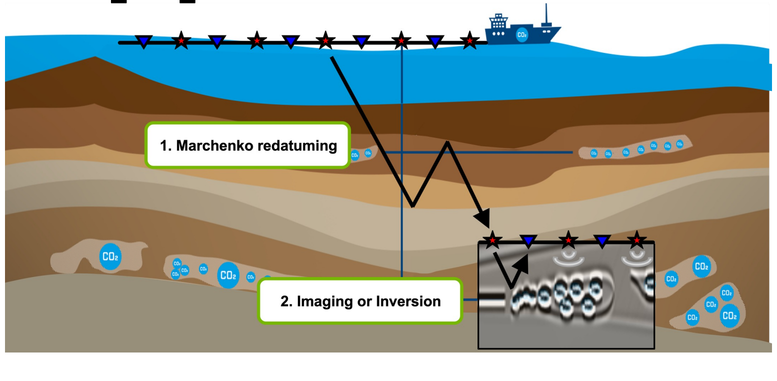

Seismic Imaging

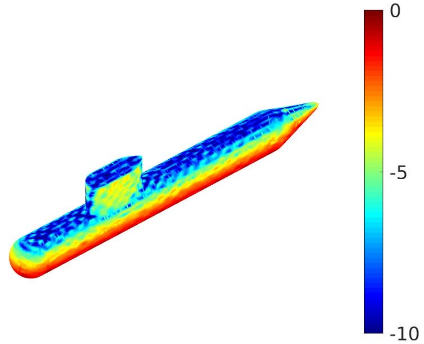

3D Computational Electromagnetics

Wireless Communications

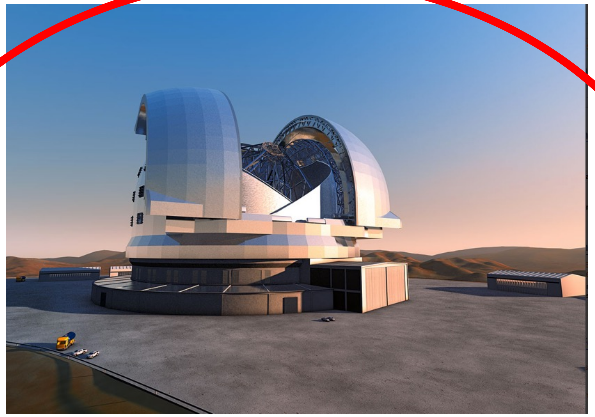

Computational Astronomy

Supporting Several Major Ground-Based Telescopes

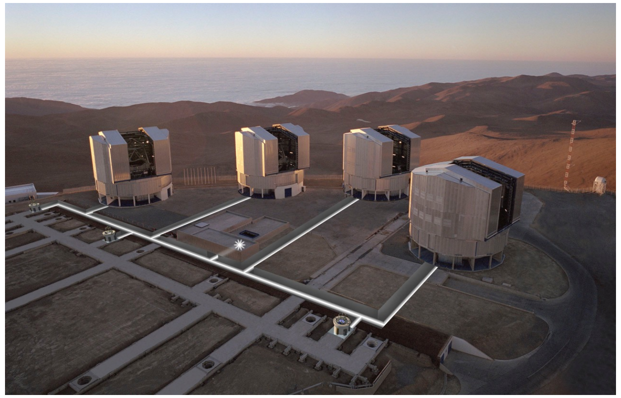

THE VERY LARGE TELESCOPE

THE SUBARU TELESCOPE

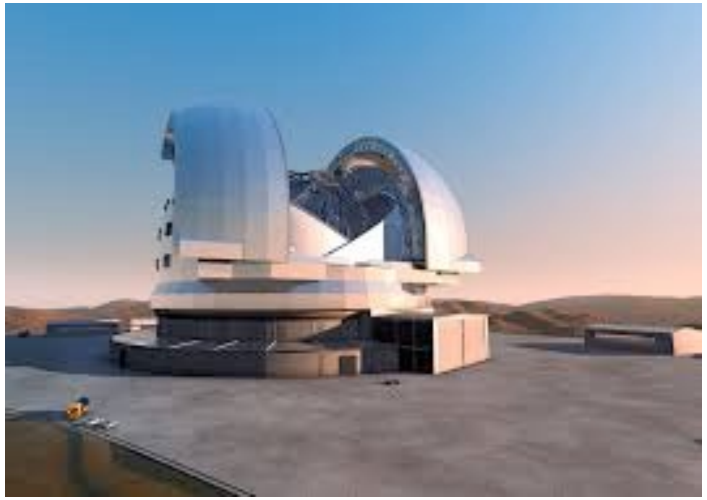

THE EUROPEAN EXTREMELY LARGE TELESCOPE

Supporting Several Major Ground-Based Telescopes

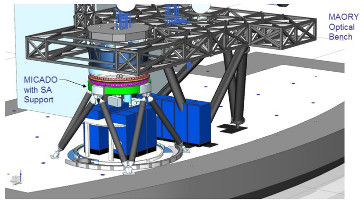

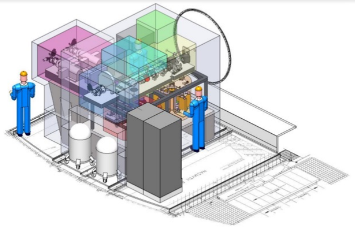

MAVIS
MCAO ASSISTED VISIBLE
IMAGER AND SPECTROGRAPH

THE VERY LARGE TELESCOPE

SCEXAO
THE SUBARU
CORONOGRAPHIC
EXTREME AO

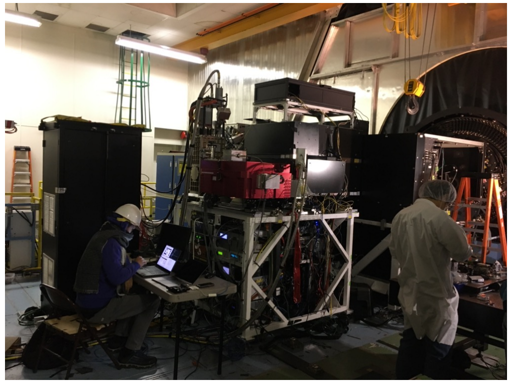

THE SUBARU TELESCOPE

EPICS
EXOPLANET IMAGING CAMERA
AND SPECTROGRAPH

MAORY
MULTI-CONJUGATE
ADAPTIVE OPTICS RELAY

MICADO
MULTI-ADAPTIVE OPTICS IMAGING
CAMERA FOR DEEP OBSERVATIONS

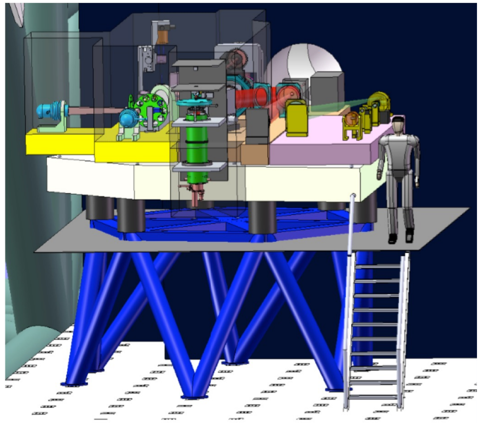

THE EUROPEAN EXTREMELY LARGE TELESCOPE

The Atmospheric Turbulence

The sun observed with a compact camera

- Disturbs the trajectory of light rays
- Reduces astronomical images quality

Adaptive Optics for Giant Telescopes

Control in real-time the shape of the incoming wavefront

- **Sensors are cameras** equipped with an optical device (lenslet array, pyramidal prism, etc…)
- **Deformable mirrors** to compensate for wavefront distortions
- **Typical rate of operation is 1kHz**
- Compute **pipeline latency below 1 millisecond**
- **Stable time-to-solution** is critical to ensure stable operations (jitter of the order of 10s of µs)
- **Matrix-Vector Multiplication (MVM)**

is the most critical computational kernel

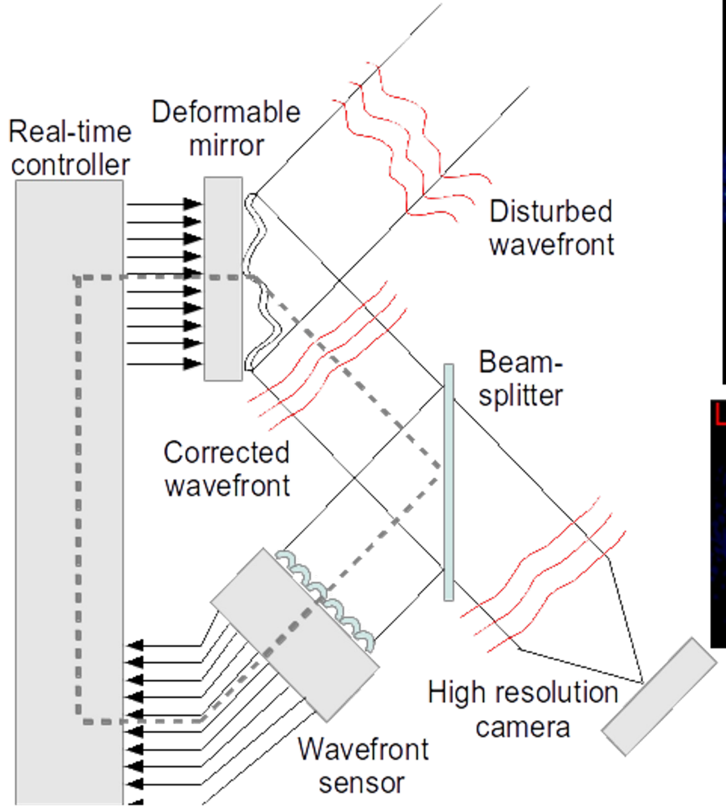

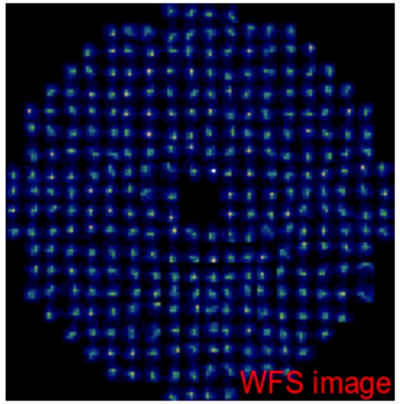

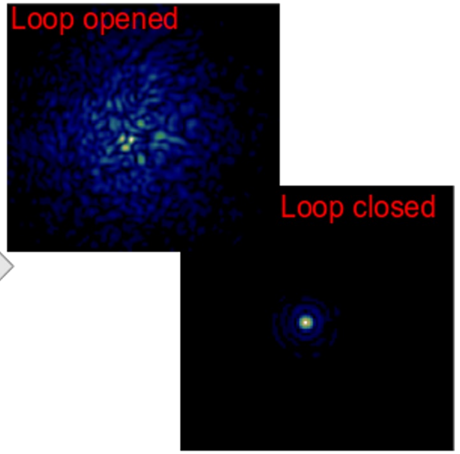

One of The Key Hardware Components: The Deformable Mirrors

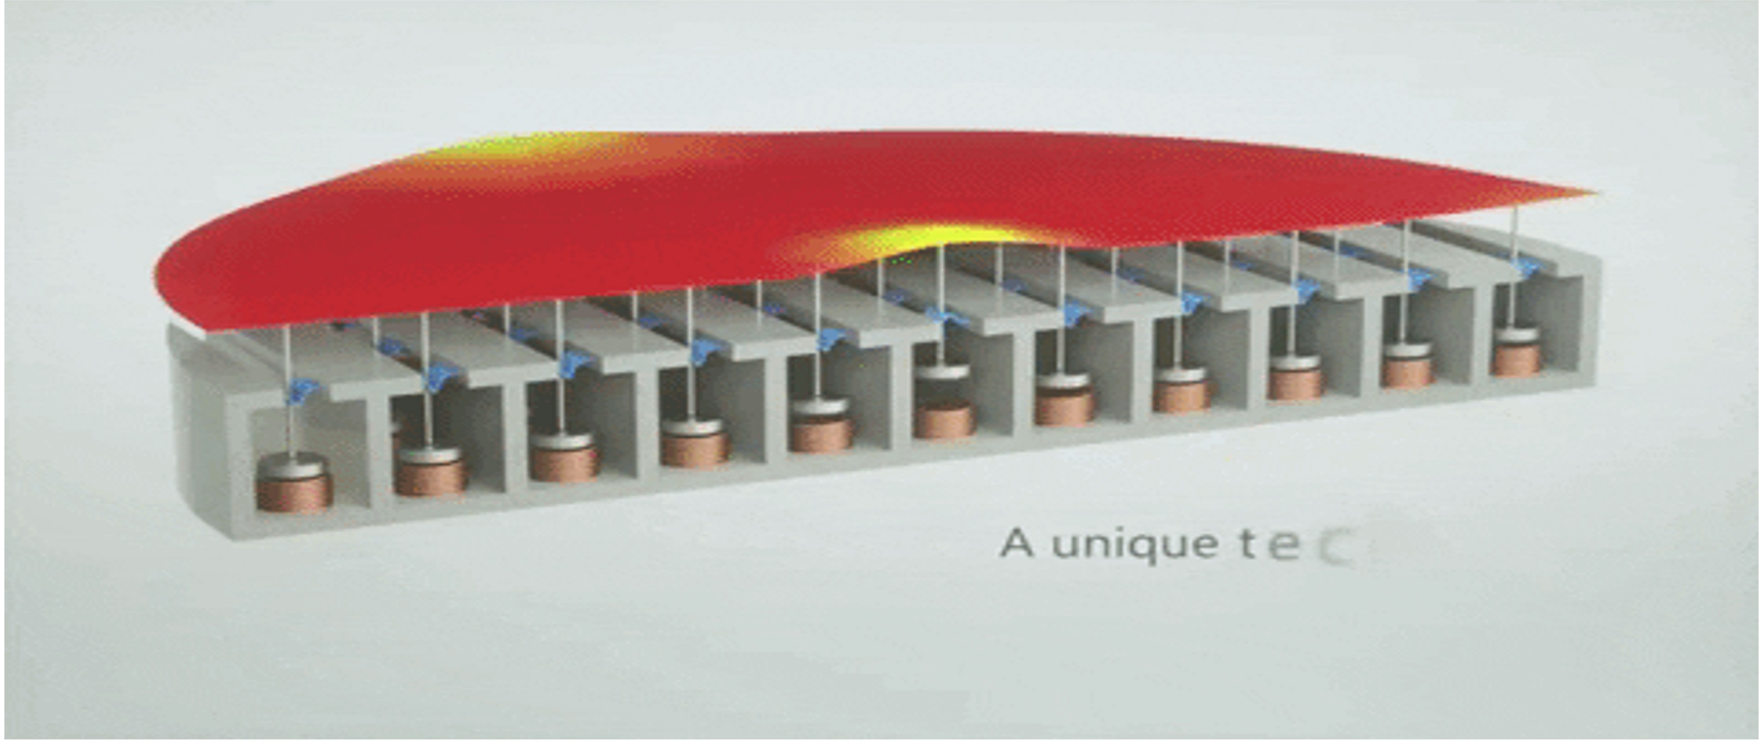

MAVIS in a nutshell

Flagship 3rd-generation instrument for the Very Large Telescope

- Imager & Spectrograph fed by **Multi-Conjugate AO**

- **Scaling up** the whole AO concept:
    - More actuators
    - Faster control system
    - Exquisite calibrations

- **Fast-track** project:
    - First-light by 2026

- **Deeper & Sharper** than any space-based instrument

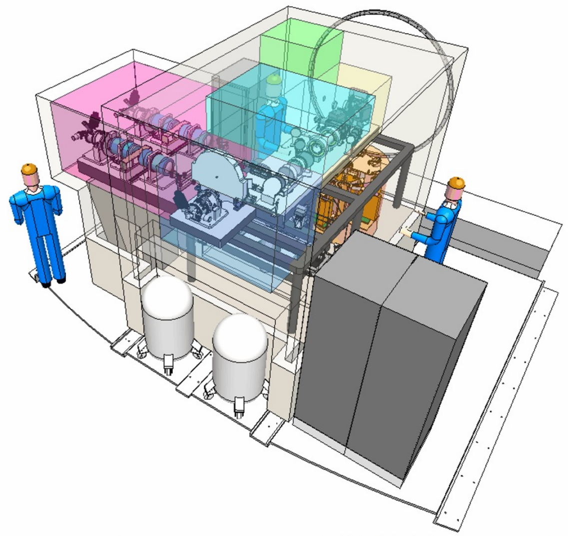

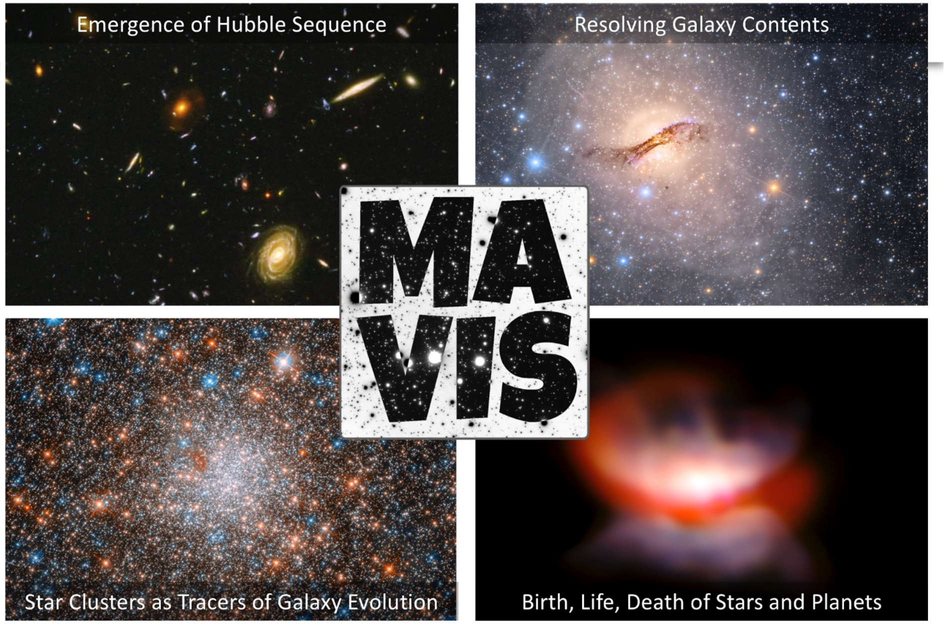

Rank Analysis of the Tomographic Reconstructor

Splitting the matrix into tiles and looking at ranks

- Tiles size aligned with system parameters
- Data sparse, opportunity for low-rank matrix approximations

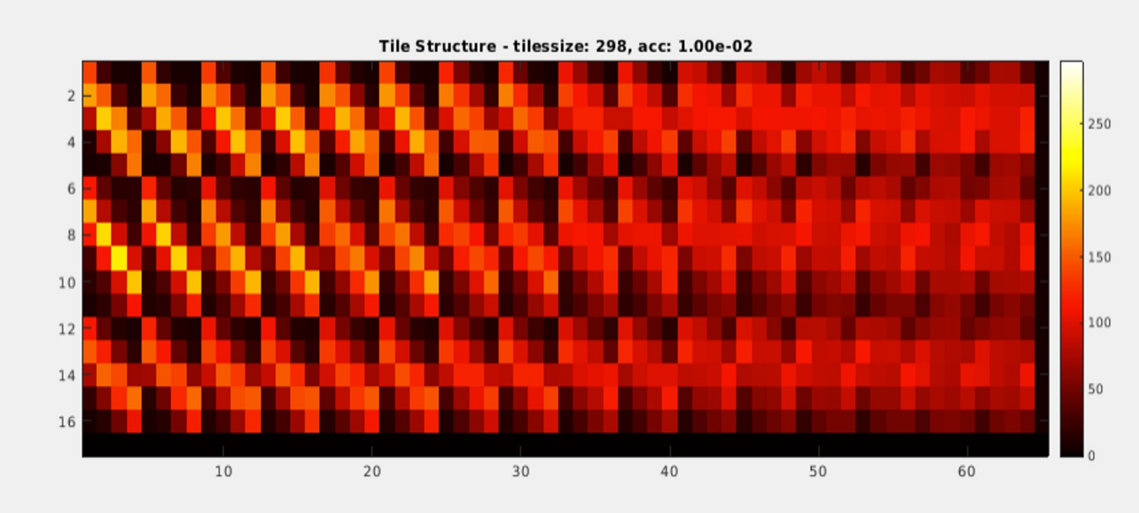

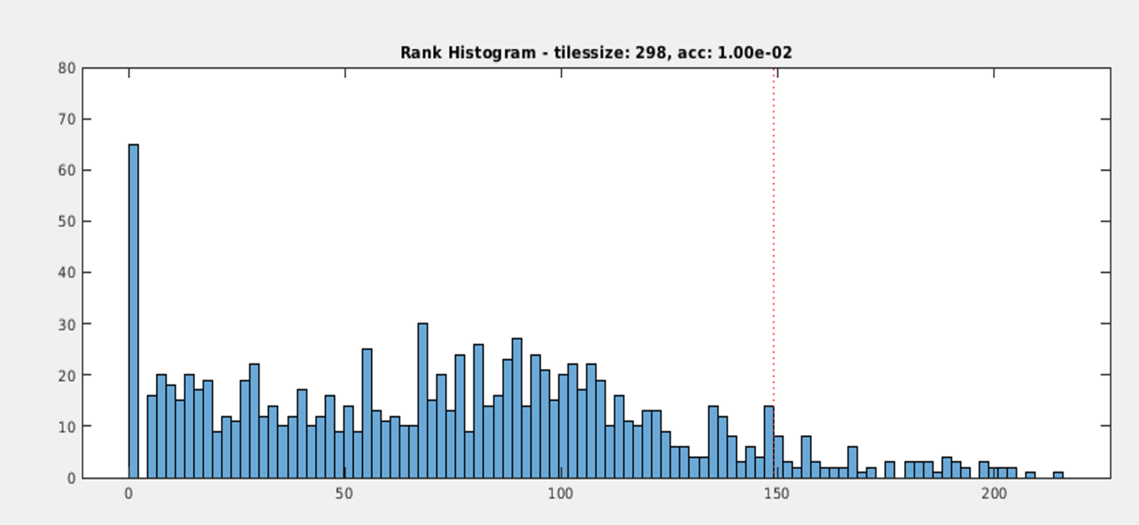

TLR-MVM

Tile Dense
Matrix-Vector Multiplication

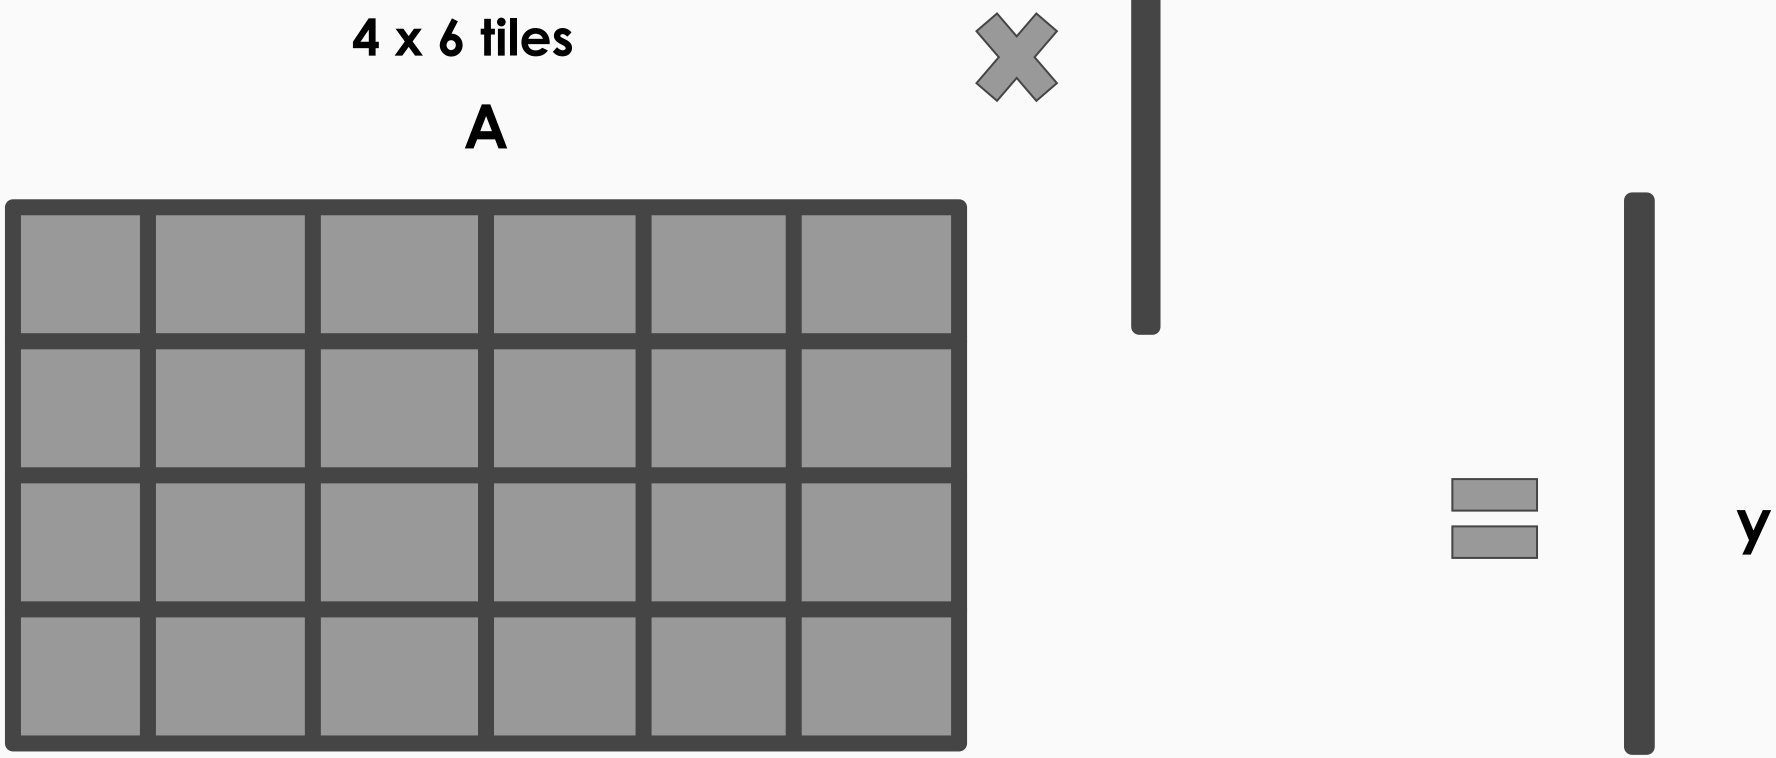

TLR-MVM

1) Compress once up-front (SVD-like algorithms)

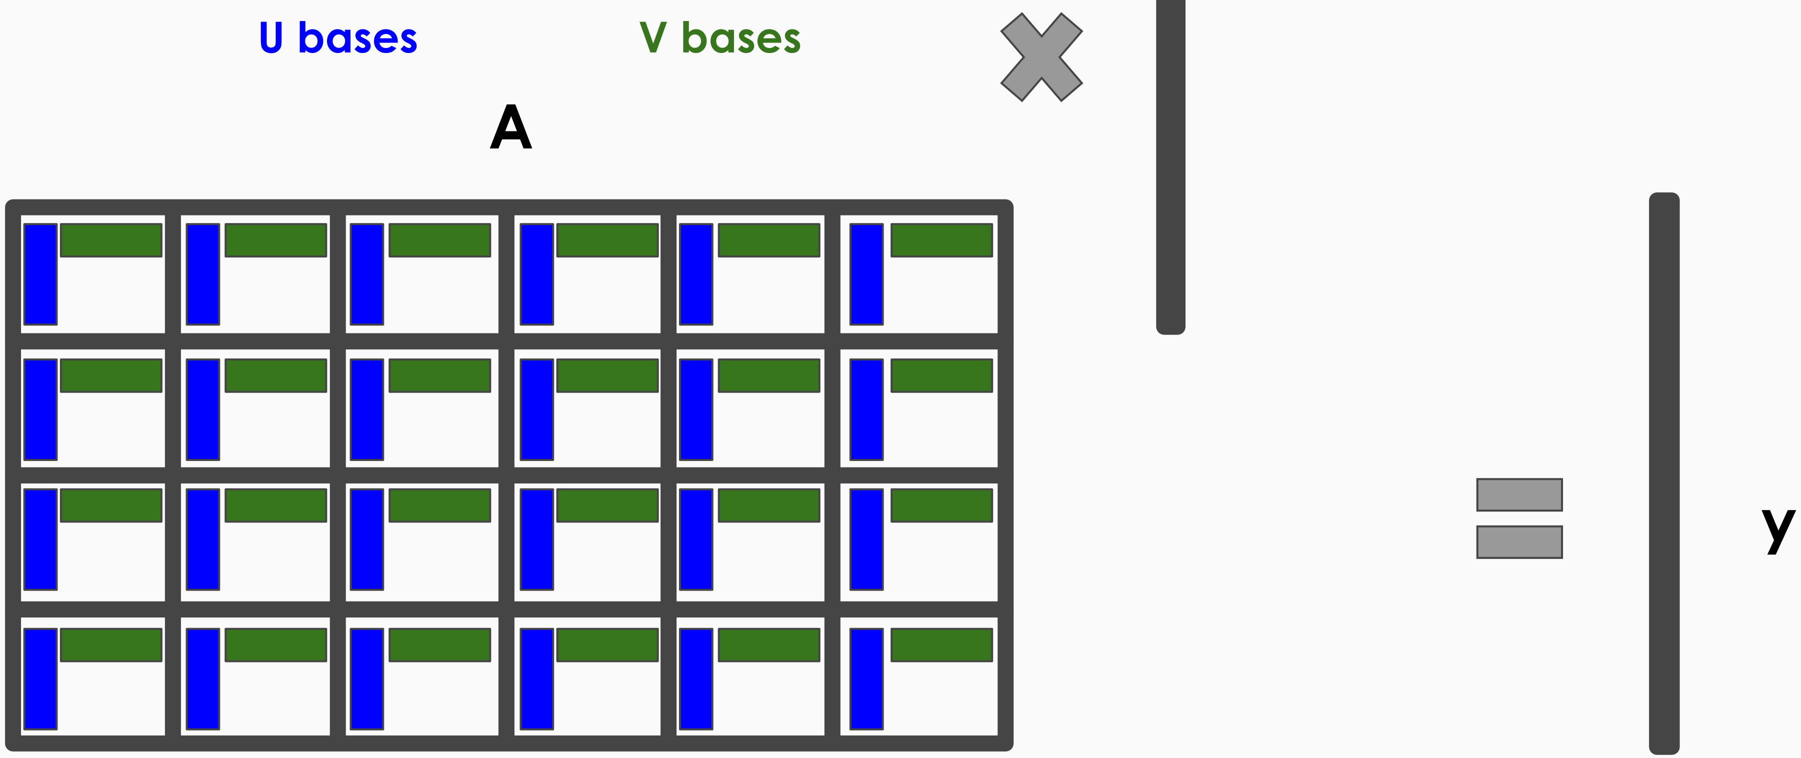

TLR-MVM

2) Stack the bases

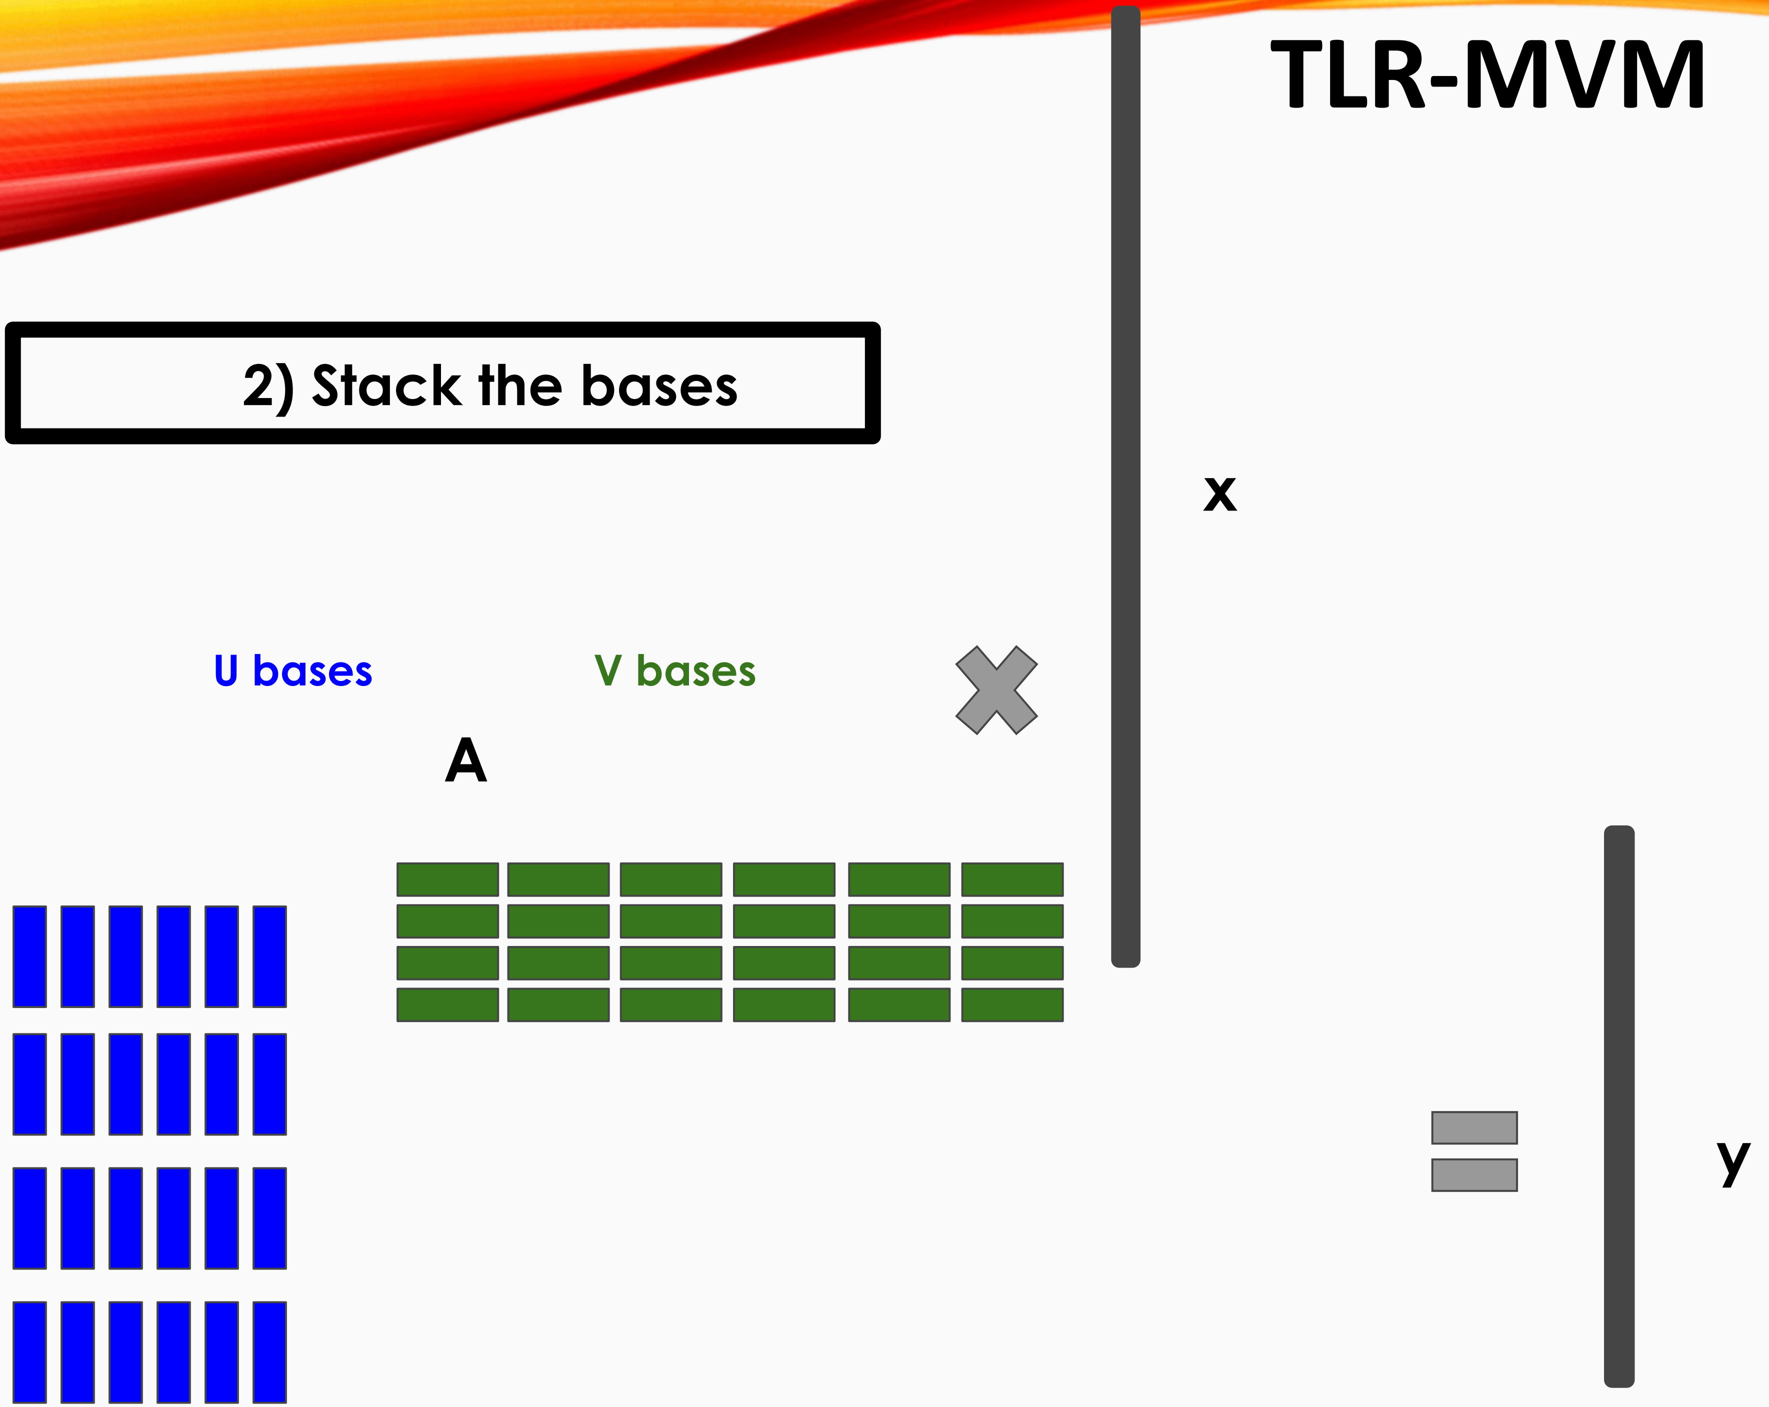

TLR-MVM

3) Calculate (per red part):
Yv = V . x

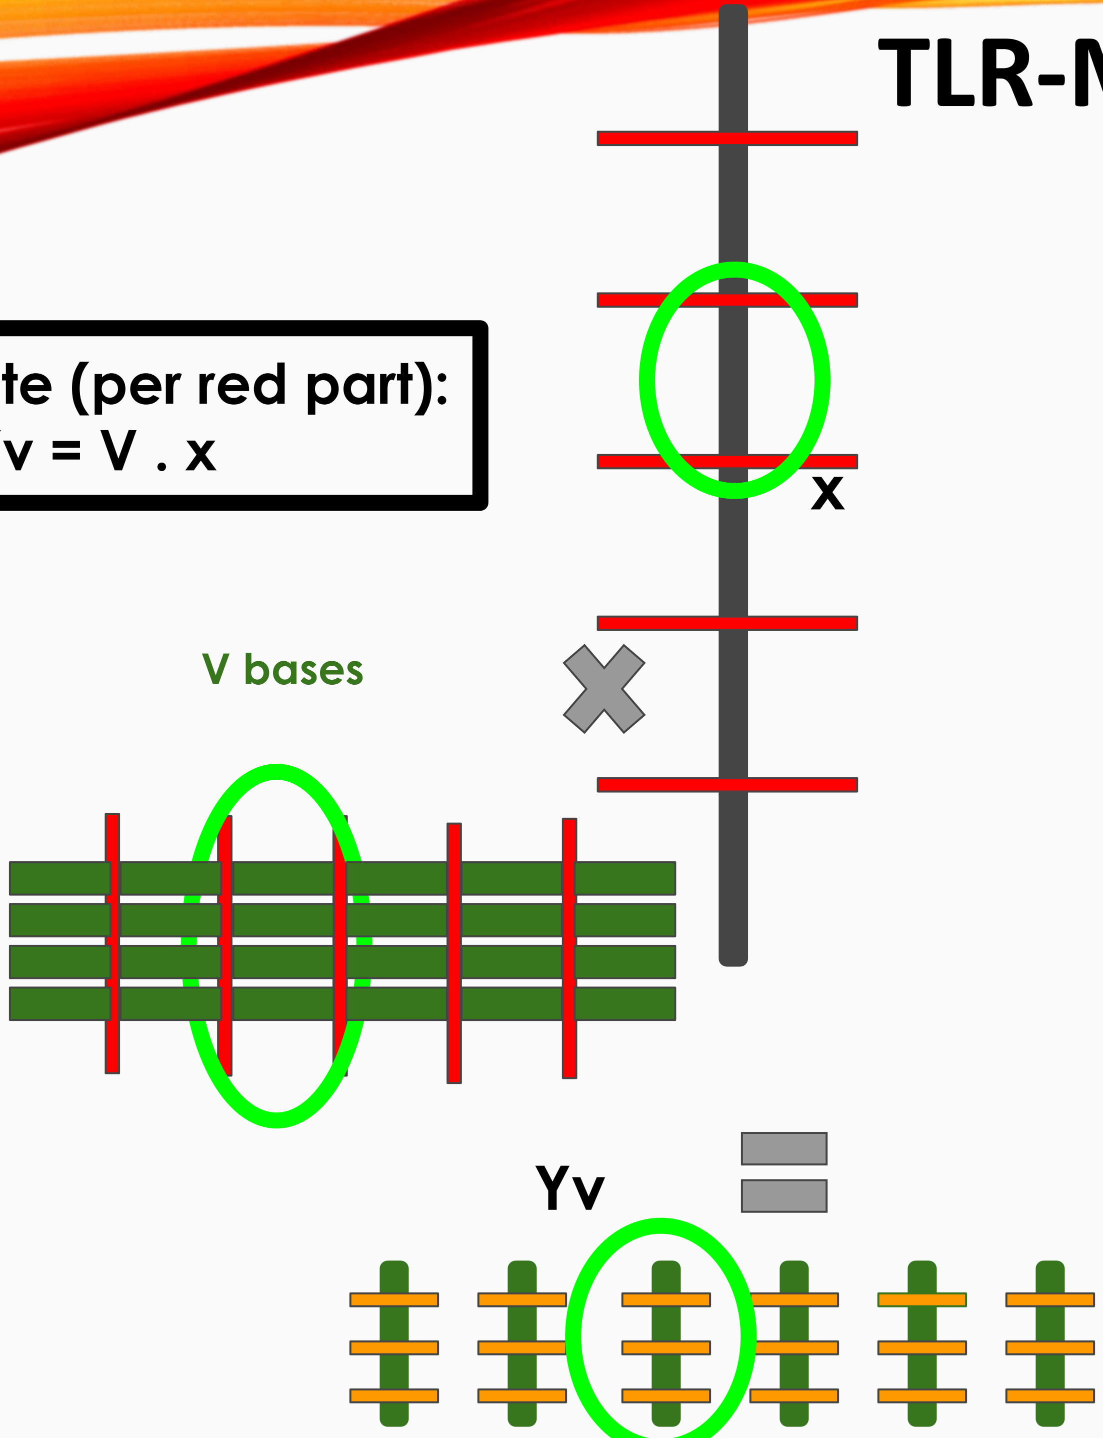

TLR-MVM

4) Translate
Yv (V bases) to Yu (U bases)

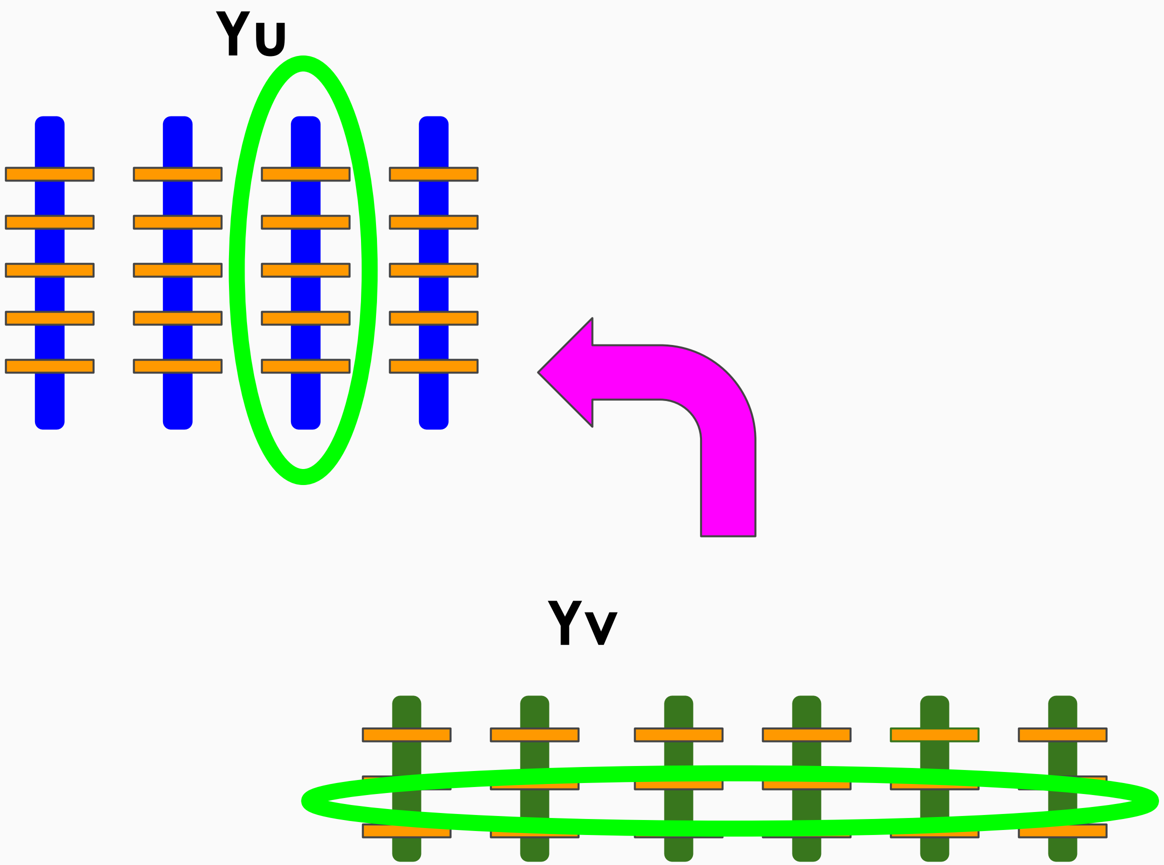

TLR-MVM

5) Calculate
y = U . Yu

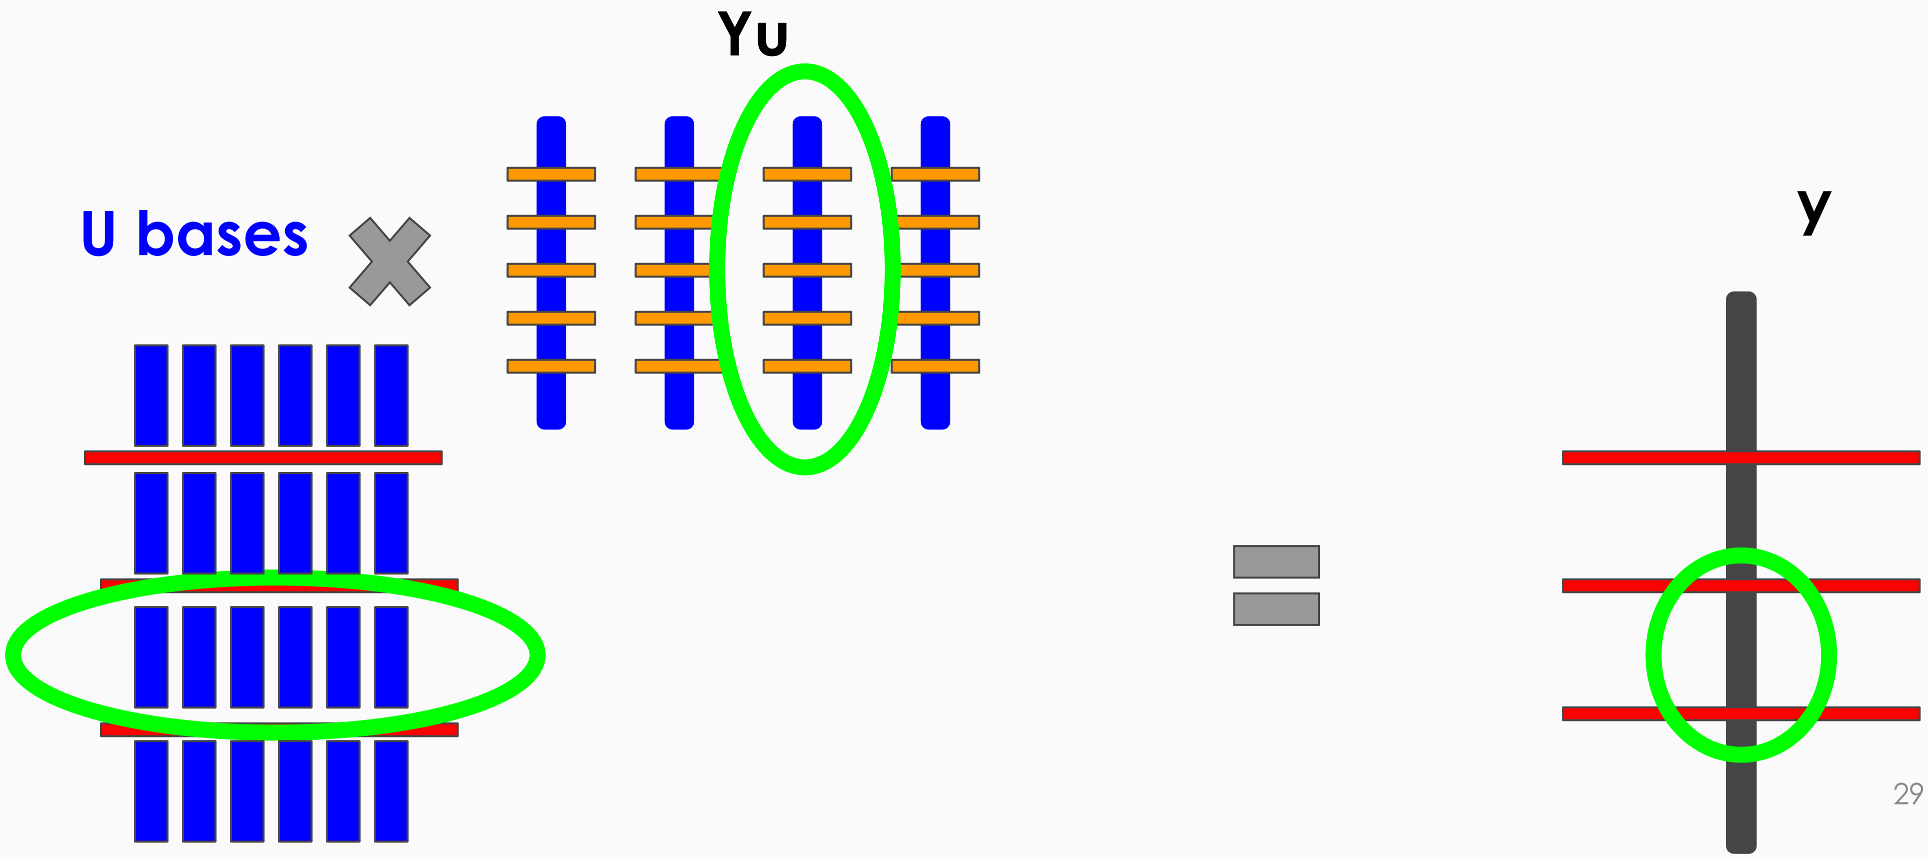

TLR-MVM

Rely on batch GEMV calls
w/ variable sizes

5) Calculate
$y = U . Yu$

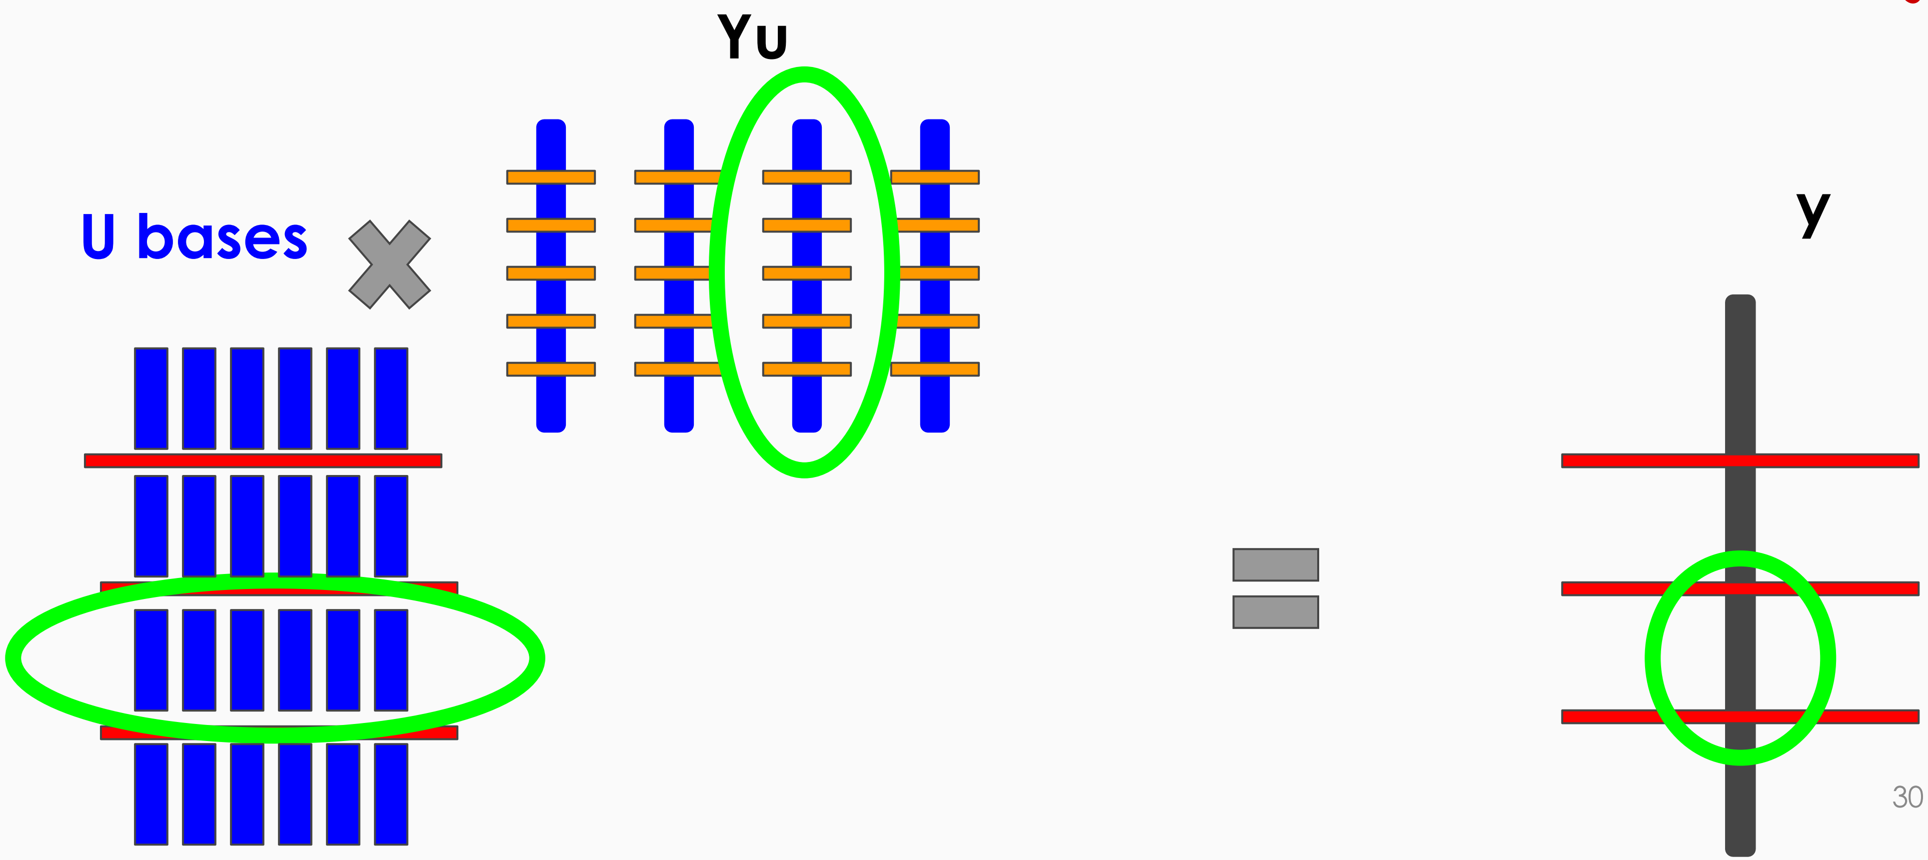

Numerical Accuracy Assessment on MAVIS Datasets

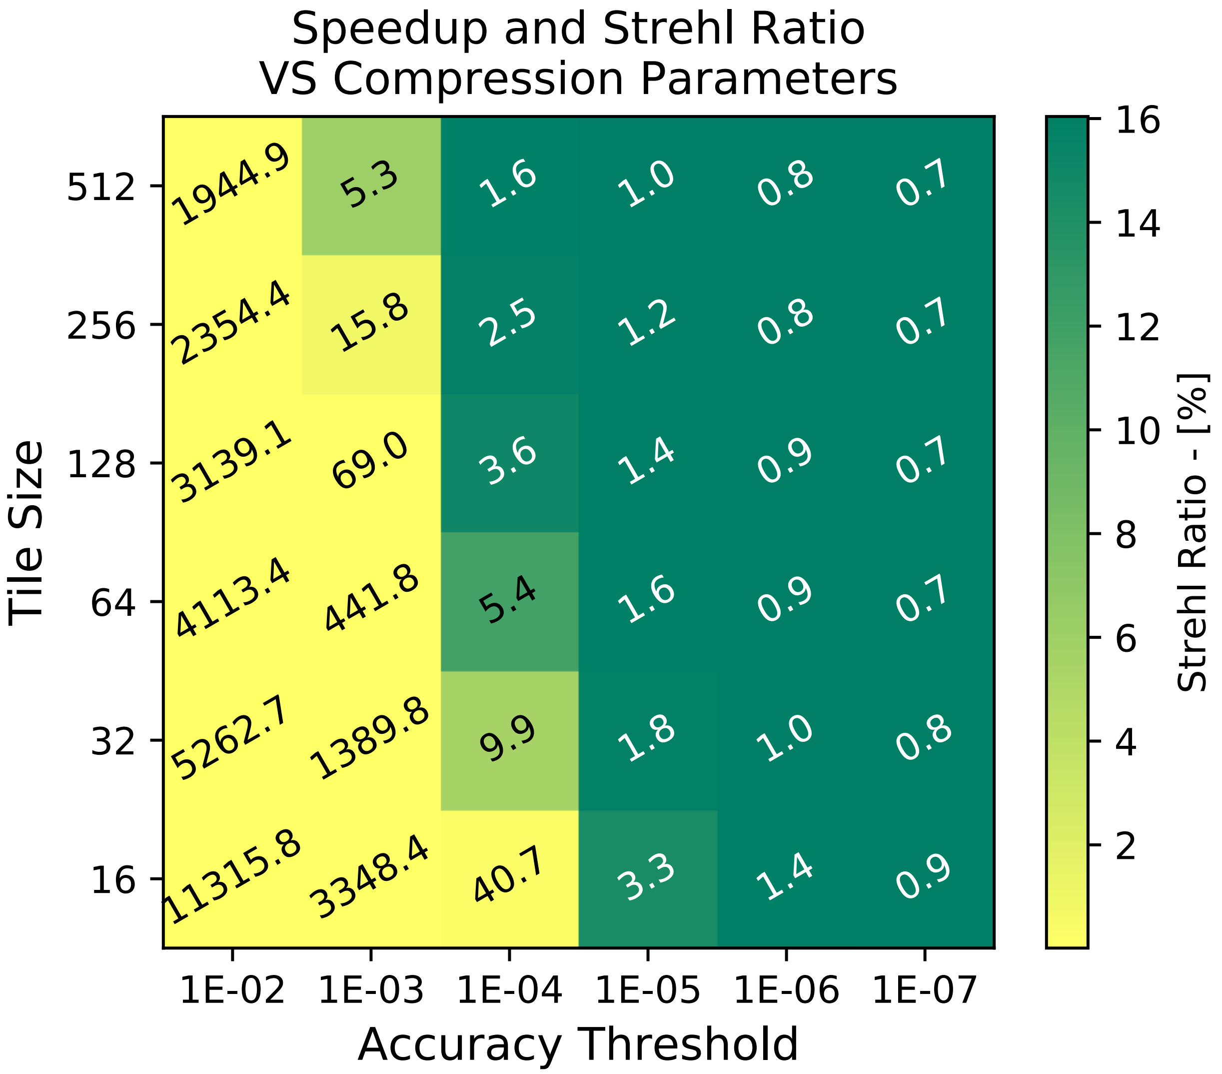

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

x86

MPI + OpenMP

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

ARM

MPI +

OpenMP

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

Accelerators

ROCm / CUDA

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

Vector

MPI +

OpenMP

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

HBM

Hardware/Software Specifications

Vendor	Intel	AMD		Fujitsu	NVIDIA	NEC
Family	Cascade Lake	EPYC Rome	Instinct	Primergy A64FX	Ampere GPU	SX-Aurora TSUBASA
Model	6248	7702	MI100	FX1000	A100	B300-8
Node(s)/Card(s)	1	1	1	16	1	8
Socket(s)	2	2	N/A	4	N/A	N/A
Cores	40	128	7680	48	6912	8
GHz	2.5	2.2	1.5	2.2	2.6	1.6
Memory Sustained BW	384GB DDR4 232GB/s	512GB DDR4 330GB/s	32GB HBM2 1.2TB/s	32GB HBM2 800GB/s	40GB HBM2e 1.5TB/s	48GB HBM2 1.5TB/s
LLC Sustained BW	27.5MB 1.1TB/s	**512MB** 4TB/s	8MB 3TB/s	32MB 3.6TB/s	40MB 4.8TB/s	16MB 2.1TB/s
Compiler	Intel compiler 19.1.0	GCC compiler 8.2.0		Fujitsu compiler 4.5.0	NVCC 11.0	NEC compiler 3.1.1
BLAS library	Intel MKL 2020	BLIS 3.0.0		Fujitsu SSL II	cuBLAS 11.0	NEC NLC 2.1.0
MPI library	OpenMPI 4.0.3	OpenMPI 3.1.2		Fujitsu MPI 4.0.1	NCCL 2.0	NEC MPI 2.13.0

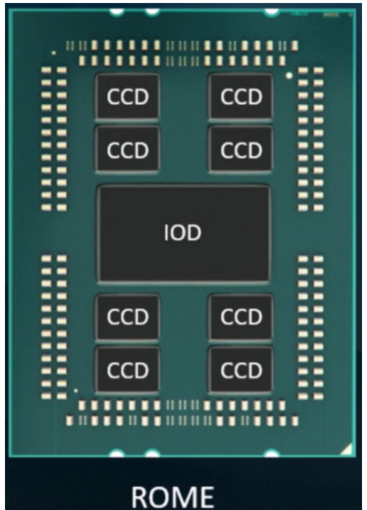

- IOD = “I/O Die” doing all the memory/PCI/other socket traffic
- CCD = “Core Compute Die”, a chiplet having compute cores only
- CCX = “Core Compute CompleX”, a set of cores sharing a L3 cache

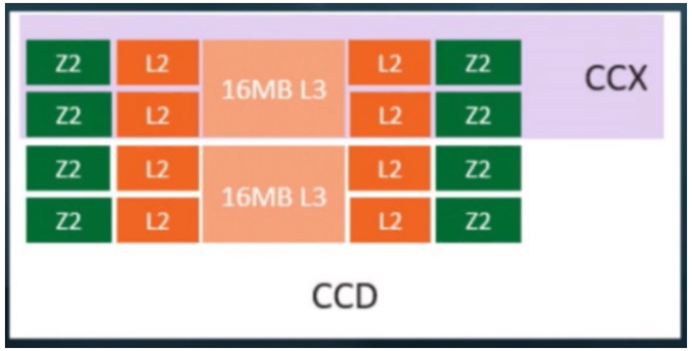

Sustained Bandwidth on Synthetic Datasets

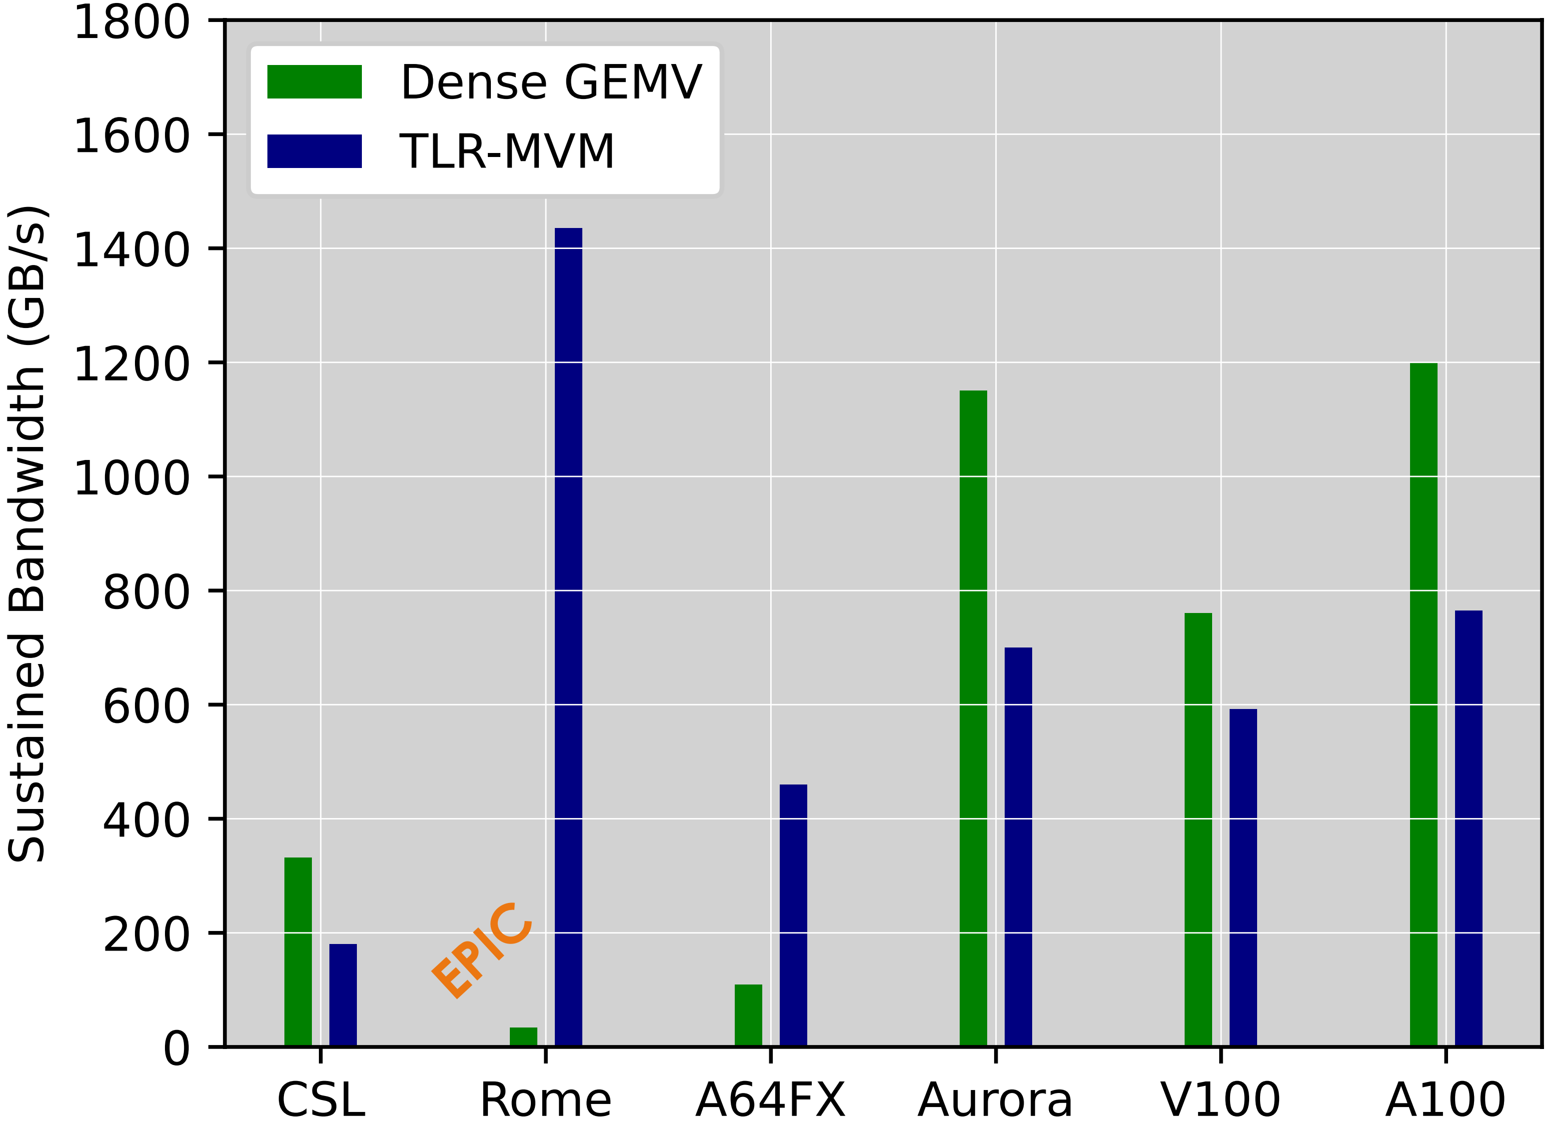

Higher is better

Rely on batch GEMV calls w/ variable sizes

Time to Solution on Synthetic Datasets

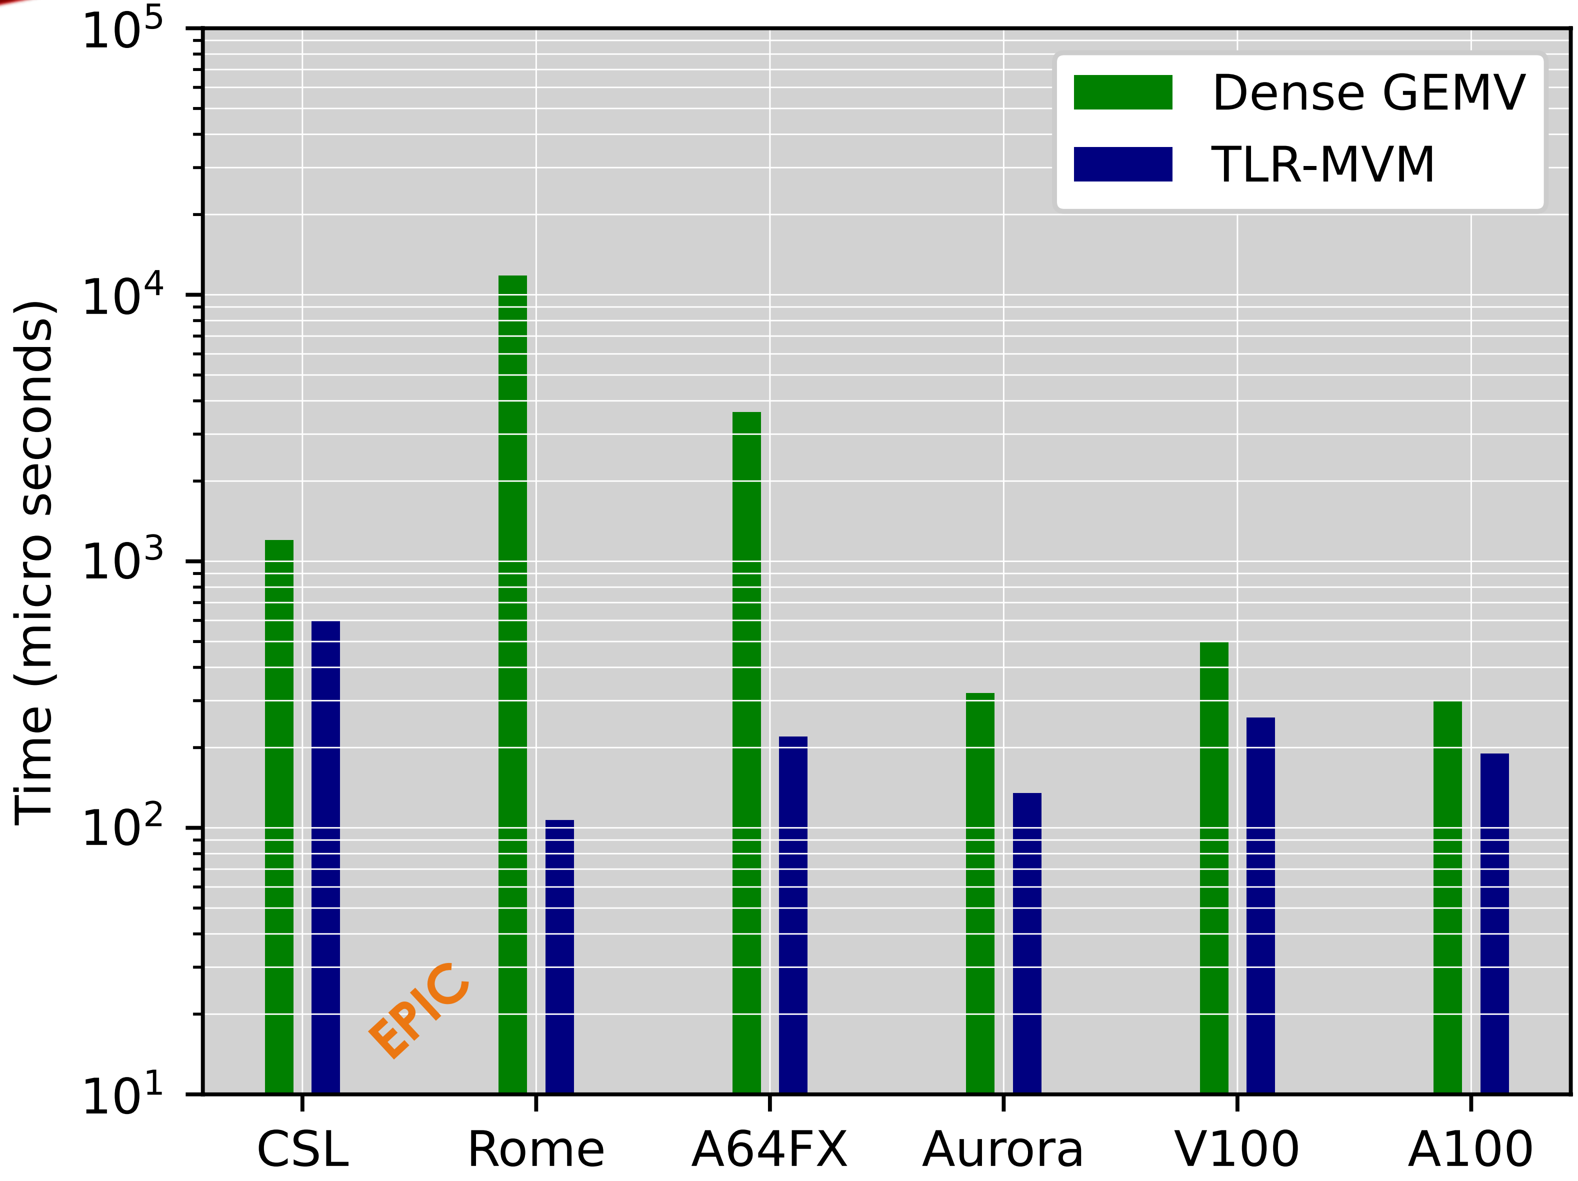

Lower is better

Rely on batch GEMV calls w/ variable sizes

Roofline Performance Model

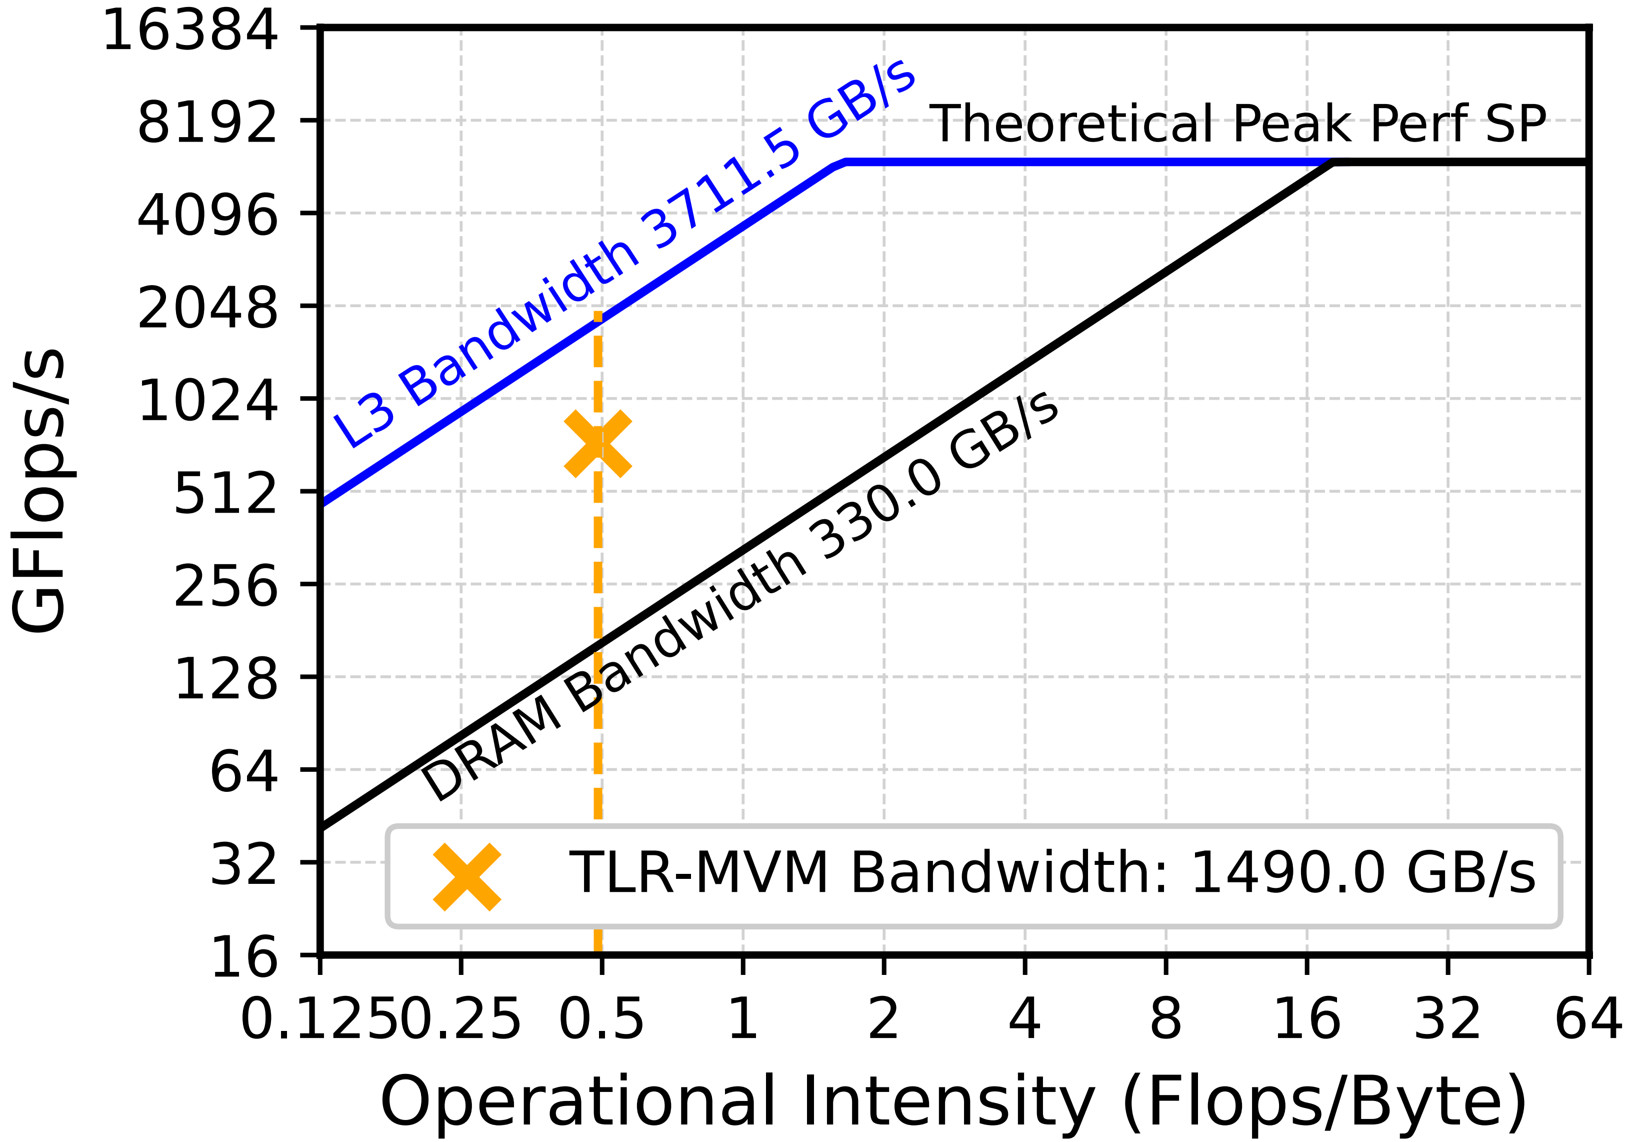

Performance Across AMD x86 Generations

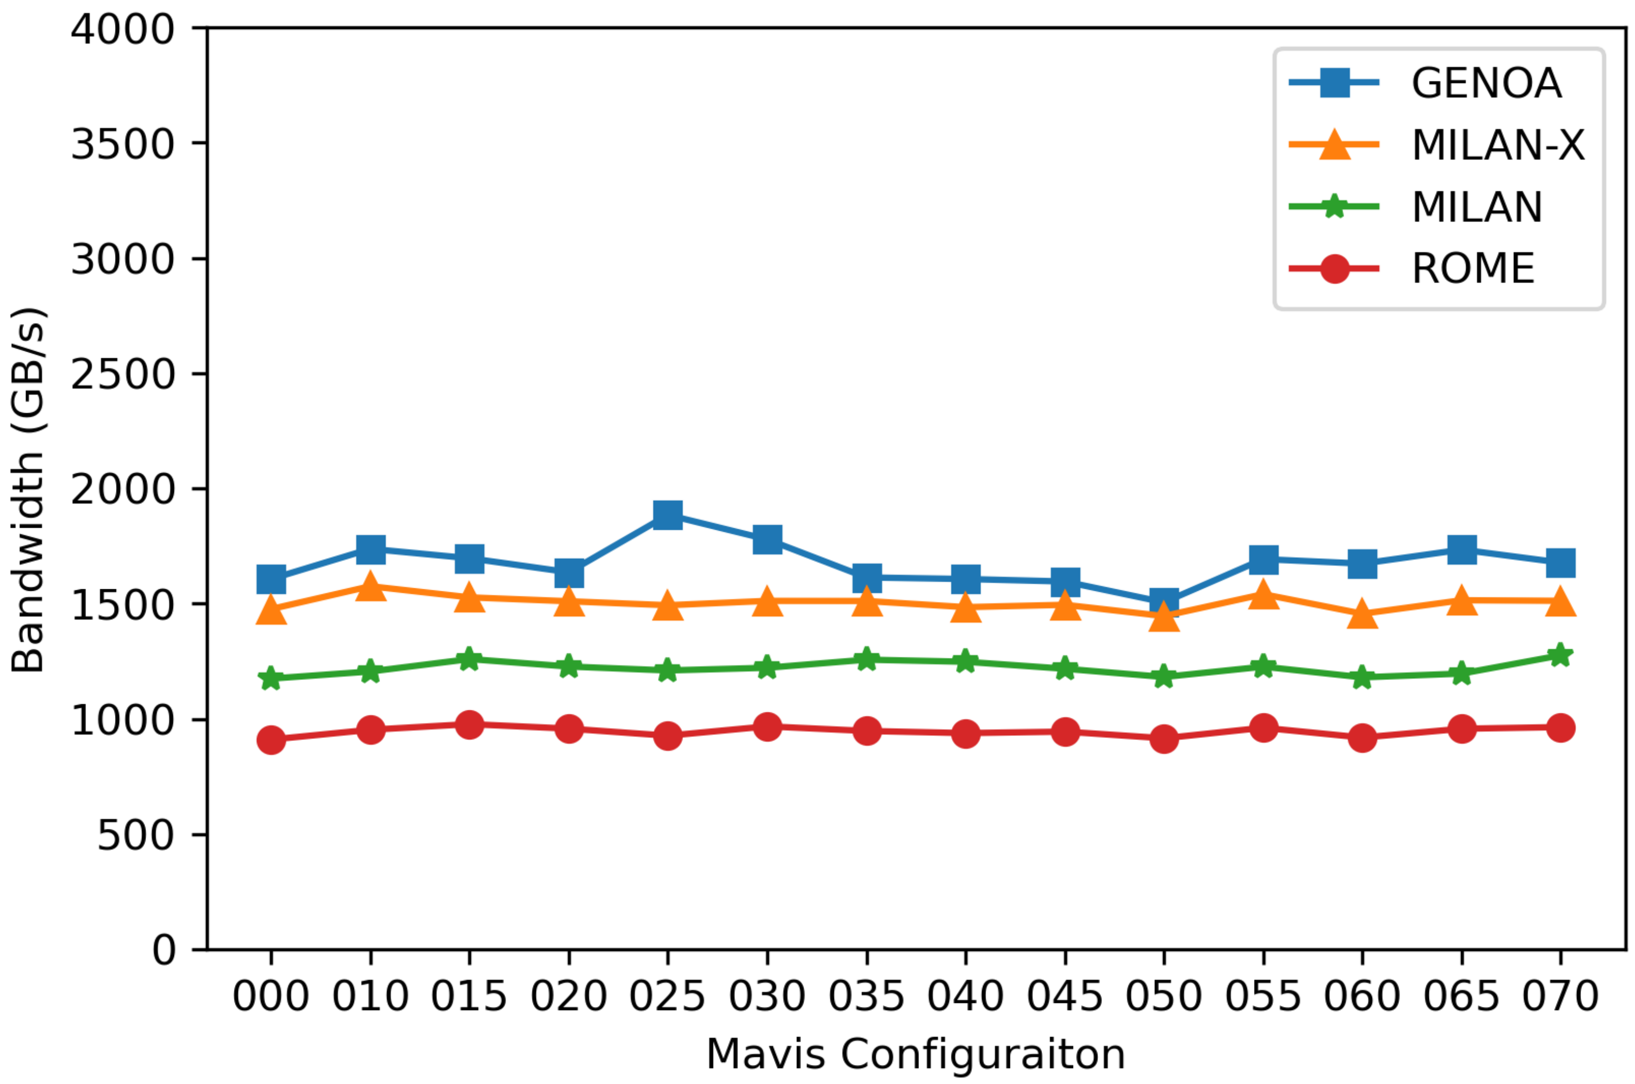

The Case for Mixed Precisions

Data transactions with AO hardware rely on UINT16

- WFS cameras provide 12-16 bits data stream as input
- Deformable mirror actuators are controlled through a set of UINT16 commands

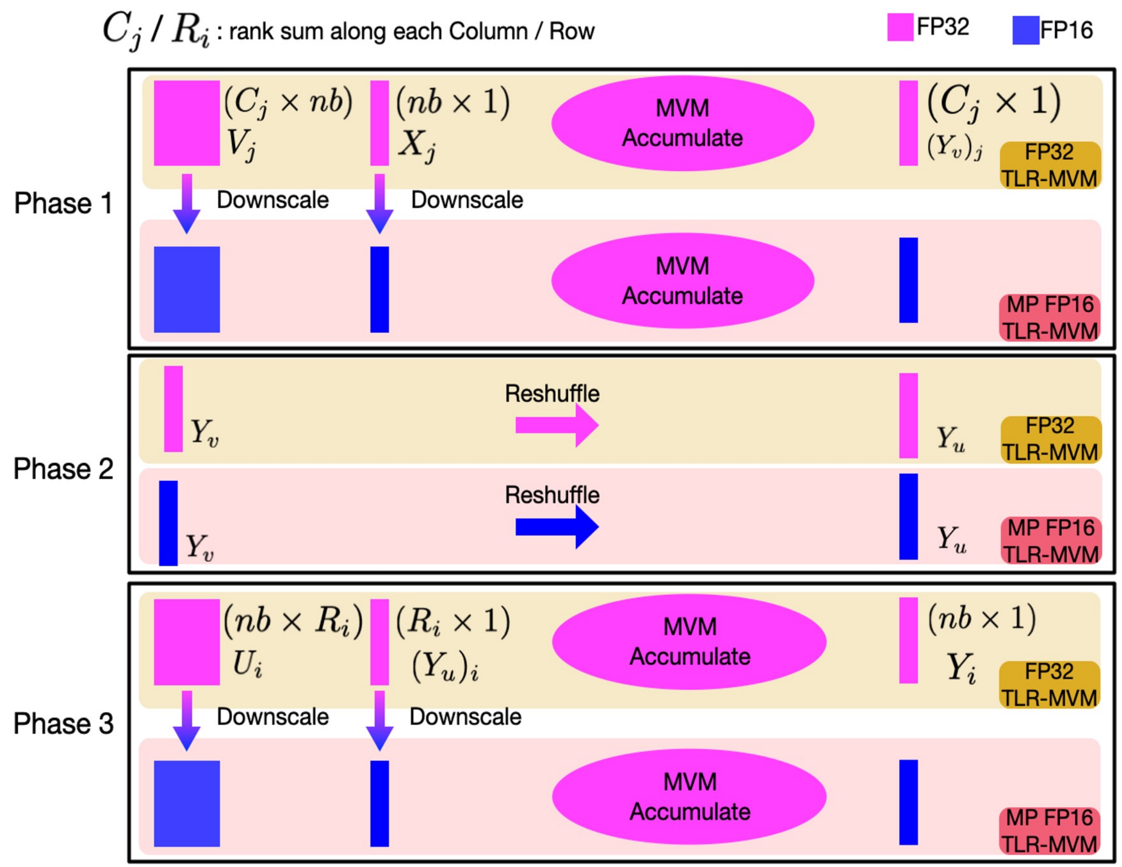

MAVIS Instrument

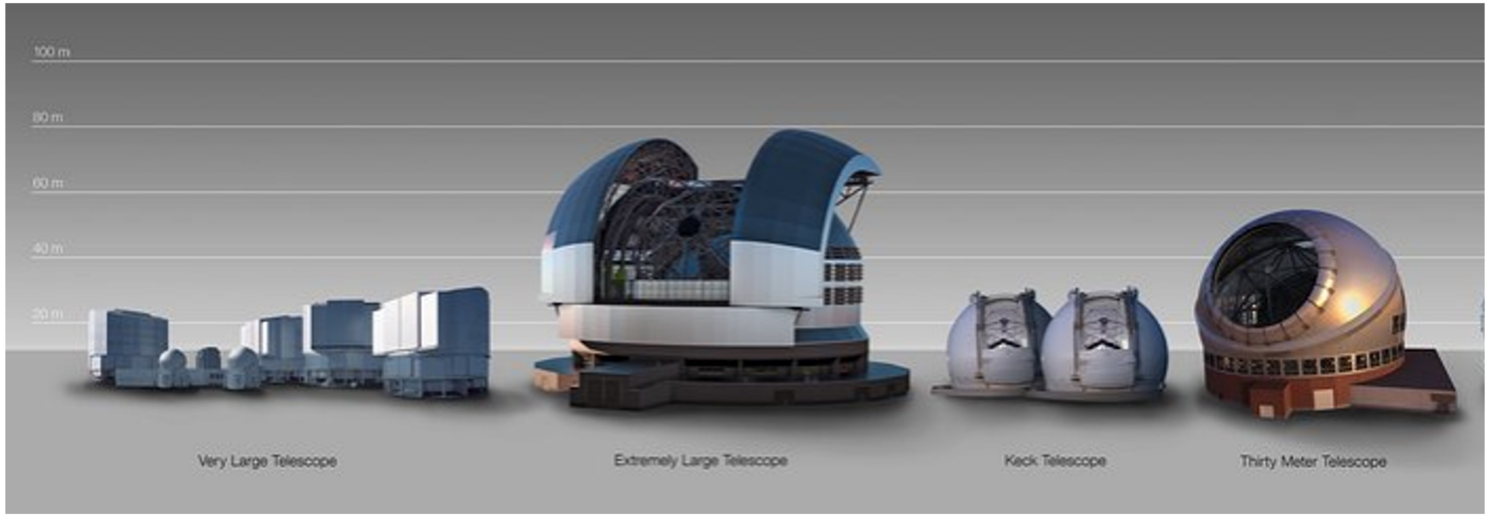

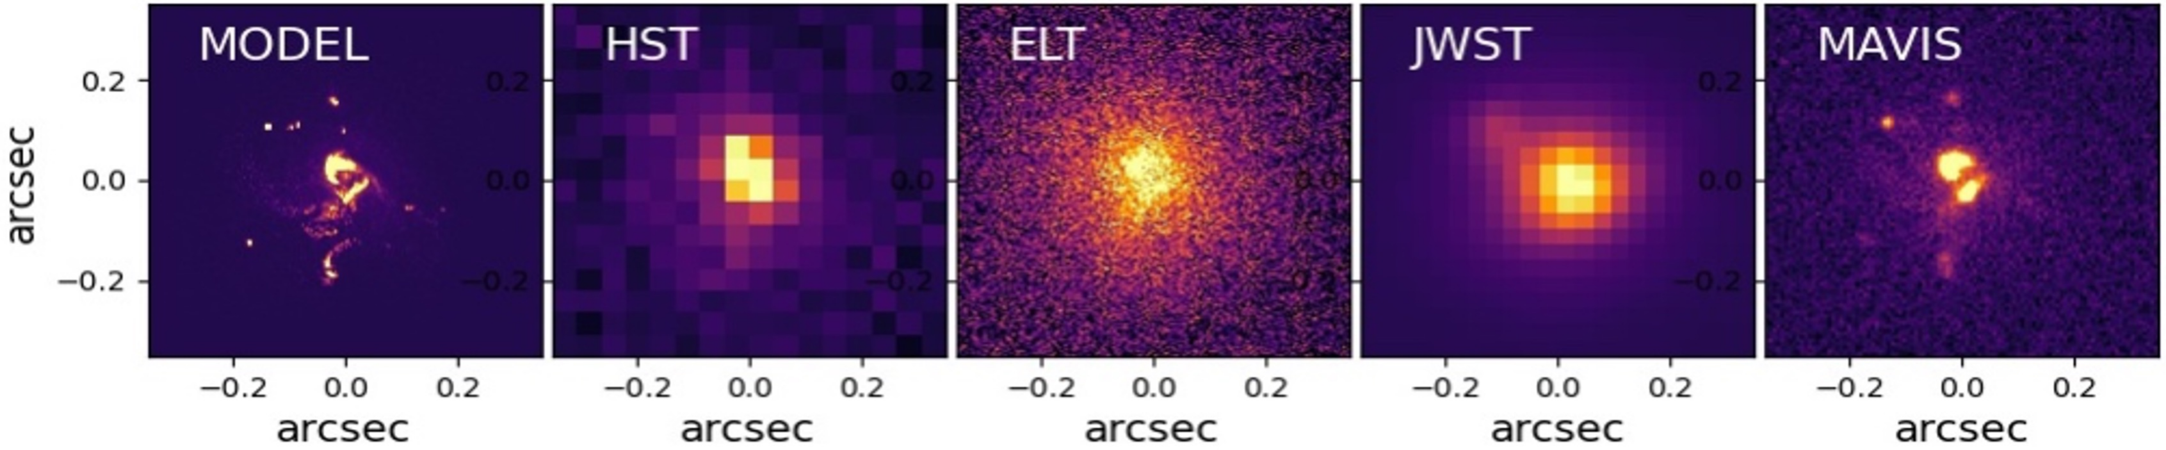

HPC Scientific Applications

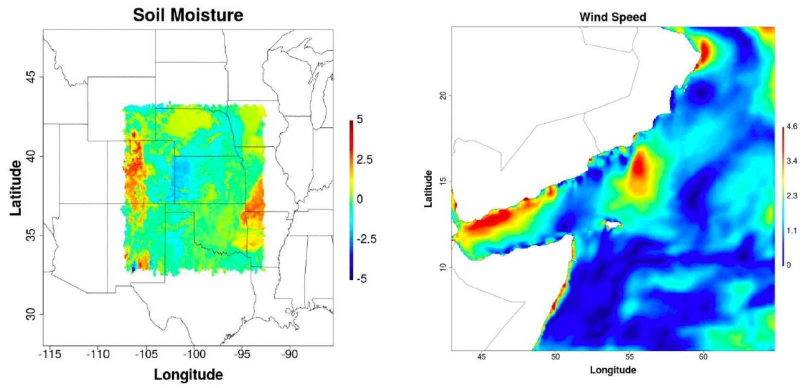

Fig. 4: **Left:** Soil moisture residuals at the topsoil of the Mississippi River basin. **Right**: Wind speed (m/s) in the Arabian Sea.

3D Geospatial Statistics

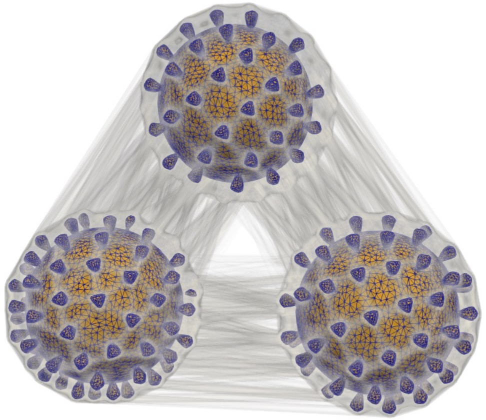

3D Mesh Deformations

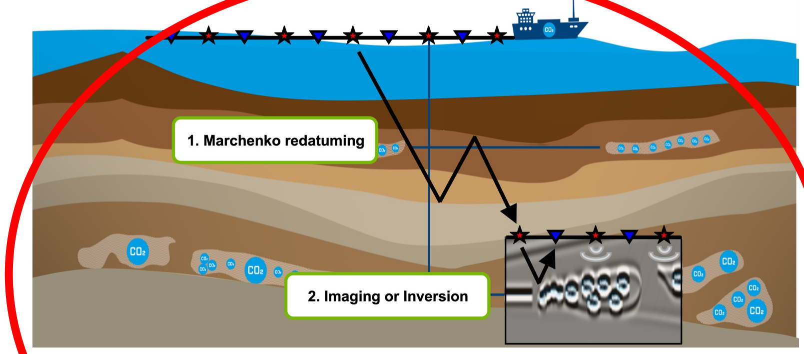

Seismic Imaging

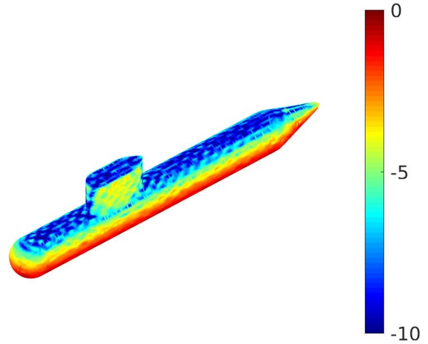

3D Computational Electromagnetics

Wireless Communications

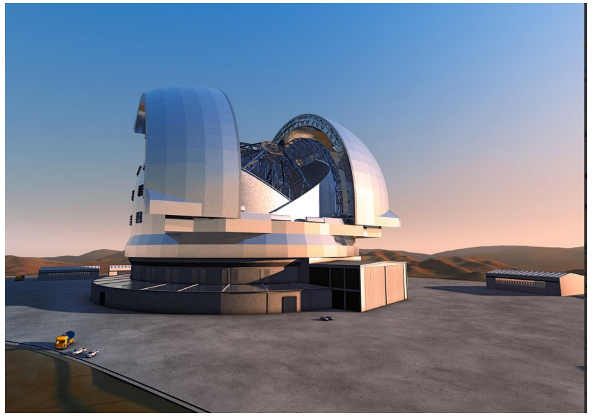

Computational Astronomy

Powering Seismic Redatuming w/ TLR-MVM

Seismic redatuming is an important technique to get insights from the Earth’s subsurface.

This requires solving an inverse problem. Traditionally, due to computational challenges, only the adjoint is applied.

Some latest research show an alternative method to improve the solution of inverse problems by using an iterative solver, e.g., conjugate gradient iterative solver. This comes at the cost of evaluating multiple expensive MVM operations, as shown in the following equations:

$$\mathbf{x} = \mathbf{R}^H \mathbf{y}: \quad x(t, \mathbf{x}_R, \mathbf{x}_A) = \mathcal{F}^{-1}_{\omega_{max}} \left(\int_{\delta\mathbb{D}} R^*(\omega, \mathbf{x}_B, \mathbf{x}_R) \mathcal{F}_{\omega_{max}}(y(t, \mathbf{x}_B, \mathbf{x}_A)) d\mathbf{x}_B \right),$$

$$\mathbf{y} = \mathbf{R}\mathbf{x}: \quad y(t, \mathbf{x}_B, \mathbf{x}_A) = \mathcal{F}^{-1}_{\omega_{max}} \left(\int_{\delta\mathbb{D}} R(\omega, \mathbf{x}_B, \mathbf{x}_R) \mathcal{F}_{\omega_{max}}(x(t, \mathbf{x}_R, \mathbf{x}_A)) d\mathbf{x}_R \right).$$

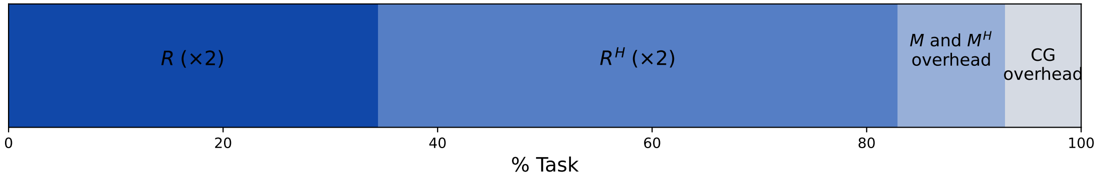

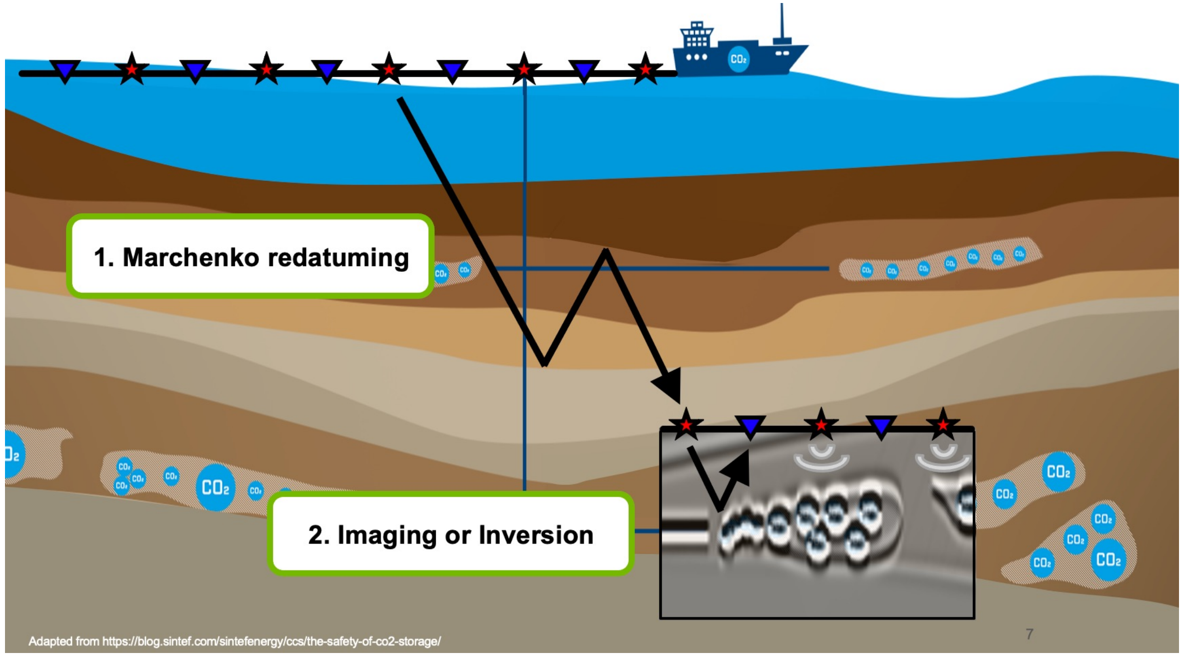

Powering Seismic Redatuming w/ TLR-MVM

We use tile low-rank matrix-vector multiplication (TLR-MVM) to address the complexity bottleneck.

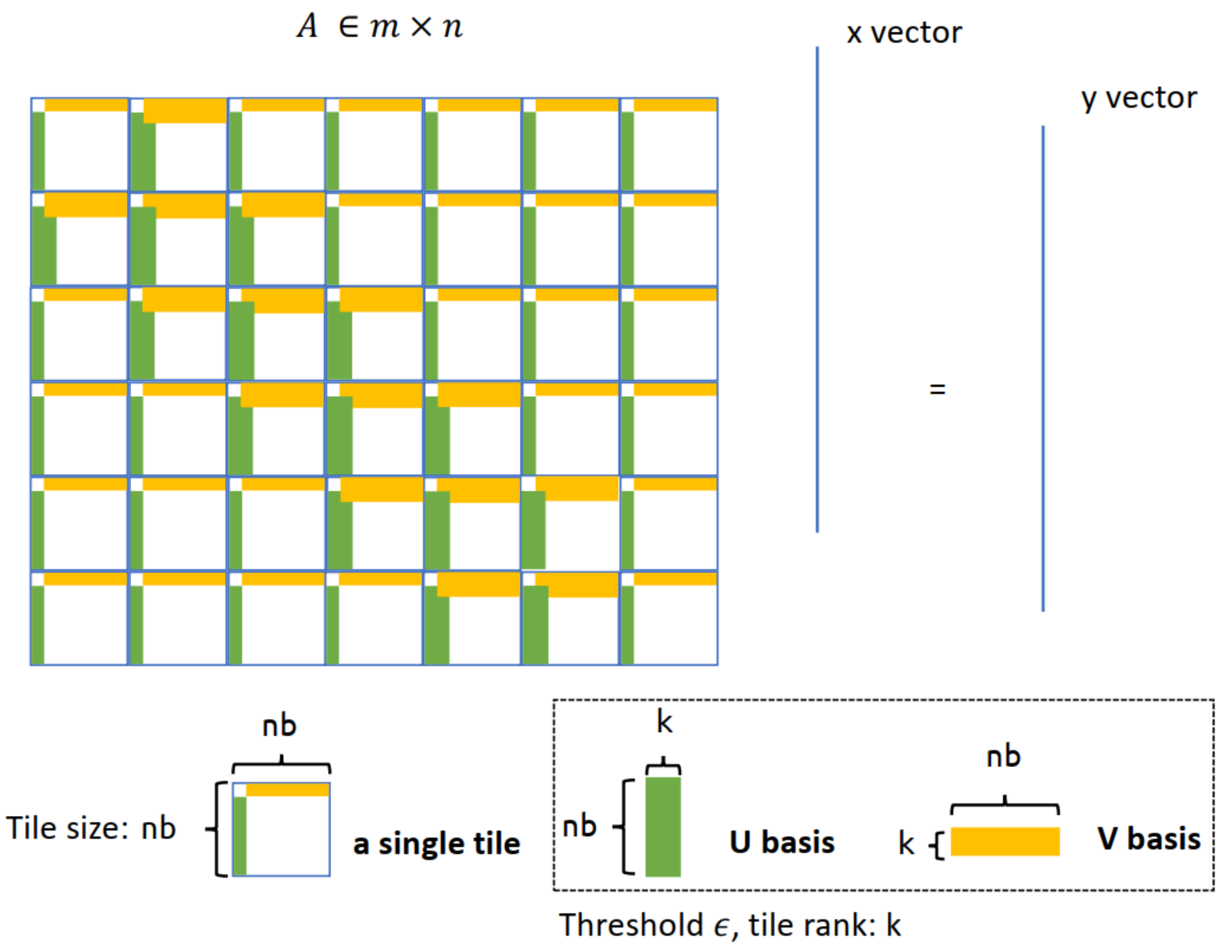

Distance-Aware Reordering

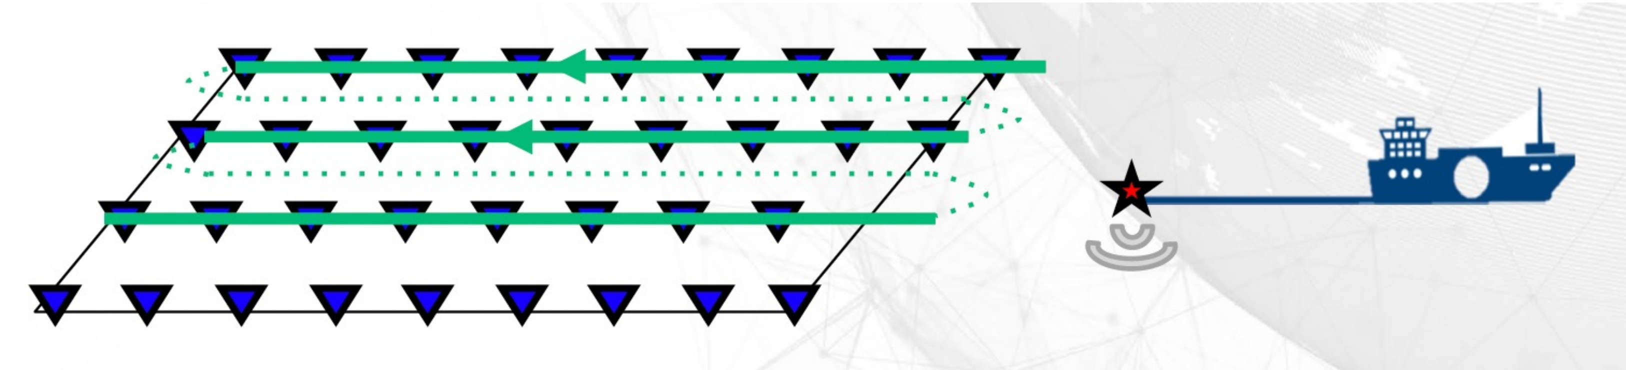

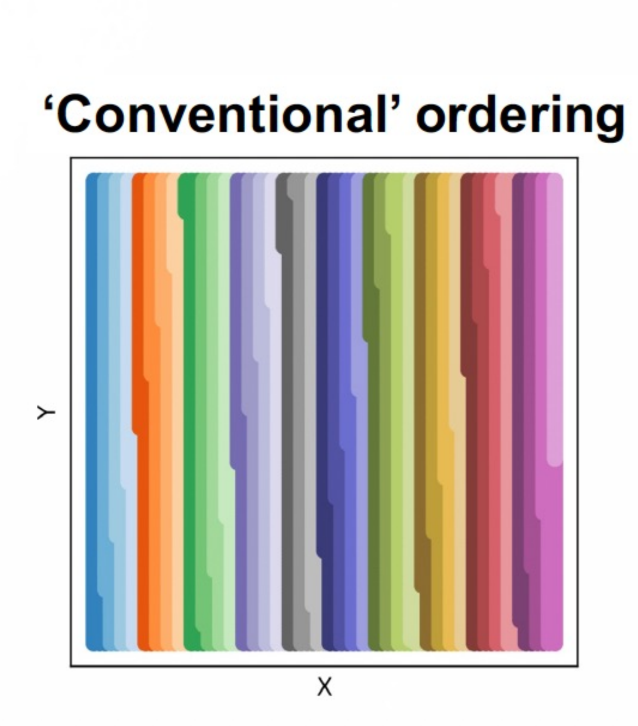

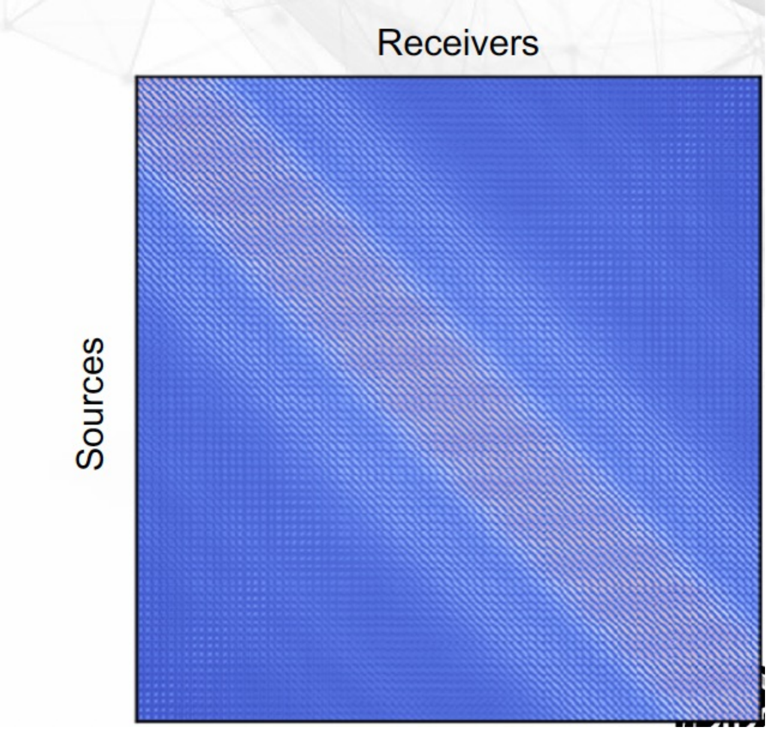

Distance-Aware Reordering

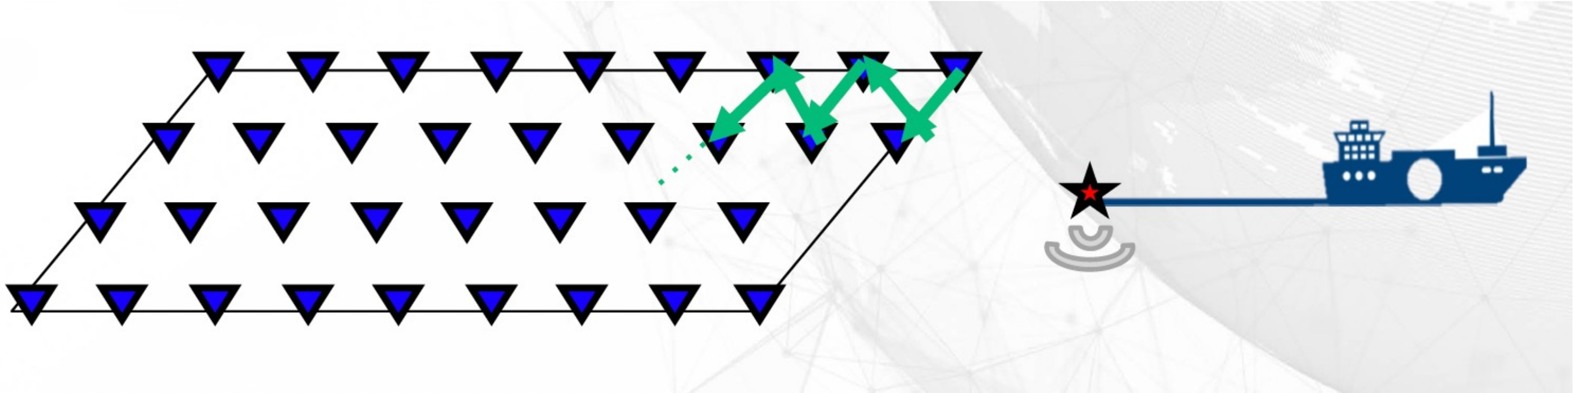

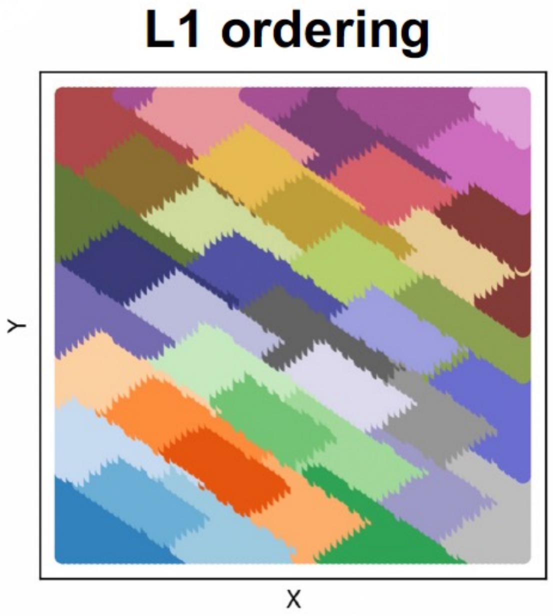

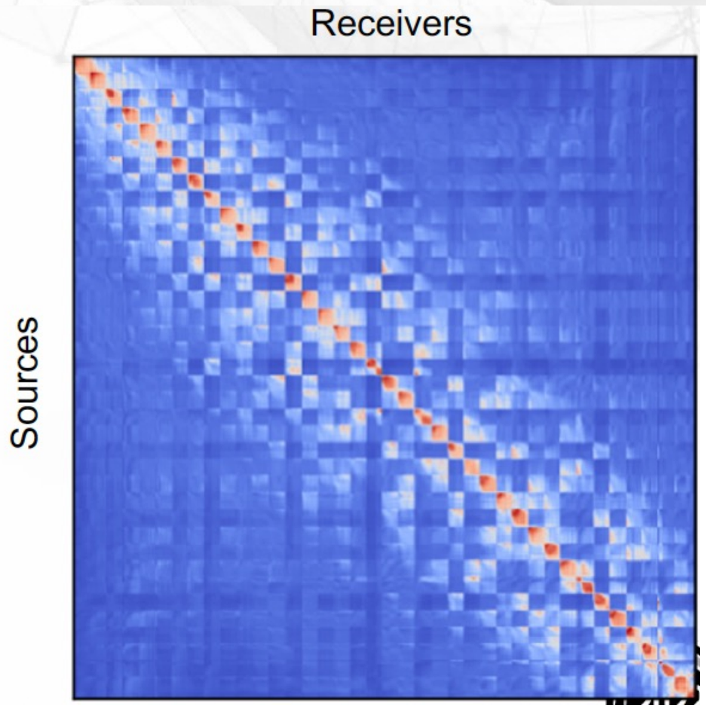

Distance-Aware Reordering

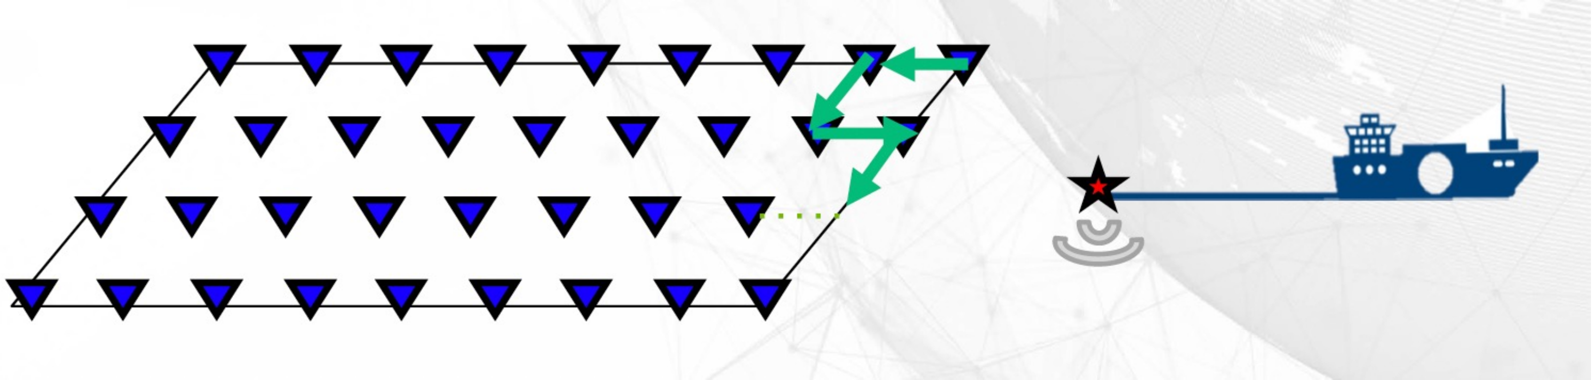

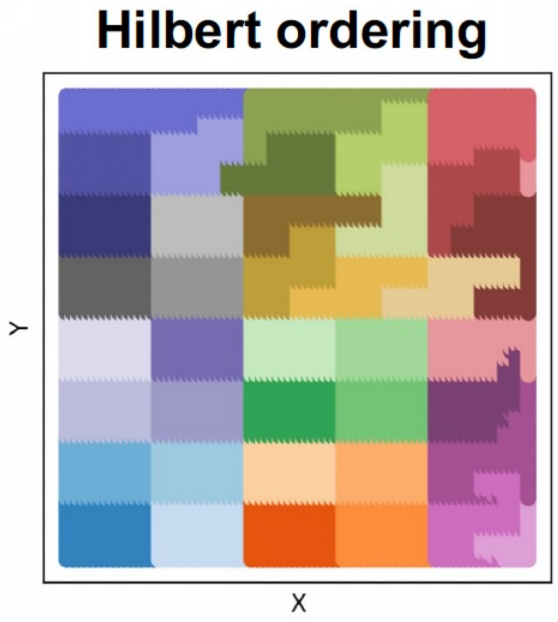

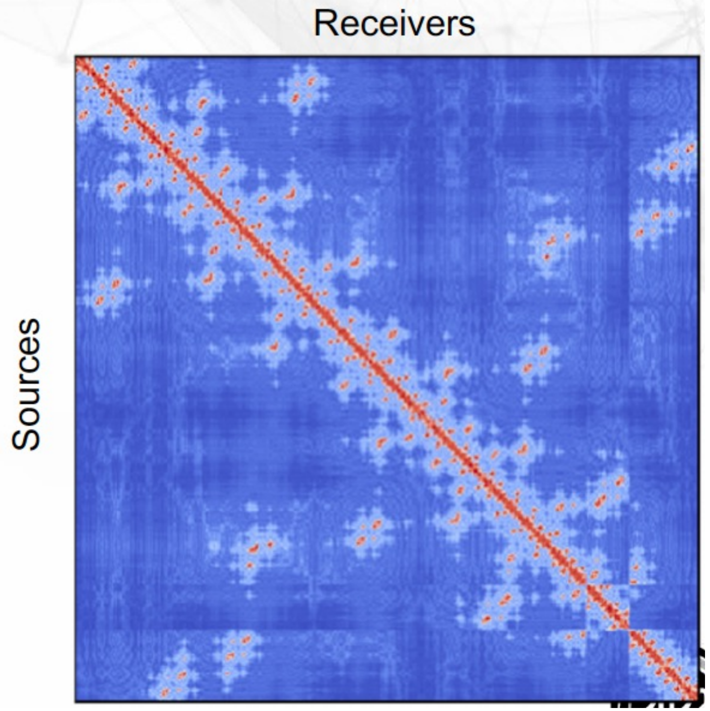

Distance-Aware Reordering

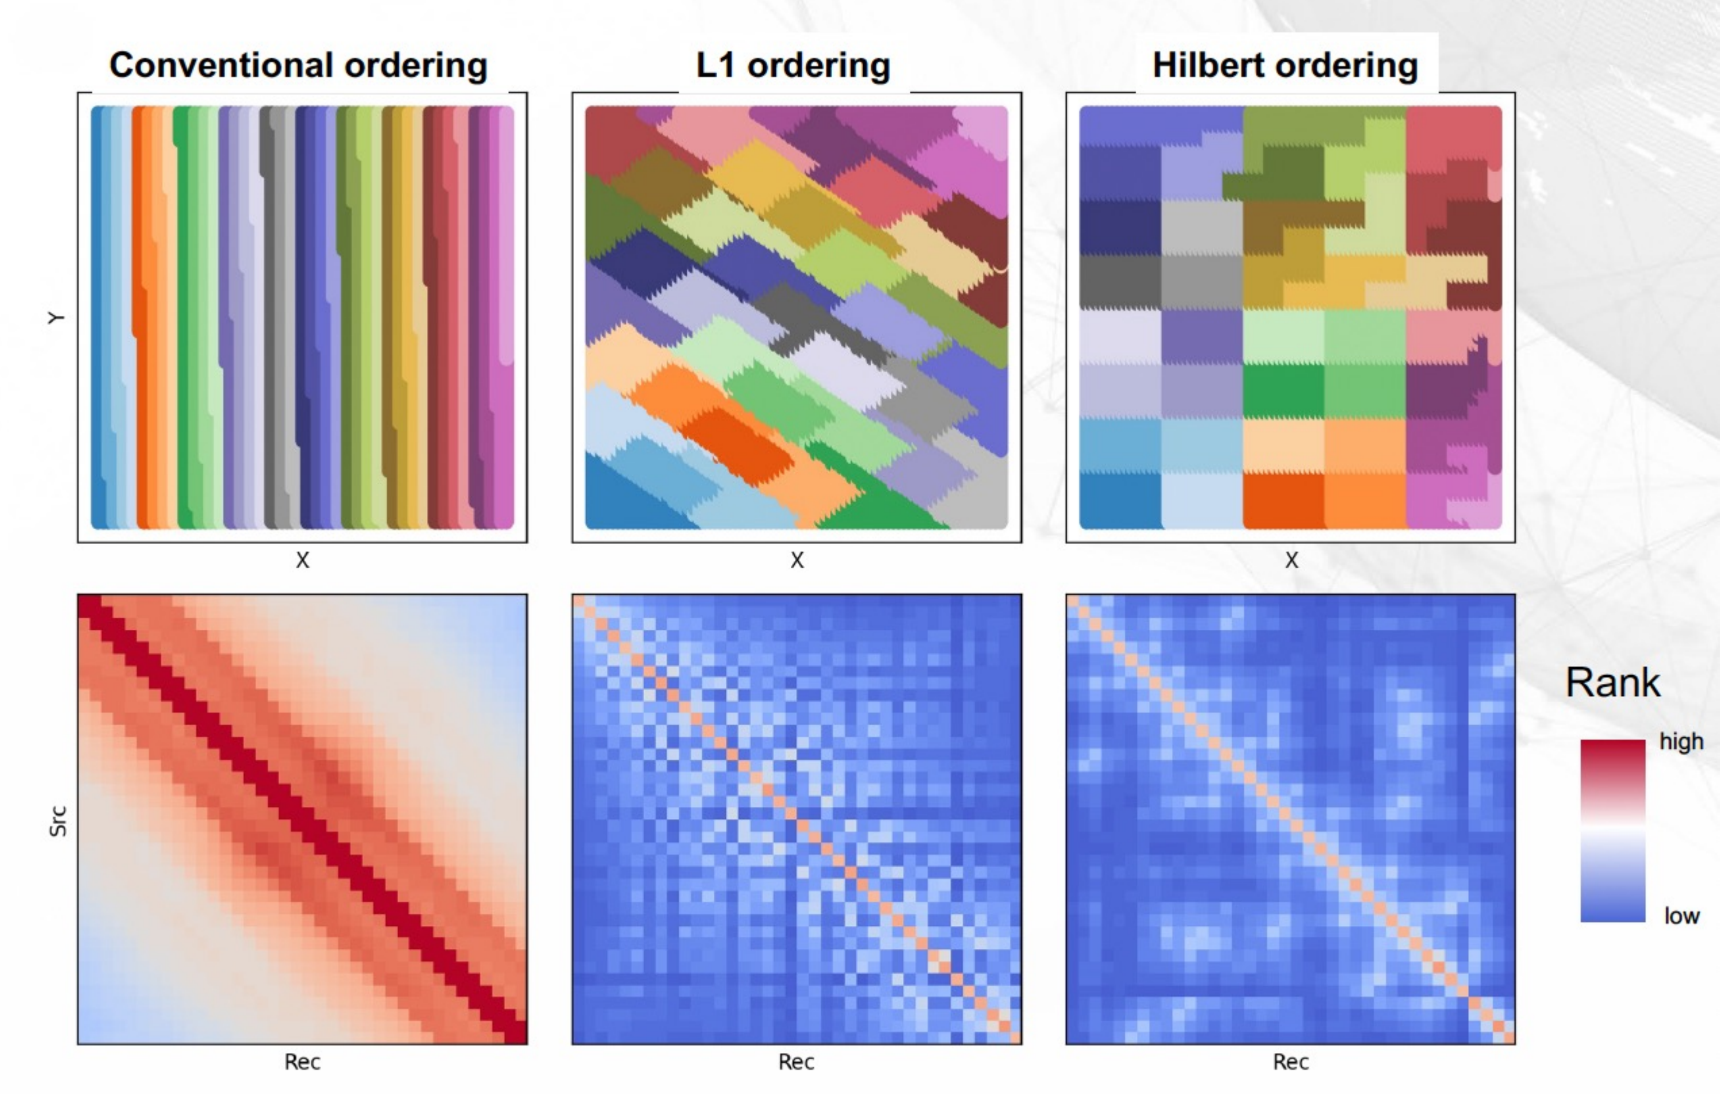

SEG/EAGE Overthrust Model

Jointly developed between the Society of Exploration Geophysicists (SEG) and the European Association of Geo- scientists and Engineers (EAGE)

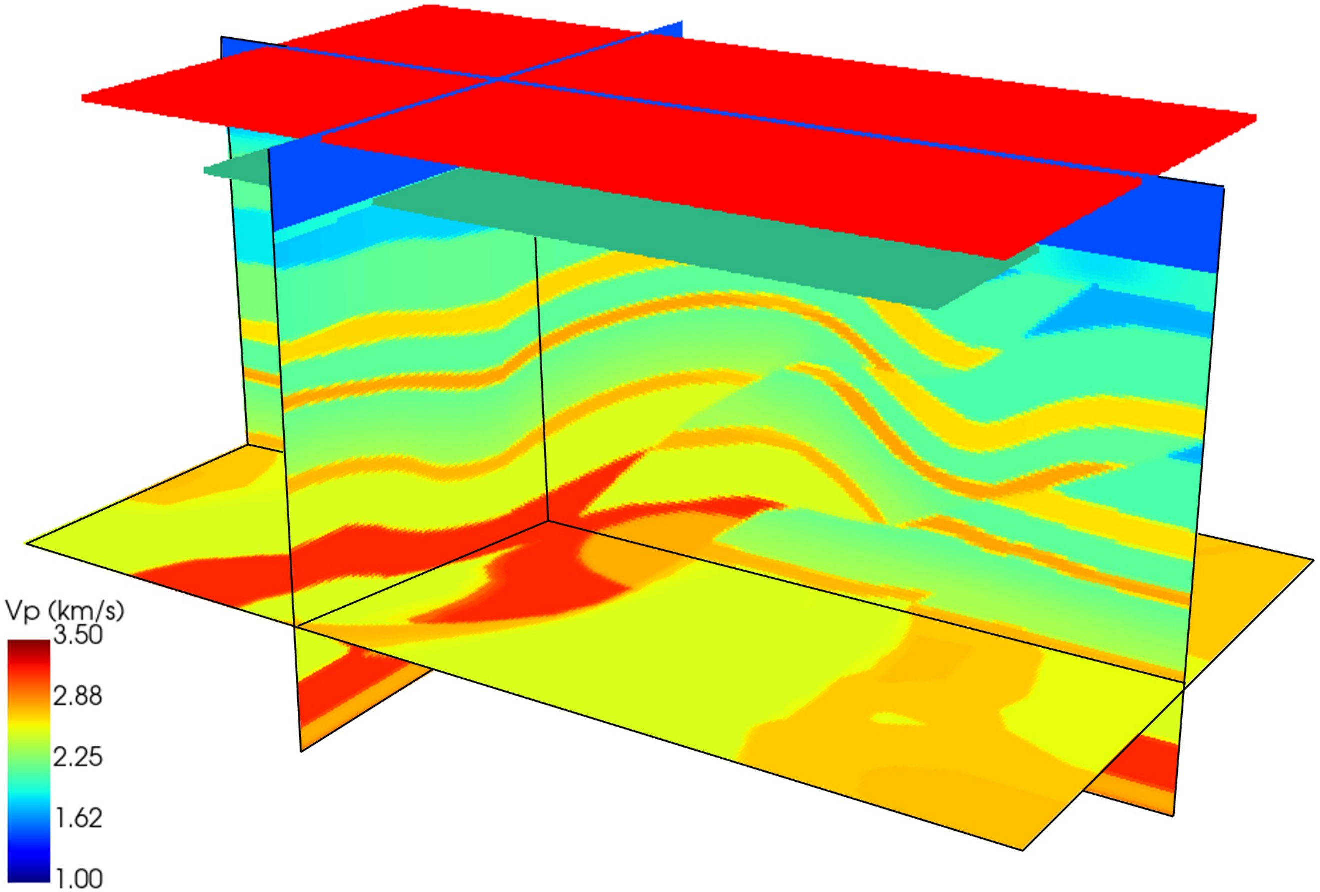

- 3D Geological open model
- 3 x 5 x 2.3 km^3
- 217 × 120 sources
- 177×90 receivers
- 230 complex-valued frequency matrices of size **26040 × 15930**

Numerical Accuracy

Checking the traces of 8 receivers

Accuracy Threshold: 1e-4

Good

Bad

Ground Truth

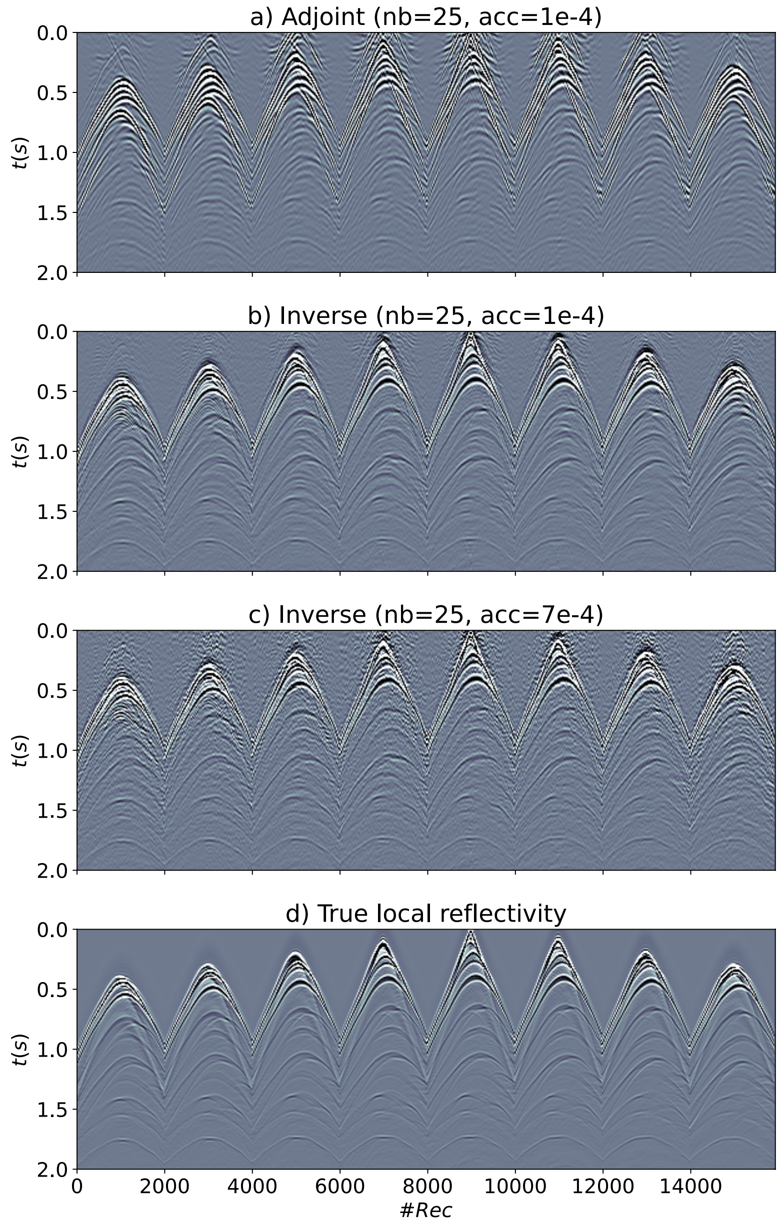

Numerical Accuracy

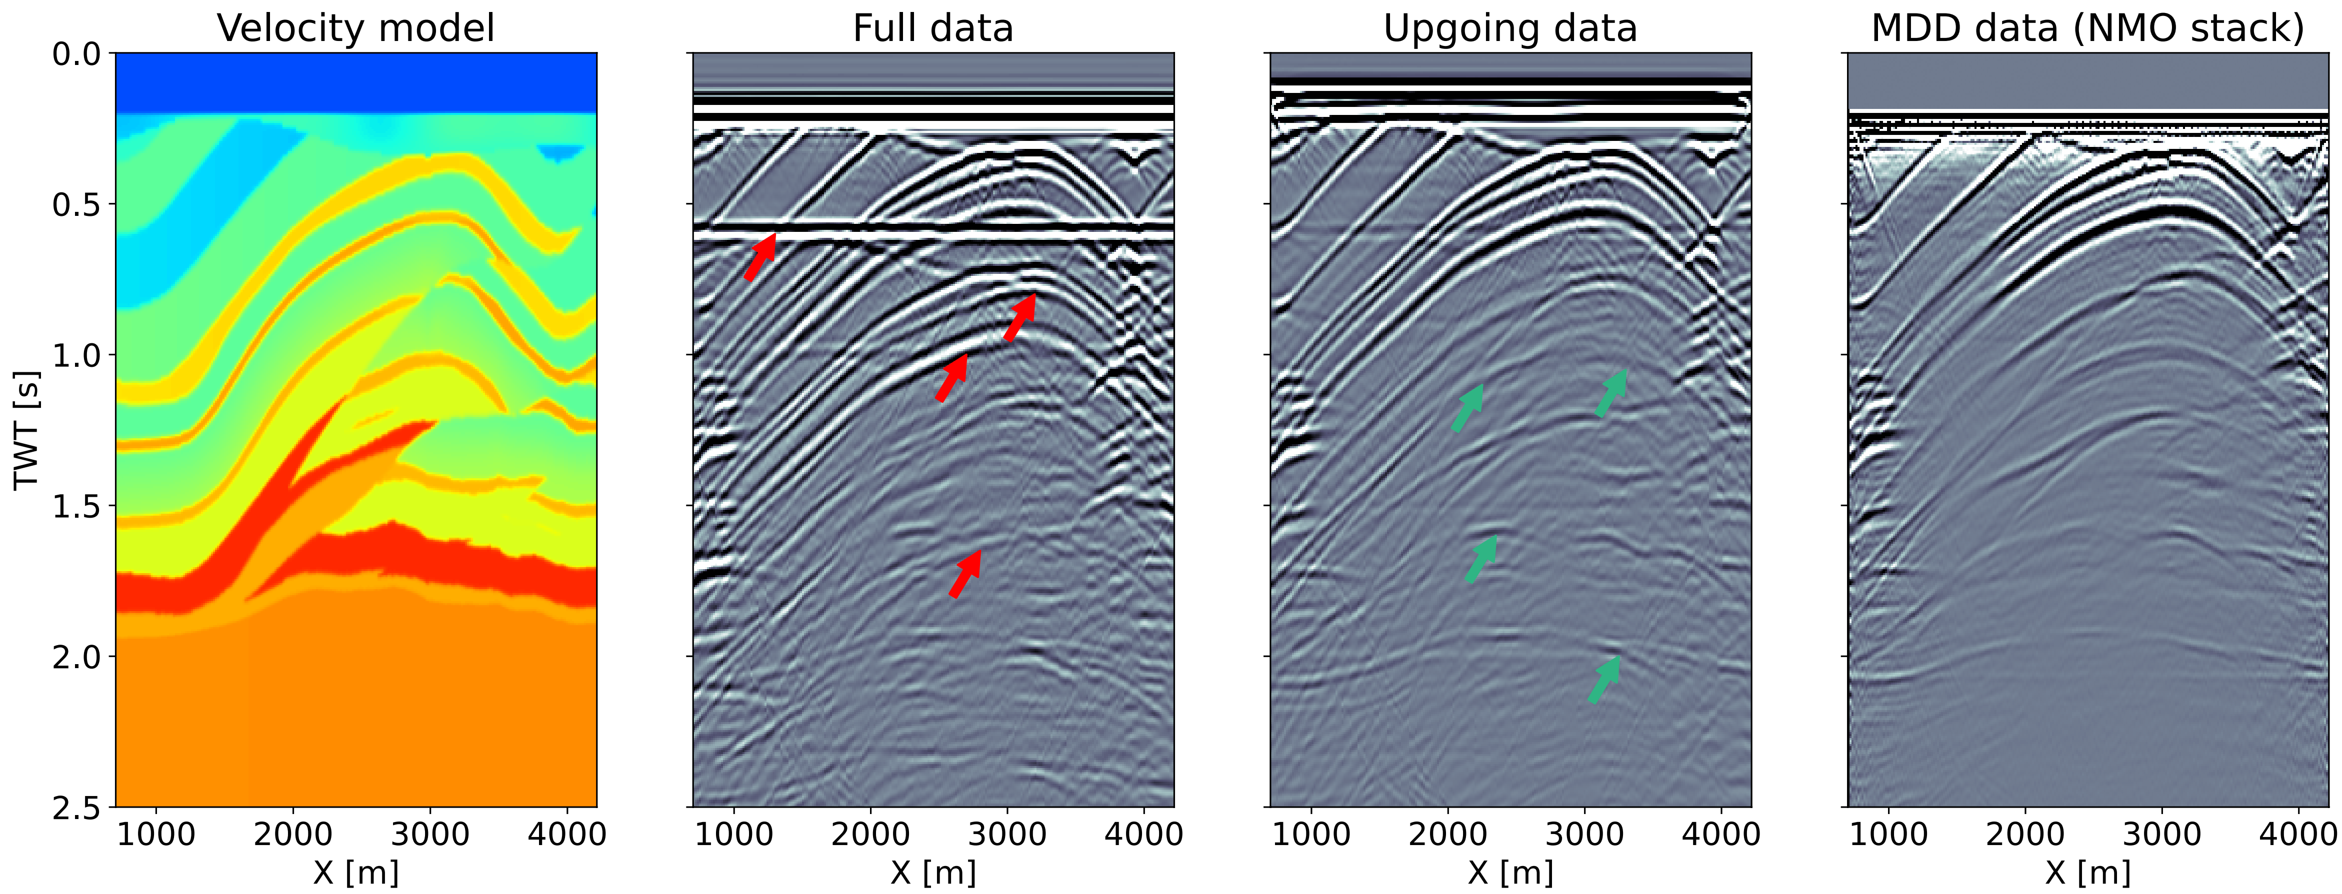

Post-acquisition processing powered by TLR-MVM to remove free-surface related effects

Performance Results

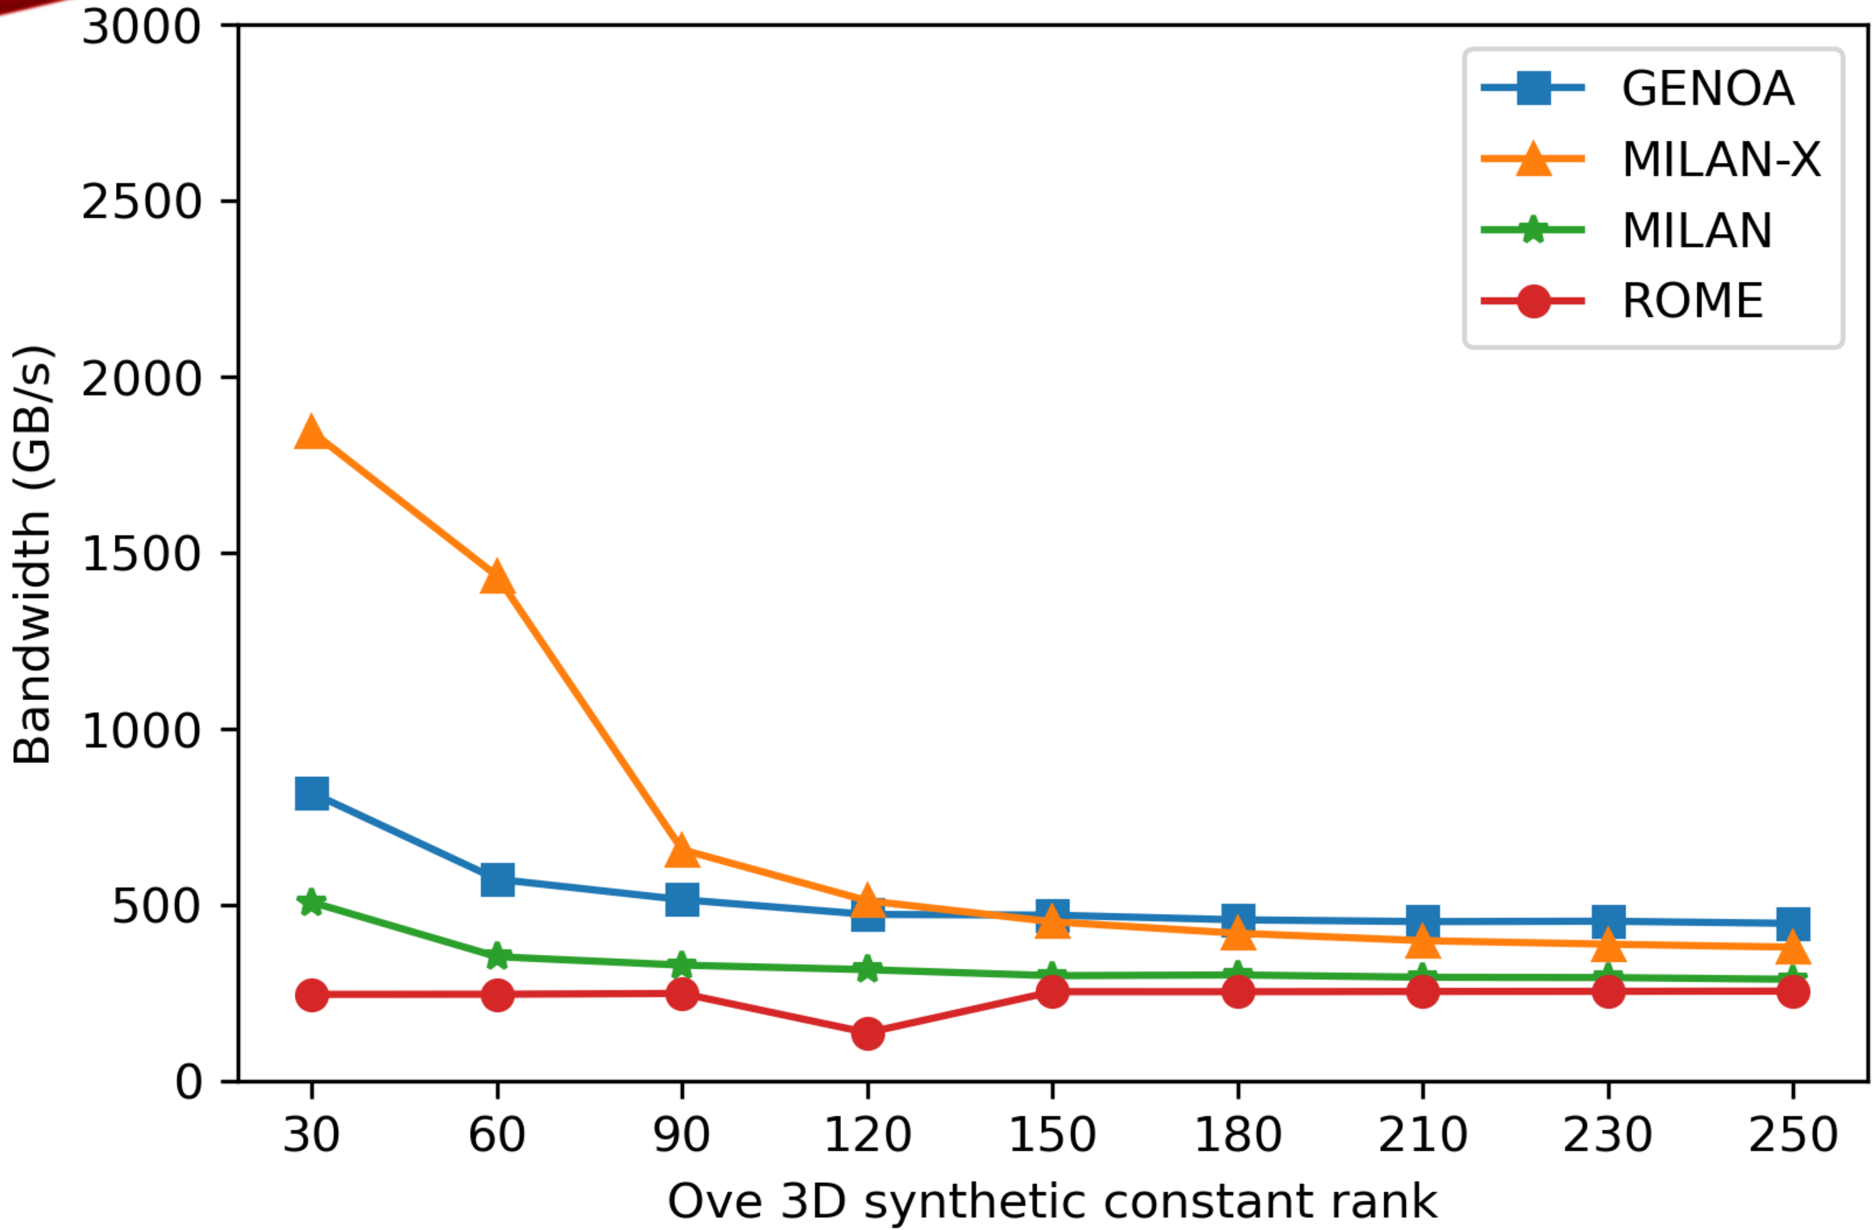

Performance Results

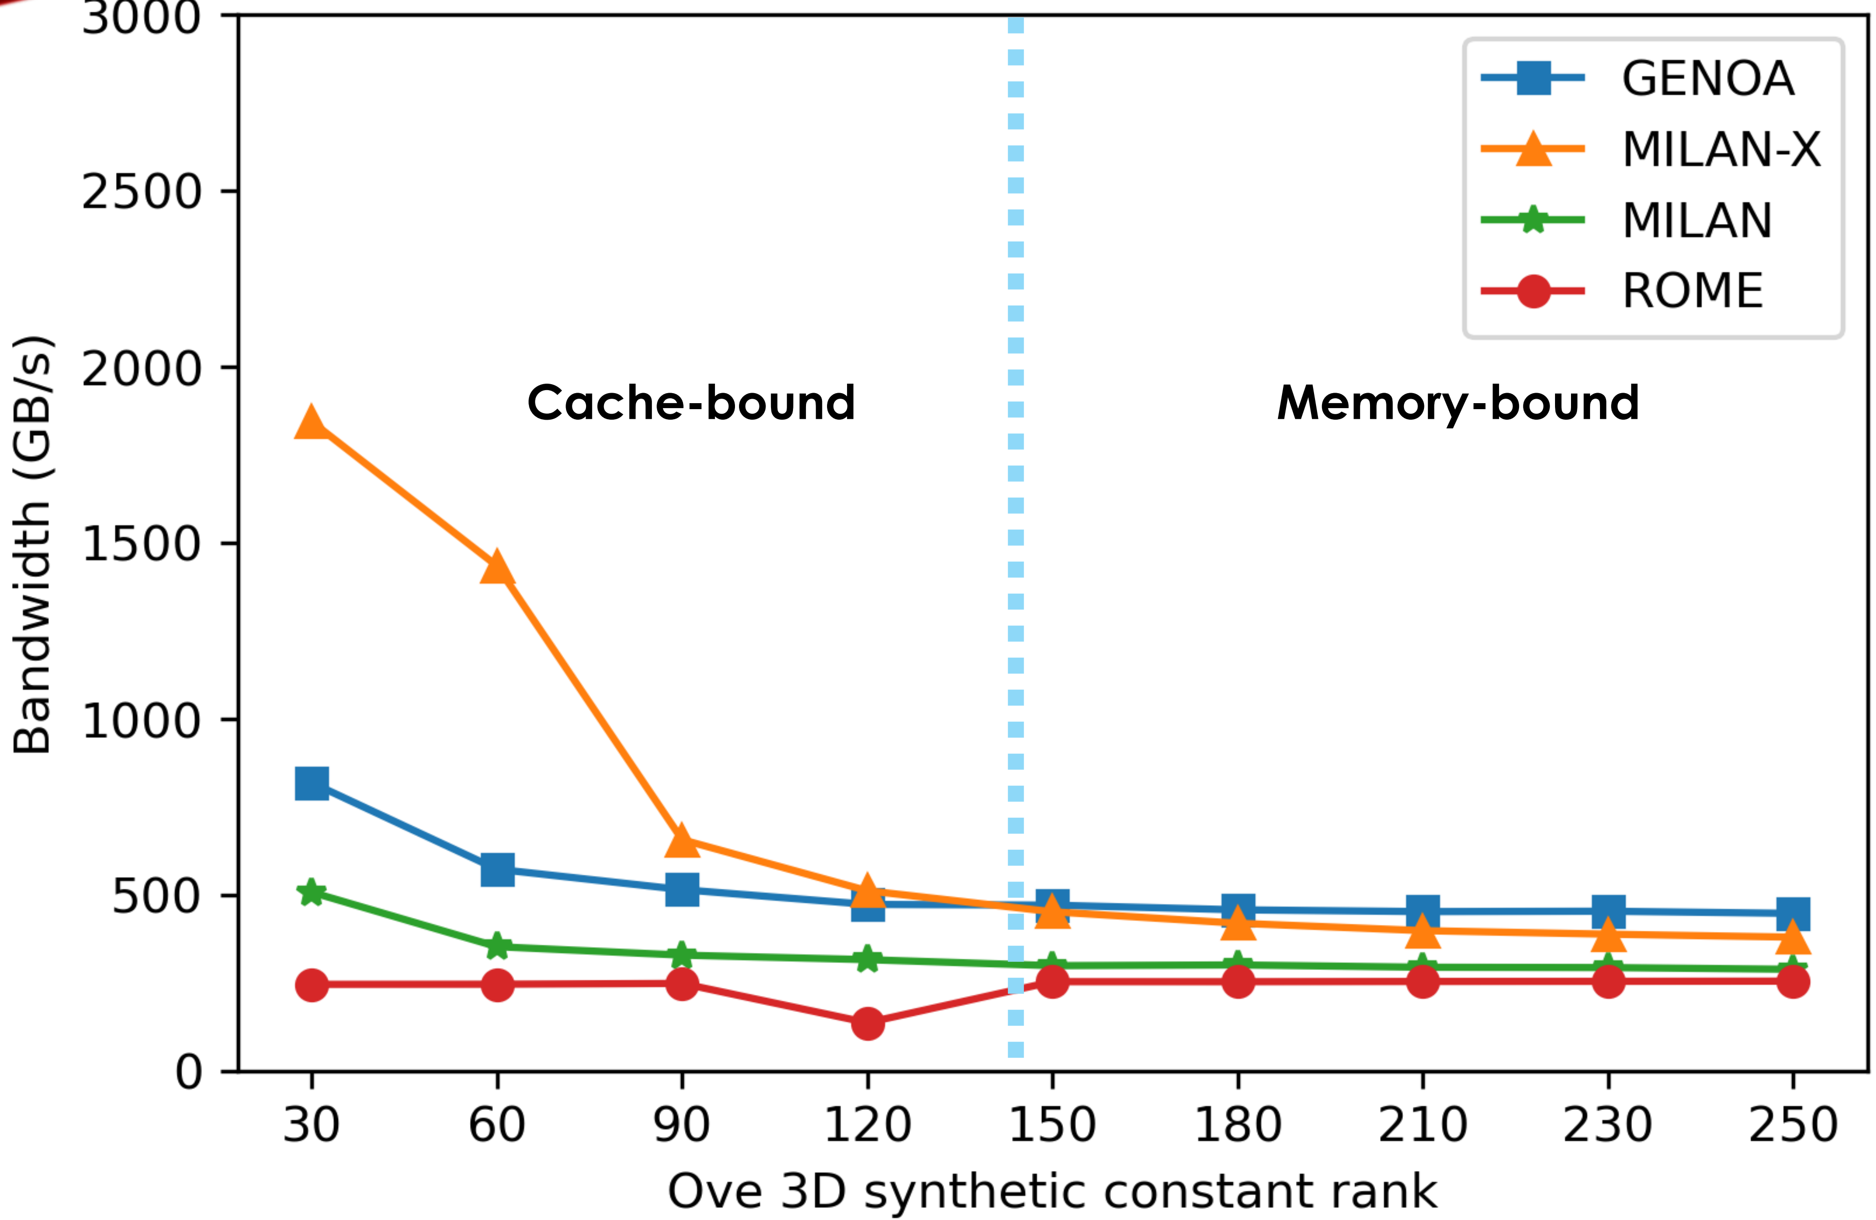

Performance Results

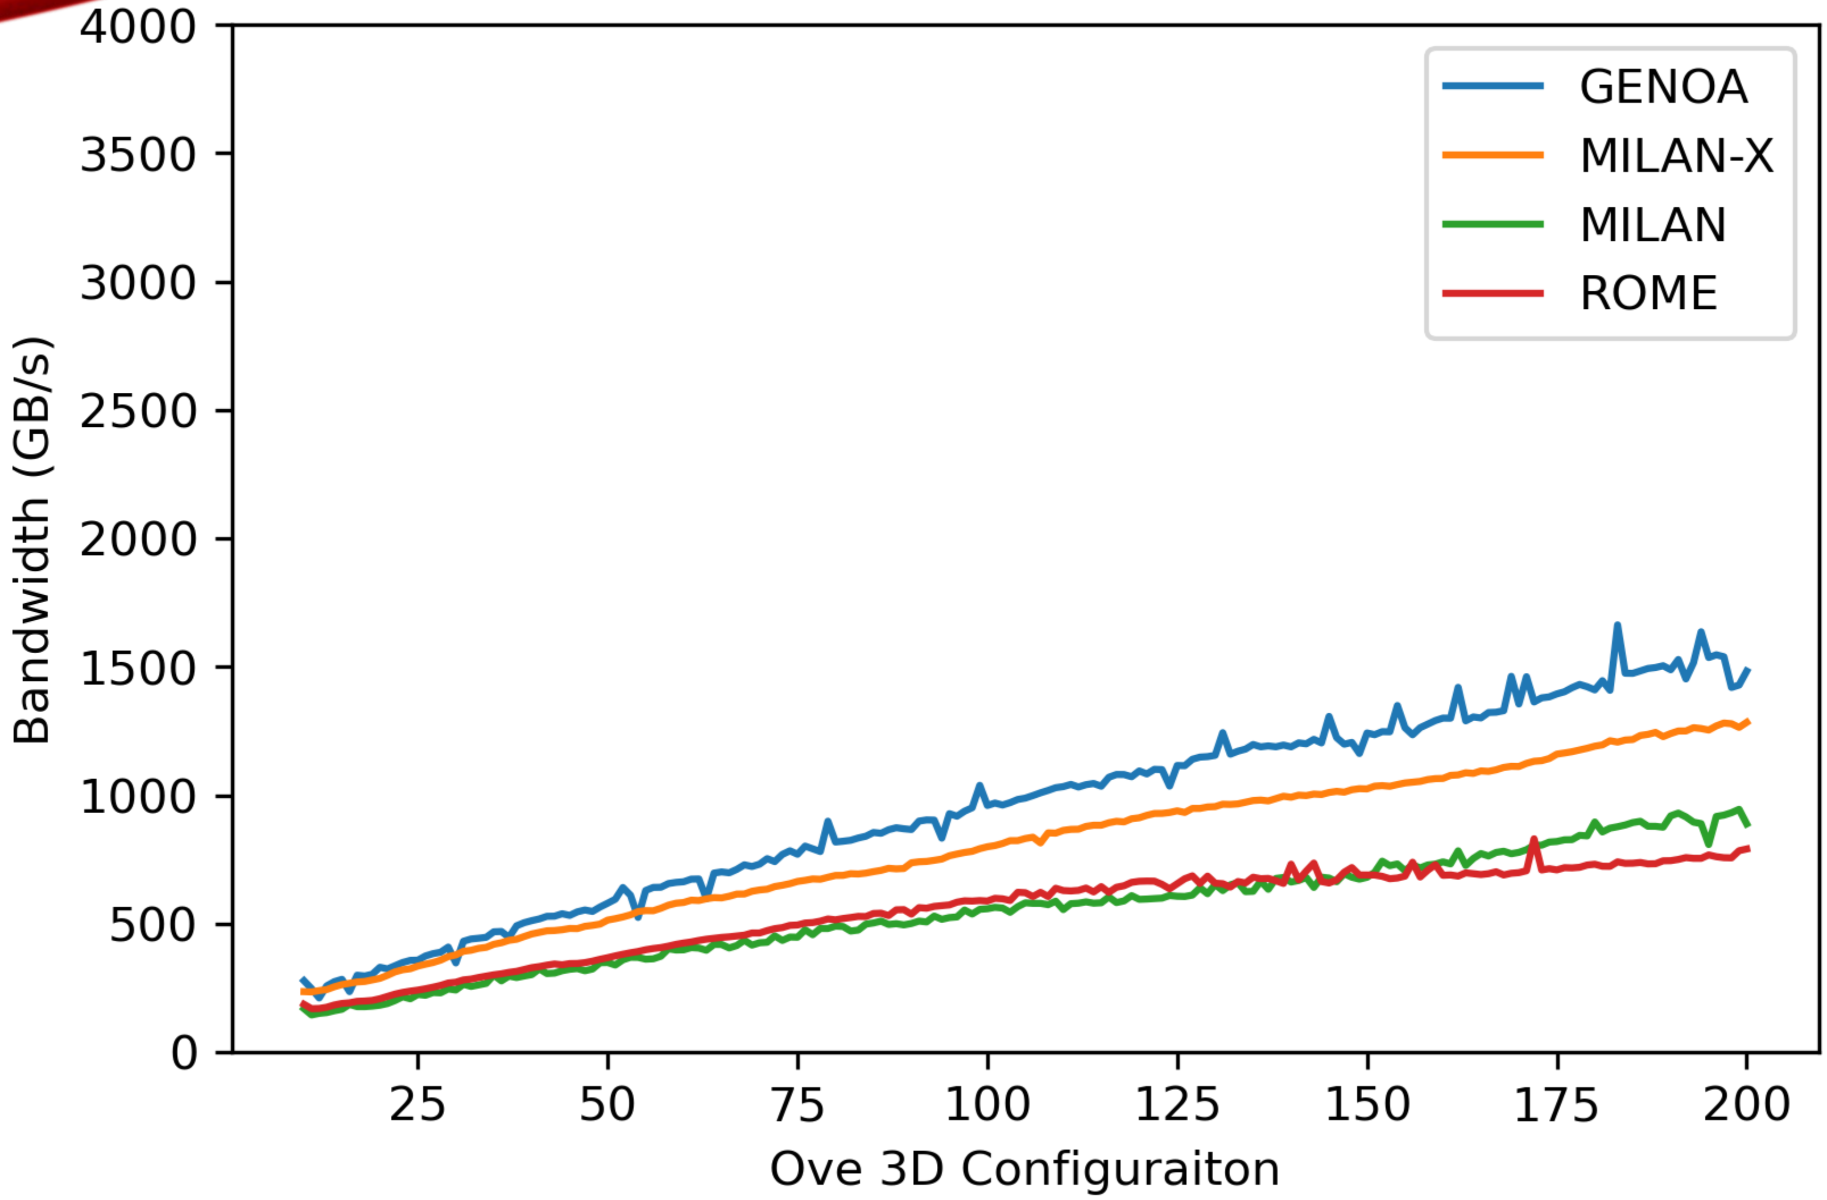

Performance Results

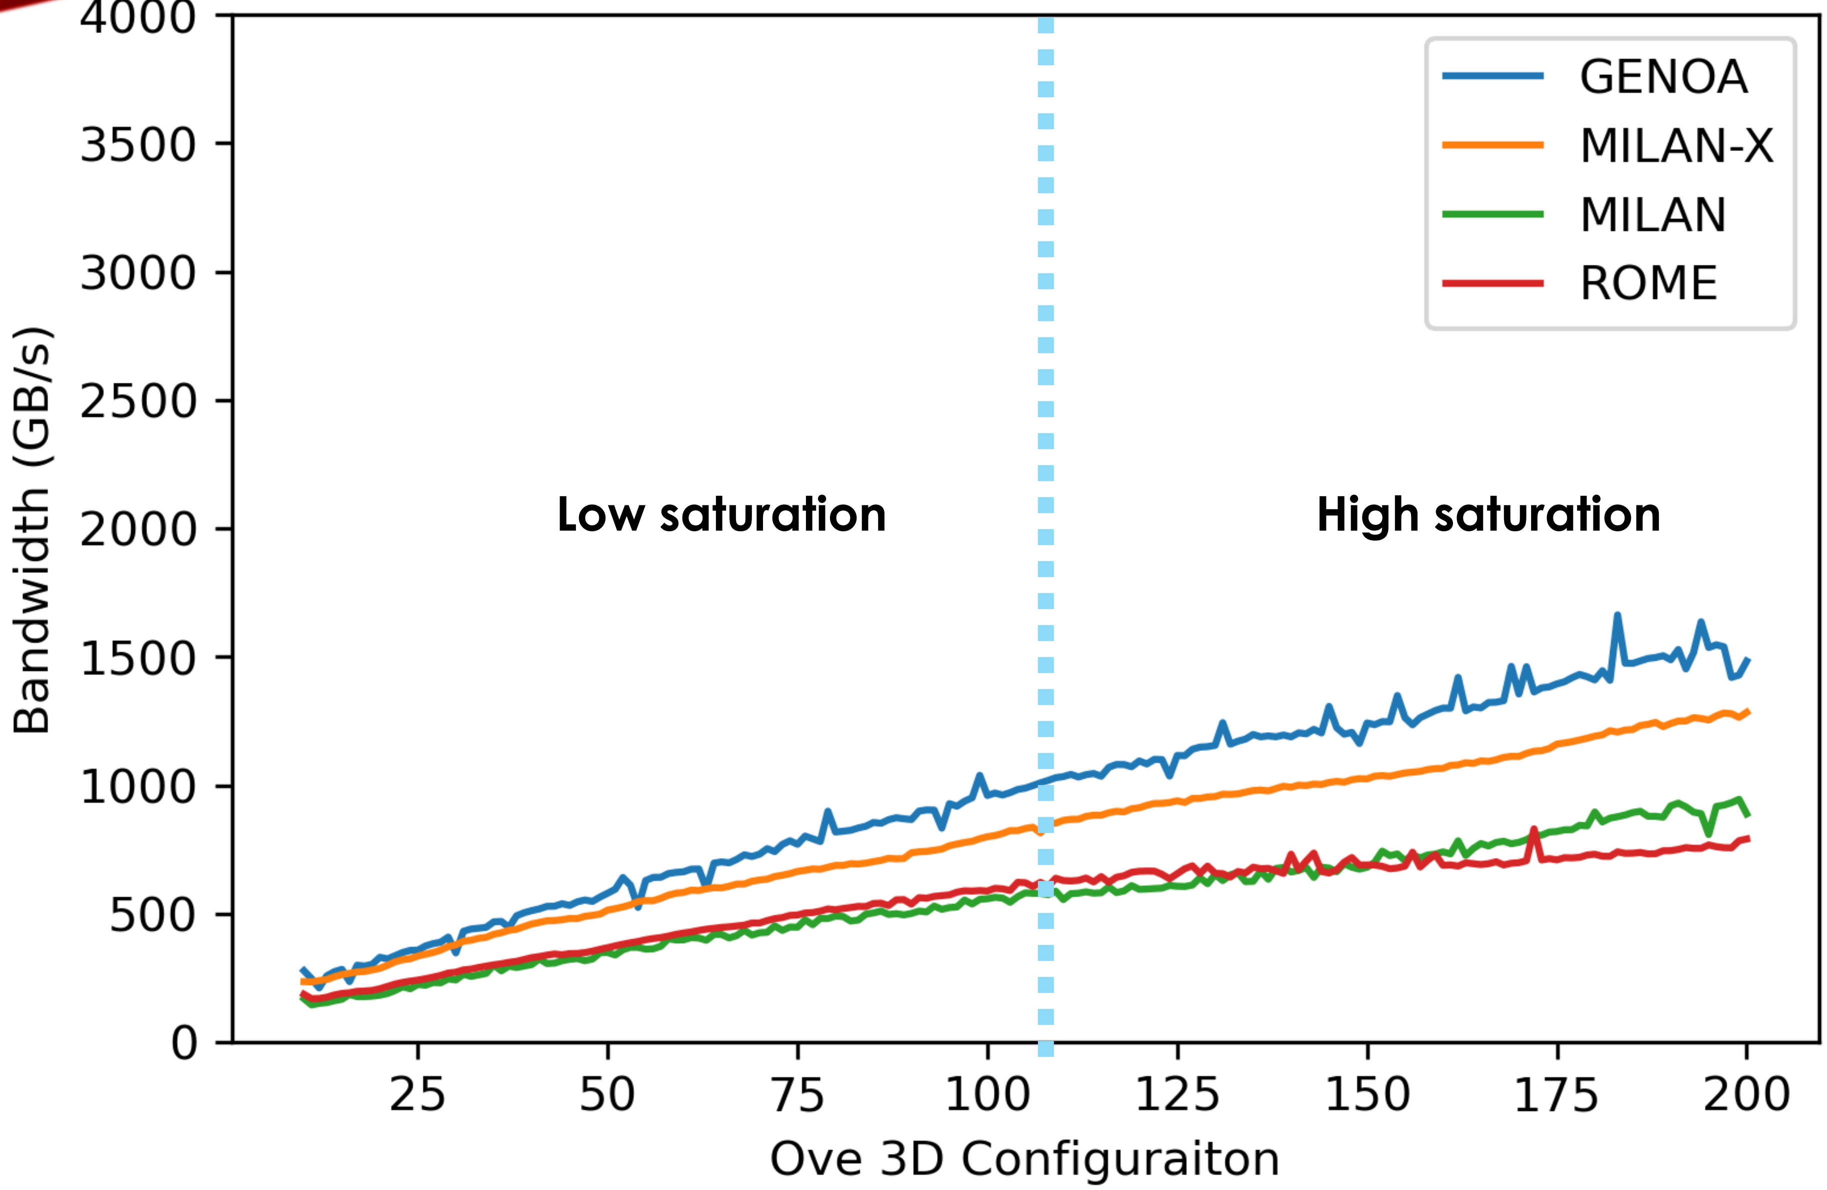

Summary

- Algorithms first!
- Low-rank matrix approximations are key for solving challenging scientific problems at scale
- Reconciling HPC workloads with the hostile hardware landscape
- Steering AI-focused hardware for HPC scientific applications is worth exploring (*ISC23 paper presentation*)
- Exploiting cache size and leveraging its high bandwidth

We are recruiting!

THANK YOU!

@HatemLtaief

Hatem.Ltaief@kaust.edu.sa

NHR PerfLab Online Seminar
May 16th 2023