国家超级计算无锡中心
National Supercomputing Center in Wuxi

Sunway TaihuLight: Designing and Tuning Scientific Applications at the Scale of 10 Million Cores

Haohuan Fu

National Supercomputing Center in Wuxi

Department of Earth System Science, Tsinghua University

September 13th 2017 @ ICAS

Outline

- Sunway Machine: the Challenges and Opportunities
- Scientific Computing with 10 Million Cores
- Long Term Plan for Sunway TaihuLight

Sunway-I:

- CMA service, 1998
- commercial chip
- 0.384 Tflops
- 48th of TOP500

Sunway BlueLight:

- NSCC-Jinan, 2011
- 16-core processor
- 1 Pflops
- 14th of TOP500

Sunway TaihuLight:

- NSCC-Wuxi, 2016
- 260-core processor
- 125 Pflops
- 1st of TOP500

The Sunway Machine Family

SW26010: Sunway 260-Core Processor

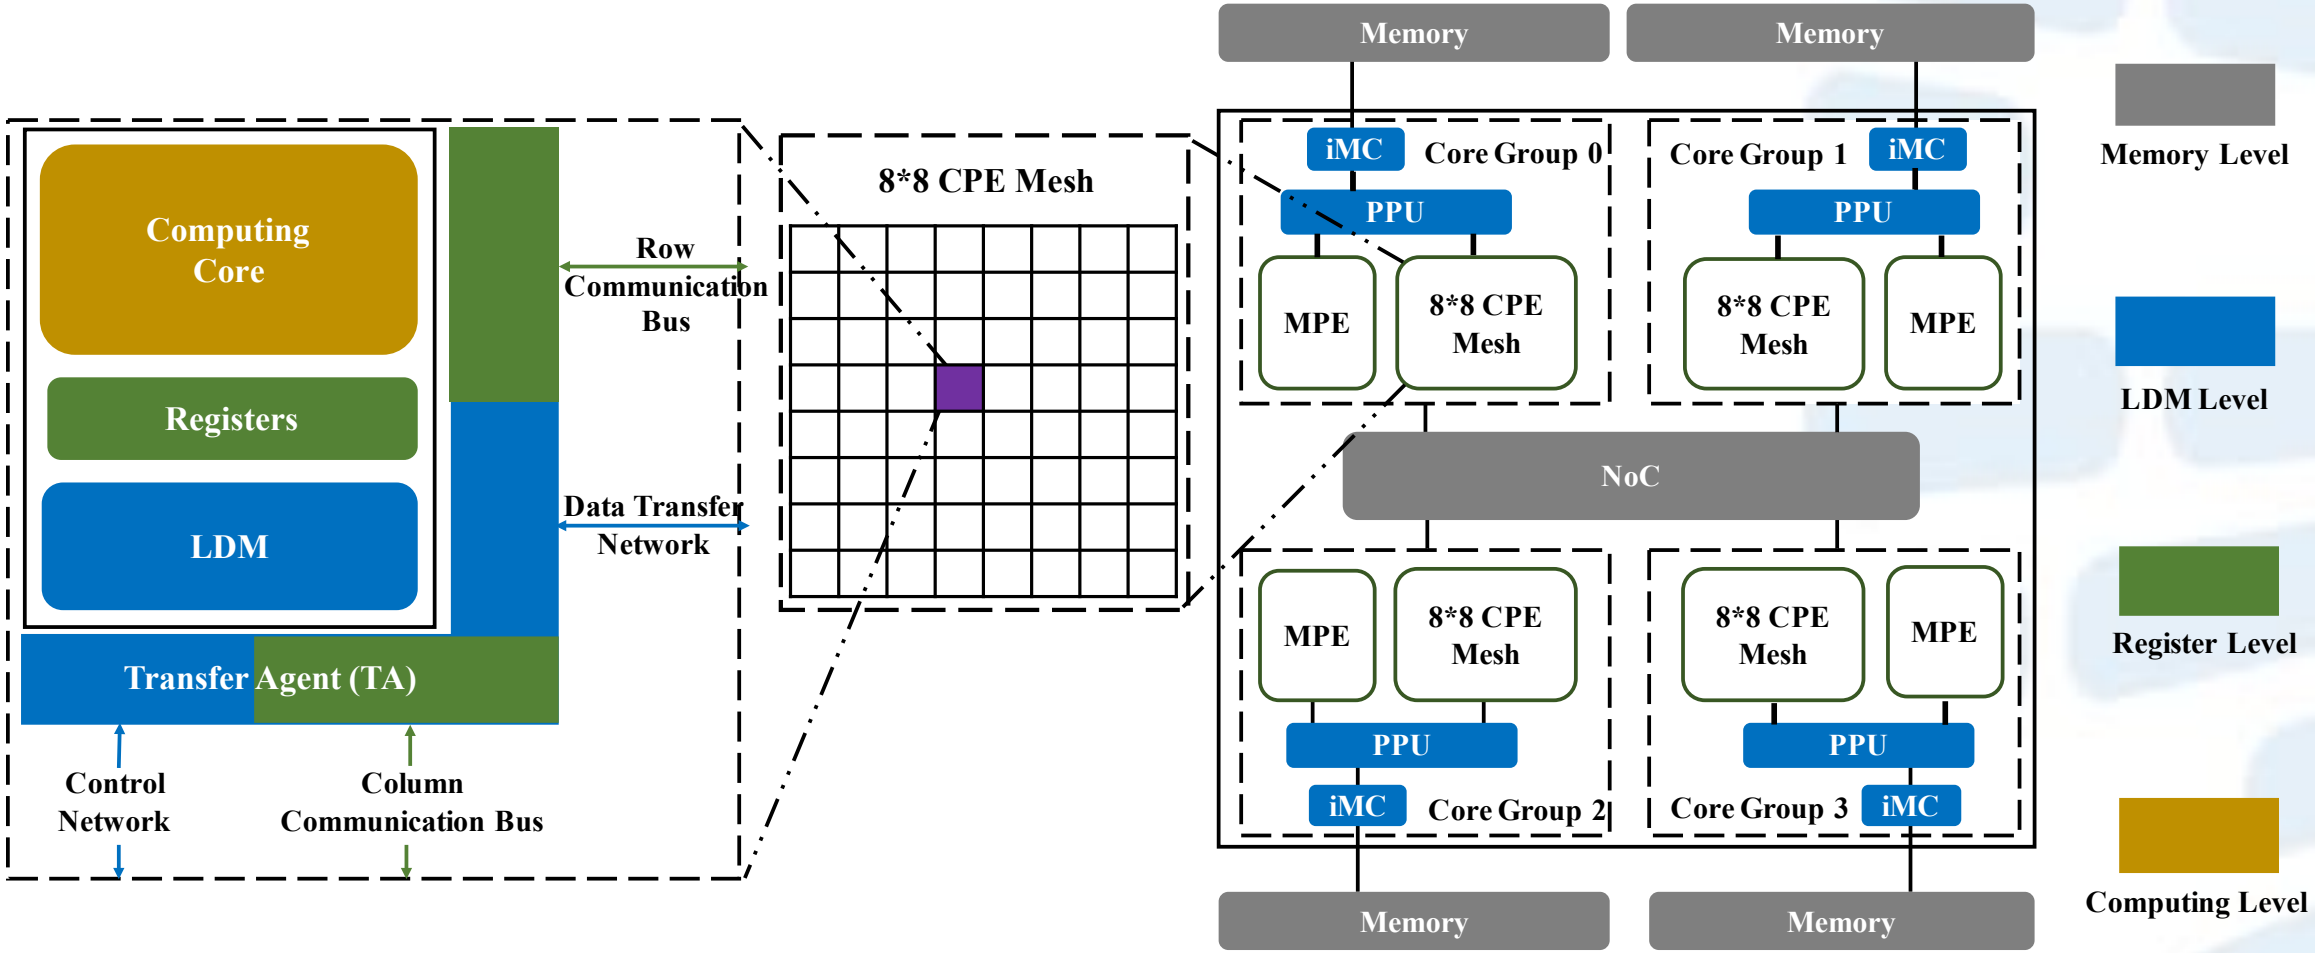

High-Density Integration of the Computing System

- A Five-Level Integration Hierarchy
  - computing node
  - computing board
  - super node
  - cabinet
  - entire computing system

High-Density Integration of the Computing System

- A Five-Level Integration Hierarchy
  - computing node
  - computing board
  - super node
  - cabinet
  - entire computing system

High-Density Integration of the Computing System

- A Five-Level Integration Hierarchy
  - computing node
  - computing board
  - super node
  - cabinet
  - entire computing system

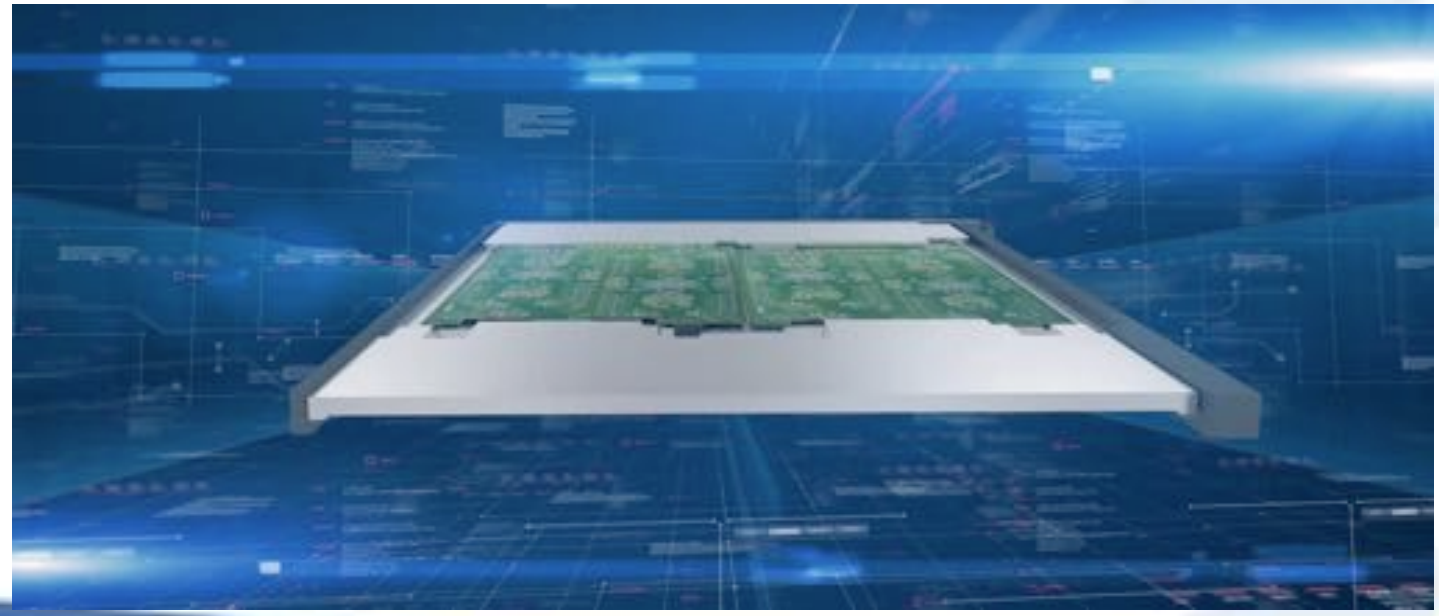

High-Density Integration of the Computing System

- A Five-Level Integration Hierarchy
  - computing node
  - computing board
  - super node
  - cabinet
  - entire computing system

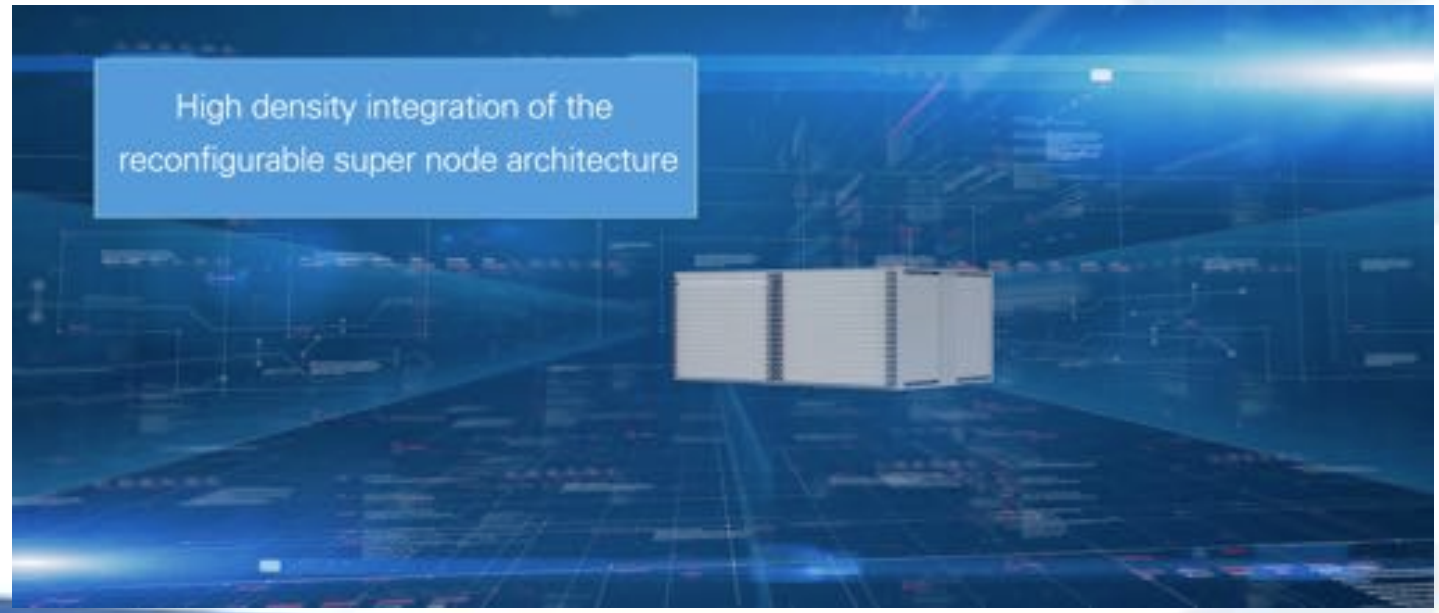

High-Density Integration of the Computing System

- A Five-Level Integration Hierarchy
  - computing node
  - computing board
  - super node
  - cabinet
  - entire computing system

How to Connect the 10 Million Cores?

$$40 \times 4 \times 256 \times 4 \times (1+8 \times 8) = 10,649,600$$

How to Connect the 10 Million Cores?

$$40 \times 4 \times 256 \times 4 \times (1 + 8 \times 8) = 10{,}649{,}600$$

2D core array with row and column buses

How to Connect the 10 Million Cores?

$$40 \times 4 \times 256 \times 4 \times (1+8 \times 8) = 10{,}649{,}600$$

2D core array with row and column buses

Network on Chip

How to Connect the 10 Million Cores?

$$40 \times 4 \times 256 \times 4 \times (1+8 \times 8) = 10{,}649{,}600$$

2D core array with row and column buses

Network on Chip

Customized Network Board to Fully Connect 256 Nodes

How to Connect the 10 Million Cores?

$$40 \times 4 \times 256 \times 4 \times (1+8 \times 8) = 10{,}649{,}600$$

2D core array with row and column buses

Network on Chip

Customized Network Board to Fully Connect 256 Nodes

Sunway Net

Tweet Comments from Prof. Satoshi Matsuoka

Satoshi Matsuoka
@ProfMatsuoka

I was quite impressed with the engineering quality of TaihuLight, different from previous Chinese machines; now truly rivals US, Japan in SC twitter.com/profmatsuoka/s…

下午4:40 - 2016年11月3日 发自 東京 目黒区

Tweet Comments from Prof. Satoshi Matsuoka

Tweet Comments from Prof. Satoshi Matsuoka

Tweet Comments from Prof. Satoshi Matsuoka

Satoshi Matsuoka
@ProfMatsuoka

Finally their design was cost&utility conscious. No expensive parts, quacky architecture, etc. Sunway apparently plans to sell the machine.

下午6:08 - 2016年11月3日

Tweet Comments from Prof. Satoshi Matsuoka

Satoshi Matsuoka
@ProfMatsuoka

Finally their design was cost&utility conscious. No expensive parts, quacky architecture, etc. Sunway apparently plans to sell the machine.

下午6:08 - 2016年11月3日

Sunway Micro

Outline

- Sunway Machine: the Challenges and Opportunities
- Scientific Computing with 10 Million Cores
- Long Term Plan for Sunway TaihuLight

Machine Capability Comparison

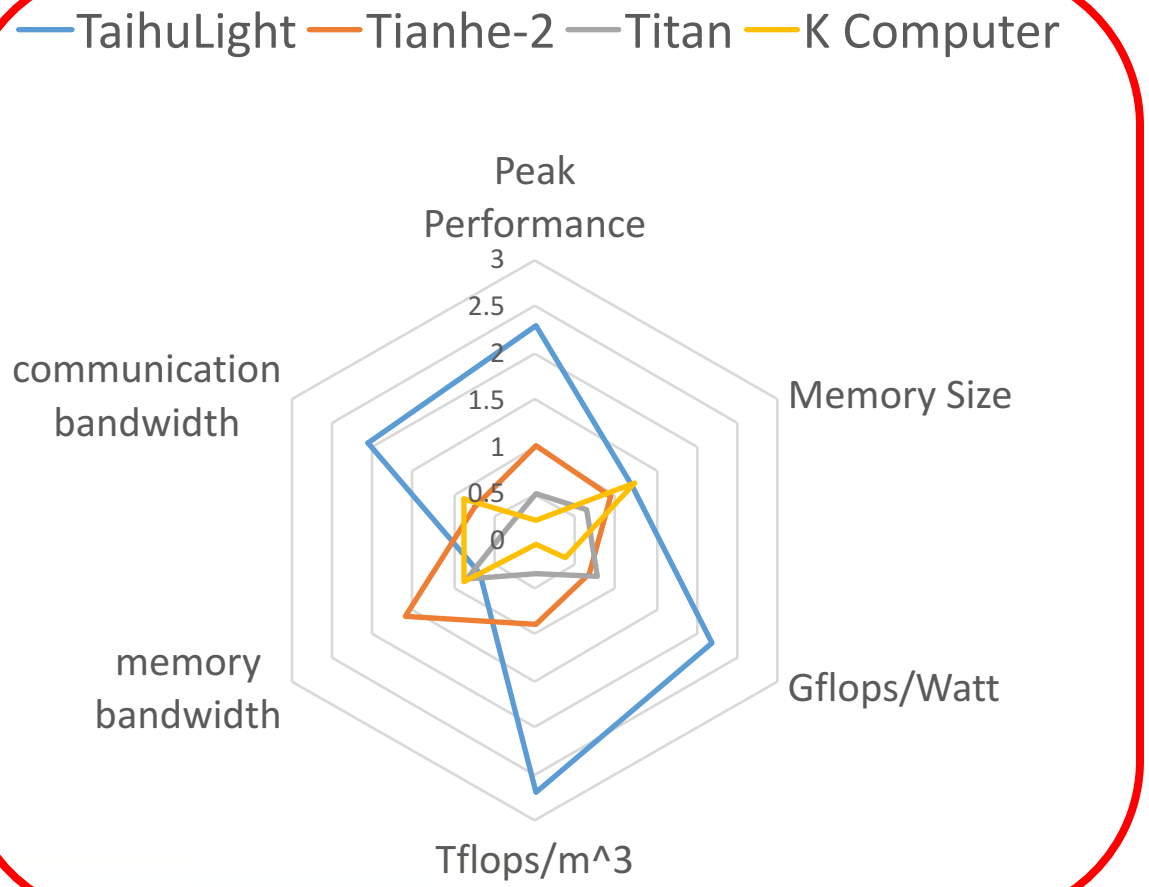

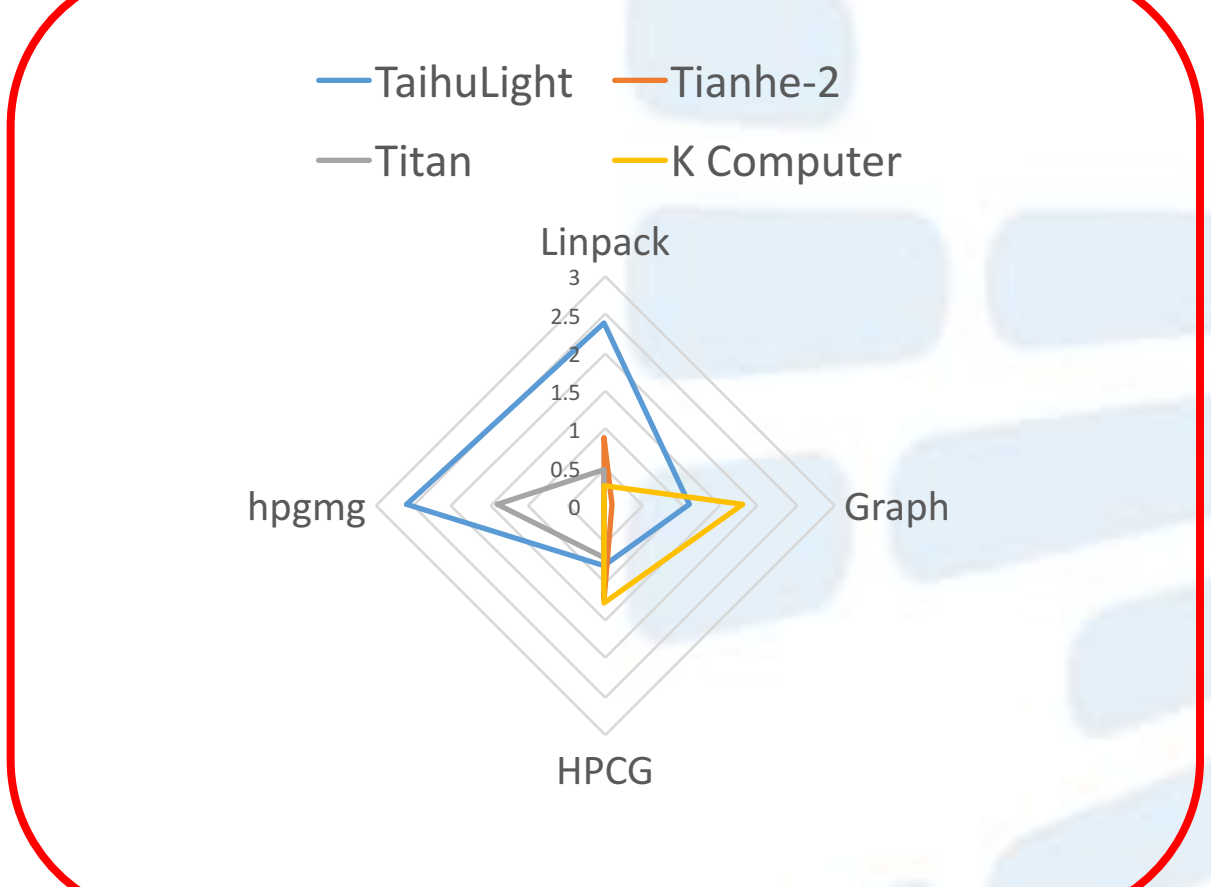

Major Features to Consider

Sunway TaihuLight

125 Pflops		10 million cores	user-controlled 64 KB LDM
32 GB and 136GB/s per node	22 flops/byte	MPE + CPE	register communication among CPEs

Major Features to Consider

Sunway TaihuLight

125 Pflops		10 million cores	user-controlled 64 KB LDM
32 GB and 136GB/s per node	22 flops/byte	MPE + CPE	register communication among CPEs

Intel KNL 7250 of Cori: 6.5 flops/byte
NVIDIA P100 of Piz Daint: 7.2 flops/byte

Major Challenge #1: Scaling

Sunway TaihuLight

125 Pflops		10 million cores	user-controlled 64 KB LDM
32 GB and 136GB/s per node	22 flops/byte	MPE + CPE	register communication among CPEs

Major Challenge #2: Memory Wall

Sunway TaihuLight

125 Pflops	10 million cores	user-controlled 64 KB LDM
32 GB and 136GB/s per node	22 flops/byte	MPE + CPE	register communication among CPEs

Major Challenge #2: Memory Wall

Sunway TaihuLight

125 Pflops

10 million cores

user-controlled 64 KB LDM

32 GB and 136GB/s per node

22 flops/byte

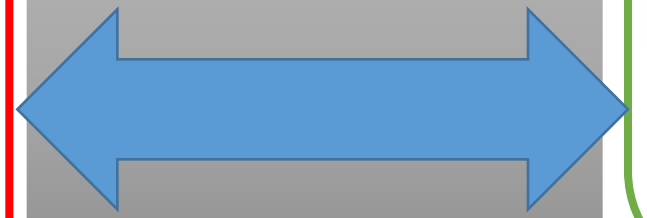

register communication among CPEs

Refactoring and Redesigning

An (Incomplete) List of Full-Scale Applications

2016	2017
Fully Implicit Solver for Atmospheric Dynamics	Extreme-scale Graph Processing Framework
Surface Wave Modeling	Simulation of Planetary Rings
Phase Field Simulations of Coarsening Dynamics	Simulations of Quantum Spin Liquid States via PEPS++
Atomistic Simulation of Silicon Nanowires	Molecular Dynamics Simulation of Condensed Covalent Materials
Run-away Electron Trajectory Simulation	cryo-EM Macromolecule Structure Determination
Genome Functional Annotation and Homeotic Gene Building	Redesigning CAM-SE
Spacecraft CFD Numerical Simulation	Nonlinear Earthquake Simulation

An (Incomplete) List of Full-Scale Applications

2016 Gordon Bell Finalists	2017
Fully Implicit Solver for Atmospheric Dynamics	Extreme-scale Graph Processing Framework
Surface Wave Modeling	Simulation of Planetary Rings
Phase Field Simulations of Coarsening Dynamics	Simulations of Quantum Spin Liquid States via PEPS++
Atomistic Simulation of Silicon Nanowires	Molecular Dynamics Simulation of Condensed Covalent Materials
Run-away Electron Trajectory Simulation	cryo-EM Macromolecule Structure Determination
Genome Functional Annotation and Homeotic Gene Building	Redesigning CAM-SE
Spacecraft CFD Numerical Simulation	Nonlinear Earthquake Simulation

An (Incomplete) List of Full-Scale Applications

2016 Gordon Bell Prize	2017
Fully Implicit Solver for Atmospheric Dynamics	Extreme-scale Graph Processing Framework
Surface Wave Modeling	Simulation of Planetary Rings
Phase Field Simulations of Coarsening Dynamics	Simulations of Quantum Spin Liquid States via PEPS++
Atomistic Simulation of Silicon Nanowires	Molecular Dynamics Simulation of Condensed Covalent Materials
Run-away Electron Trajectory Simulation	cryo-EM Macromolecule Structure Determination
Genome Functional Annotation and Homeotic Gene Building	Redesigning CAM-SE
Spacecraft CFD Numerical Simulation	Nonlinear Earthquake Simulation

An (Incomplete) List of Full-Scale Applications

2016 Gordon Bell Prize

- Fully Implicit Solver for Atmospheric Dynamics
- Surface Wave Modeling
- Phase Field Simulations of Coarsening Dynamics
- Atomistic Simulation of Silicon Nanowires
- Run-away Electron Trajectory Simulation
- Genome Functional Annotation and Homeotic Gene Building
- Spacecraft CFD Numerical Simulation

2017 Gordon Bell Finalists

- Extreme-scale Graph Processing Framework
- Simulation of Planetary Rings
- Simulations of Quantum Spin Liquid States via PEPS++
- Molecular Dynamics Simulation of Condensed Covalent Materials
- cryo-EM Macromolecule Structure Determination
- Redesigning CAM-SE
- Nonlinear Earthquake Simulation

163,840 processes | 65 threads

racks | *chips* | *core-groups* | *cores* | *total number of cores*

$$40 \times 1{,}024 \times 4 \times 65 = 10{,}649{,}600$$

DD-MG $\mathcal{K}$-cycle

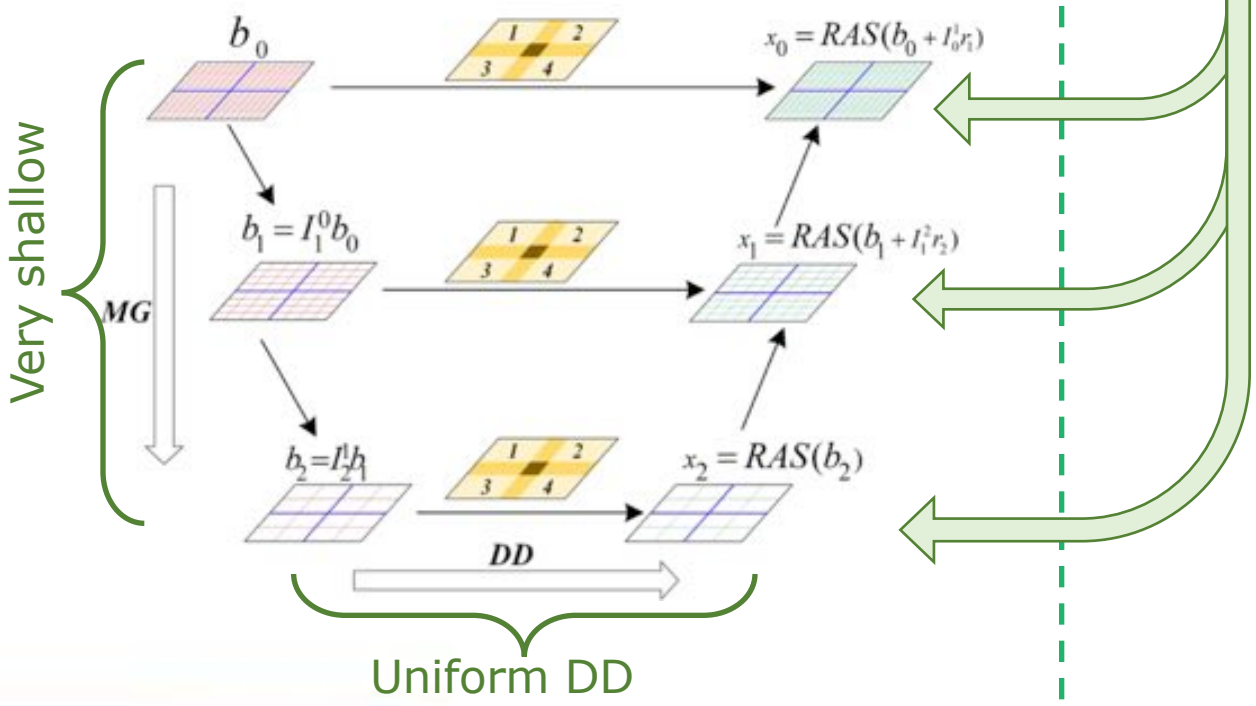

Plug & Play

Now let's find a way to design a subdomain solver.

163,840 processes

65 threads

racks *chips* *core-groups* *cores* *total number of cores*

$$40 \times 1{,}024 \times 4 \times 65 = 10{,}649{,}600$$

DD-MG $\mathcal{K}$-cycle

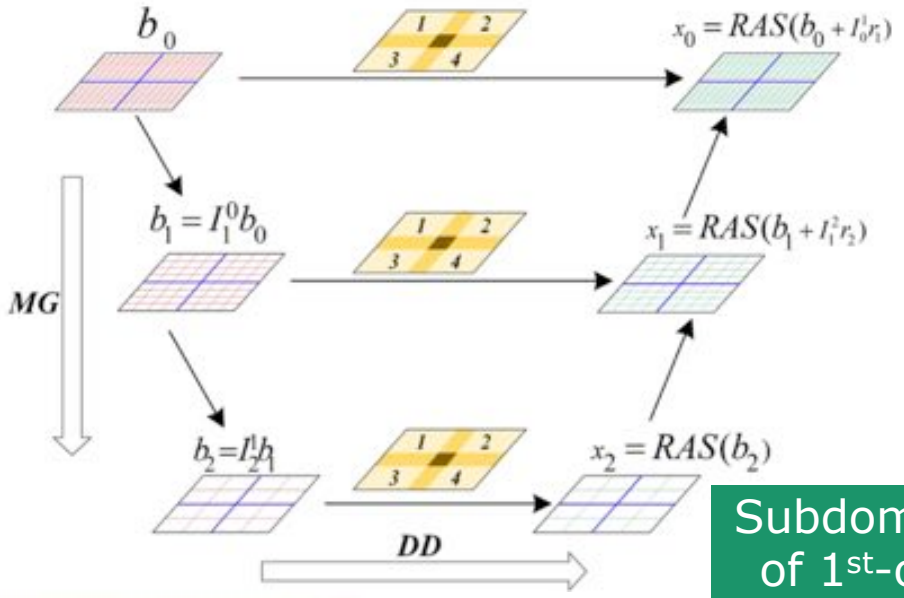

Geometry-based pipelined ILU (GP-ILU)

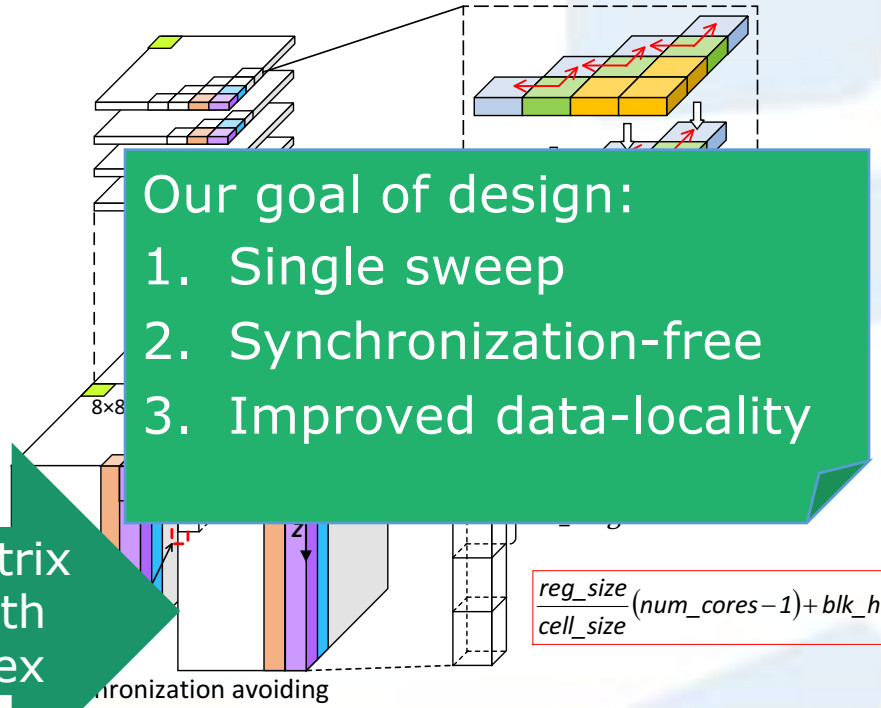

Our goal of design:

1. Single sweep
2. Synchronization-free
3. Improved data-locality

Subdomain matrix of 1st-order with geometric index

$$\frac{reg\_size}{cell\_size}(num\_cores-1)+blk\_height<dim\_z$$

Strong-scaling results

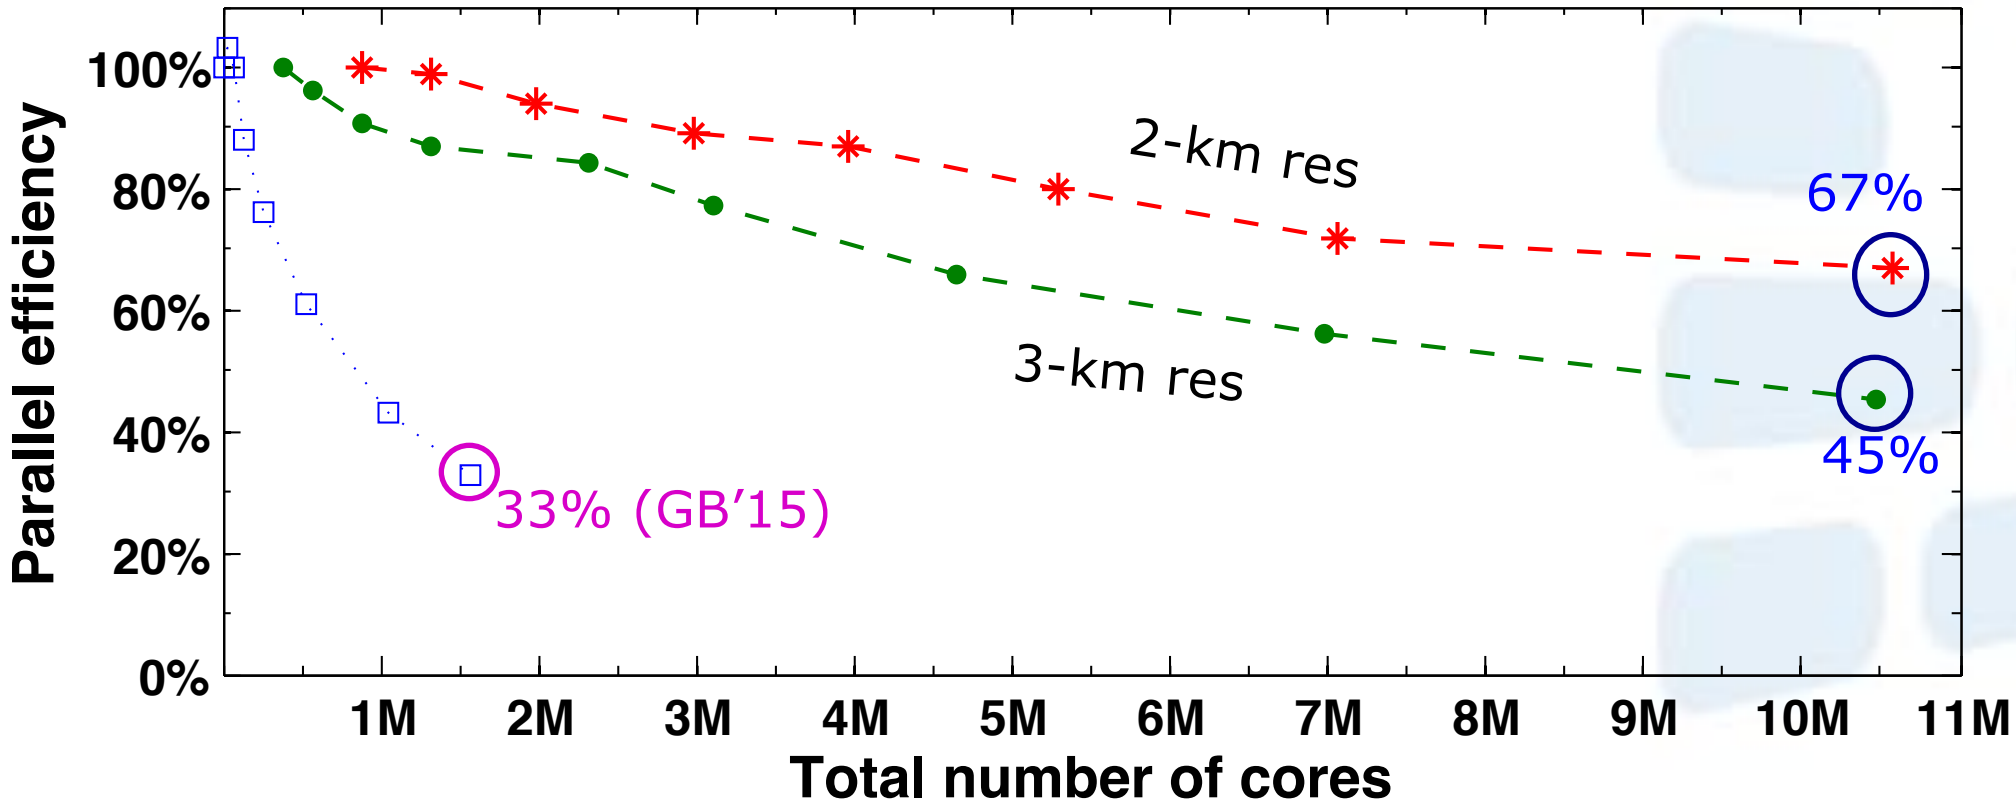

The 3-km res run: 1.01 SYPD with 10.6M cores, dt=240s, I/O penalty <5%

Weak-scaling results

The 488-m res run: 0.07 SYPD, 10.6M cores, dt=240s, 89.5X speedup over explicit

Application (II): Porting CESM and Redesigning CAM-SE for Sunway TaihuLight

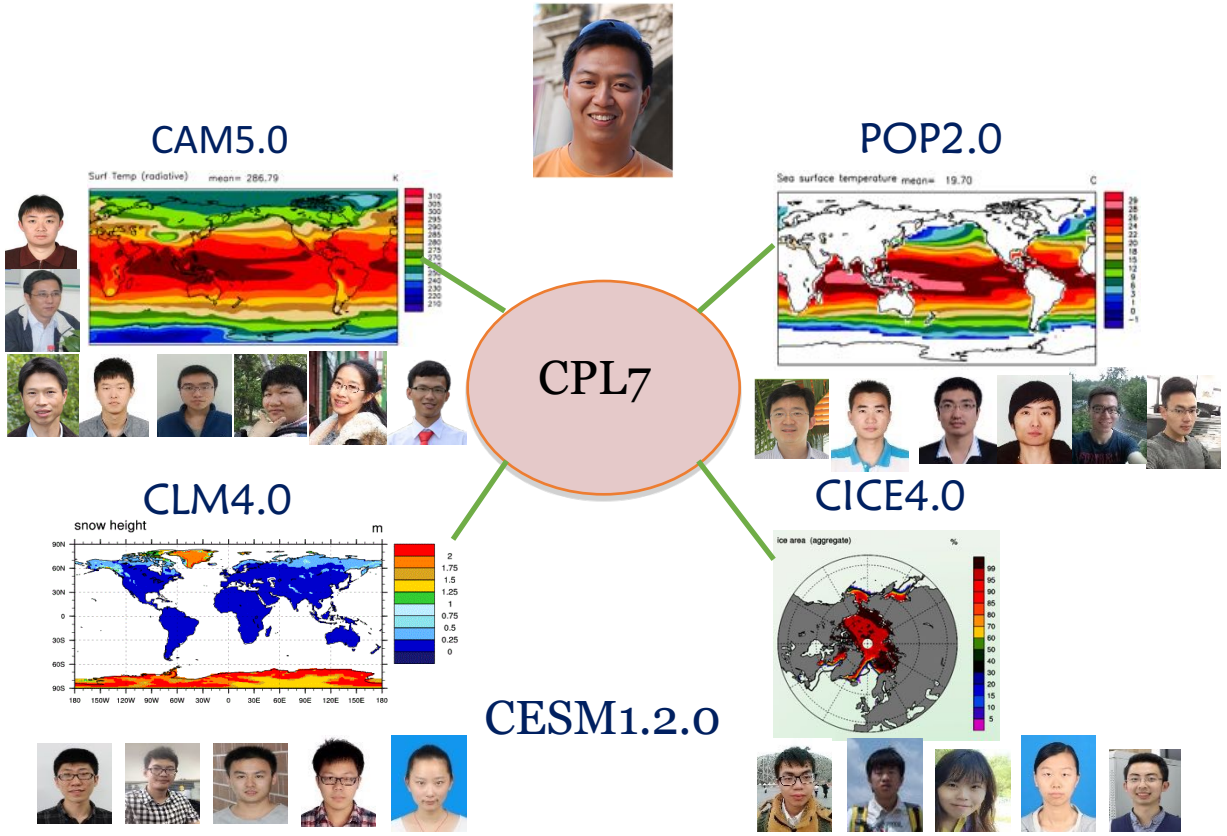

Tsinghua + BNU 30+ Professors and Students

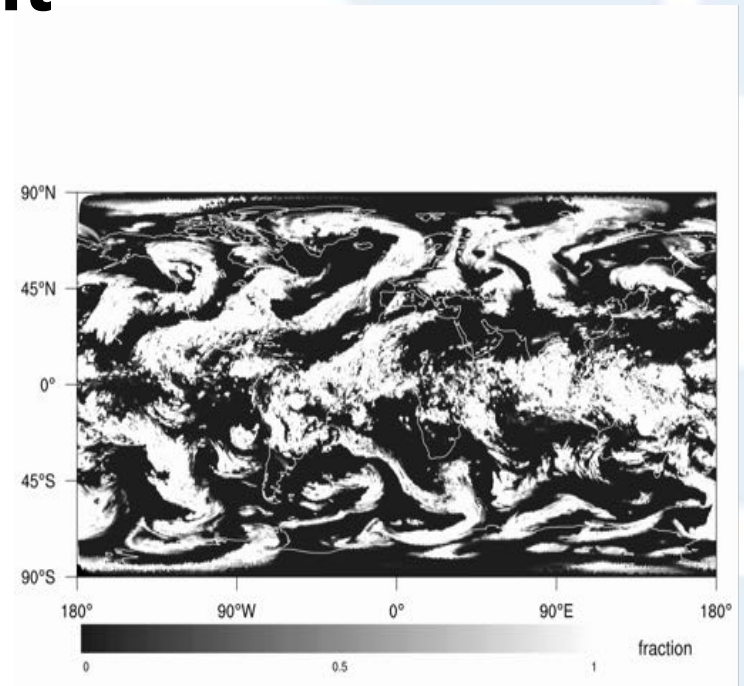

- Four component models, millions lines of code
- Large-scale run on Sunway TaihuLight
  - 24,000 MPI processes
  - Over one million cores
- 10-20x speedup for kernels
- 2-3x speedup for the entire model

"Refactoring and Optimizing the Community Atmosphere Model (CAM) on the Sunway TaihuLight Supercomputer", in Proceedings of SC 2016.

Application (II): Porting CESM and Redesigning CAM-SE for Sunway TaihuLight

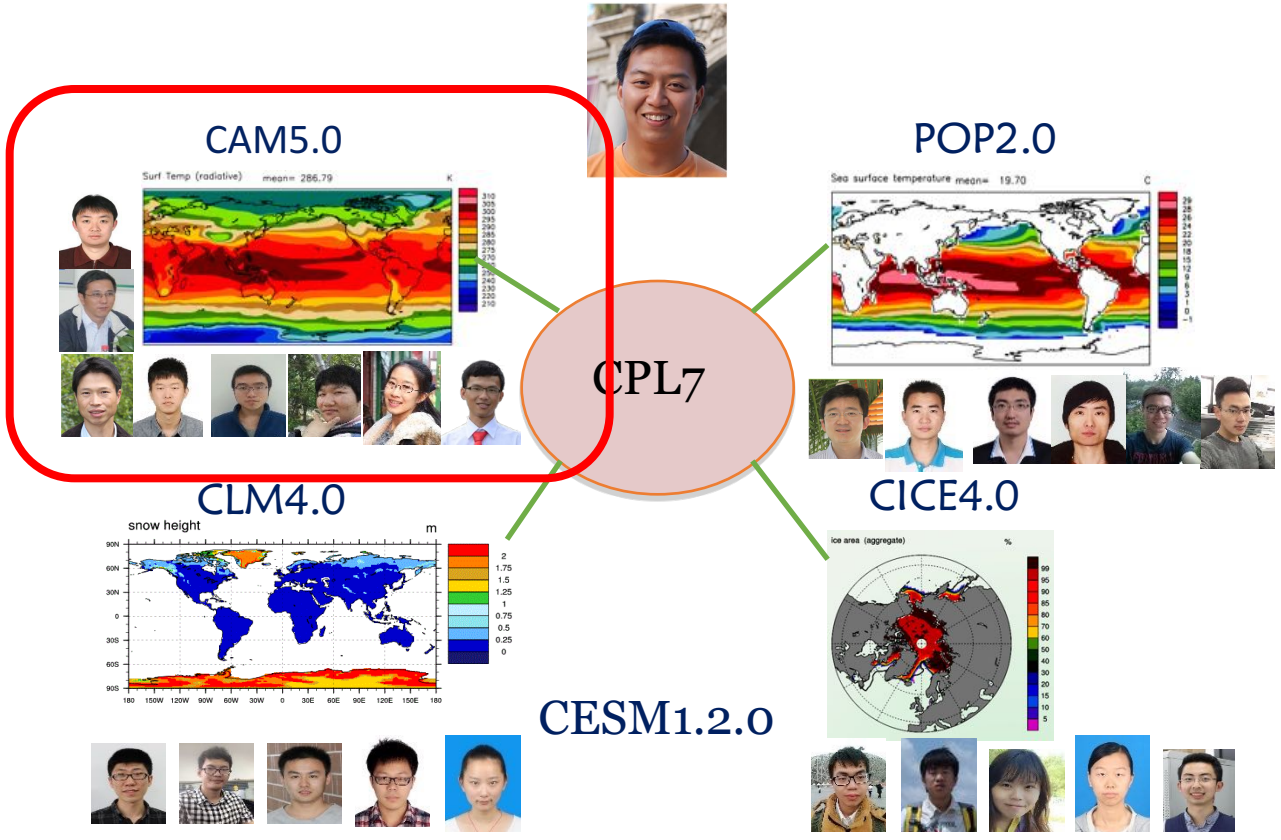

Tsinghua + BNU 30+ Professors and Students

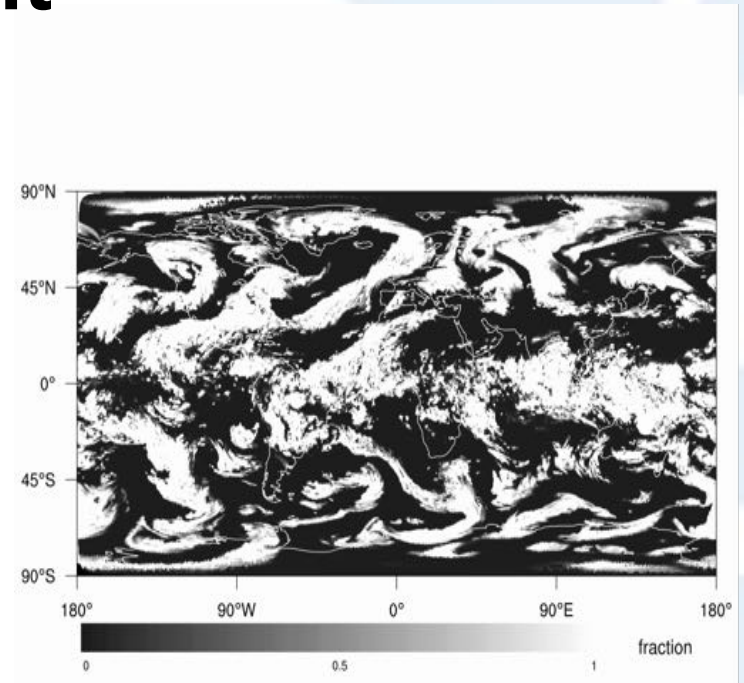

- Four component models, millions lines of code
- Large-scale run on Sunway TaihuLight
  - 24,000 MPI processes
  - Over one million cores
- 10-20x speedup for kernels
- 2-3x speedup for the entire model

"Refactoring and Optimizing the Community Atmosphere Model (CAM) on the Sunway TaihuLight Supercomputer", in Proceedings of SC 2016.

Major Challenges

a high complexity in application, and a heavy legacy in the code base (millions lines of code)

an extremely complicated MPMD program with no hotspots (or hundreds of hotspots)

misfit between the in-place design philosophy and the new architecture

lack of people with interdisciplinary knowledge and experience

OpenACC-based Refactoring of CAM

Pass tracers (u, v) to dynamics

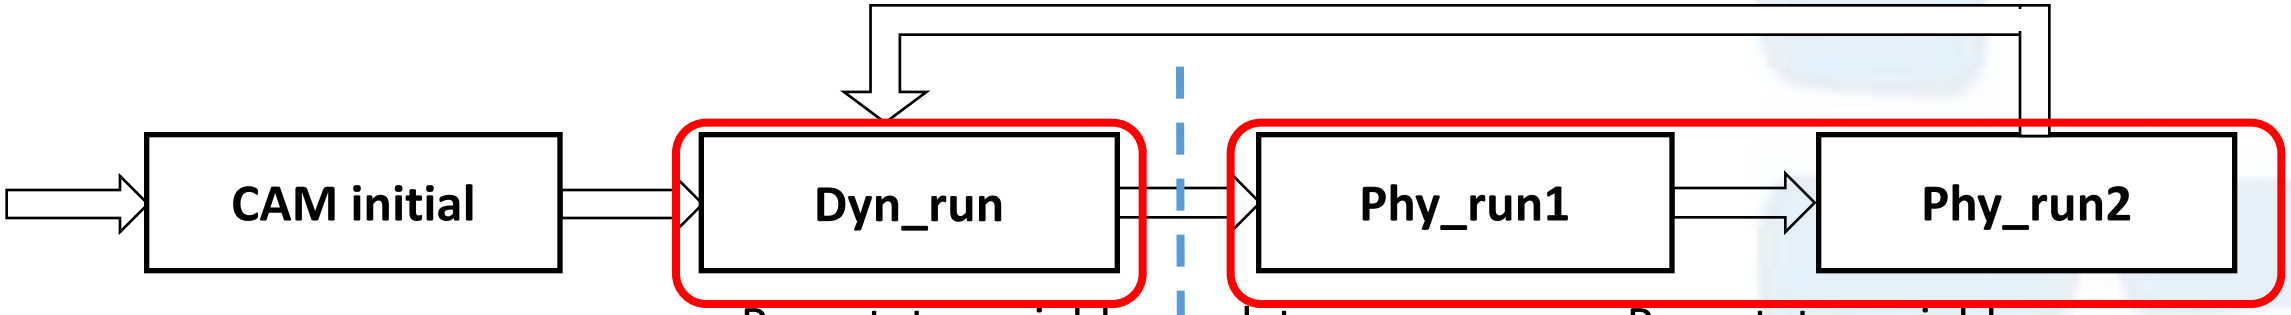

Pass state variables and tracers

Pass state variables

- manual transformation of loops
- manual OpenACC parallelization and optimization on code and data structures

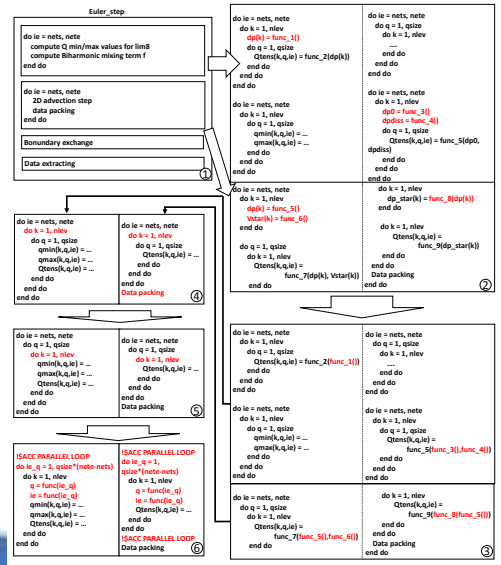

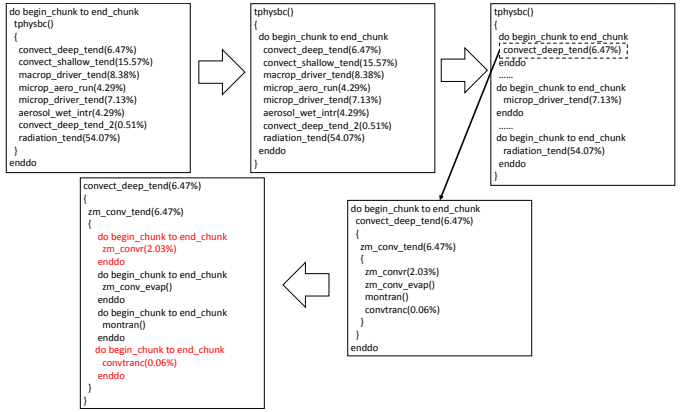

- tool based transformation of loops

CAM model: scalability and speedup

- million core scale, 2.81 SYPD
- many-core refactoring for the entire model
- competitive simulation speed to the same model on NCAR Yellowstone

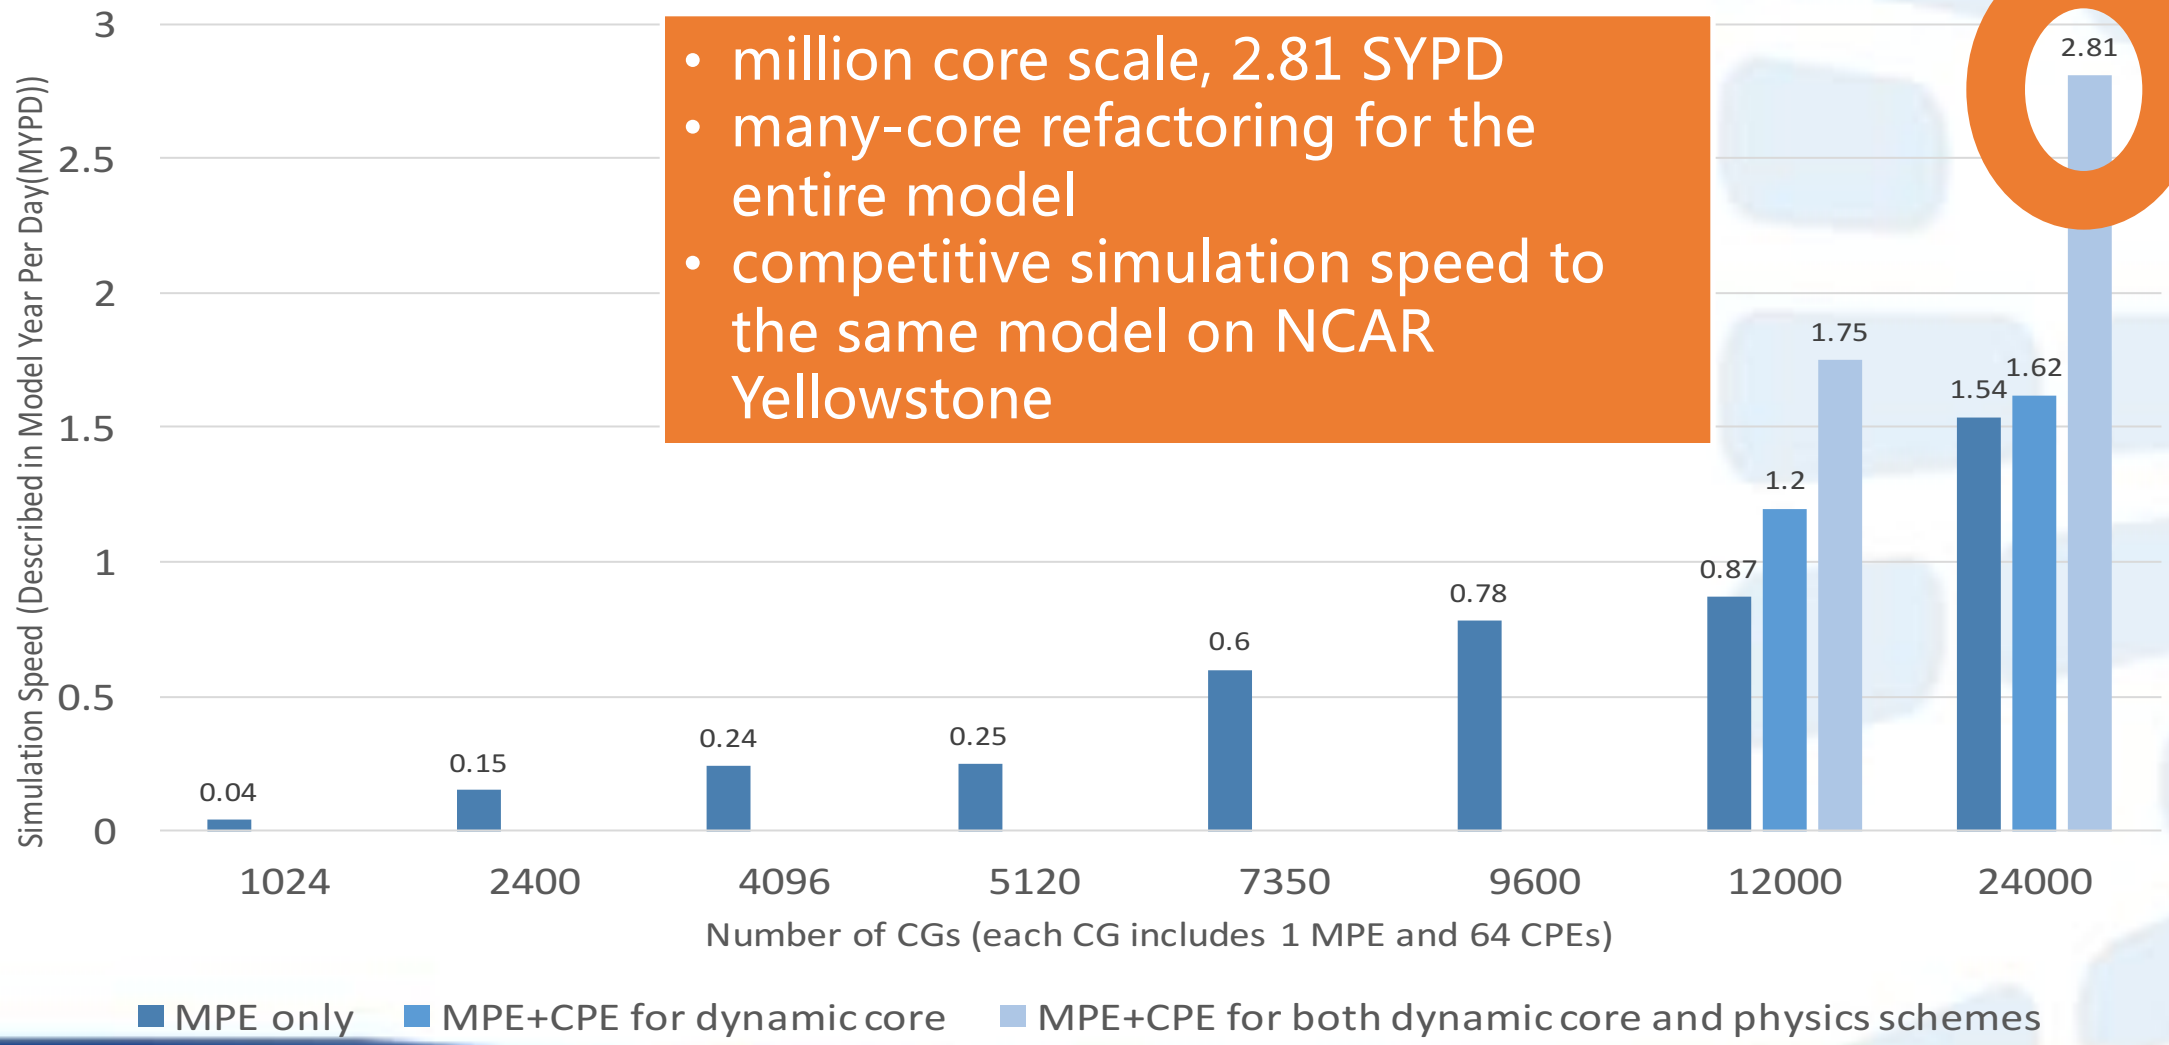

Athread-based Fine-grained Redesign

- Step 1: rewrite of Fortran OpenACC code to Athread C code
  - finer memory control through a specific DMA scheme
  - more efficient vectorization

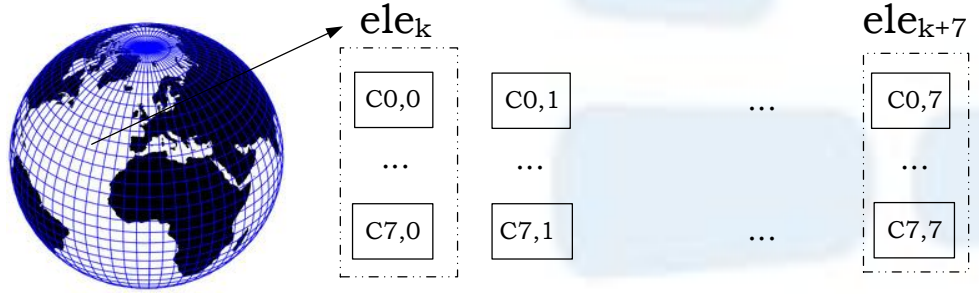

- Step 2: register-communication based redesign
  - remove data dependency
  - expose more parallelism

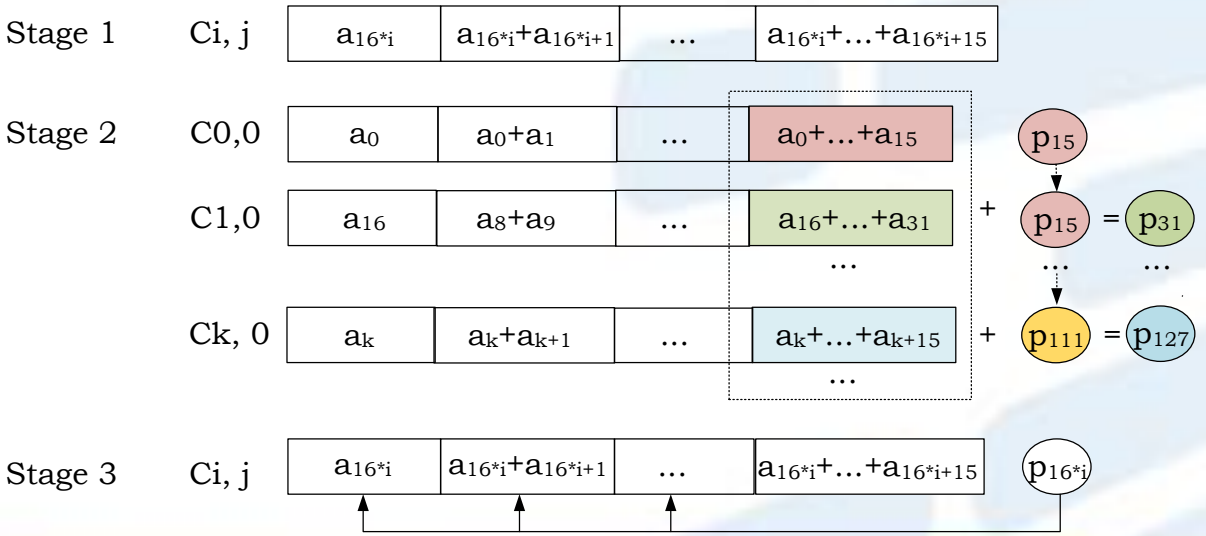

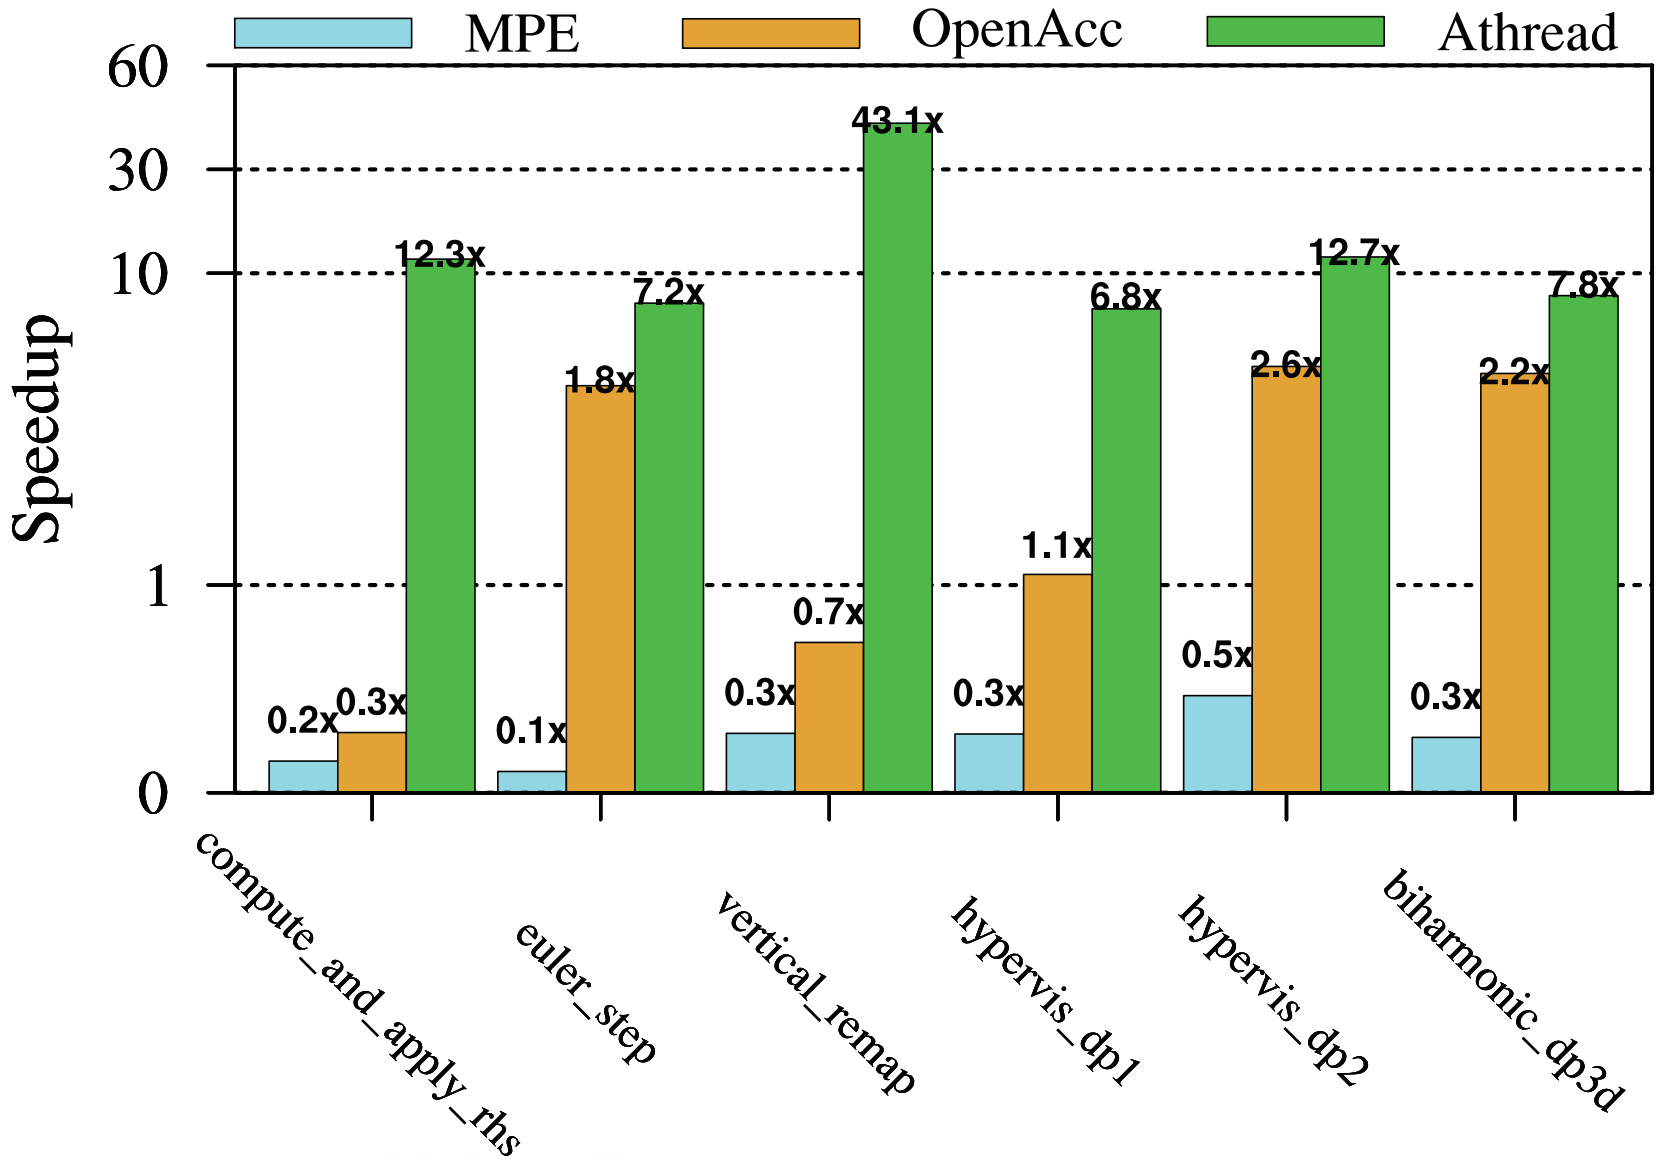

1 Sunway CG (64 CPEs) could be equivalent to
0.1x Intel Core
or
1.8x Intel Core
or
7.2x Intel Core
or in certain cases
43.1x Intel Core

Performance Improvement through Redesign

Scaling the Dynamic Core to Millions of Cores

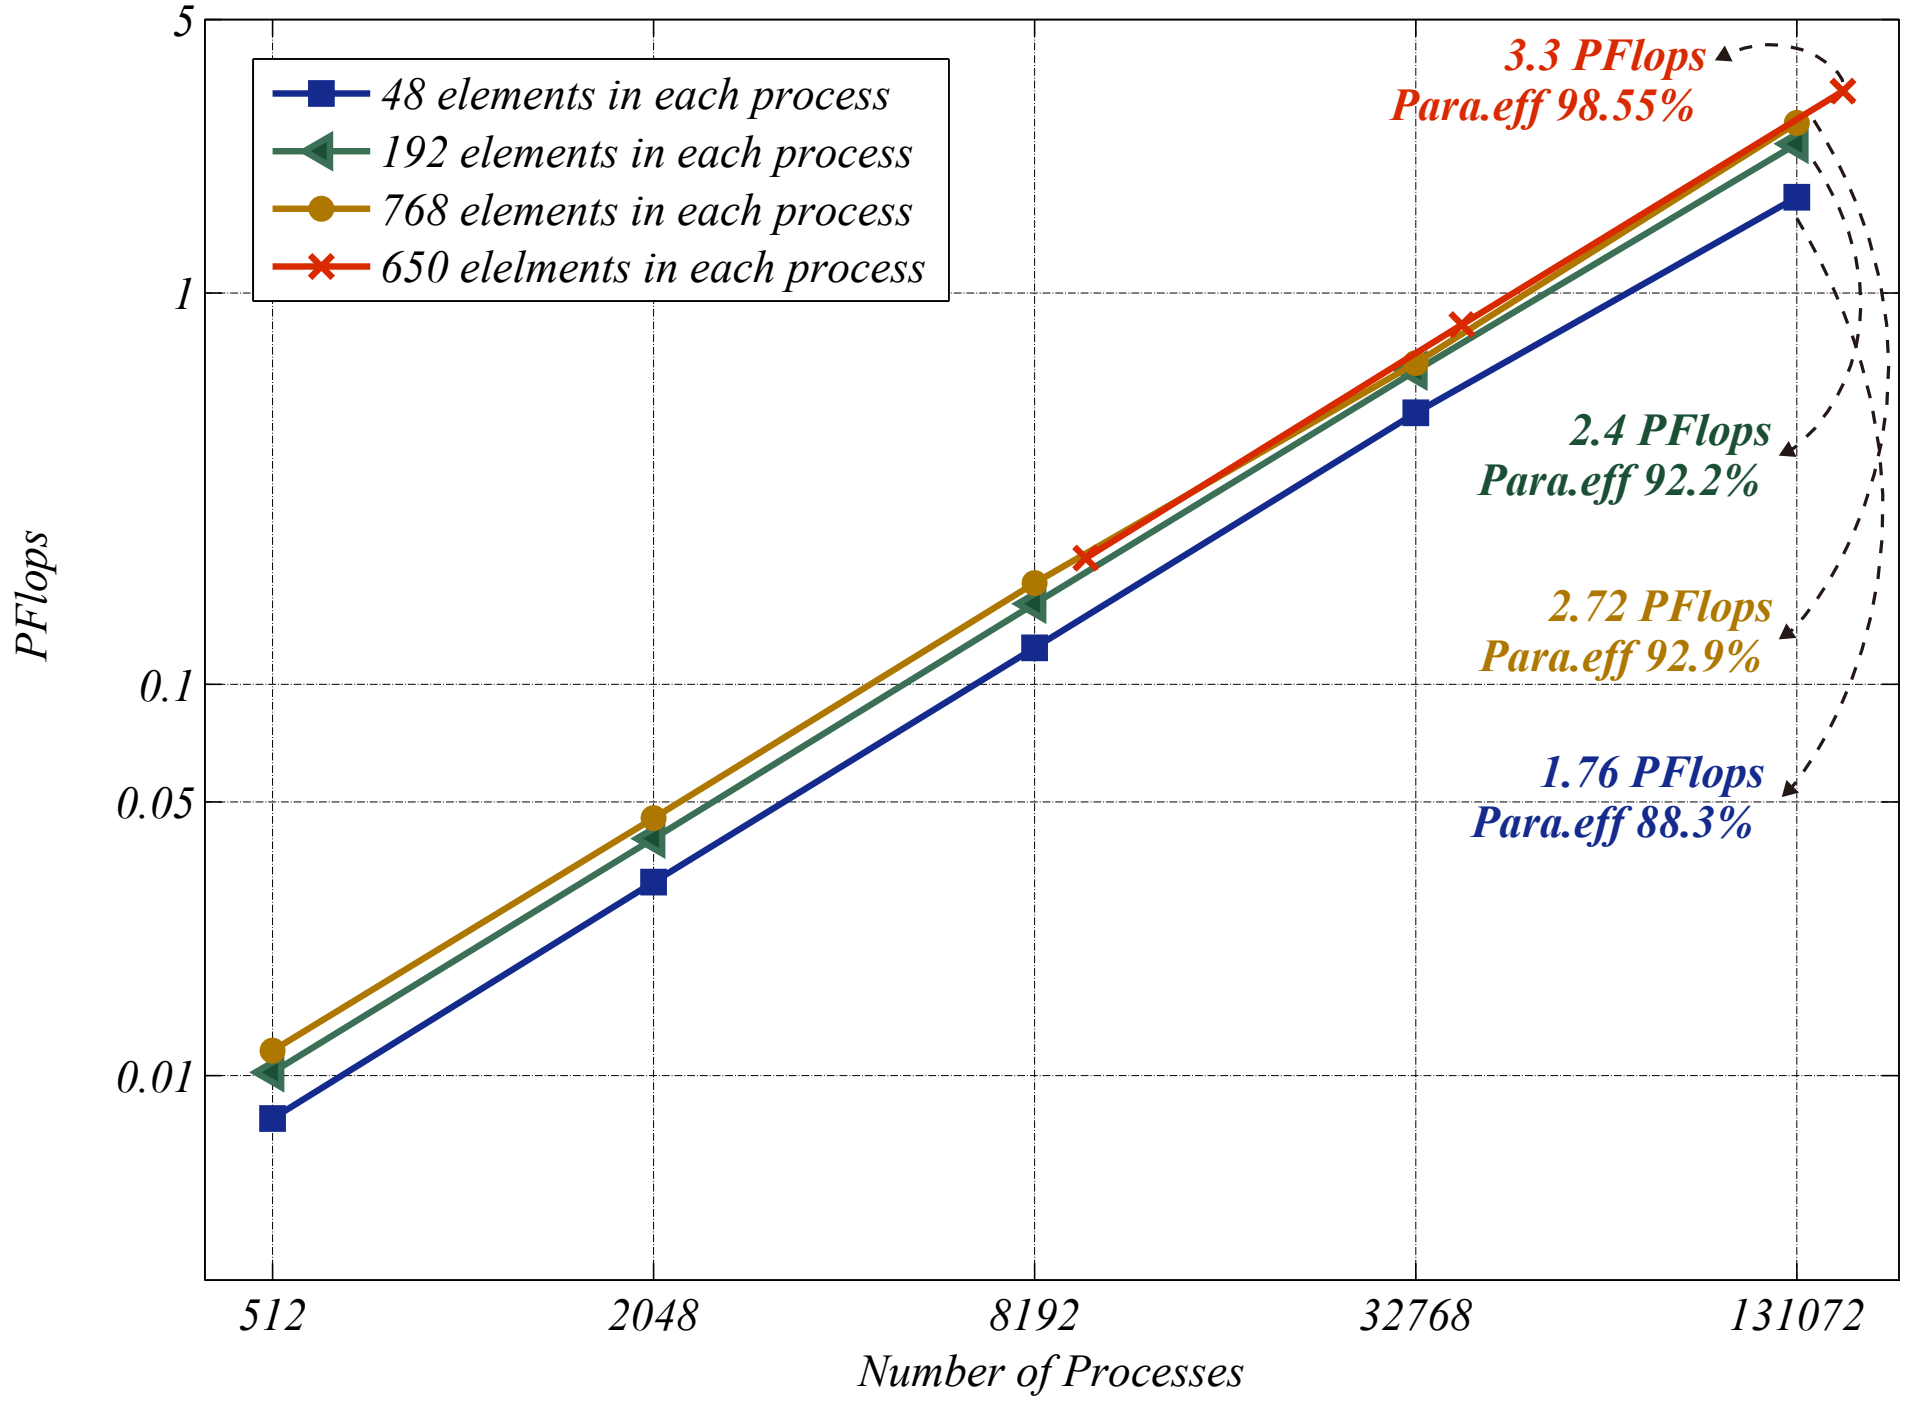

Simulation of Hurricane Katrina

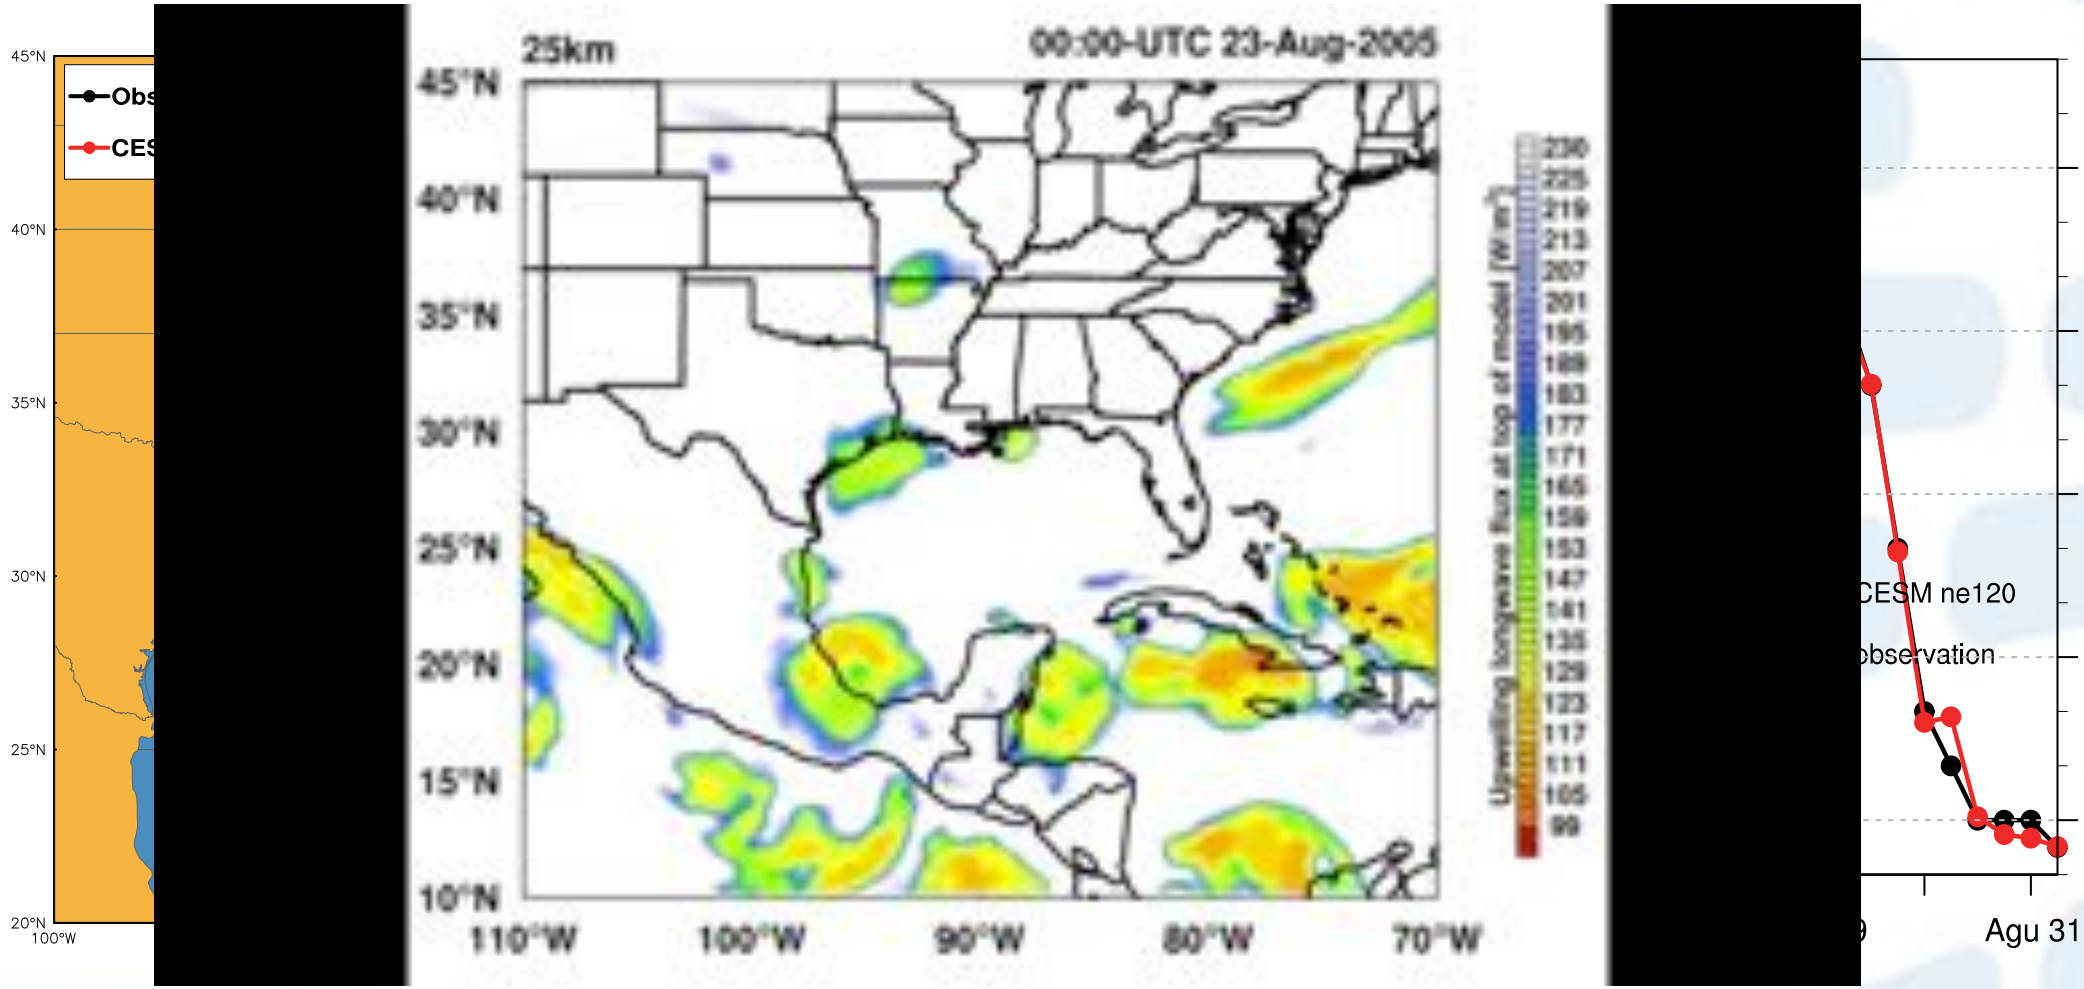

Application (III): Nonlinear Earthquake Simulation on Sunway TaihuLight

- Dynamic rupture source generator (originated from CG-FDM)
- Seismic wave propagation (originated from AWP-ODC)
- Other utilities:
  - source partitioner
  - 3D Model Interpolator
  - Restart controller

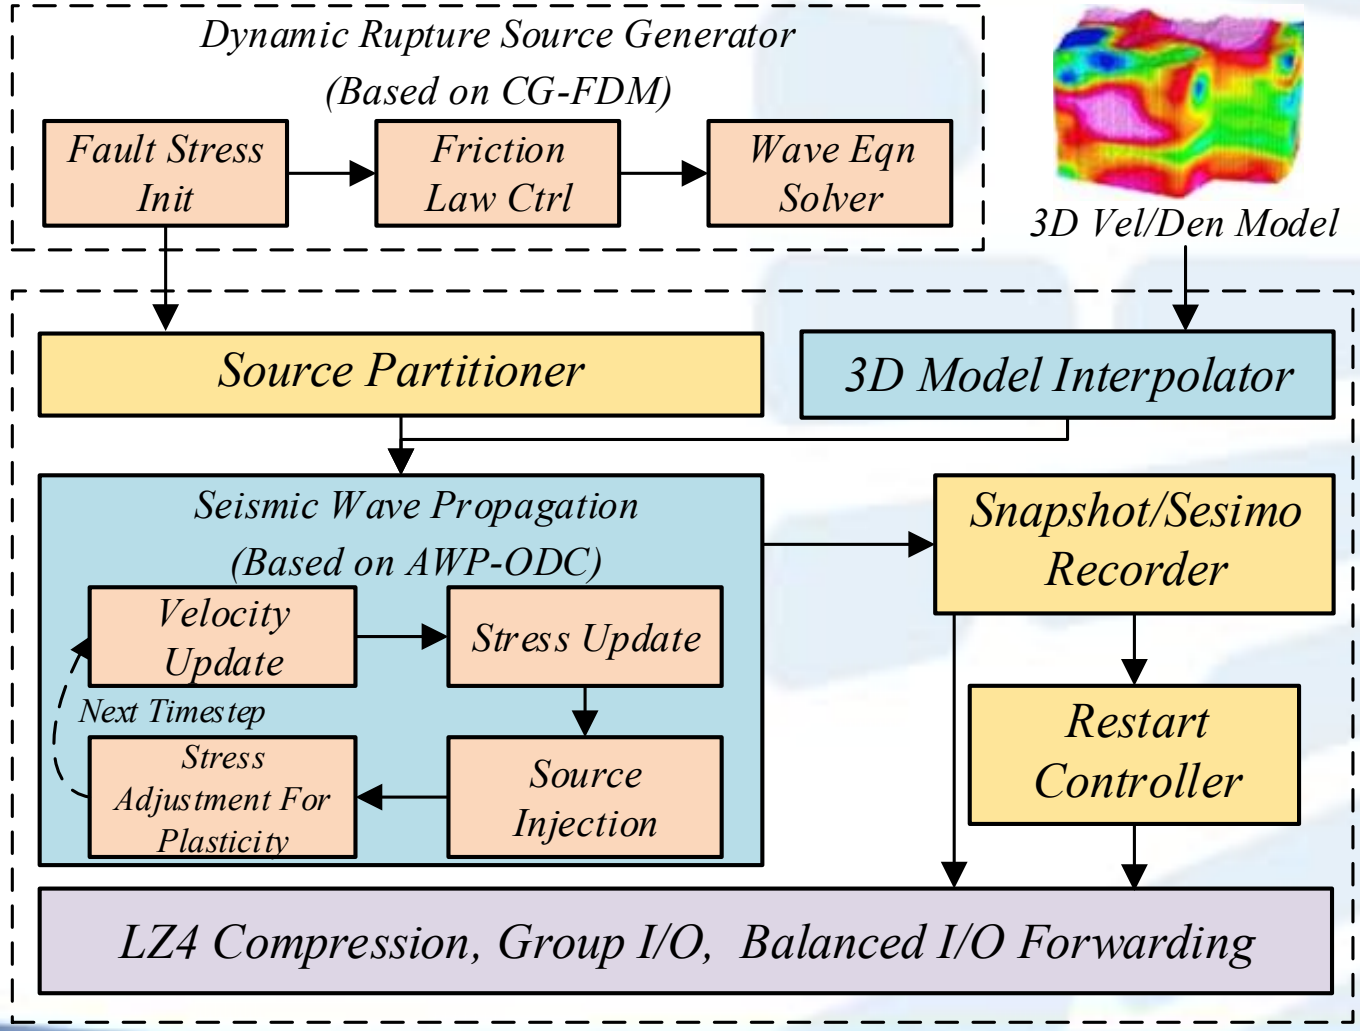

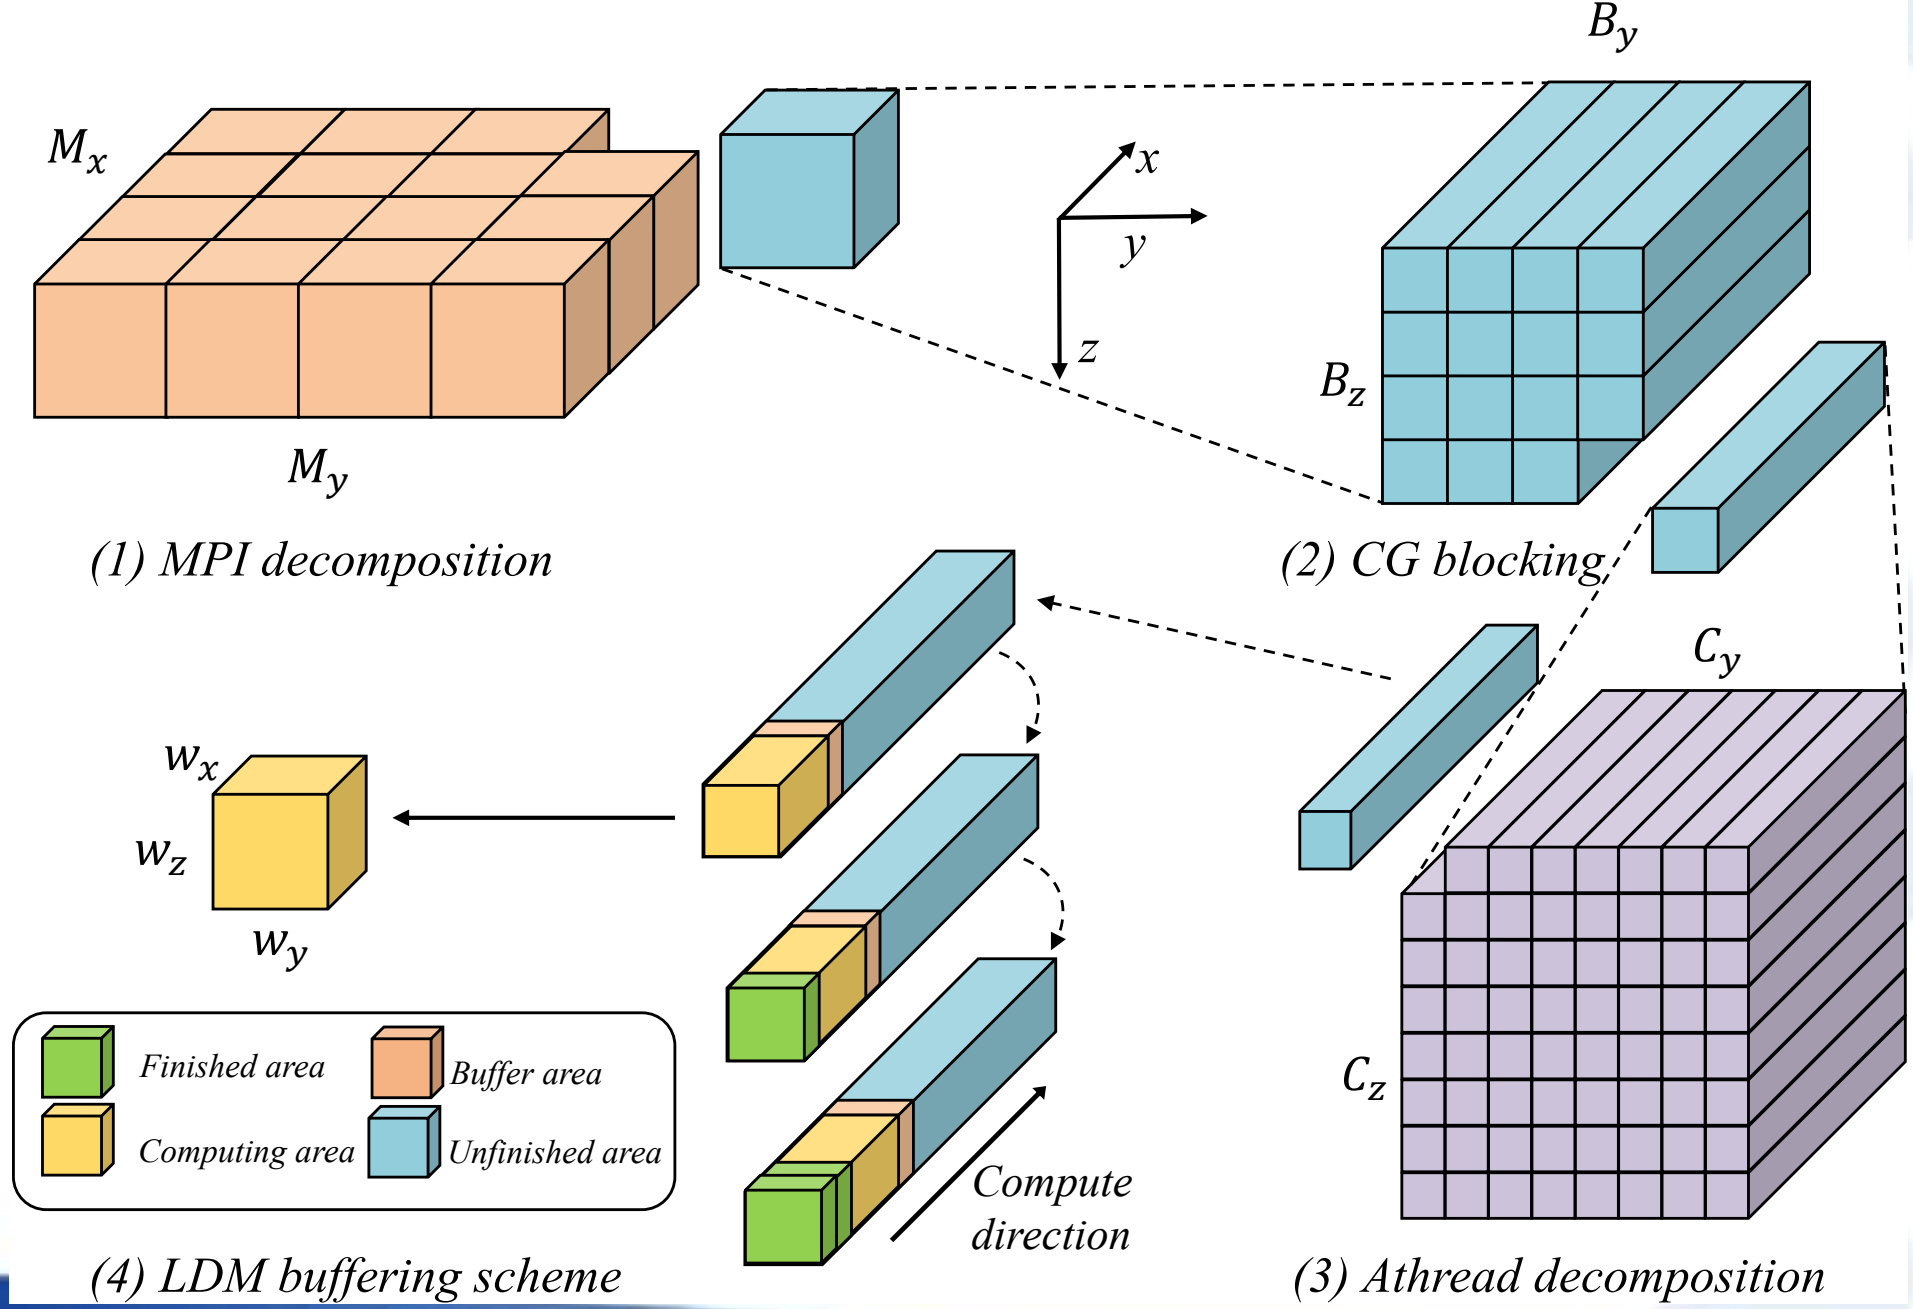

Multi-Level Domain Decomposition

A Balanced Memory Scheme

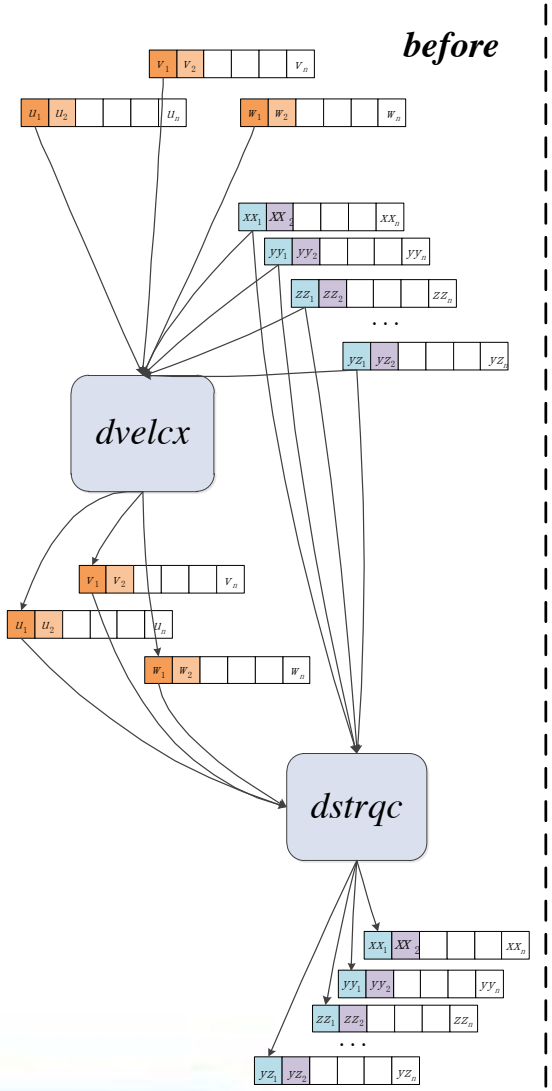

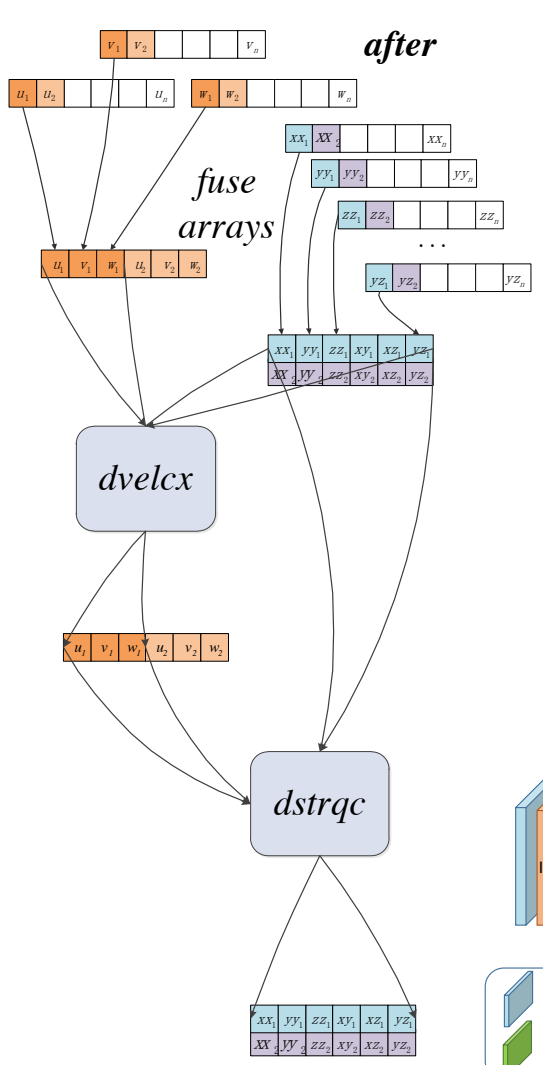

(1) array fusion,

(2) halo exchange through register communication,

(3) and optimized blocking configuration guided by an analytical model

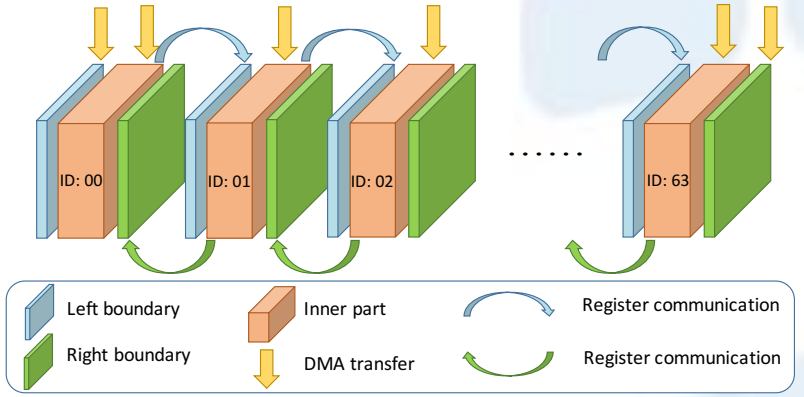

On-the-fly Compression

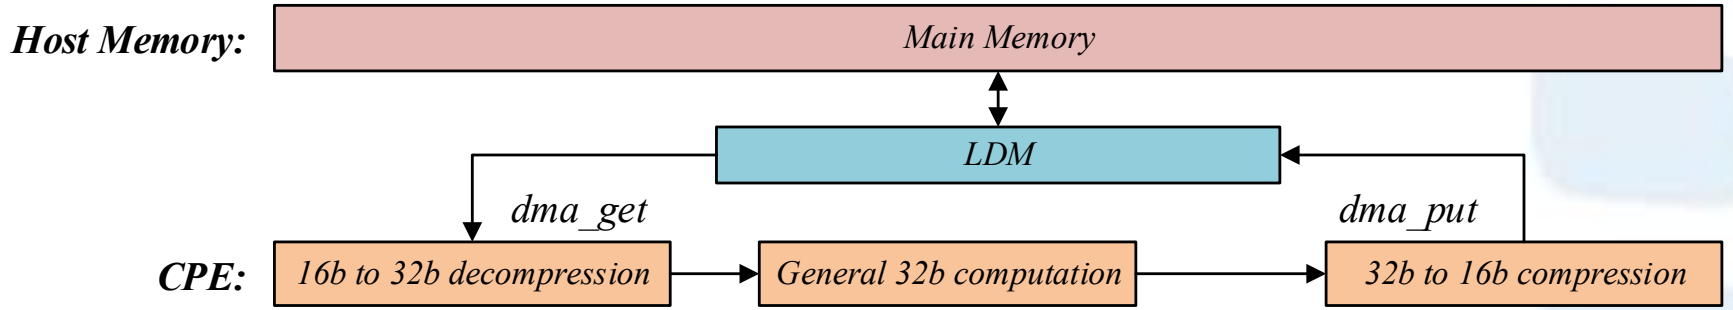

(b) Computation workflow

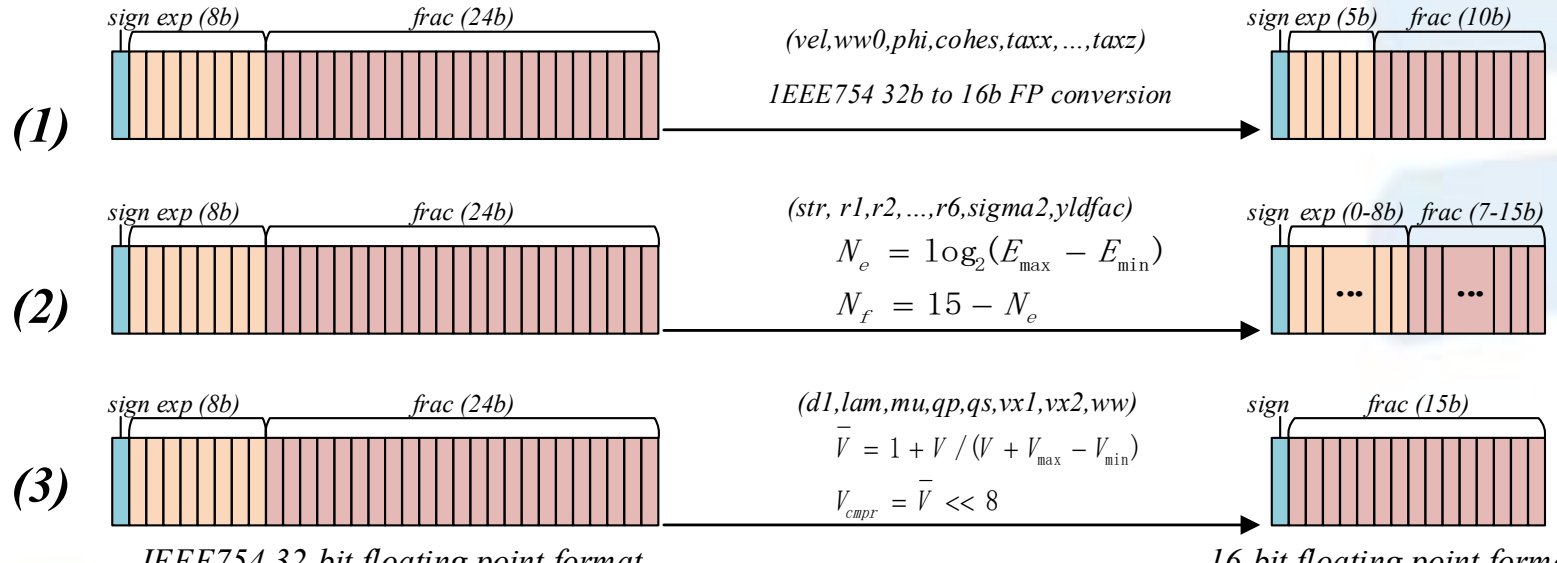

IEEE754 32-bit floating point format

16-bit floating point formats

(d) Compression algorithms

Speedup: 64 CPE vs 1 MPE

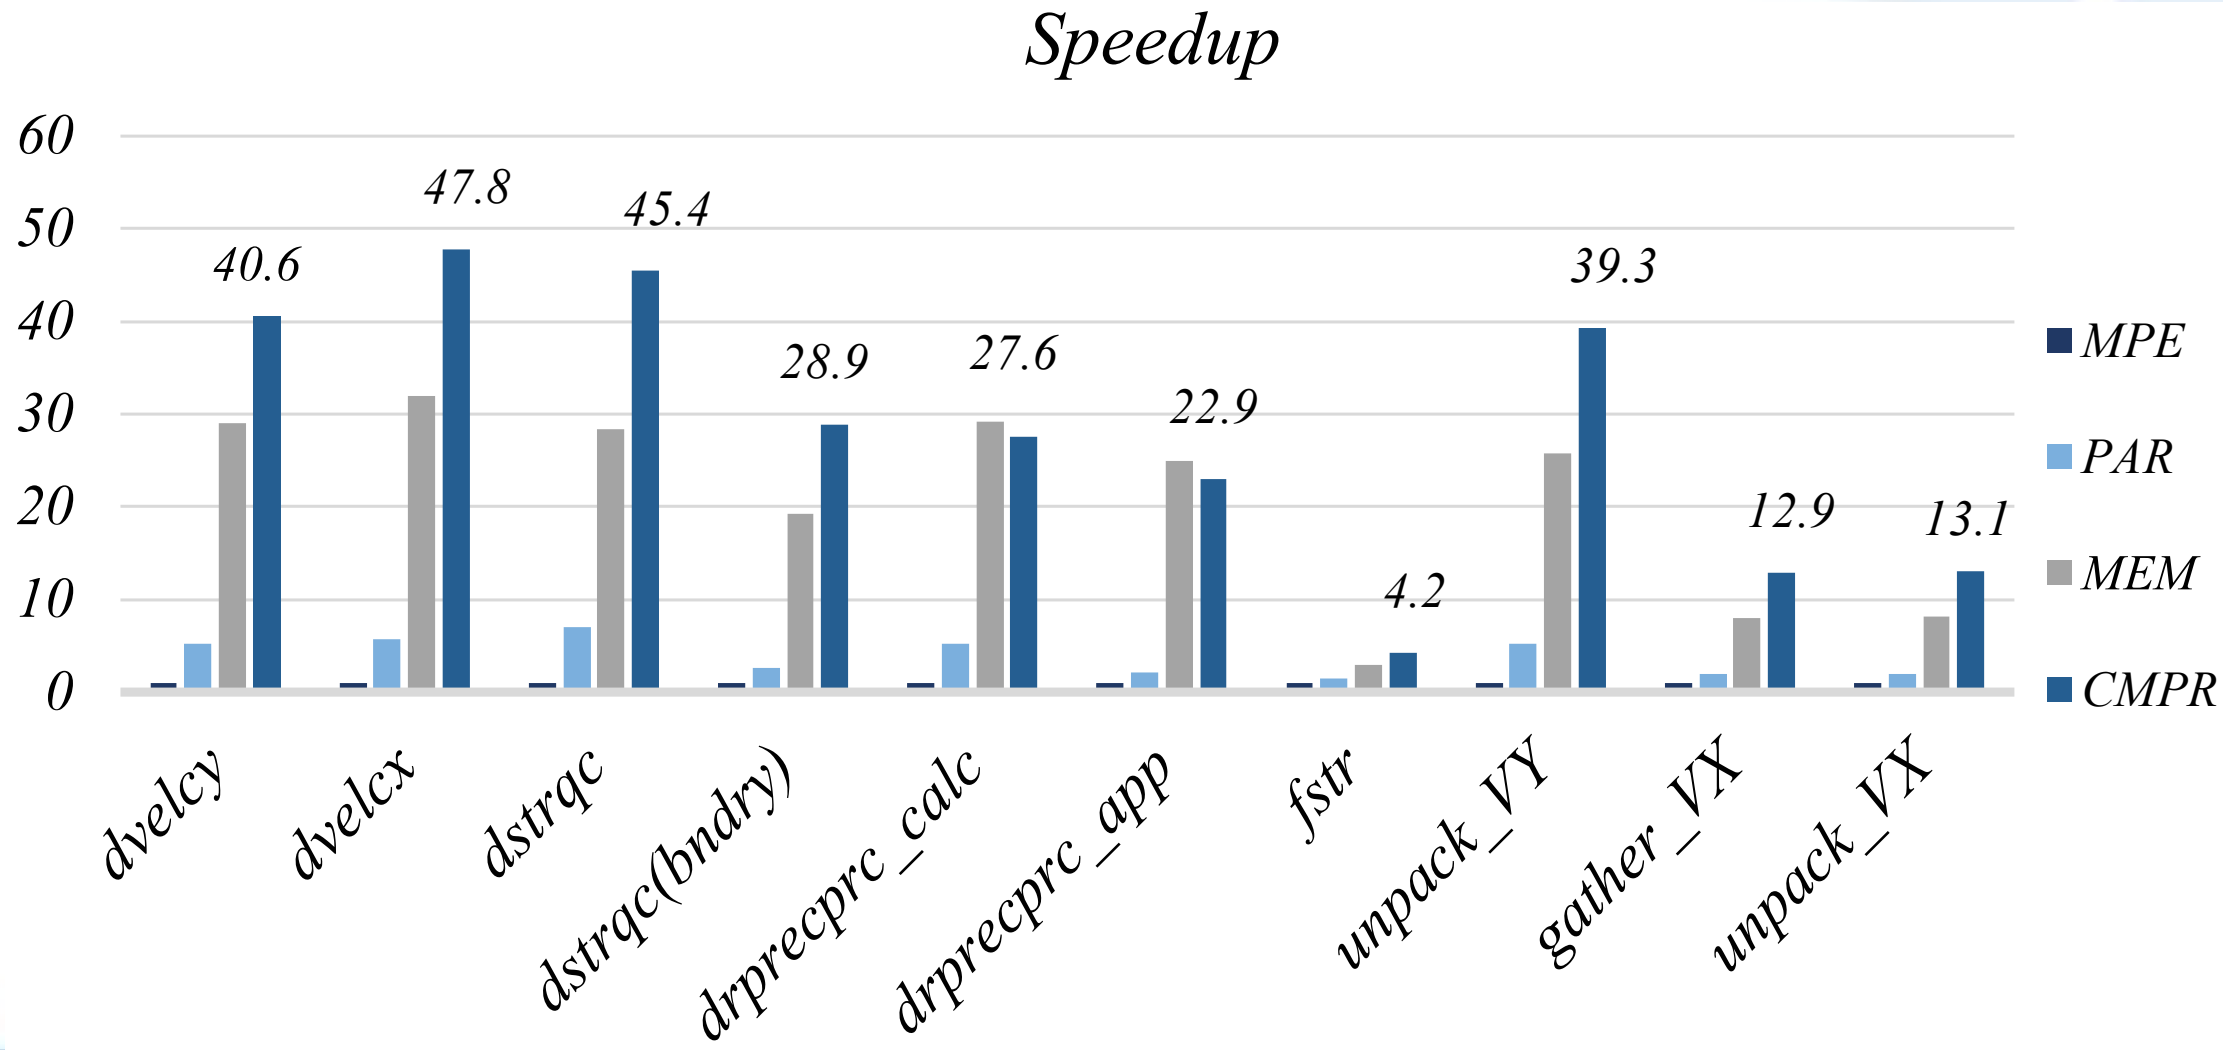

Memory Bandwidth Utilization

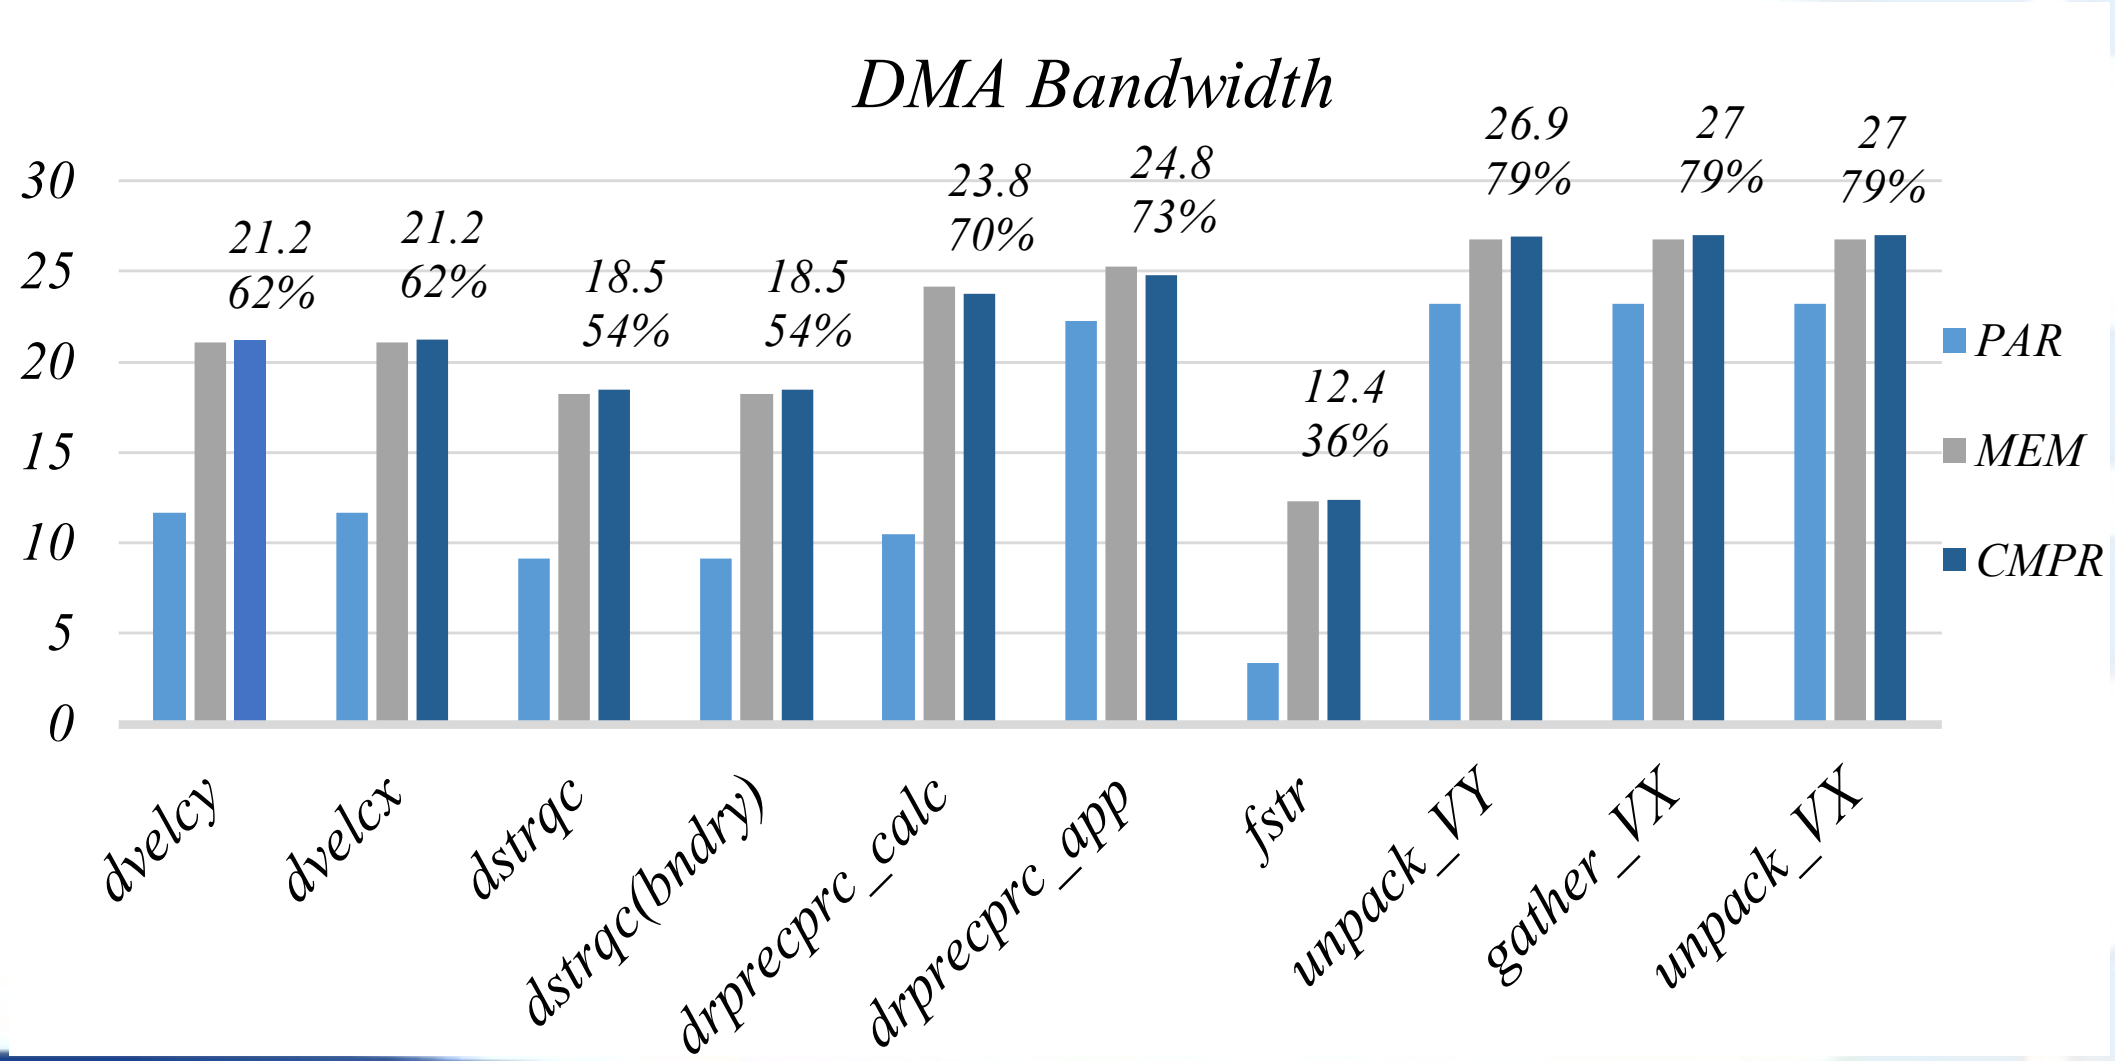

Weak Scaling

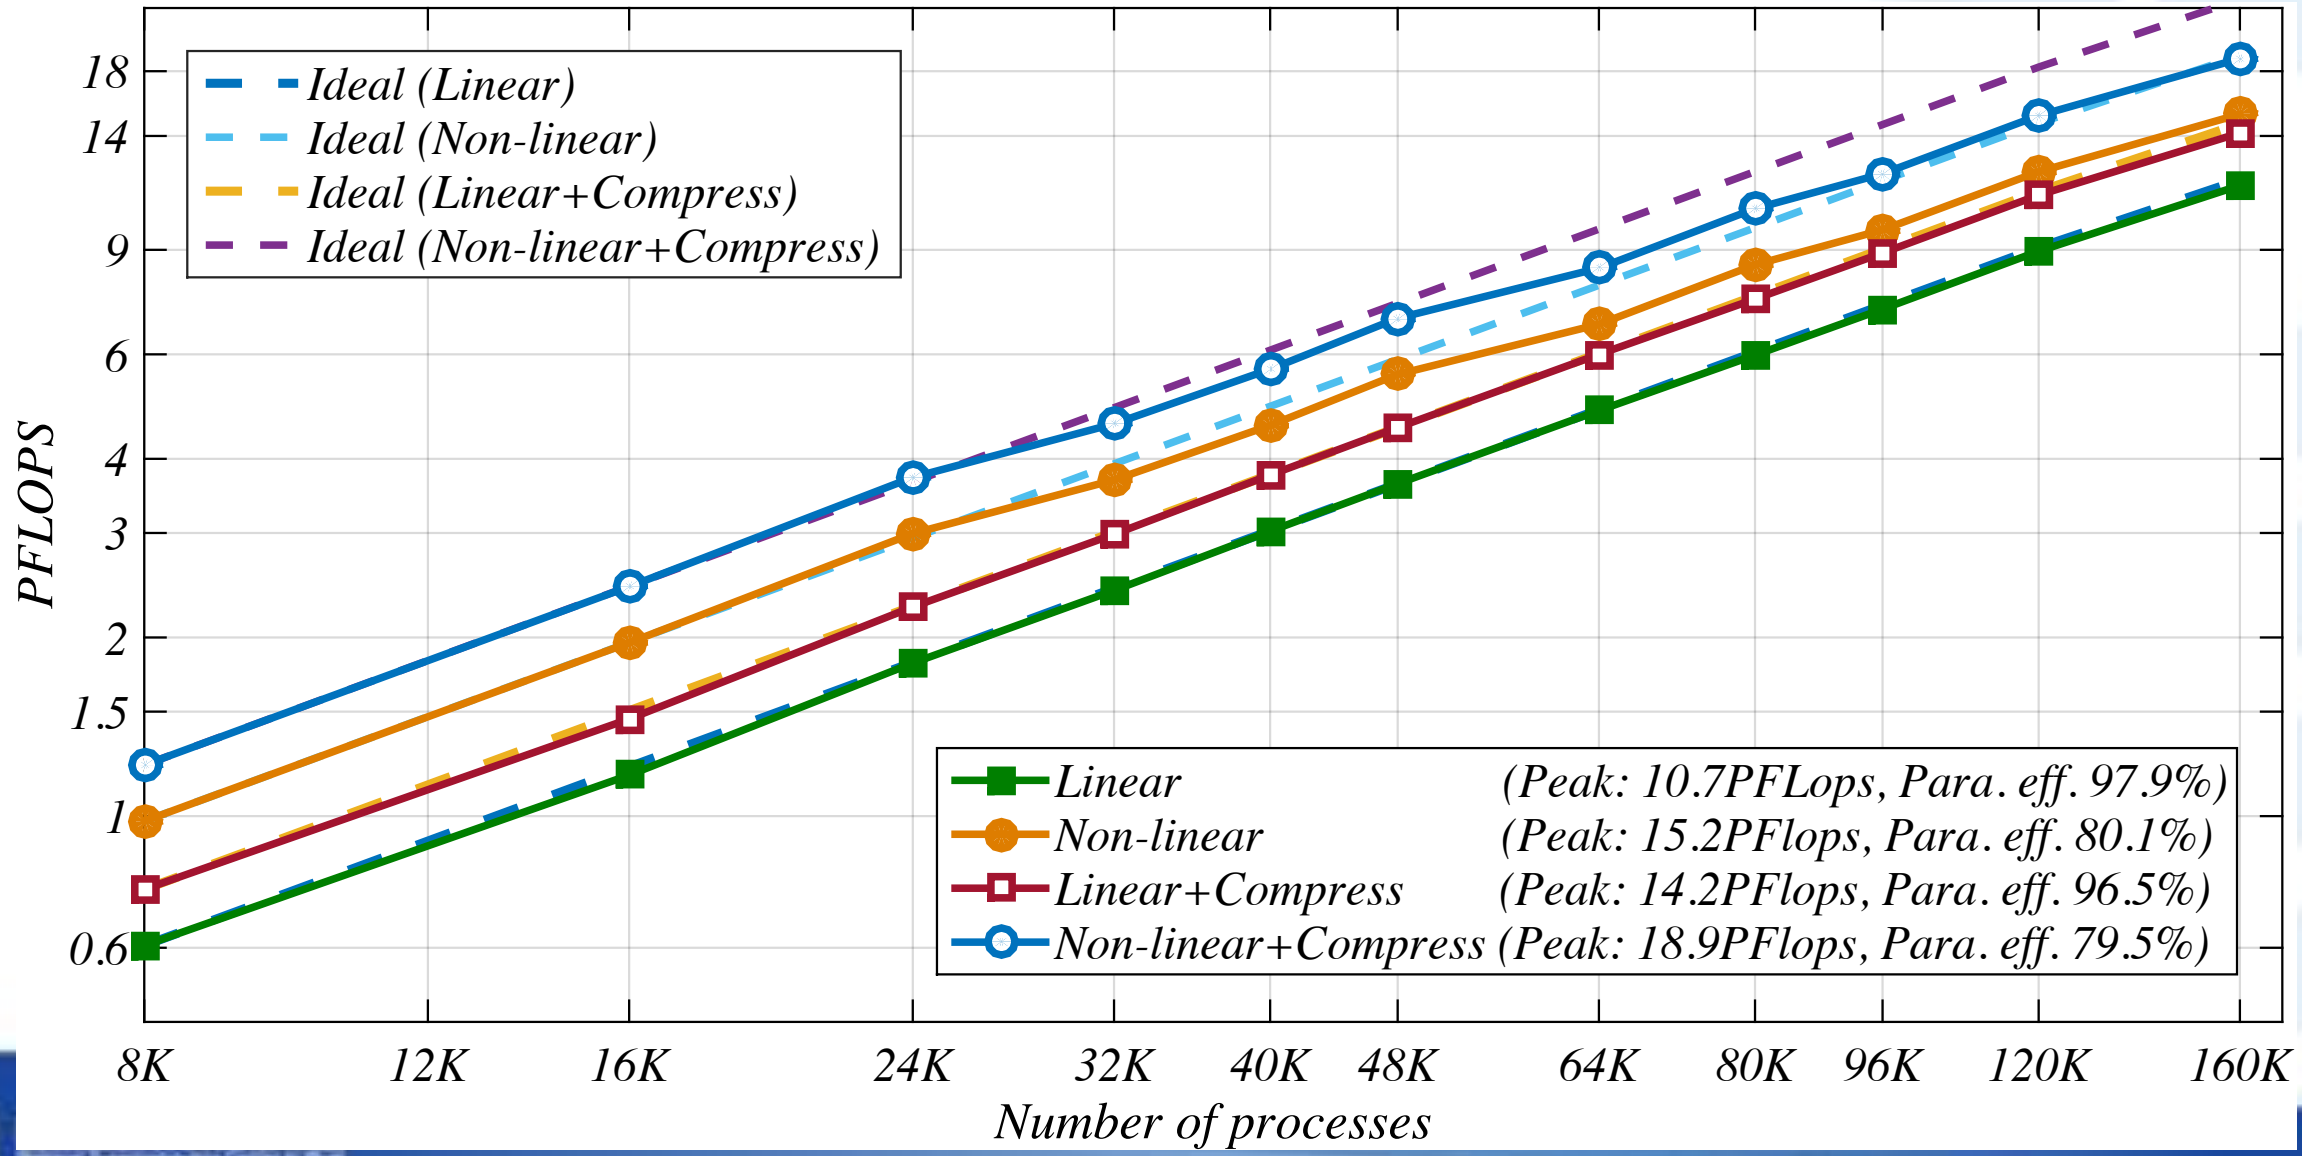

Strong Scaling

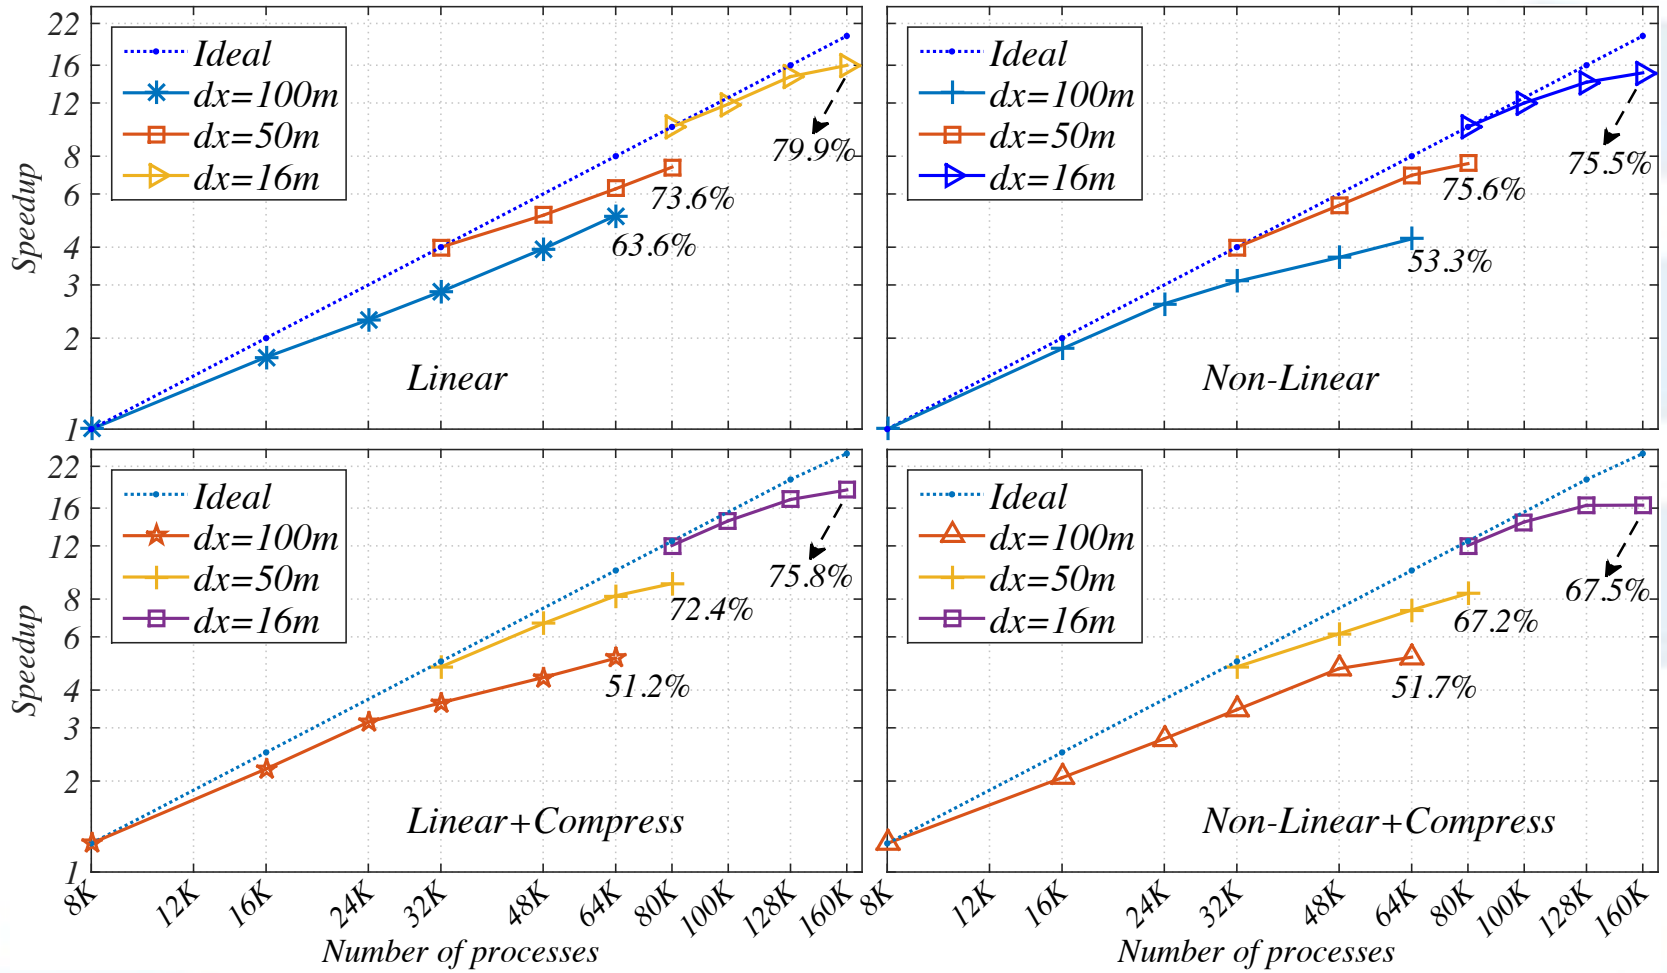

Simulation Results: 200m vs 16m

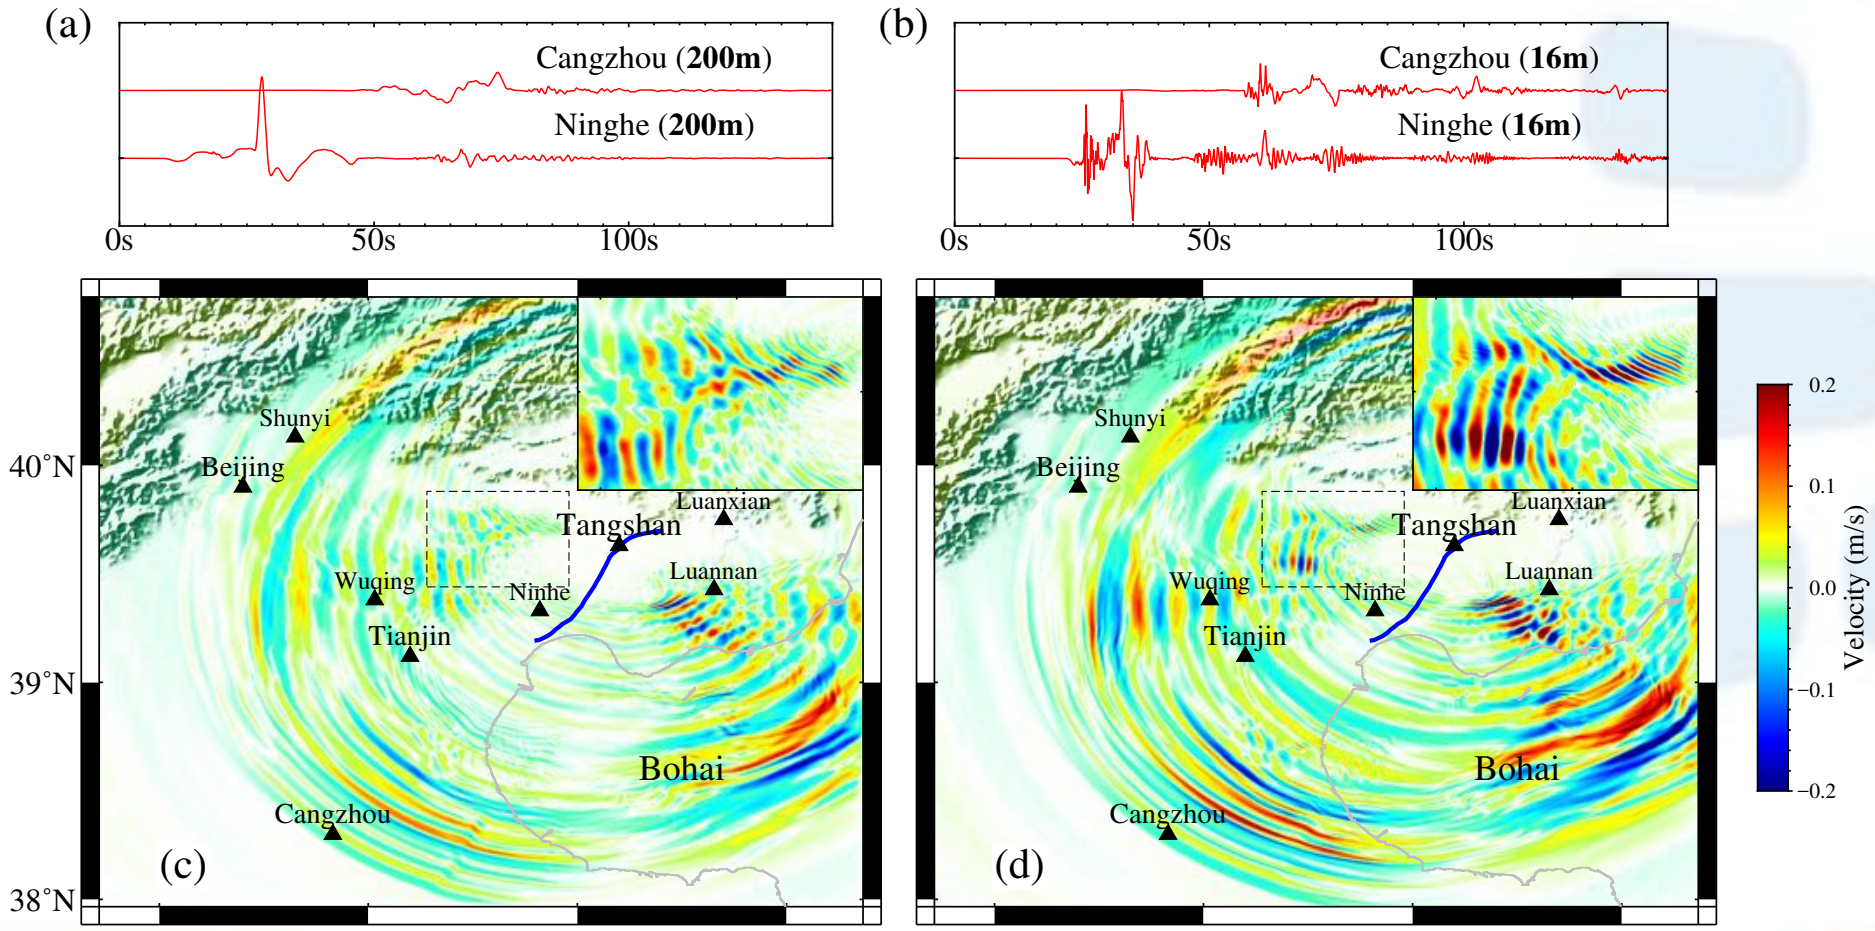

Outline

- Sunway Machine: the Challenges and Opportunities
- Scientific Computing with 10 Million Cores
- Long Term Plan for Sunway TaihuLight

Long Term Plan

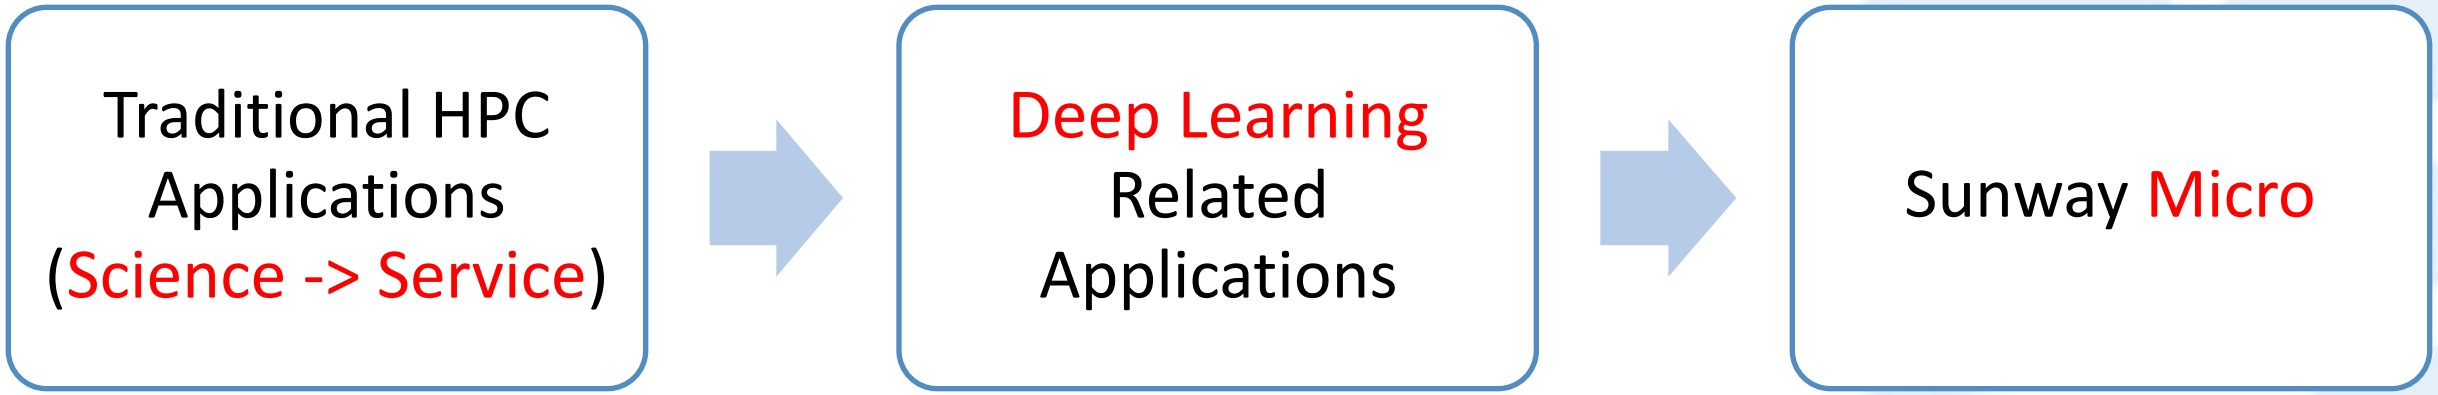

Long Term Plan

Traditional HPC Applications
(Science -> Service)

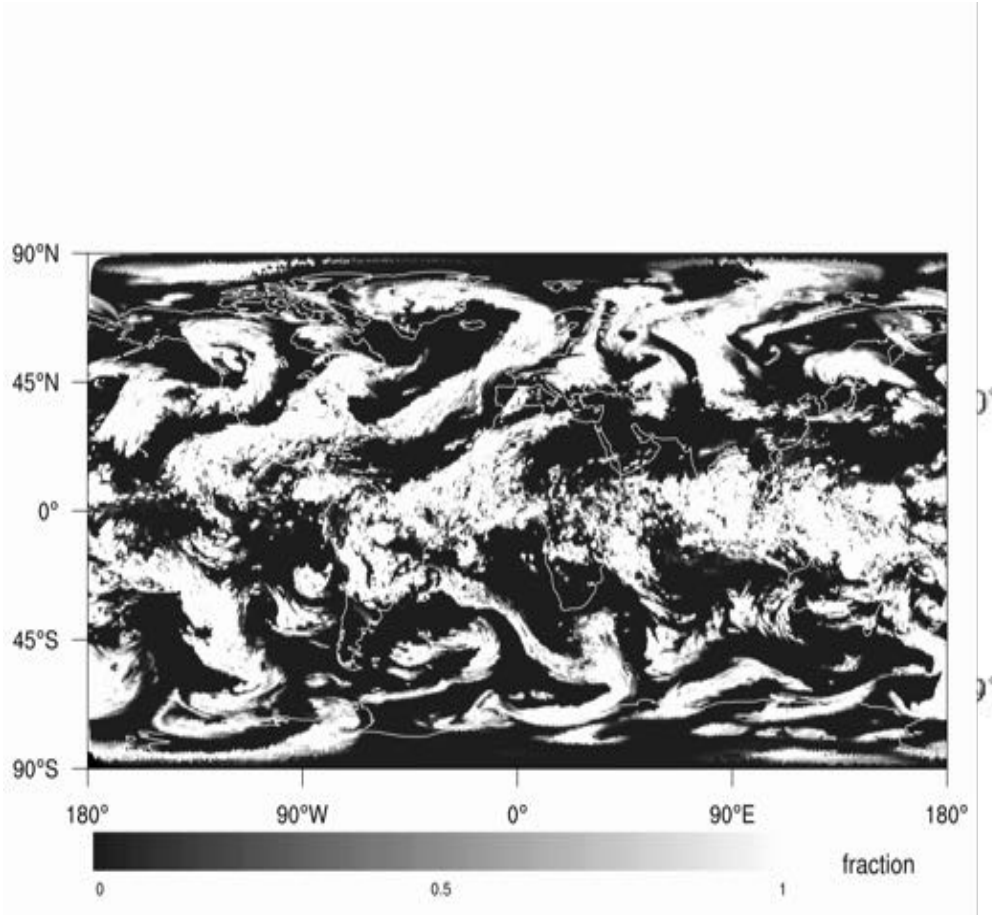

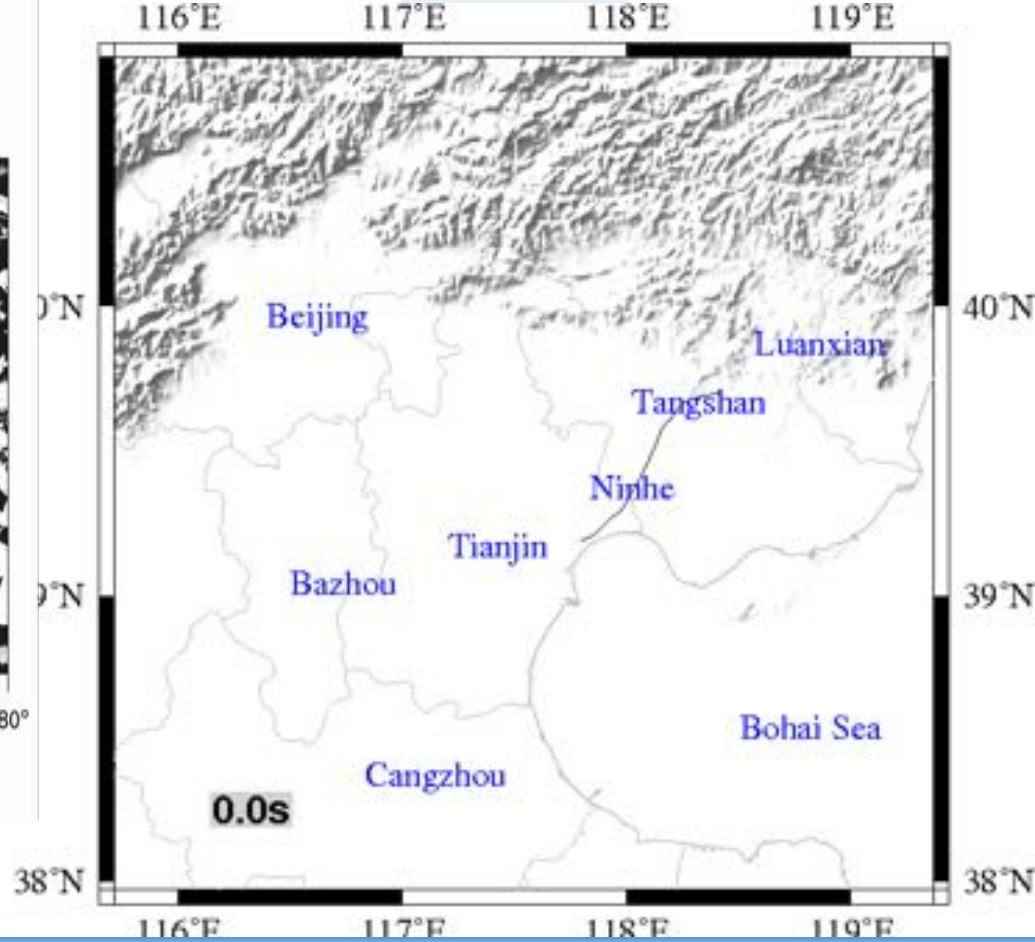

“15-Pflops Nonlinear Earthquake Simulation on Sunway TaihuLight: Enabling Depiction of Realistic 10 Hz Scenarios”, Gordon Bell Prize Finalist, SC 2017.

Long Term Plan

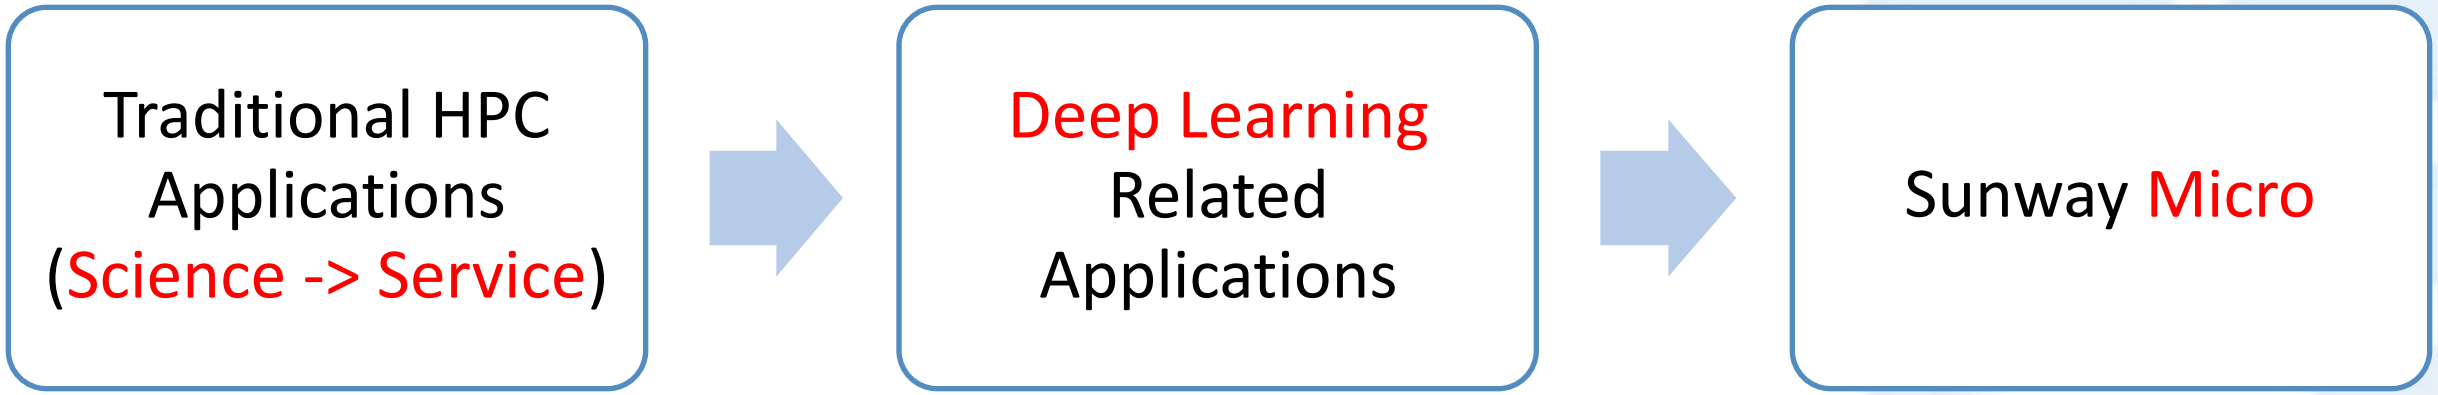

Long Term Plan

Traditional
Applicati
(Science -> S

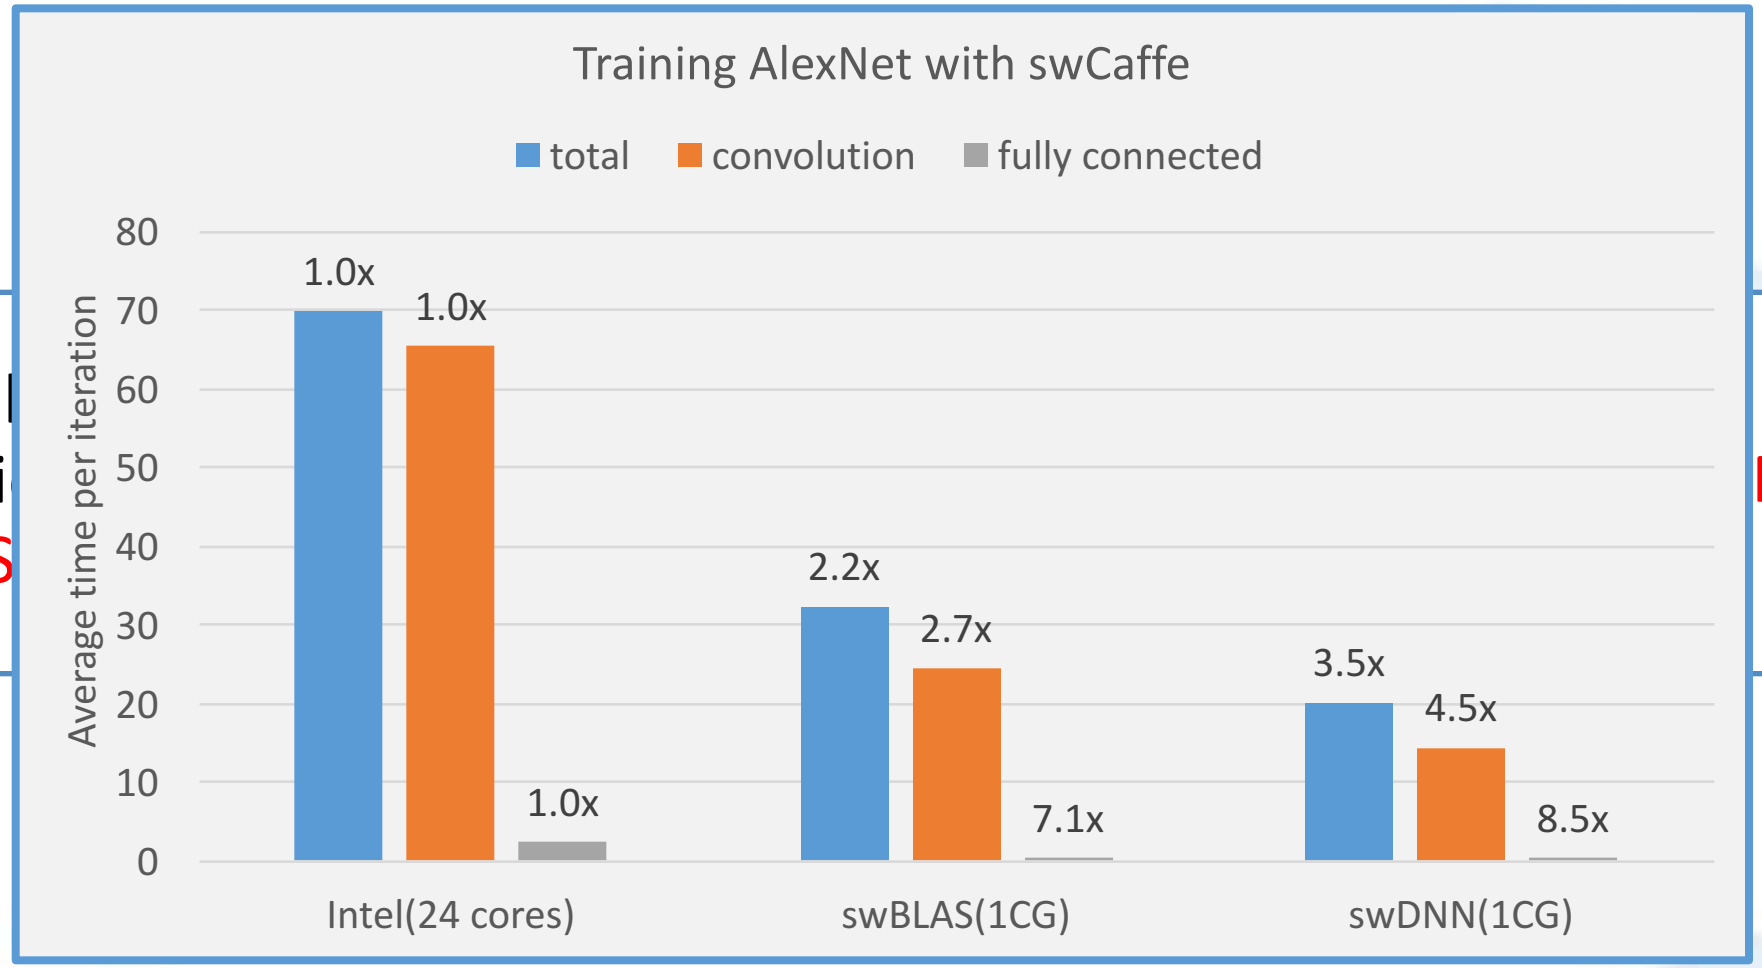

Micro

Long Term Plan

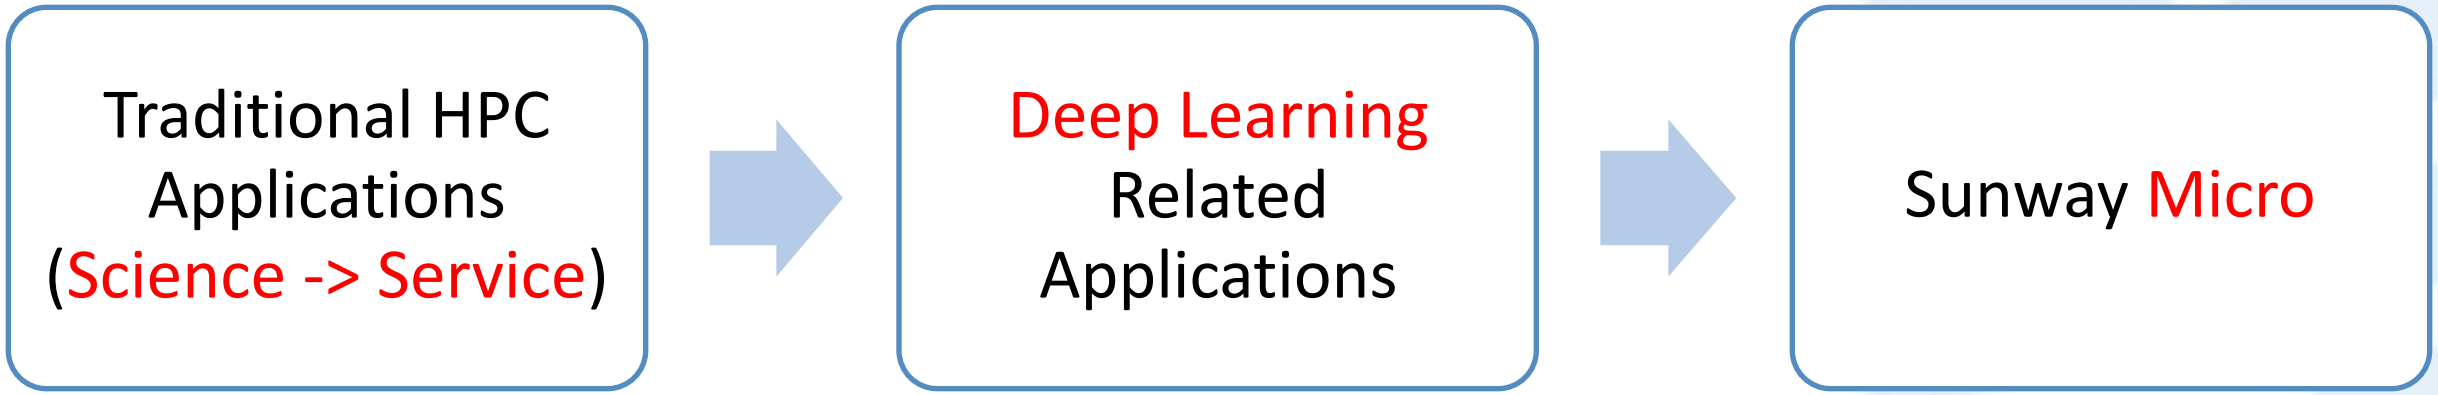

Acknowledgements

- MOST, China: major sponsors of the HPC hardware and software development
- NRCPC: vendor of the machine
- NCAR: Rich Loft, John Dennis, Allison Baker, Haiying Xu (support and advice on the CAM-SE work)
- SCEC: Yifeng Cui, Steve Day, Daniel Roten, Kim Olsen, Josh Tobin, Alex Breuer, and Dawei Mu (discussion and advice on the earthquake simulation work)

THANK YOU