Universe Sciences

SPACE

CENTRE OF EXCELLENCE FOR HPC ASTROPHYSICAL APPLICATIONS

Towards Exascale Computing for Astrophysical SimulationLeveraging the Leonardo EuroHPC System

Nitin Shukla | HPC Engineer CINECA | October 1st, 2025 | Copenhagen, Denmark

Co-funded by the European Union

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441.

EuroHPC Joint Undertaking

Rise of GPUs

HPC systems: from Giga to Exa and beyond …

Parallel scaling has replaced frequency as primary way to improve computational performance

Multi-core architecture perform more operations than single-core processor at same clock frequency by executing tasks concurrently.

GPU architecture handles massive parallel computations by leveraging thousands of small, efficient **cores**.

Organized in **Streaming Multiprocessors (SMs)** ⇒

execute thousands of threads simultaneously.

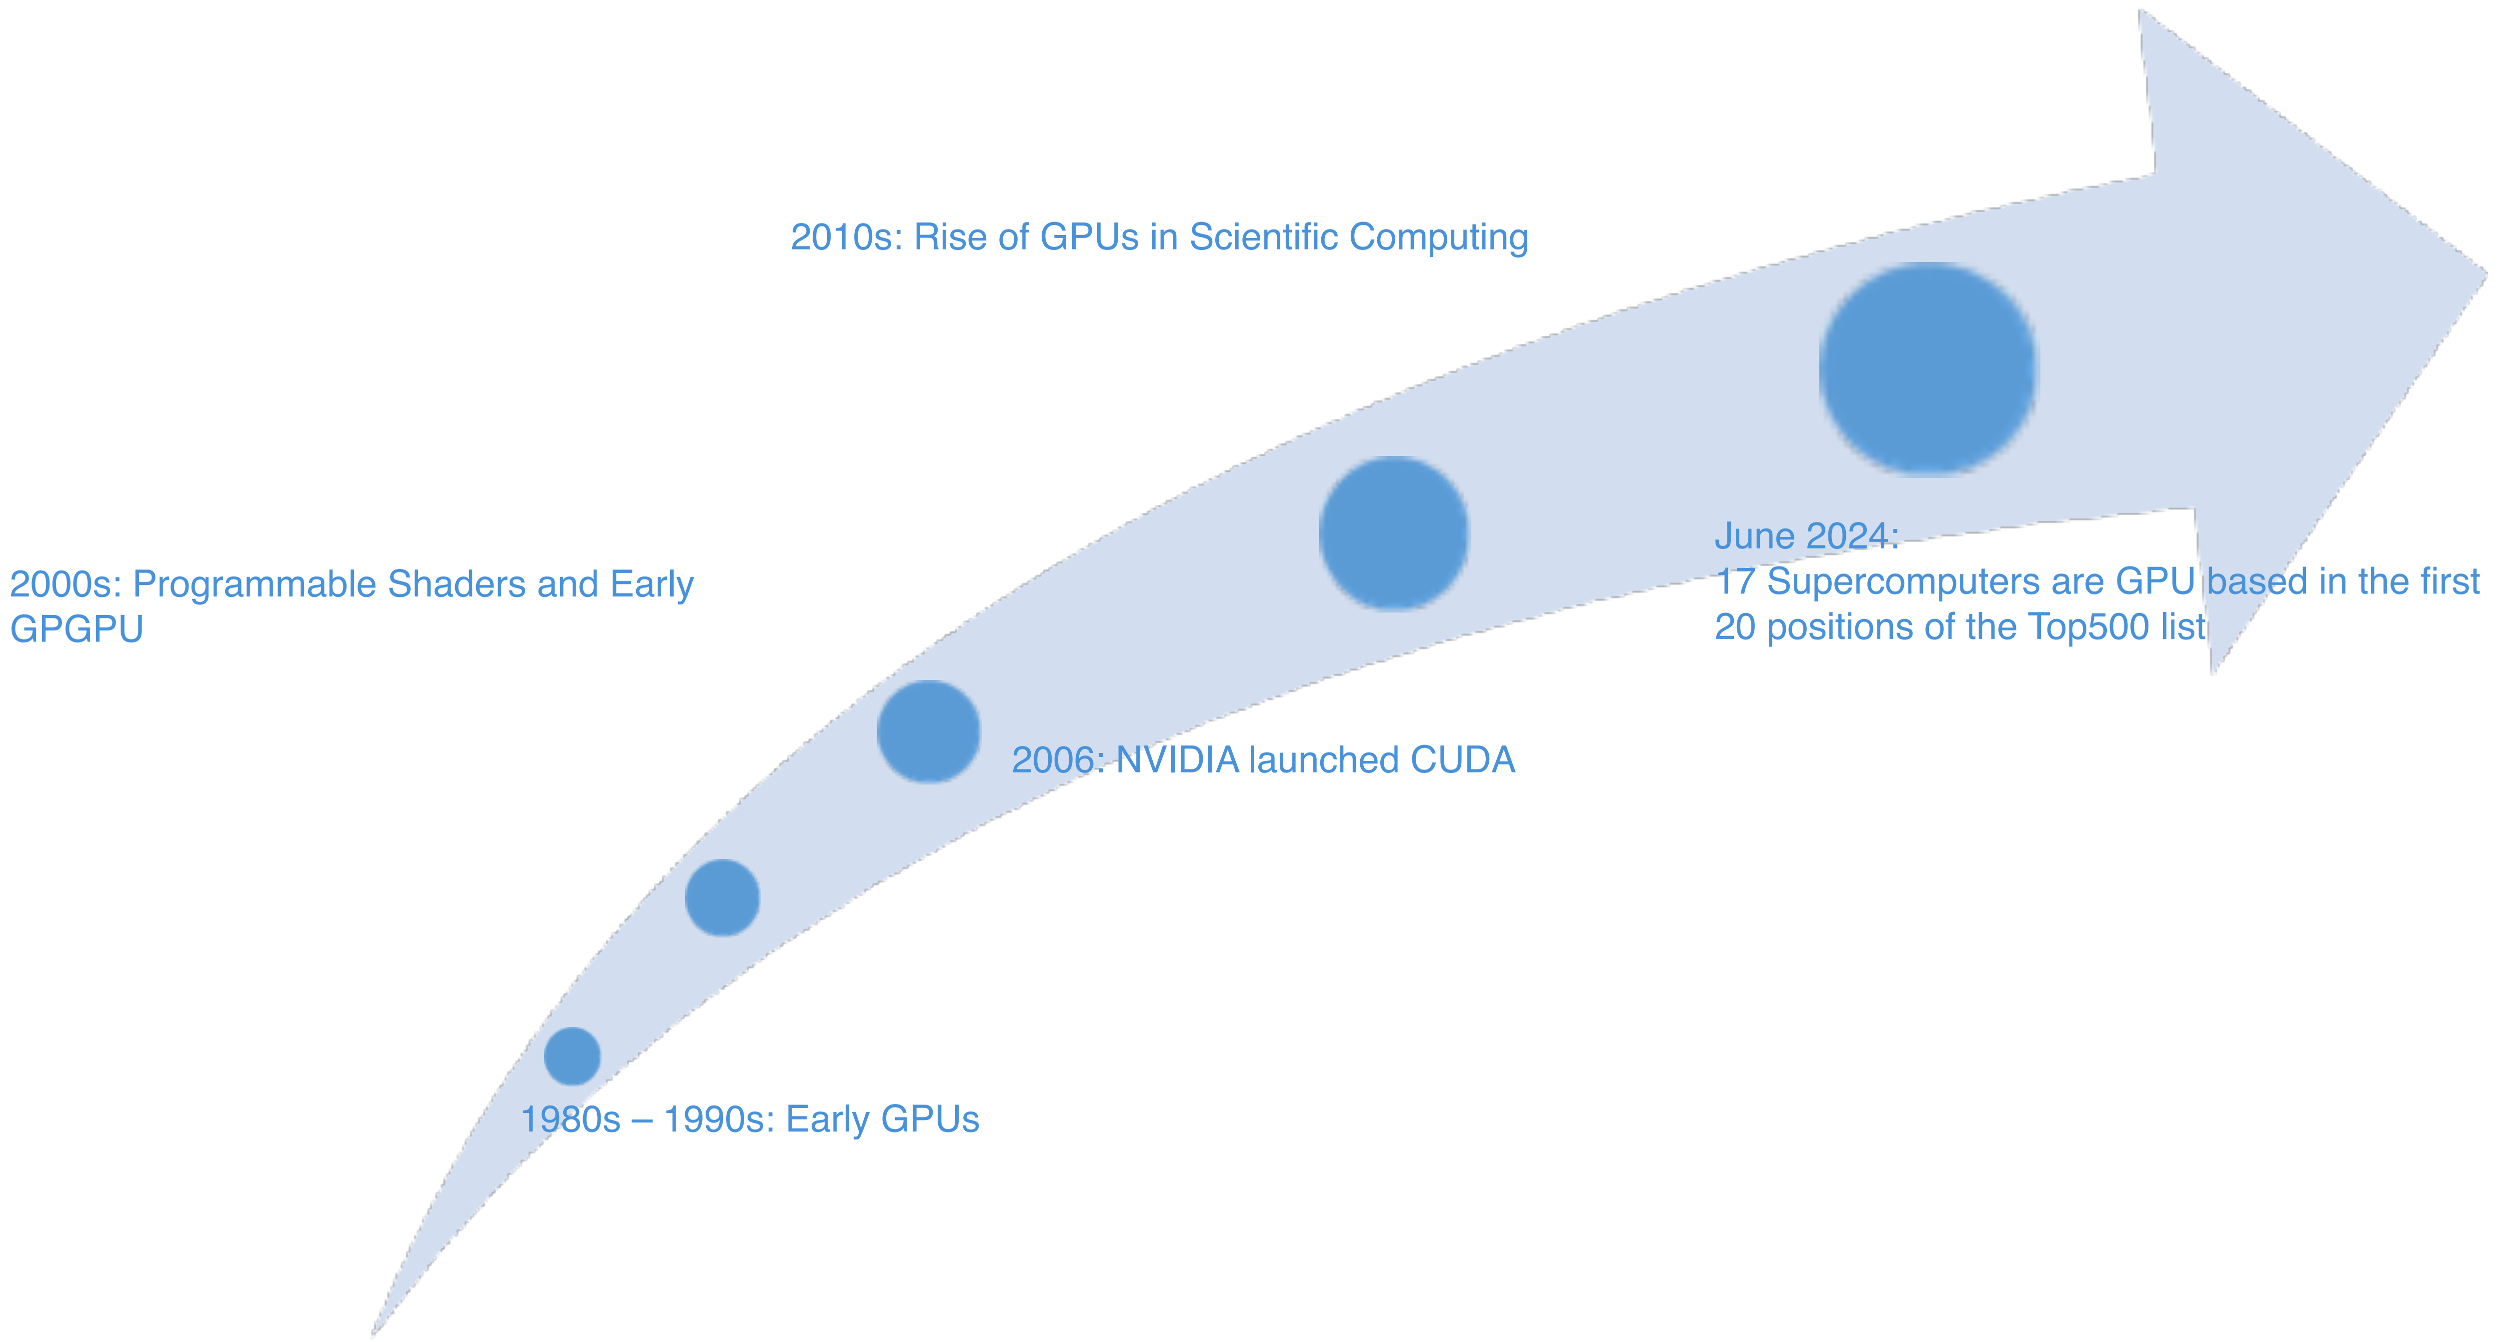

EURO HPC SYSTEM

Exascale

Pre-exascale

Petascale / mid-range

Quantum system / simulator

https://www.space-coe.eu/

Scalable Parallel Astrophysical Codes for Exascale

1st January 2023 → 31st December 2026

SPACE

Scalable Parallel Astrophysical Codes for Exascale

the first project at CIENCA for Astrophysics and Plasma simulations

Mission

To evolve seven widely used European codes into exascale paradigms through co-design activities involving scientists, HPC experts, and developers

Objectives

To enable widely used European A&C HPC codes to efficiently exploit pre-exascale systems and prepare them for exascale and beyond

Impacts

To adopt a user-driven, co-design approach for application development, with a sustainable software strategy

To deploy selected applications, promote their use through outreach and training, and create a skilled talent pool in Europe to drive high-performance solutions in academia and facilitate the transition to exascale technologies

SPACE: a Center of Excellence for Research–Industry Synergy

fifteen partners from eight countries

Scientific Applications

7

Tenure: 48 months

Budget: 8 Millions

European Universities

9

Research Institute:

Italy, Germany, France, Norway, Belgium

EU-HPC

6

HPC centres: CINECA, IT4I BSC

Companies: EnginSoft, E4, Atos

CINECA’s Contribution to Accelerating SPACE-COE Flagship Codes i.e PLUTO, OpenGadget, and iPIC3D

Mesh-Based Codes Approach

Multidimensional compressible plasma with high mach numbers

Fluid model

- Variable (e.g density momentum etc ..) live on a static/adaptive spatial grid
- Finite volume / Finite difference methods
- Multidimensional compressible plasma with high mach numbers
  - Pluto [UniTo]
- In a Finite Difference (FD) approach, fluid variables (density, momentum, etc..) are stored, e.g., at the cell center. Equations are then discretized using, e.g

$$\frac{\partial U}{\partial t} + \nabla \cdot \mathbf{F} = 0 \rightarrow \frac{dU_{ij}}{dt} + \frac{F^x_{i+\frac{1}{2},j} - F^x_{i-\frac{1}{2},j}}{\Delta x} + \frac{F^y_{i,j+\frac{1}{2}} - F^y_{i,j-\frac{1}{2}}}{\Delta y} = 0$$

- In a Finite Volume (FV) approach, fluid are stored by their volume average, and temporal evolution is carried out using Gauss's law:

$$\langle U \rangle_{ij} = \frac{1}{\Delta V} \int_{ij} U(x,y)dxdy$$

$$\frac{\partial U}{\partial t} + \nabla \cdot \mathbf{F} = 0 \quad \rightarrow \quad \frac{d\langle U \rangle_{ij}}{dt} + \int_{ij} \mathbf{F} \cdot d\mathbf{S} = 0$$

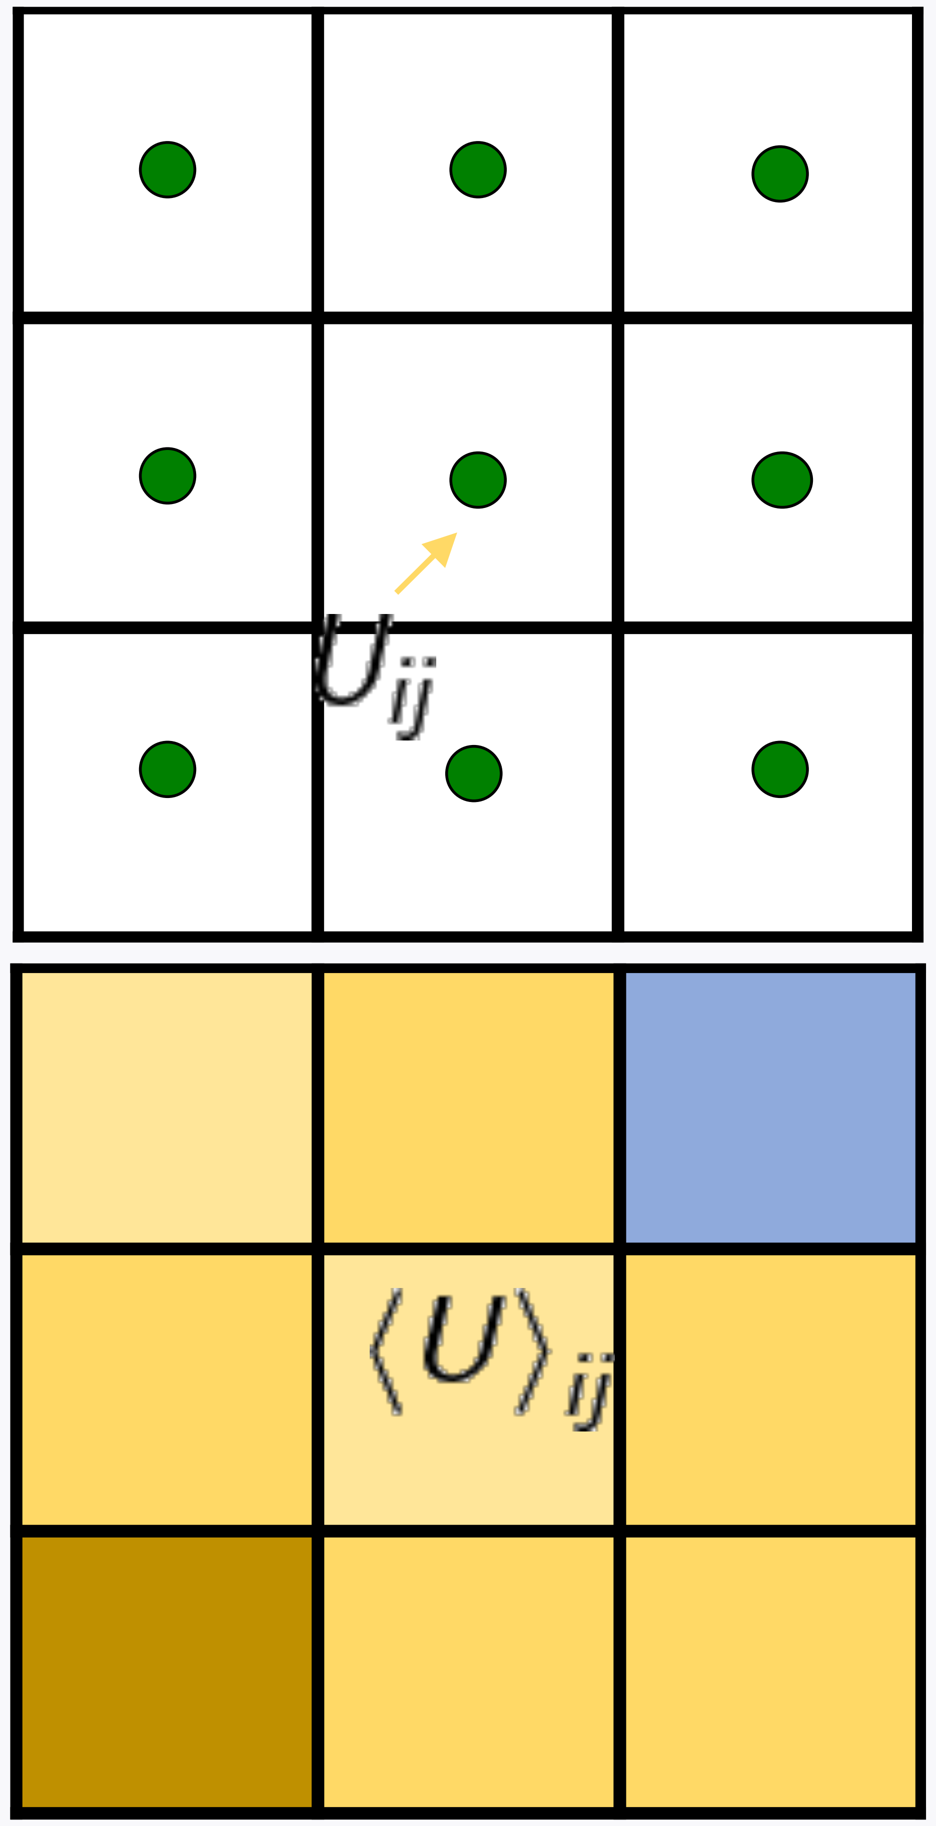

Particle-Based Codes Approach

Multidimensional compressible plasma with high mach numbers

Particle-Based Methods

- Particles act as samples of the "continuum" and carry the physical properties of the simulated system
  - N-Body codes -> OpenGadget [INAF]
  - Particle-in-cell -> iPIC3D [KU-Leuven]

Key Distinctions

- **OpenGadget** are both N-body + SPH codes

  → They do **not** rely on any mesh
  → Use **smoothing kernels** to interpolate field values at any point in space
  → Represent the continuum without needing interpolation to a mesh

- **PIC methods like iPIC3D**

  → Require a mesh for solving field equations
  → Particles and fields are tightly coupled via interpolation between mesh and particles
  → Integration of equations of motion for each particle
  → Interpolation of charge/current to field mesh (only for PIC)
  → Computation of fields on mesh points (only for PIC)
  → Interpolation of fields back to particle locations (only for PIC)

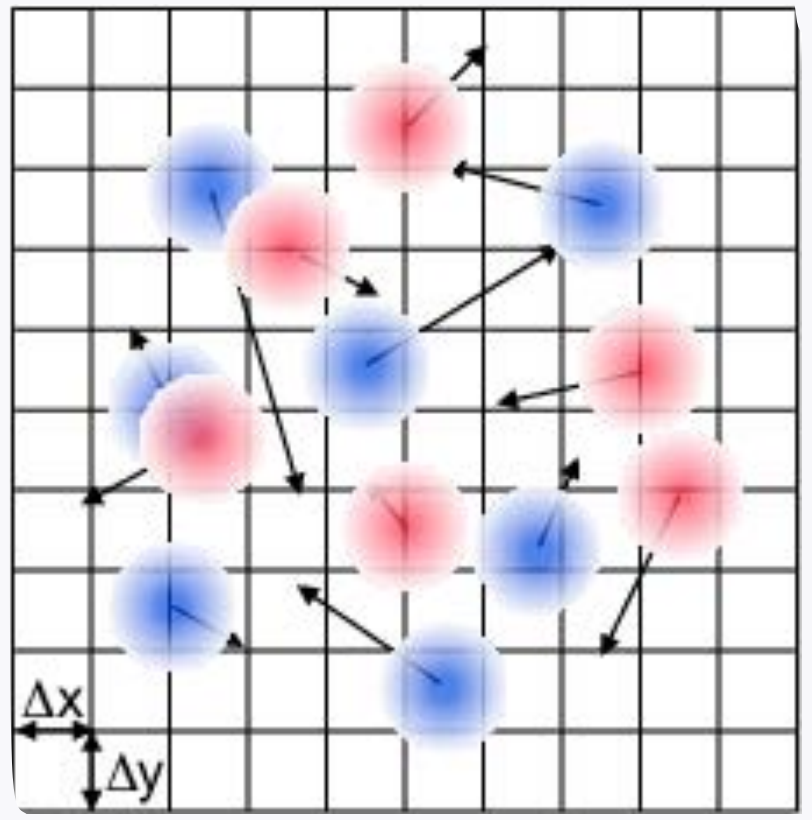

$$\frac{d^2\mathbf{x}_p}{dt^2} = \frac{d\mathbf{v}_p}{dt} = \mathbf{a}$$

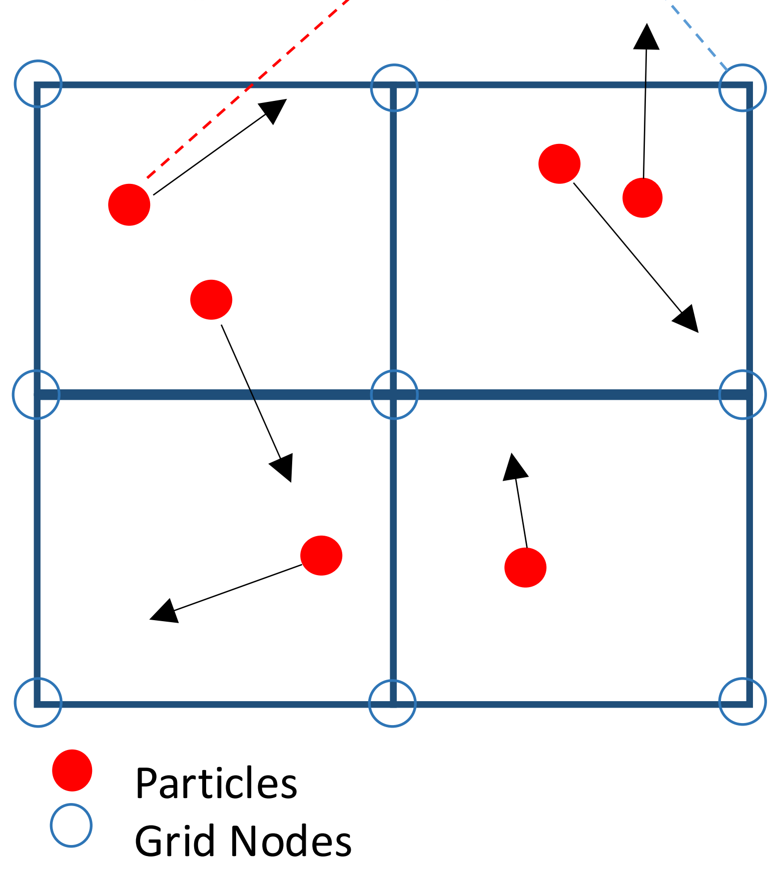

Porting Codes on the Accelerators

Programming models Overview

01 **Performance assessment**
Profiling, Benchmark test

02 **Porting code on GPU**
Offload most compute part of the code

03 **Profiling tools**
Nsight sys, Nsight compute

04 **Optimisation**
Improving data locality

01

Performance assessment

Extrae Profiling tools in collaboration with POP3/IT4I

Performance Assessment With Extrae

OpenGadget3

Simulation setup

- Cosmological box: 30 Mpc/h
- 256^3 particles
- 1–16 nodes used
- Gravity and hydrodynamics enabled
- Last 15 timesteps as Region of Interest (RoI)

Results

- Speedup diverges from optimum, efficiency <80% on 8 nodes
- Main bottlenecks: Serialization efficiency (SerE) and OpenMP Communication efficiency (OCE)
- DENS routine: 2nd longest, SerE 86%, best optimization candidate
- GRAV routine: most time-consuming, SerE 92%, low priority for optimization

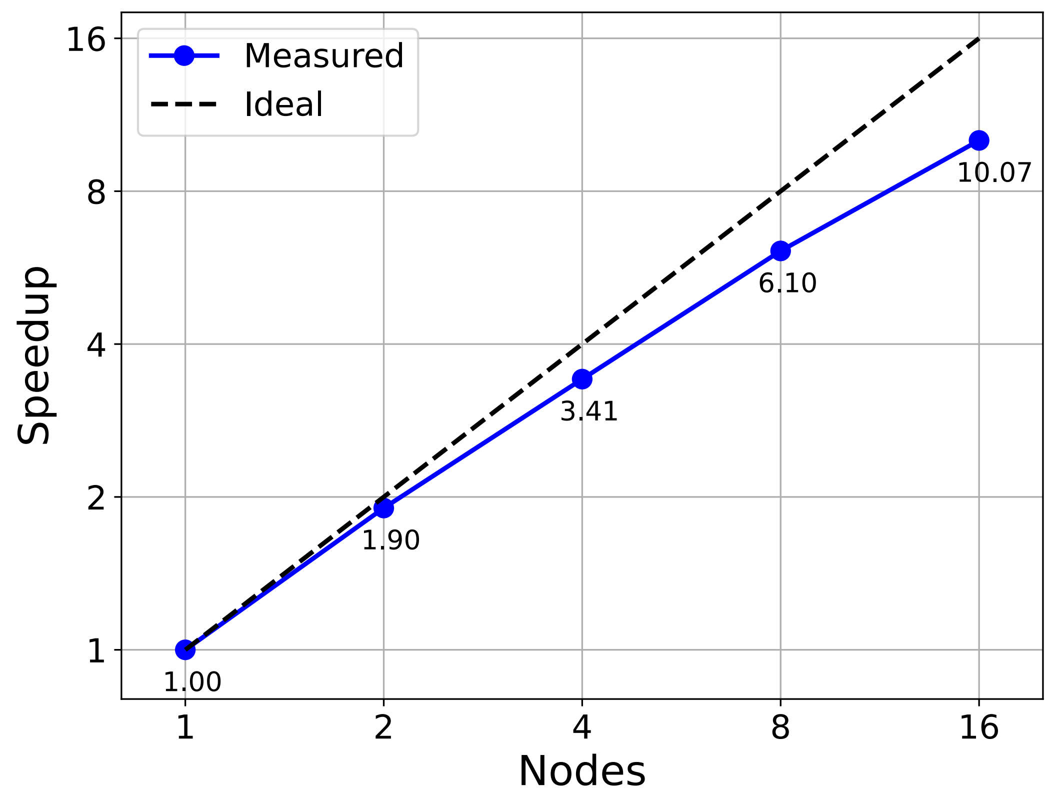

	112(16x7)[1]	224(32x7)[2]	448(64x7)[3]	896(128x7)[4]	1792(256x7)[5]
Global efficiency	83.99	79.87	71.65	64.06	52.86
-- Parallel efficiency	83.99	77.12	72.19	62.75	51.97
-- Load balance	94.49	91.07	91.33	89.74	87.46
-- Communication efficiency	88.88	84.69	79.04	69.92	59.41
-- Computation scalability	100.00	103.56	99.25	102.09	101.73
-- IPC scalability					
-- Instruction scalability					
-- Frequency scalability					

	112(16x7)[1]	224(32x7)[2]	448(64x7)[3]	896(128x7)[4]	1792(256x7)[5]
-- Hybrid Parallel efficiency	83.99	77.12	72.19	62.75	51.97
-- MPI Parallel efficiency	89.64	84.50	82.10	76.99	72.24
-- MPI Load balance	98.44	95.60	96.75	96.39	97.03
-- MPI Communication efficiency	91.07	88.39	84.86	79.87	74.45
-- Serialization efficiency	93.27	91.43	89.33	85.39	80.25
-- Transfer efficiency	97.64	96.68	94.99	93.54	92.78
-- OpenMP Parallel efficiency	93.69	91.27	87.93	81.51	71.93
-- OpenMP Load Balance	96.00	95.26	94.40	93.10	90.14
-- OpenMP Communication efficiency	97.60	95.81	93.15	87.55	79.80

Main Findings and Recommendations

OpenGadget3

Current code shows improved scaling capability

- refactoring of communication patterns
- refactoring of multi-threading support in loops
- 2x enhancement in “small steps”, i.e. in steps where very few particles are processed (and not offloaded to GPU)

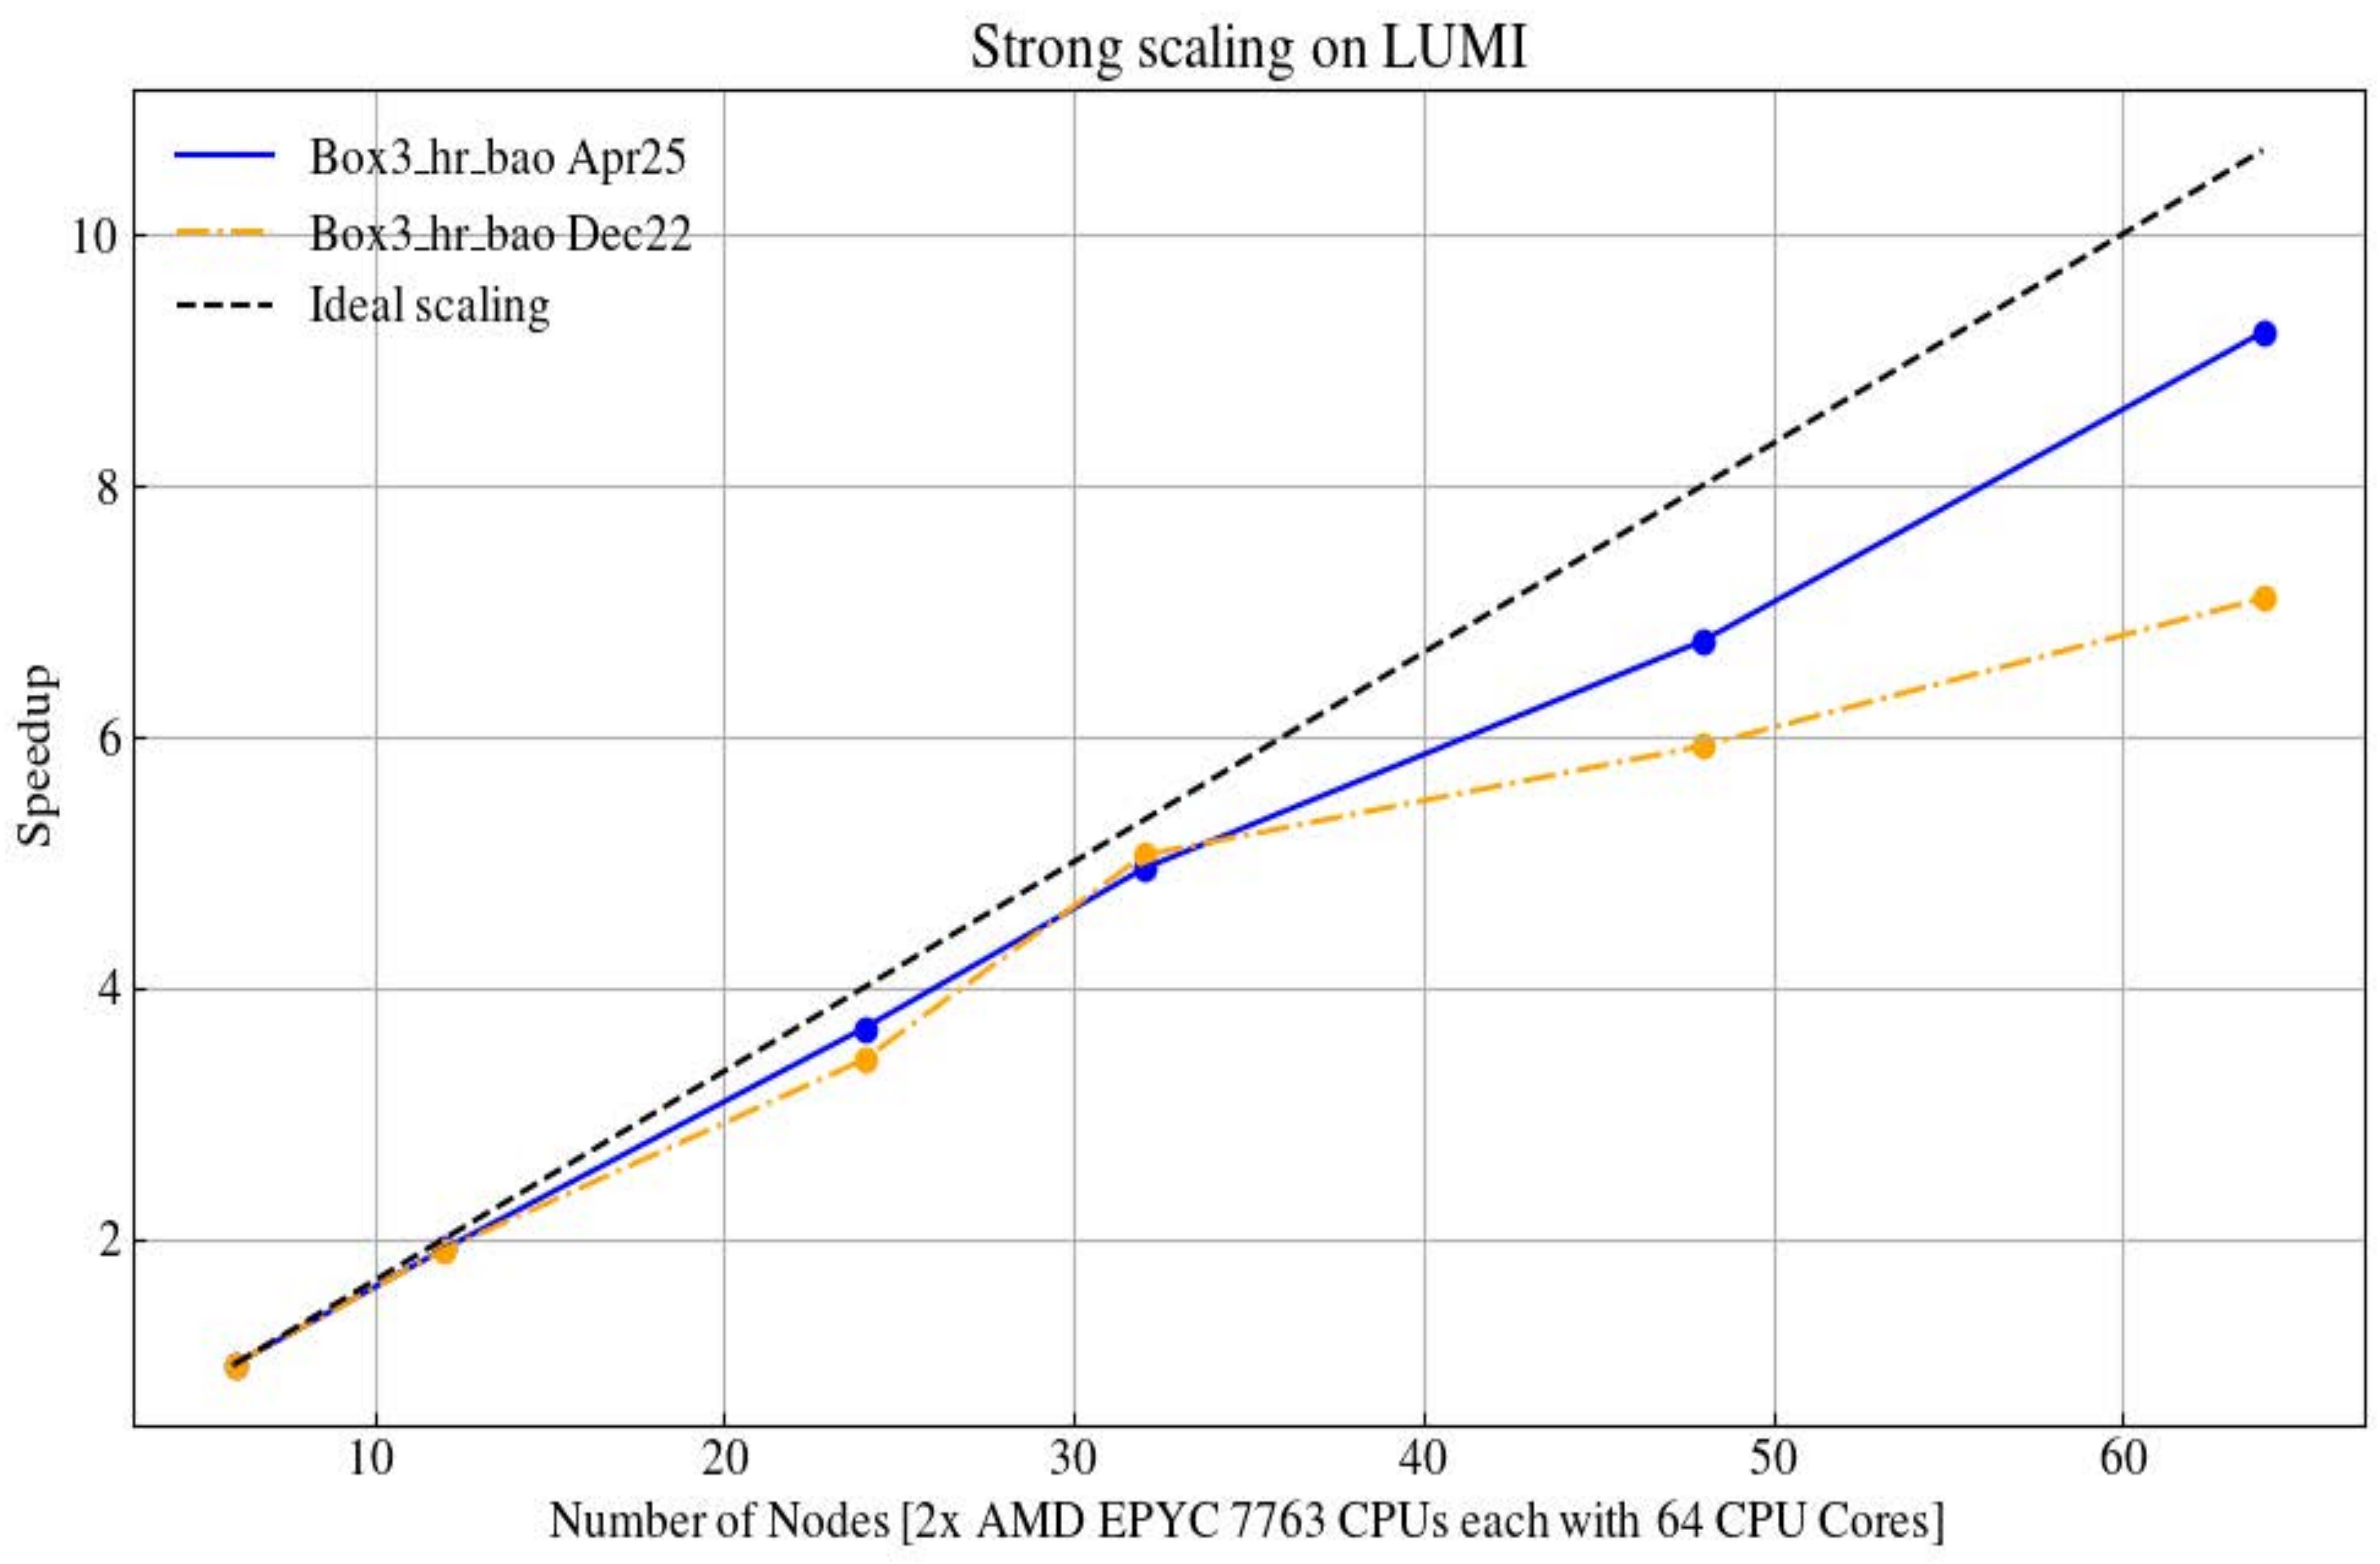

Main Findings and Recommendations

iPIC3D

Main findings

- Excessive use of MPI_Barriers - might cause unnecessary serialization
- Sequences of many small P2P messages
- Most of the time in computations
- Identify data dependencies in communication phases
- Fuse small MPI messages -> lower latency, higher transfer rate
- Reduce number of MPI_Barriers -> less synchronization
- Focus on computations, e.g. vectorization, GPU acceleration

Results

- Speedup degradation at 64 nodes (7168 processes)
- Transfer efficiency (TE) drops due to MPI latency
- CalculateField: TE 2%, severe slowdown
- GatherMoments: TE 41%
- ParticlesMover: very good efficiencies, near-perfect scaling

Simulation setup

Maxwellian distribution, 952 particles per cell, 2 species
2D grid resolution: 2240 → 1120 cells
Simulations on 2–64 nodes
RoI: 2nd cycle, 4 total cycles

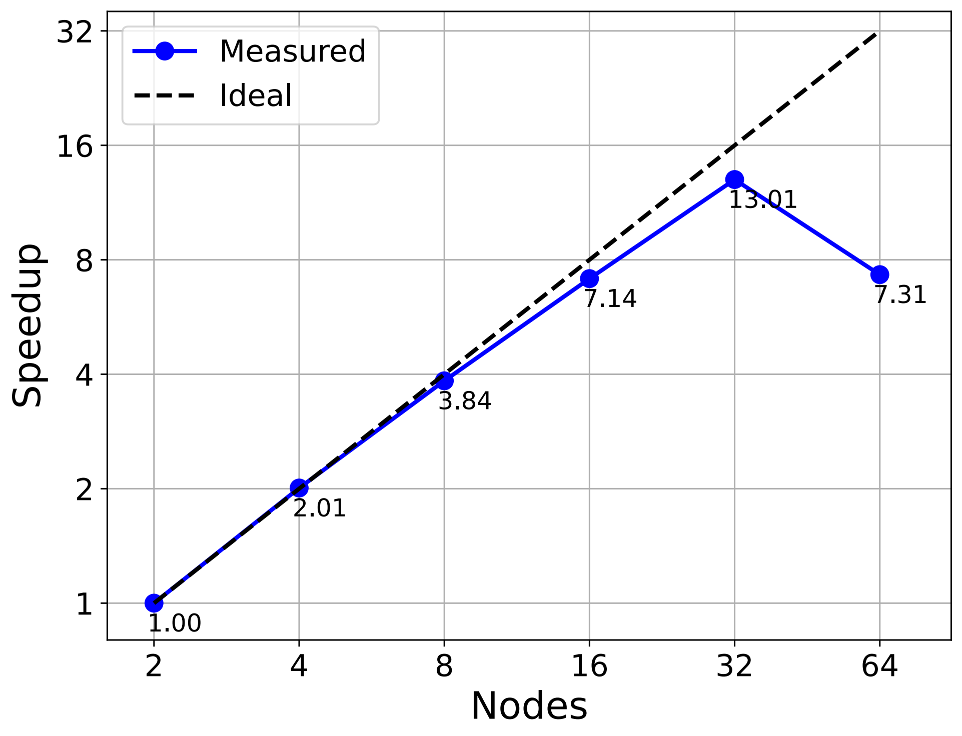

Programming Approaches That Each Code Uses To Port the Code

Keeping **performance** and **portability** in mind

- Each code uses distinct numerical techniques and extract mini-apps for kernel analysis
- OpenGadget3: memory-bound, needs new algorithms
- gPLUTO and iPIC3D have better GPU utilization

Code	NVIDIA	AMD	Portability	**Supporting technology partners**
gPluto	OpenACC	OpenMP target directives	*CPU + GPU*	CINECA
OpenGadget	OpenACC	OpenMP target directives	*CPU + GPU*	CINECA
iPIC	CUDA	HIP	*CPU+GPU*	CINECA (main), E4

02

Transition Pluto -> gPluto

Utilising the capabilities of accelerators

What Is gPLUTO?

Multidimensional compressible plasma with high mach numbers

gPLUTO: the new **GPU-enabled** version of the PLUTO code **targeting exascale facilities** and new generation hardware

- **C++** & **OpenACC, OpenMP** ⇒ high-level directive-based programming models
- Current Status:
  - **80%** of the **static-grid code** ported to GPUs ;
  - Deployed on **all JU Systems** with both CPUs and GPUs;
  - Extensive code revision;

- Public GitLab: https://gitlab.com/PLUTO-code/gPLUTO.git
- Dev GitLab: https://gitlab.com/PLUTO-code/gPLUTO-dev

What Is gPLUTO?

Main gPLUTO kernels

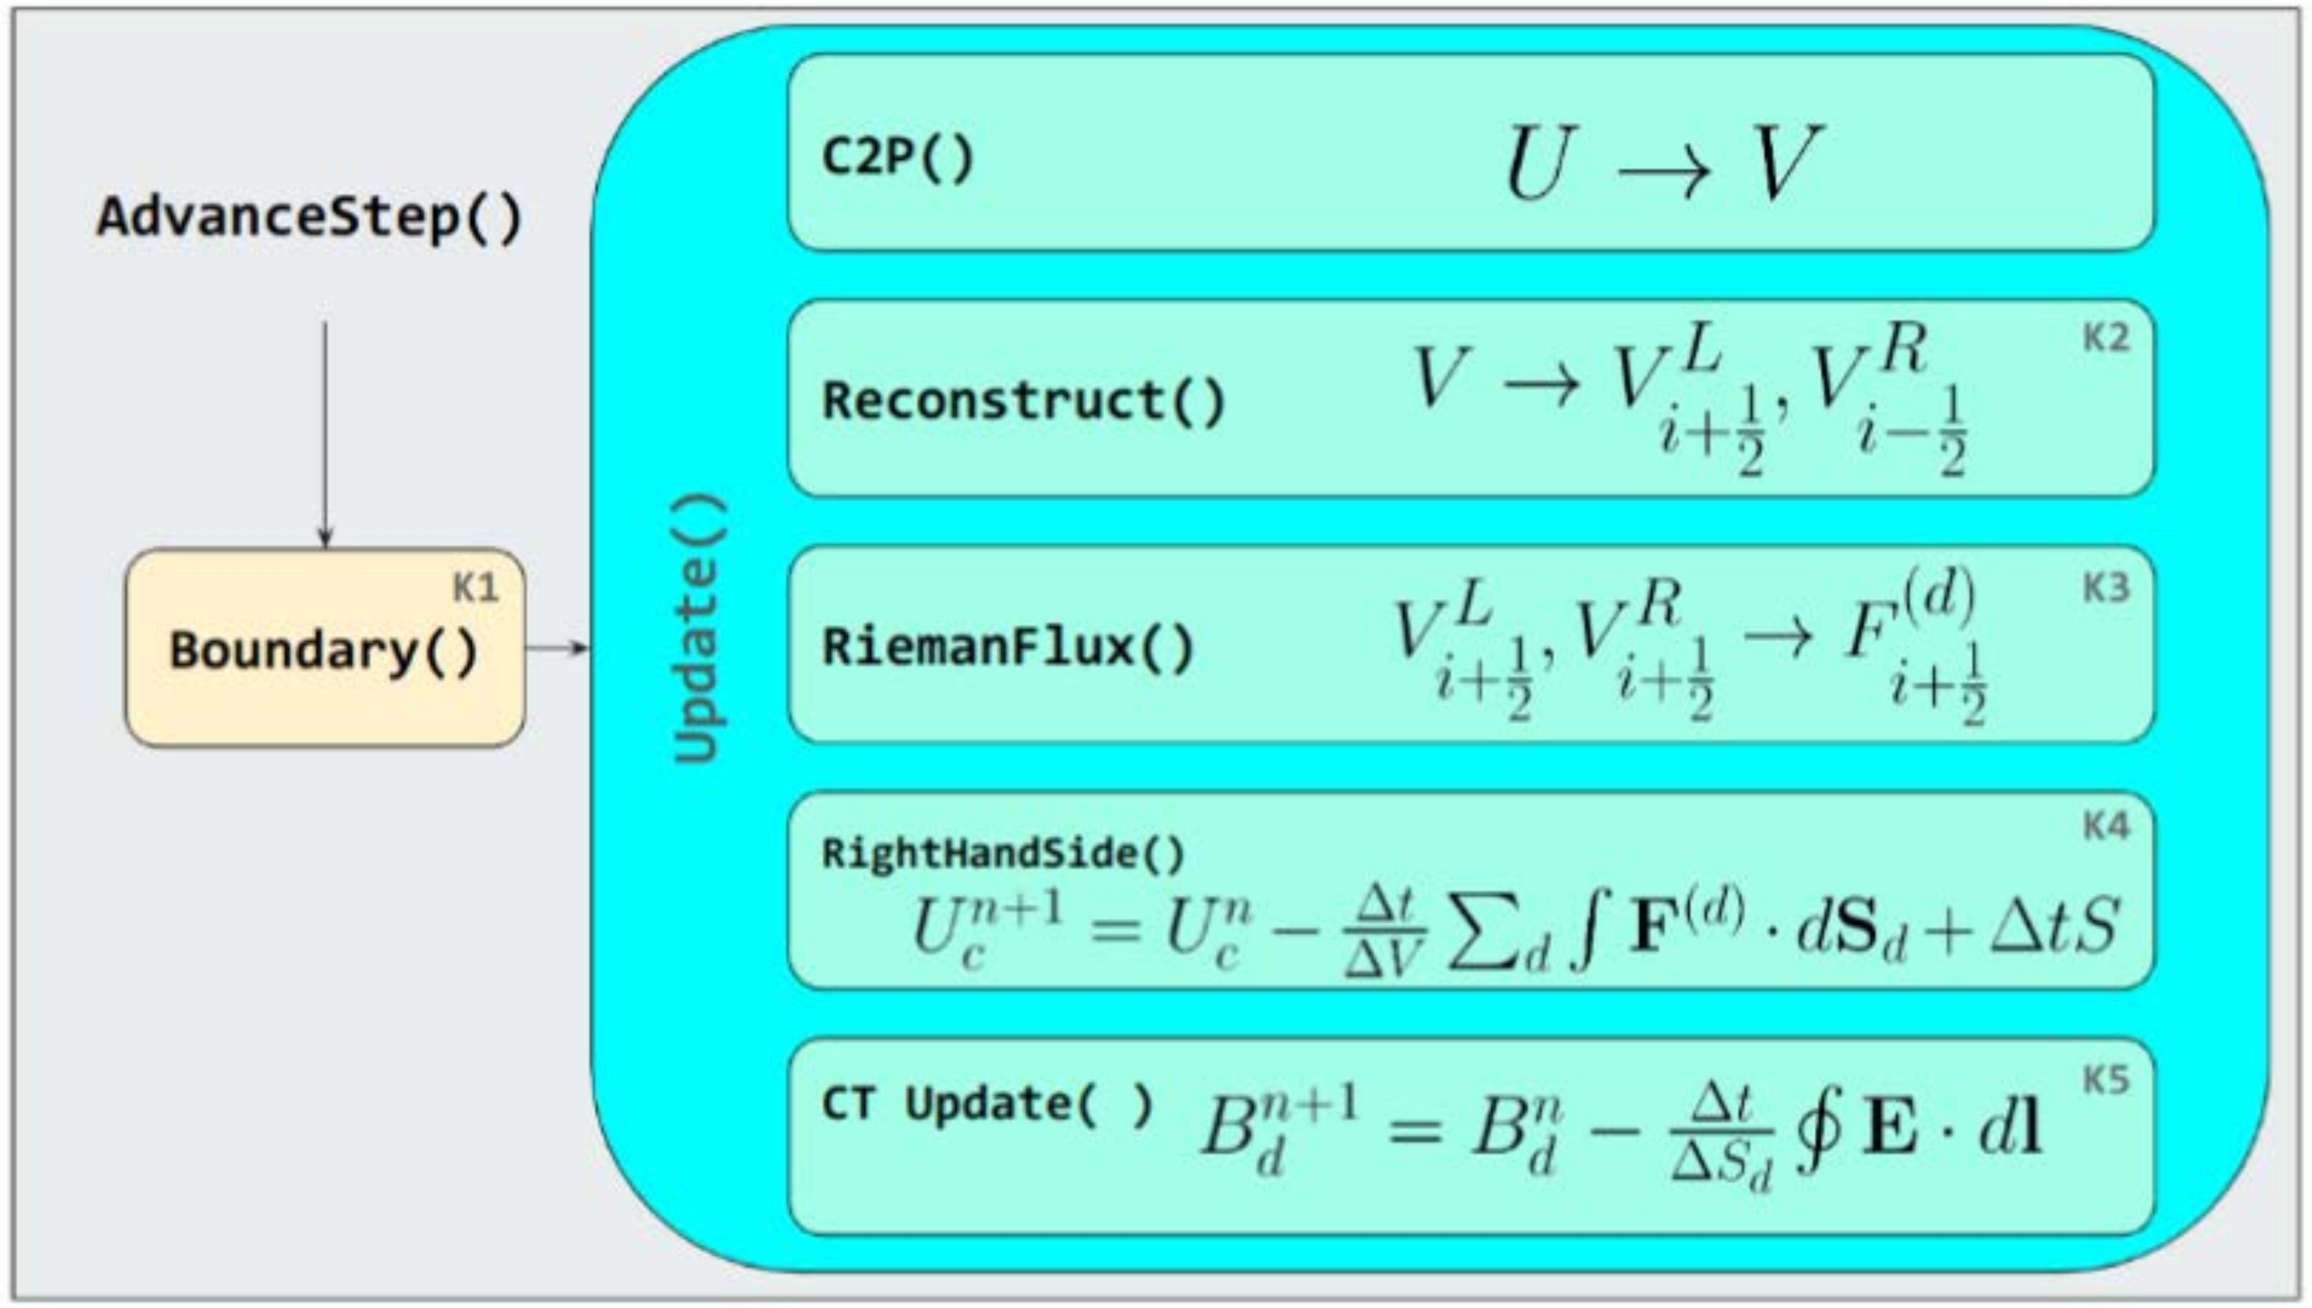

Figure 4: Diagram of the Reconstruct-solve-average (RSA) strategy

Kernel example

```
#pragma acc parallel loop collapse(2)
for (k = KBEG; k <= KEND; k++){
for (j = JBEG; j <= JEND; j++){
  #pragma acc loop
  for (i = 0; i < NX1_TOT; i++){
  for (nv = 0; nv < NVAR; nv++){
    vc[i][nv] = Vc[nv][k][j][i];
  }}

  Reconstruct (...);
  Riemann (...);
  Powell  (...);

  #pragma acc loop
  for (i = IBEG; i <=
  #pragma acc loop
  for (nv = 0; nv < N
    dU[k][j][i][nv] =          ][nv];
  }
  Cdt[k][j][i] = 0.5*
  }
}}
```

Kernel markers

```
nvtxRangePush("Routine_1");
Routine_1();
nvtxRangePop();

nvtxRangePush("Routine_2");
Routine_2();
nvtxRangePop();
```

Code owner: M. Rossazza, A. Mignone

Memory Layout

Optimal memory access = having fastest index be the index of the most inner loop.

PLUTO

- Array ****Uc defined as: `Uc[k][j][i][nv]` with nv = fastest index;
- Many for loops are over spatial indexes k, j, i + inner loop on nv;
- Inner loops <u>executed sequentially</u> ⇒ <u>threads coalesced memory access NOT achieved</u>!

gPLUTO

- Array now defined as C++ template classes:

```
Uc(i,j,k,nv);
```

- Allows for <u>optimal coalesced memory access</u>;
-

Code owner: M. Rossazza, A. Mignone

Coalesced Memory Access

Achieving Highly Optimized Code on Heterogeneous Systems: A Long and Iterative Process

- Coalesced Memory Access: technique allowing optimal usage of global memory bandwidth ⇒ consecutive threads access consecutive memory addresses.

```
#pragma acc parallel loop vector
for (i = ibeg; i <= iend; i++){
#pragma acc loop seq
for (nv = 0; nv < NVAR; nv++) {

  // MUST REVERSE INDICES HERE:

  v[i][nv] *= 2;  ➡  v[nv][i] *= 2;
}}

Using C++ templates:   v[nv][i]  ➡  v(i,nv)
```

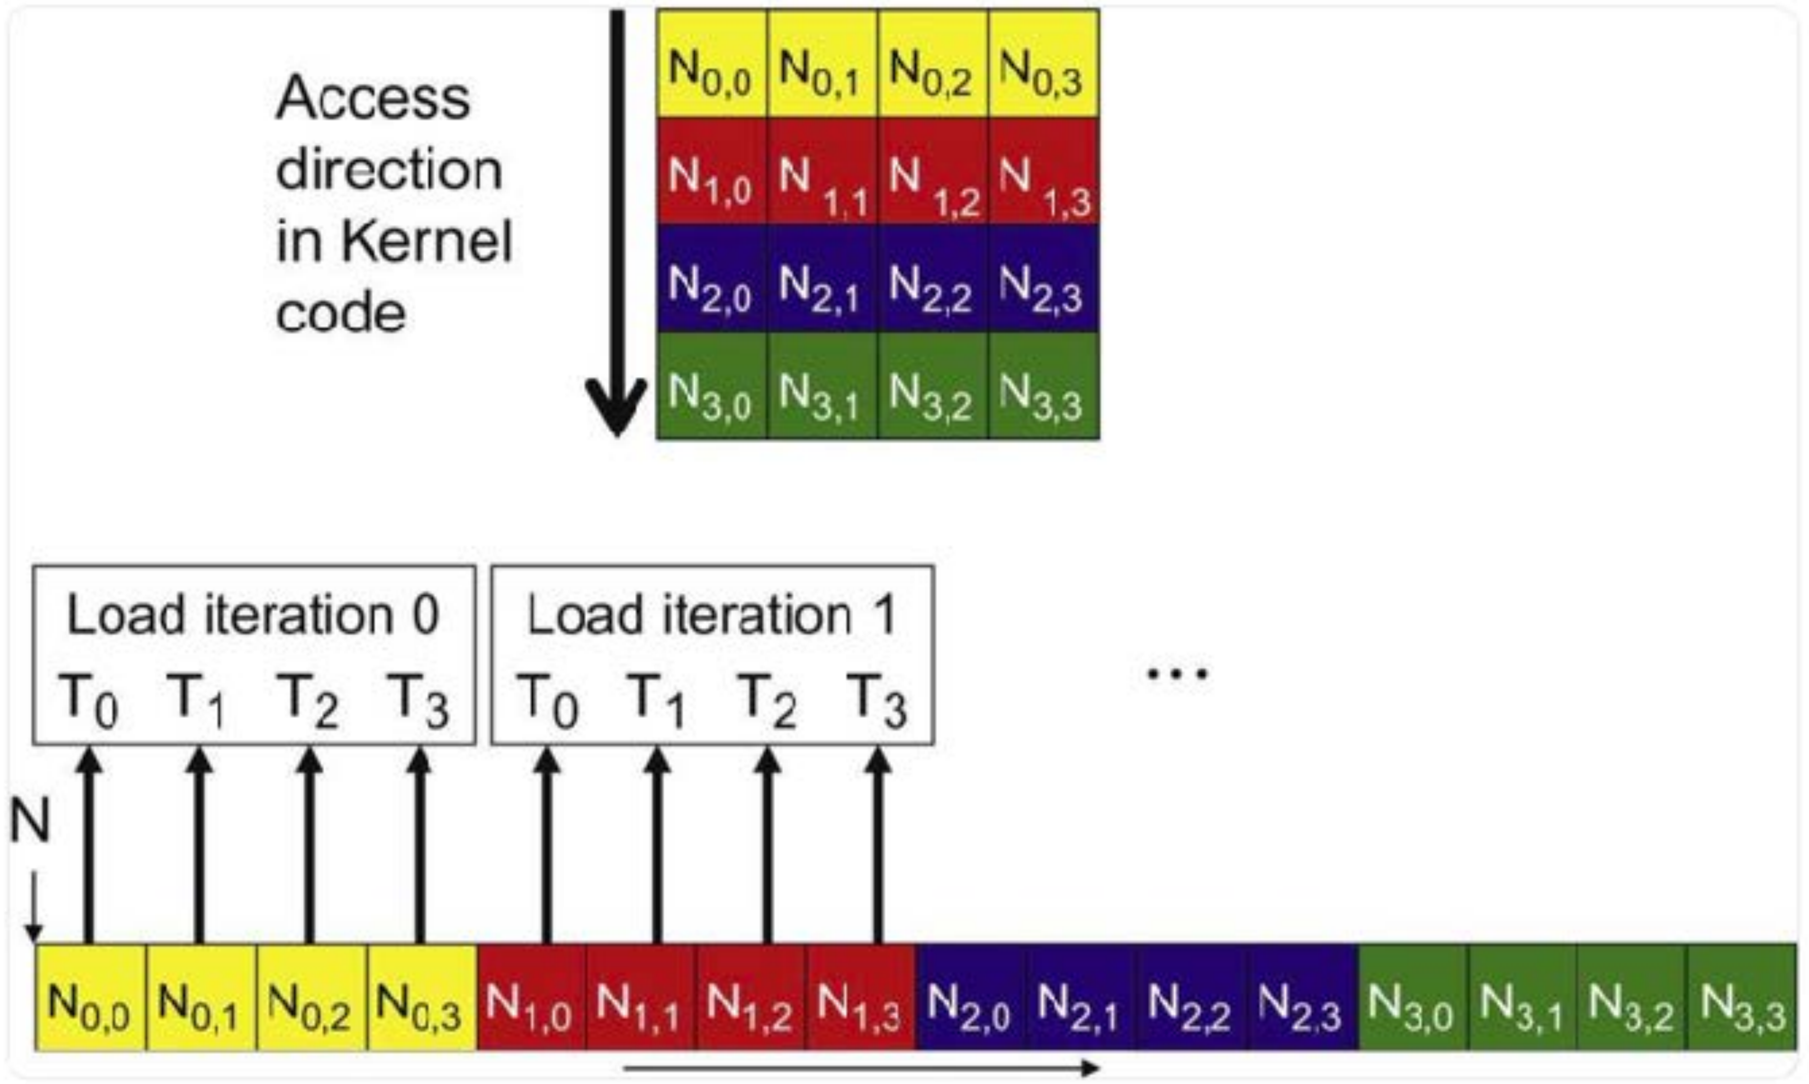

✦ Requires *different* array ordering ⇒ inner loop accelerated is the *fastest index* of the multidimensional array:

Code owner: M. Rossazza, A. Mignone

C++ Features: Classes

- Single memory block rather than multi-pointer approach;
- Order index can be easily modified for optimal memory access;
- Complexity is hidden in classes definition.

```
data->Uc = Array4D(ntot[KDIR], ntot[JDIR], ntot[IDIR], NVAR, double);

char ****Array4D (int nx, int ny, int nz, int nv, size_t dsize)
{

m = (char ****) malloc ((size_t) nx*sizeof (char ***));
m[0] = (char ***) malloc ((size_t) nx*ny*sizeof (char **));
m[0][0] = (char **) malloc ((size_t) nx*ny*nz*sizeof (char *));
m[0][0][0] = (char *) malloc ((size_t) nx*ny*nz*nv*dsize);
...
}

Uc[k][j][i][nv];
```

PLUTO: arrays.c

```
data->Uc.create(ntot[IDIR],ntot[JDIR],ntot[KDIR],NVAR);

void create(int n0, int n1, int n2, int n3) {

  n0_=n0;n1_=n1;n2_=n2;n3_=n3;

  stride0_ = 1;
  stride1_ = n0_;
  stride2_ = n1_*n0_;
  stride3_ = n2_*n1_*n0_;

  total_size_ = n0_*n1_*n2_*n3_;
  data_ = new T[total_size_];
}

Uc(i,j,k,nv);
```

gPLUTO: Array.hpp + arrays.cpp

Code owner: M. Rossazza, A. Mignone

Data Locality and Private Variables

Achieving Highly Optimized Code on Heterogeneous Systems: A Long and Iterative Process

- **Data Locality**: reduce data movement between CPU and GPU memory as much as possible.
- Data transfer *major bottleneck* ⇒ all computational part of the program should reside in GPU memory!

✦ **Private Variables**: GPU threads should perform identical operations but on different memory addresses.

✦ Without precautions, simultaneous operations are performed at the same memory address leading to incorrect results!

✦ Private variables have local scope and are allocated individually for each thread.

```
int V[8];
int A[NX][NY][NZ];

#pragma acc parallel loop collapse(3) private(V[:8])
for (i = 0; i < NX; i++){
for (j = 0; j < NY; j++){
for (k = 0; k < NZ; k++){

  A[i][j][k] *= 2.0;

  V[0] = ...;
  V[1] = ...;
  ...
}}}
```

Code owner: M. Rossazza, A. Mignone

Numerical Benchmarks for gPLUTO (~90 % Average Efficiency)

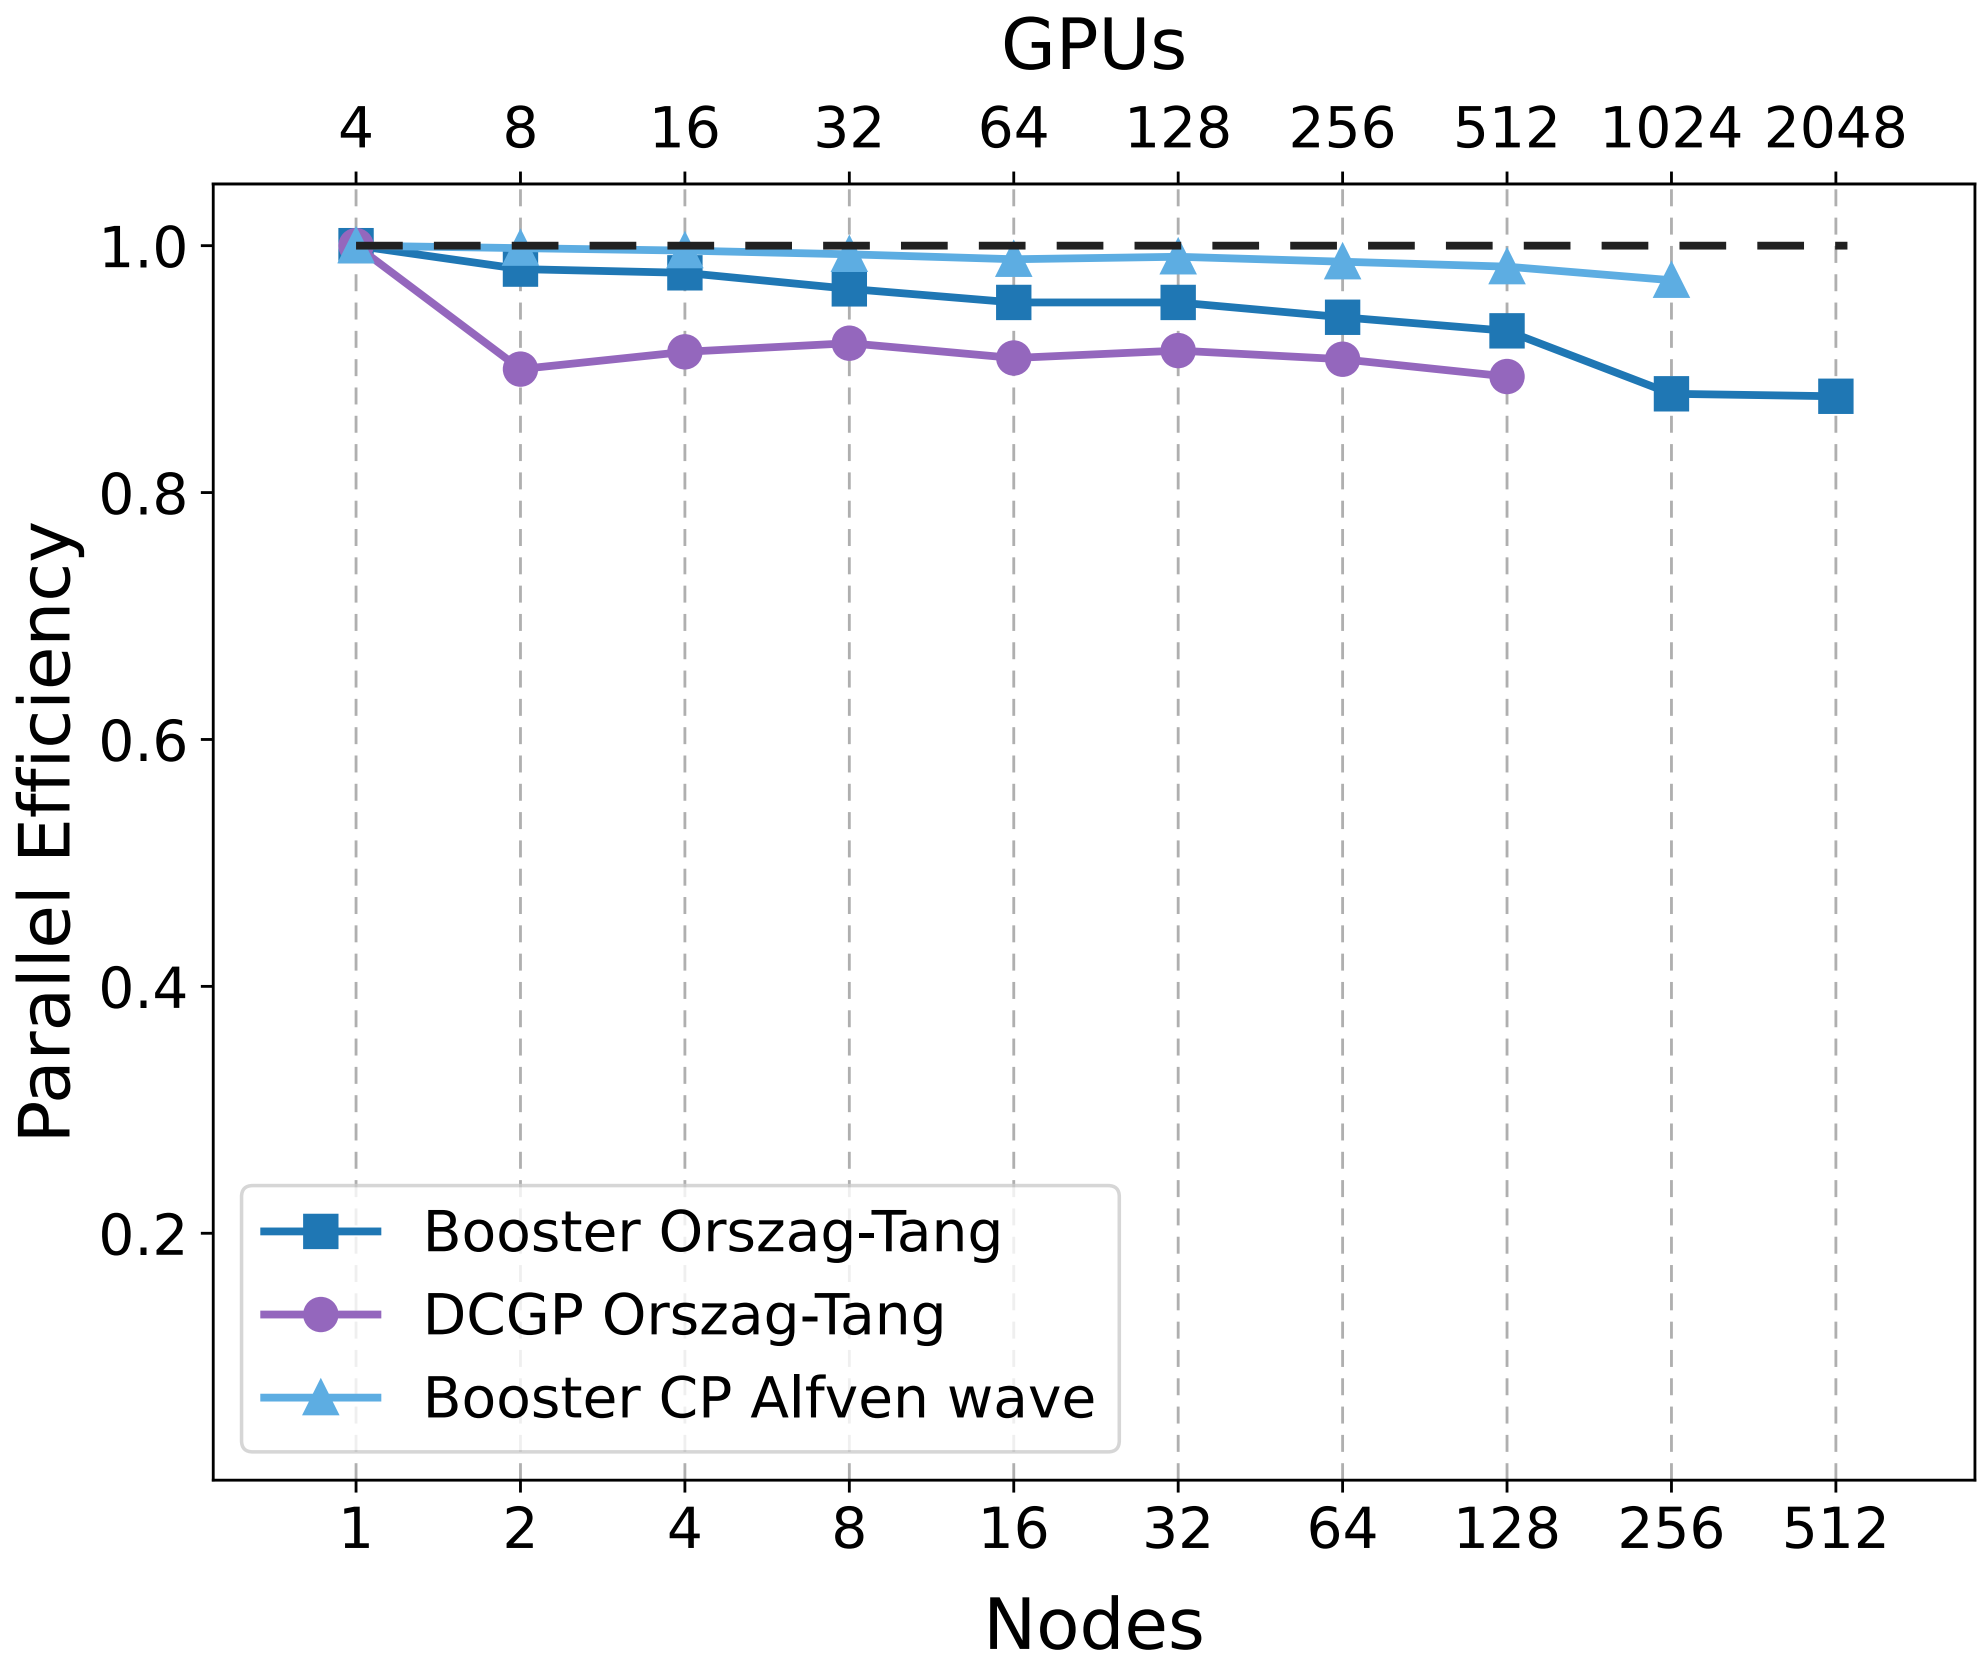

Weak scaling tests for gPLUTO on both Leonardo Booster and DCGP partitions

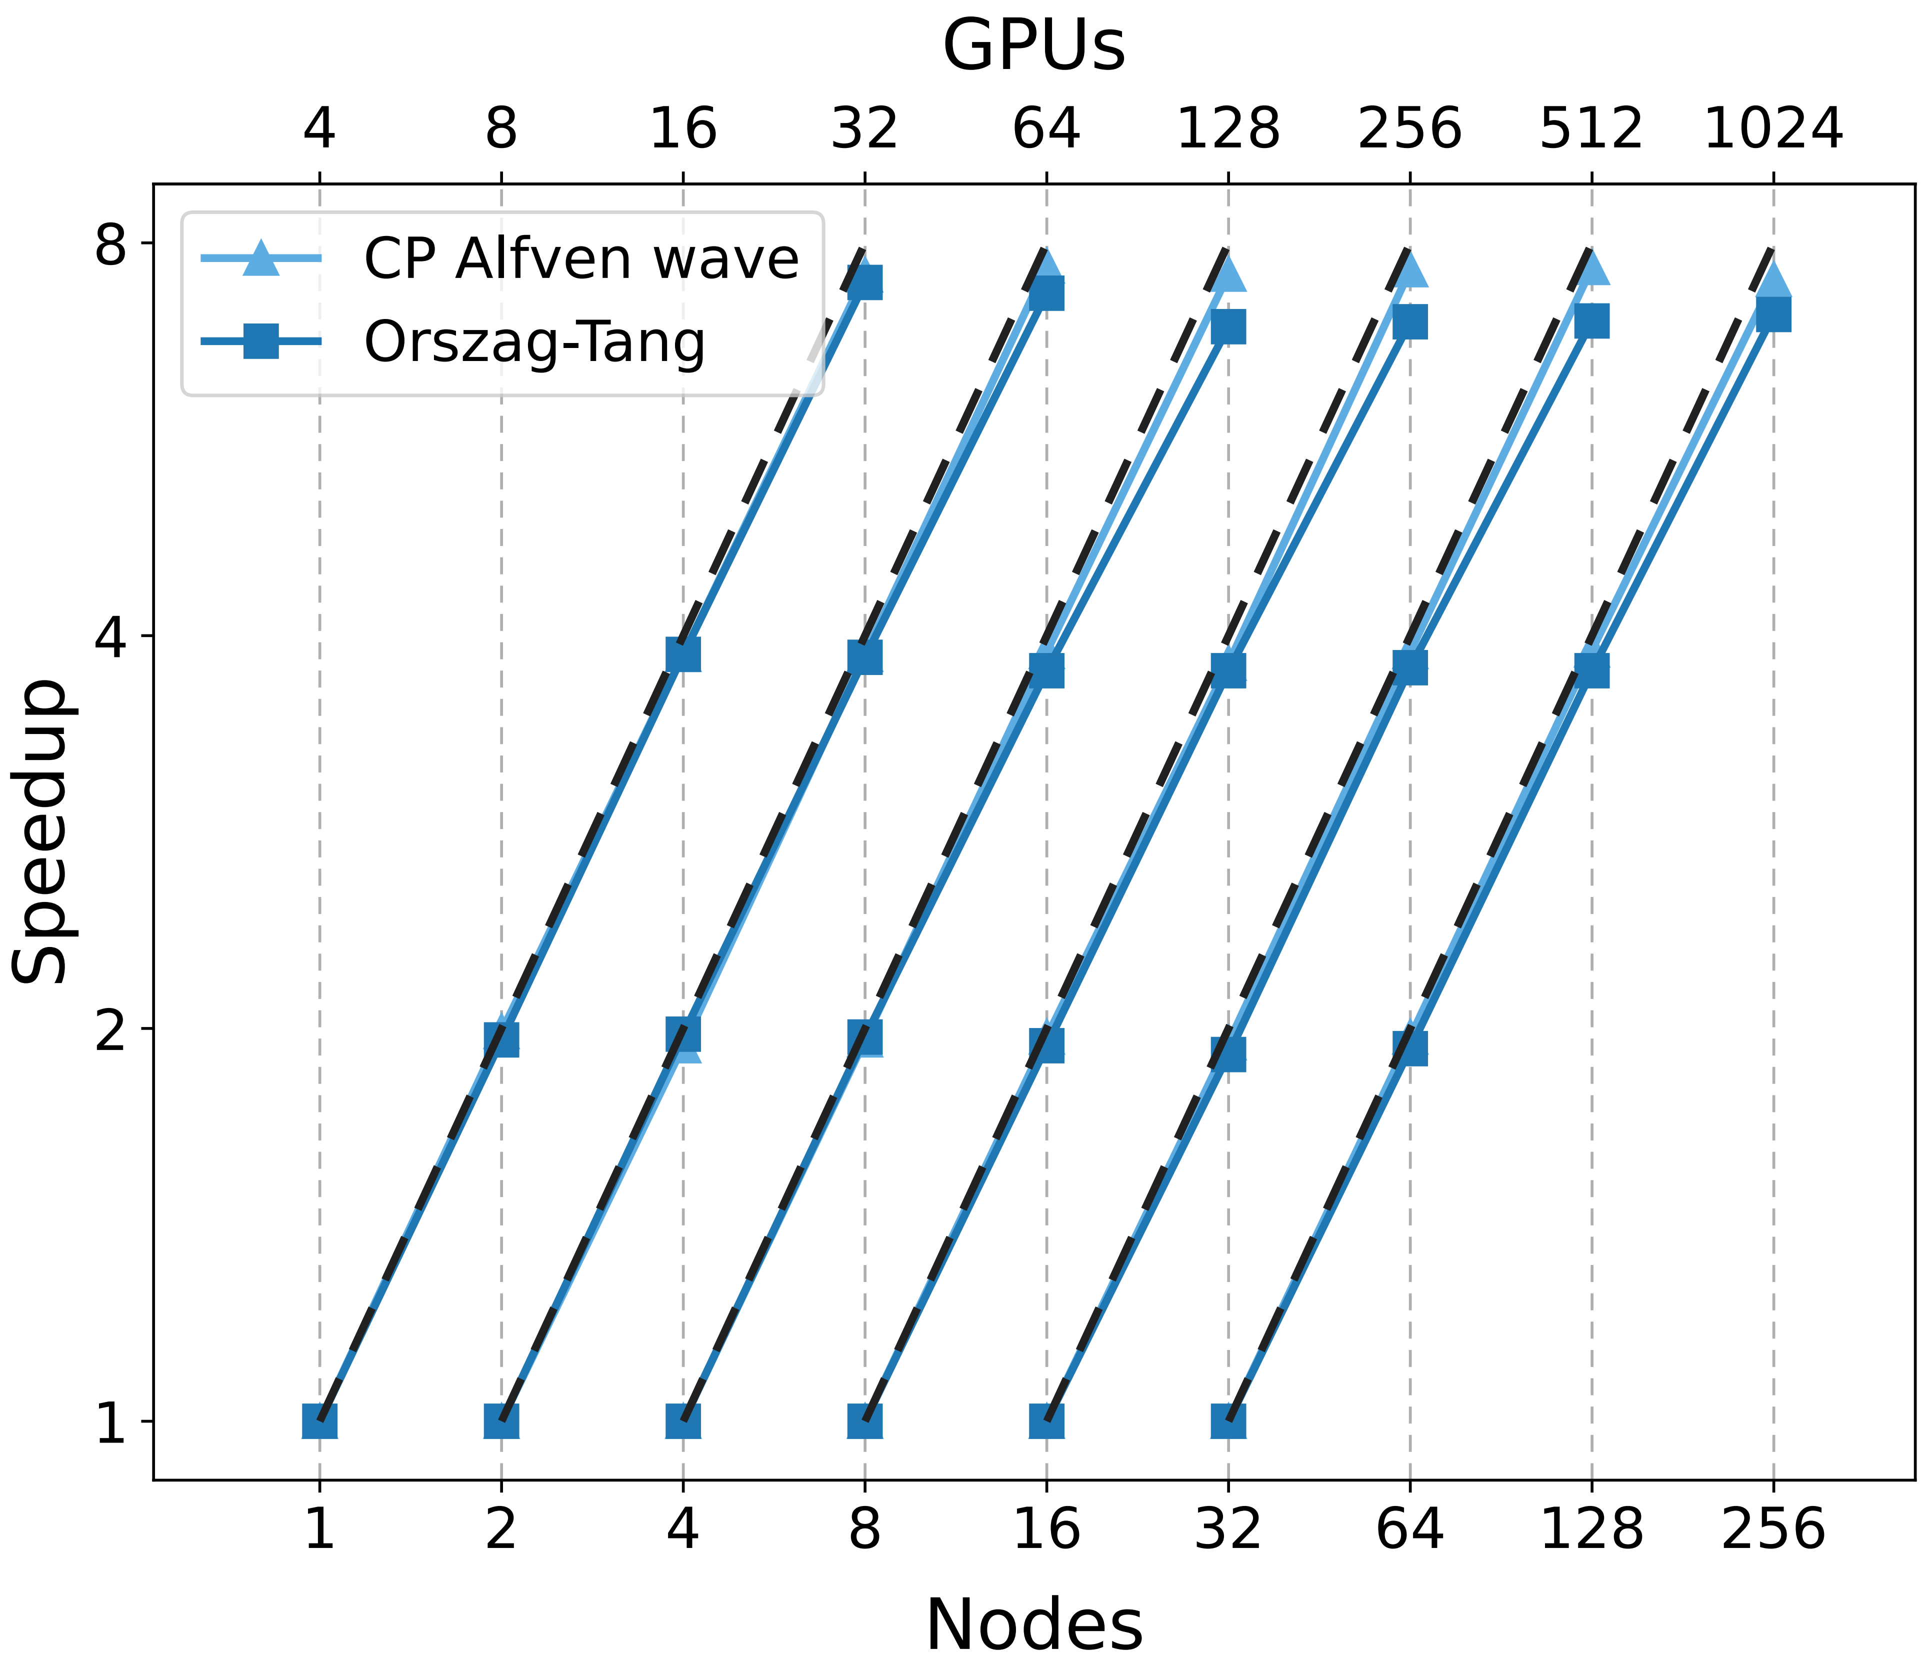

Strong scaling tests for gPLUTO on Leonardo Booster

Available Physics Modules (~60 % Ported on GPU Version)

Achieving Highly Optimized Code on Heterogeneous Systems: A Long and Iterative Process

Advection Physics (Hyperbolic PDE)

- Hydrodynamics (HD)
- Magnetohydrodynamics (MHD)
- Relativistic Hydrodynamics (RHD)
- Ideal and resistive relativistic MHD (RMHD - ResRMHD)

Source Terms

- Gravity / Body forces
- Cooling
- Heating / optically thin
- Chemical networks

Geometry

- Cartesian
- Cylindrical
- Spherical

Dissipation Physics (Parabolic PDE)

- Viscosity (Navier-Stokes)
- Thermal conduction (hydro and MHD)
- Hall MHD, Ambipolar diffusion, Magnetic resistivity
- Radiation Hydrodynamics (FLD, 2 temp)

Particle Physics

- Lagrangian particles
- Cosmic Ray
- Dust

Thermodynamics

- Ideal
- Isothermal
- Non-Constant gamma
- Synge Gas (relativistic)

LEGEND

- Ported
- in progress
- Not ported

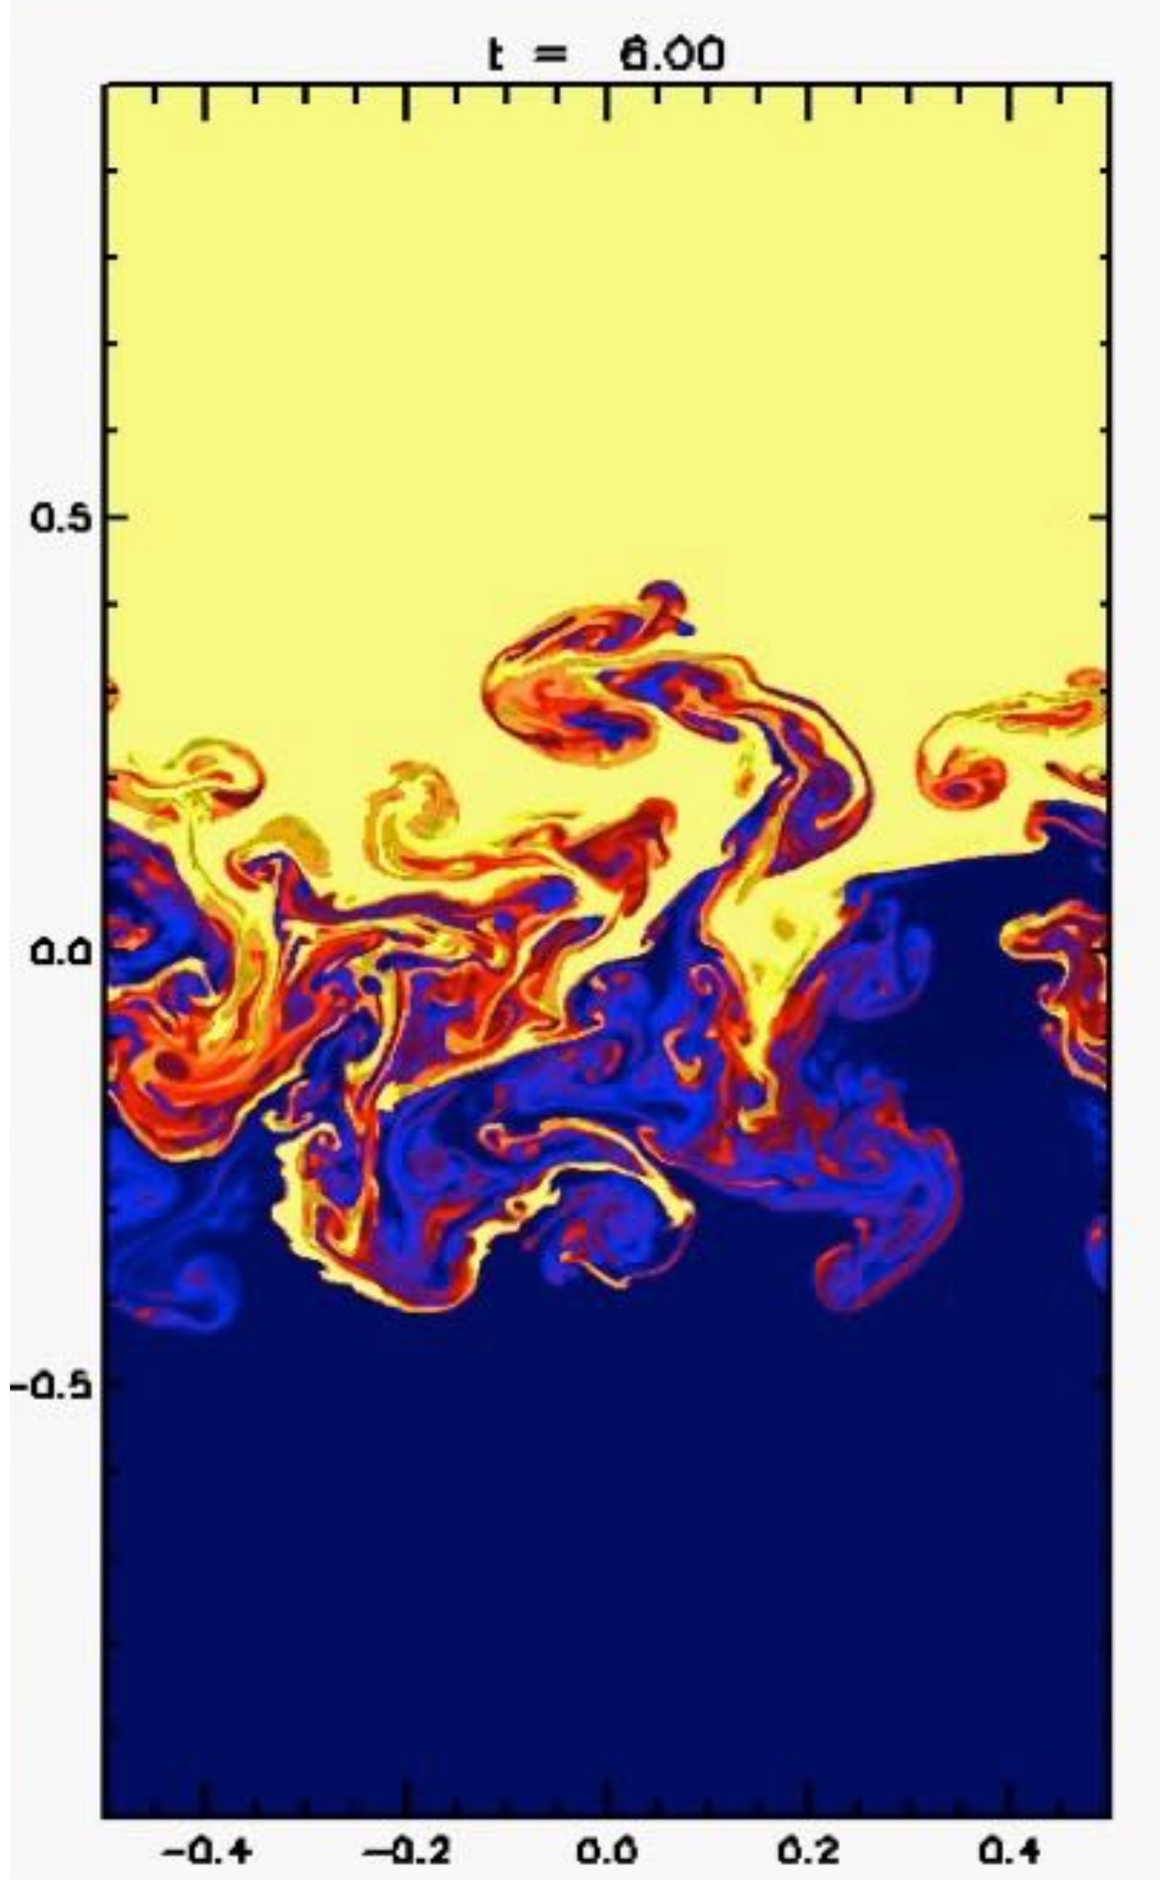

An example of a 2D Rayleigh-Taylor instability with density contrast of 2.

03

Transition OpenGaget -> OpenGaget3

Utilising the capabilities of accelerators

What Is openGadget3?

N-body cosmological simulations:

- ✦ Formation of large scale structure
- ✦ Galaxy formation and evolution
- ✦ Gas dynamics
- ✦ star formation and feedback processes
- ✦ black holes growth and AGN feedback
- ✦ Magnetic fields
- ✦ And many more…

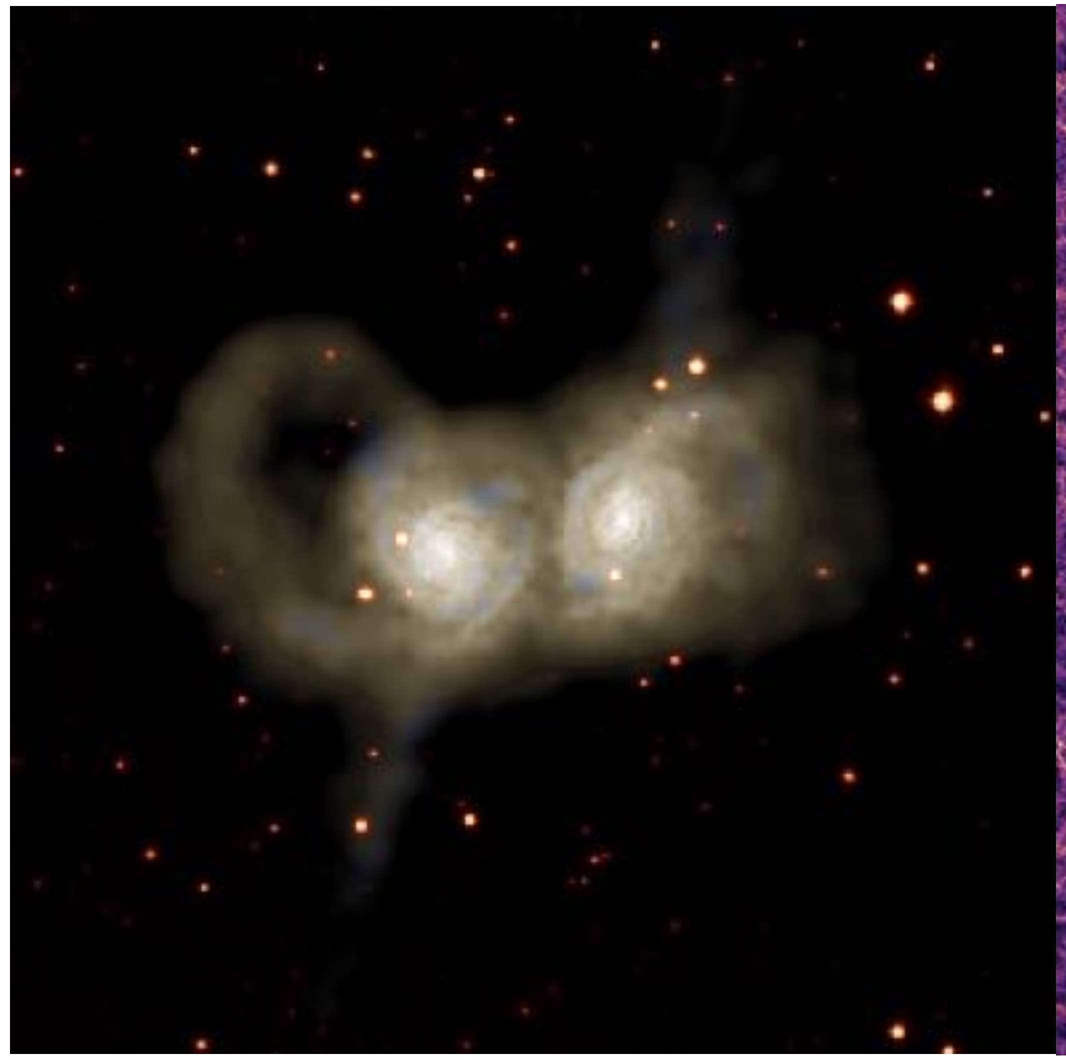

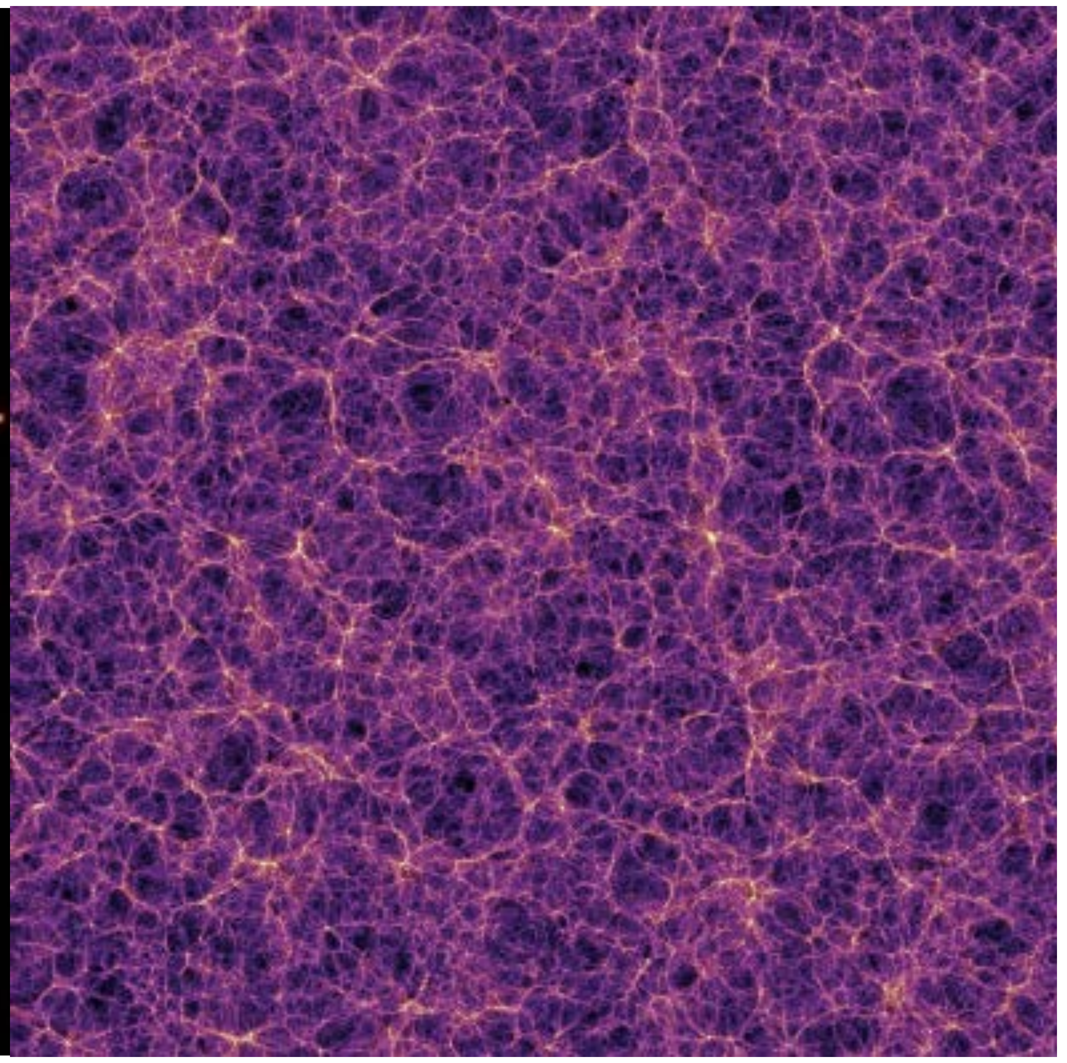

OpenGadget3 - Development

main_development | OpenGadget3 / | +

- Written in C/C++
- Parallelized with MPI and OpenMP
- I/O with custom binary files and HDF5
- OpenACC GPU implementation
- NEW: OpenMP GPU offloading

Code owner: G. S. Karademir, K. Dolag, L. Tornatore

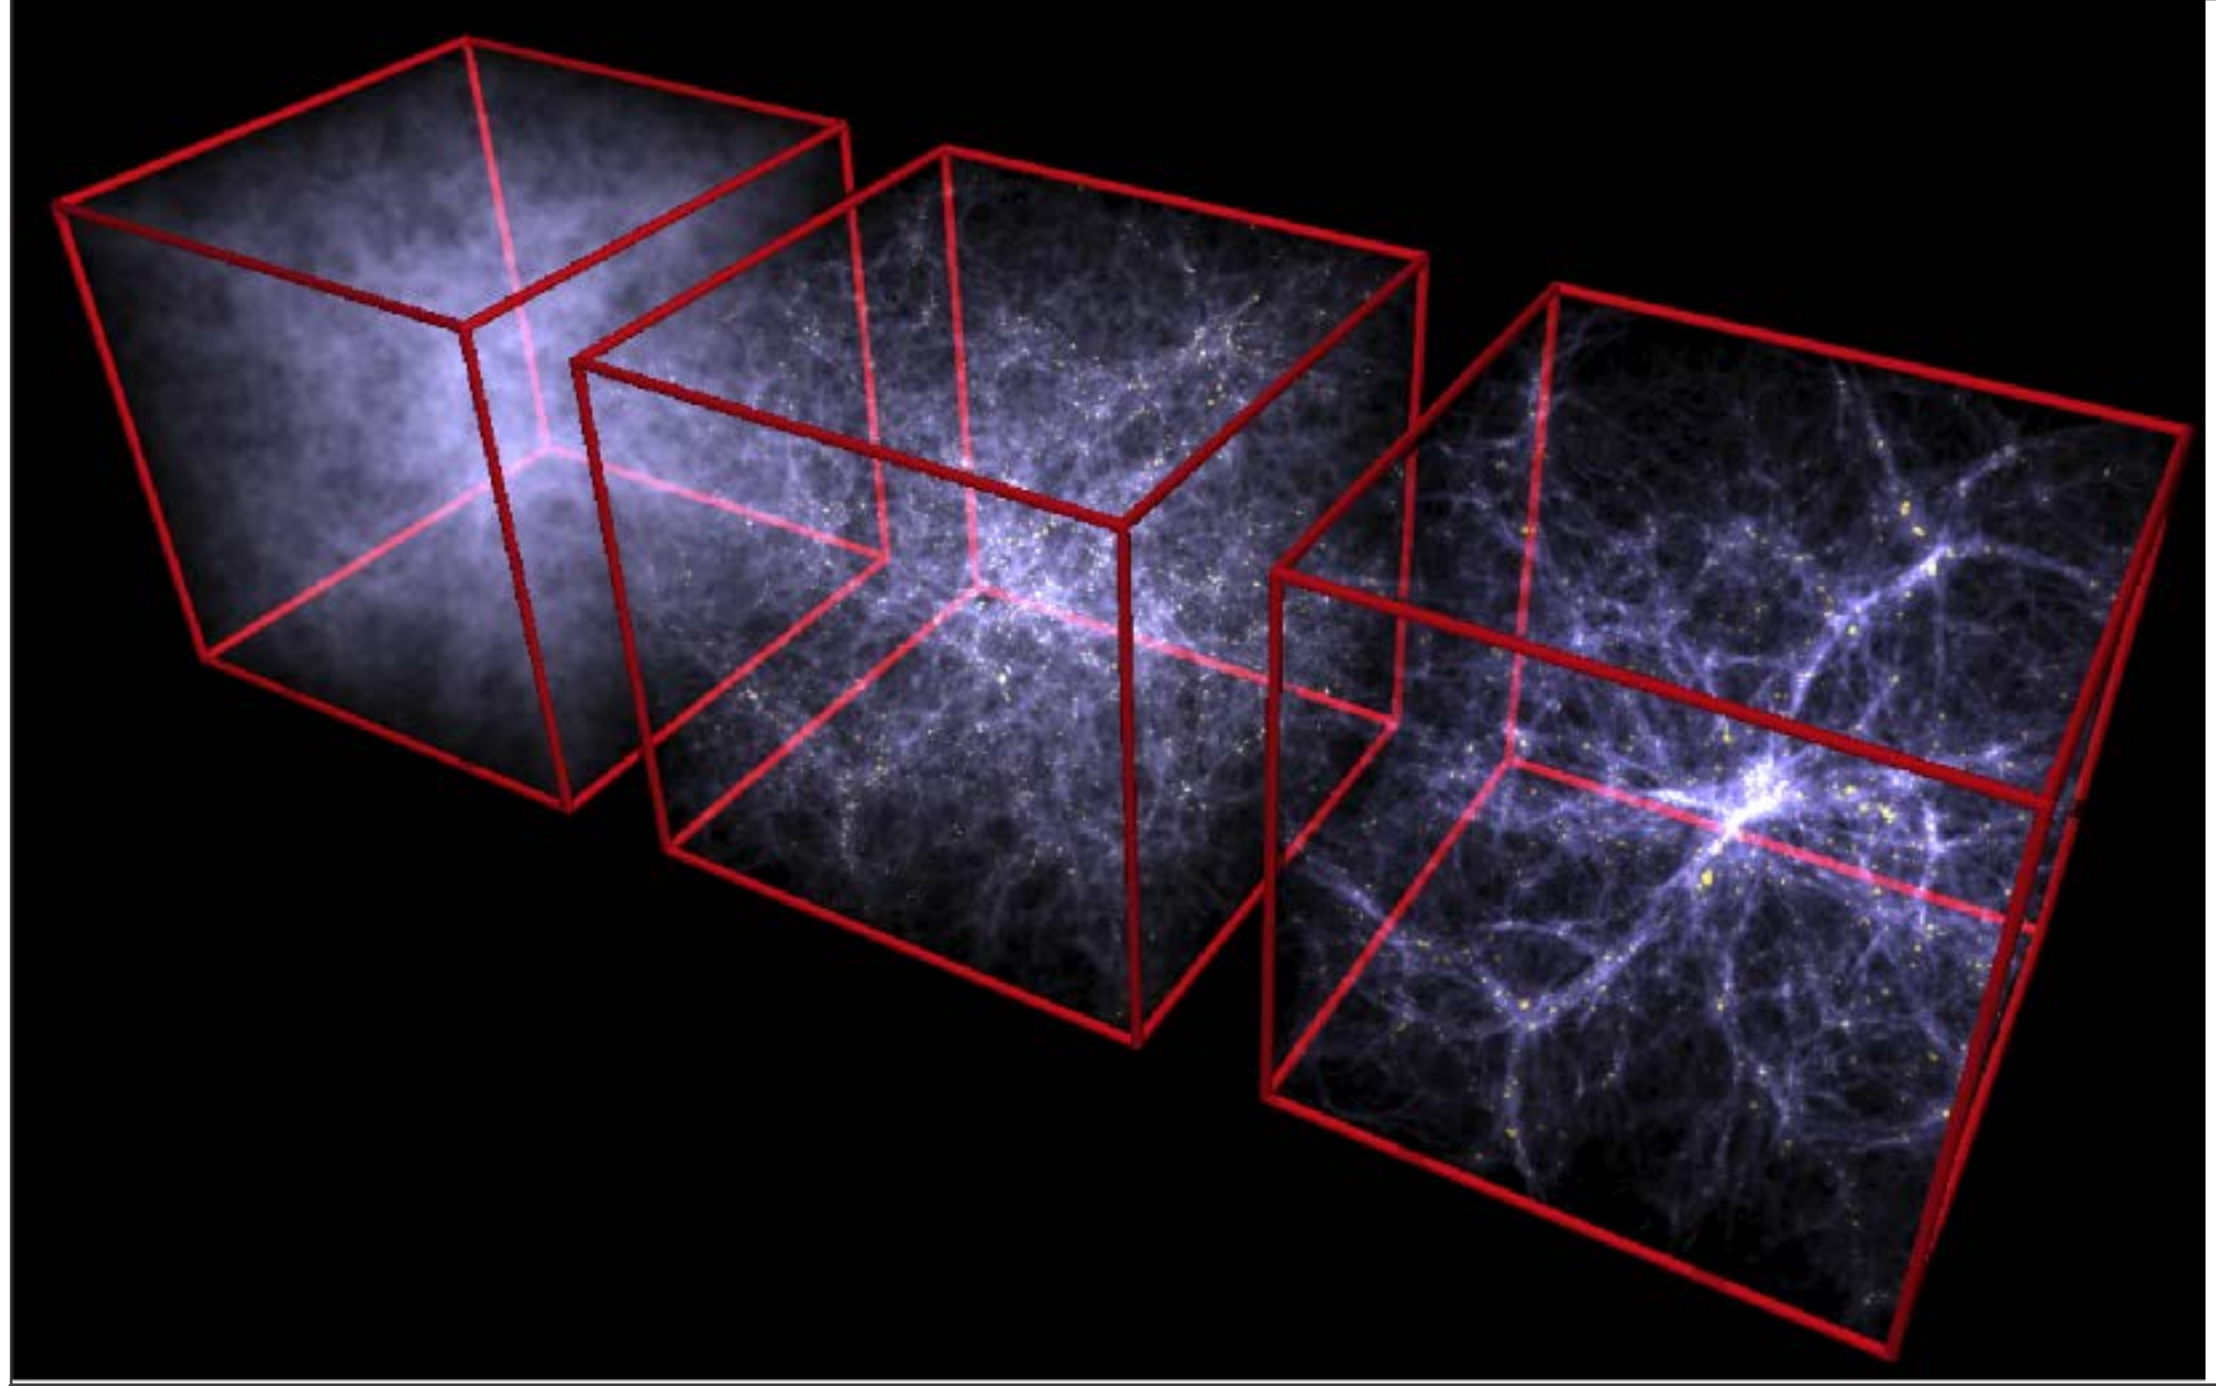

Gadget-like codes evolve astrophysical structures with adaptive time steps on Tree (Barnes-Hut) +PM for gravity and Smoothed Particle Hydrodynamics (SPH) for gas, parallelized with MPI+OpenMP (image: Springel)

OpenGadget3 GPU porting:

- GPU calculate physics interactions async while CPU search exchanges on boundaries and perform MPI com
- OpenACC allows for code reuse and run time-steps on CPU if they have few active particles

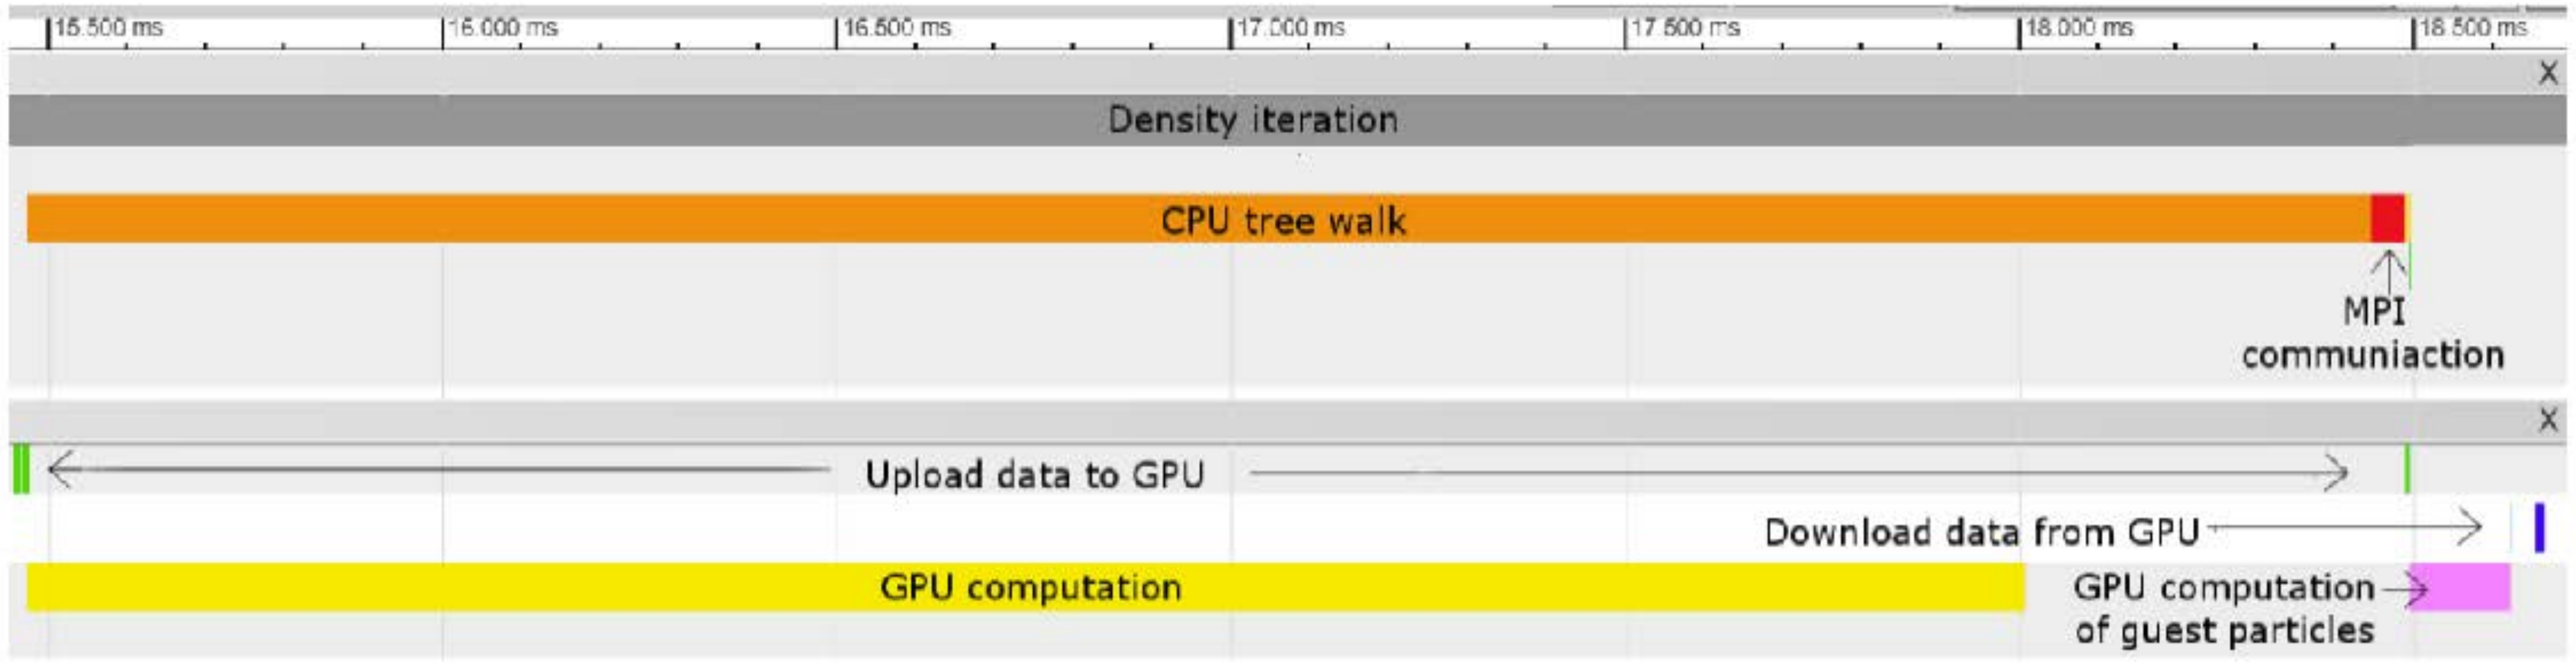

(Ragagnin et al. 2020)

Ragagnin et al.
GPU asynchronous porting:

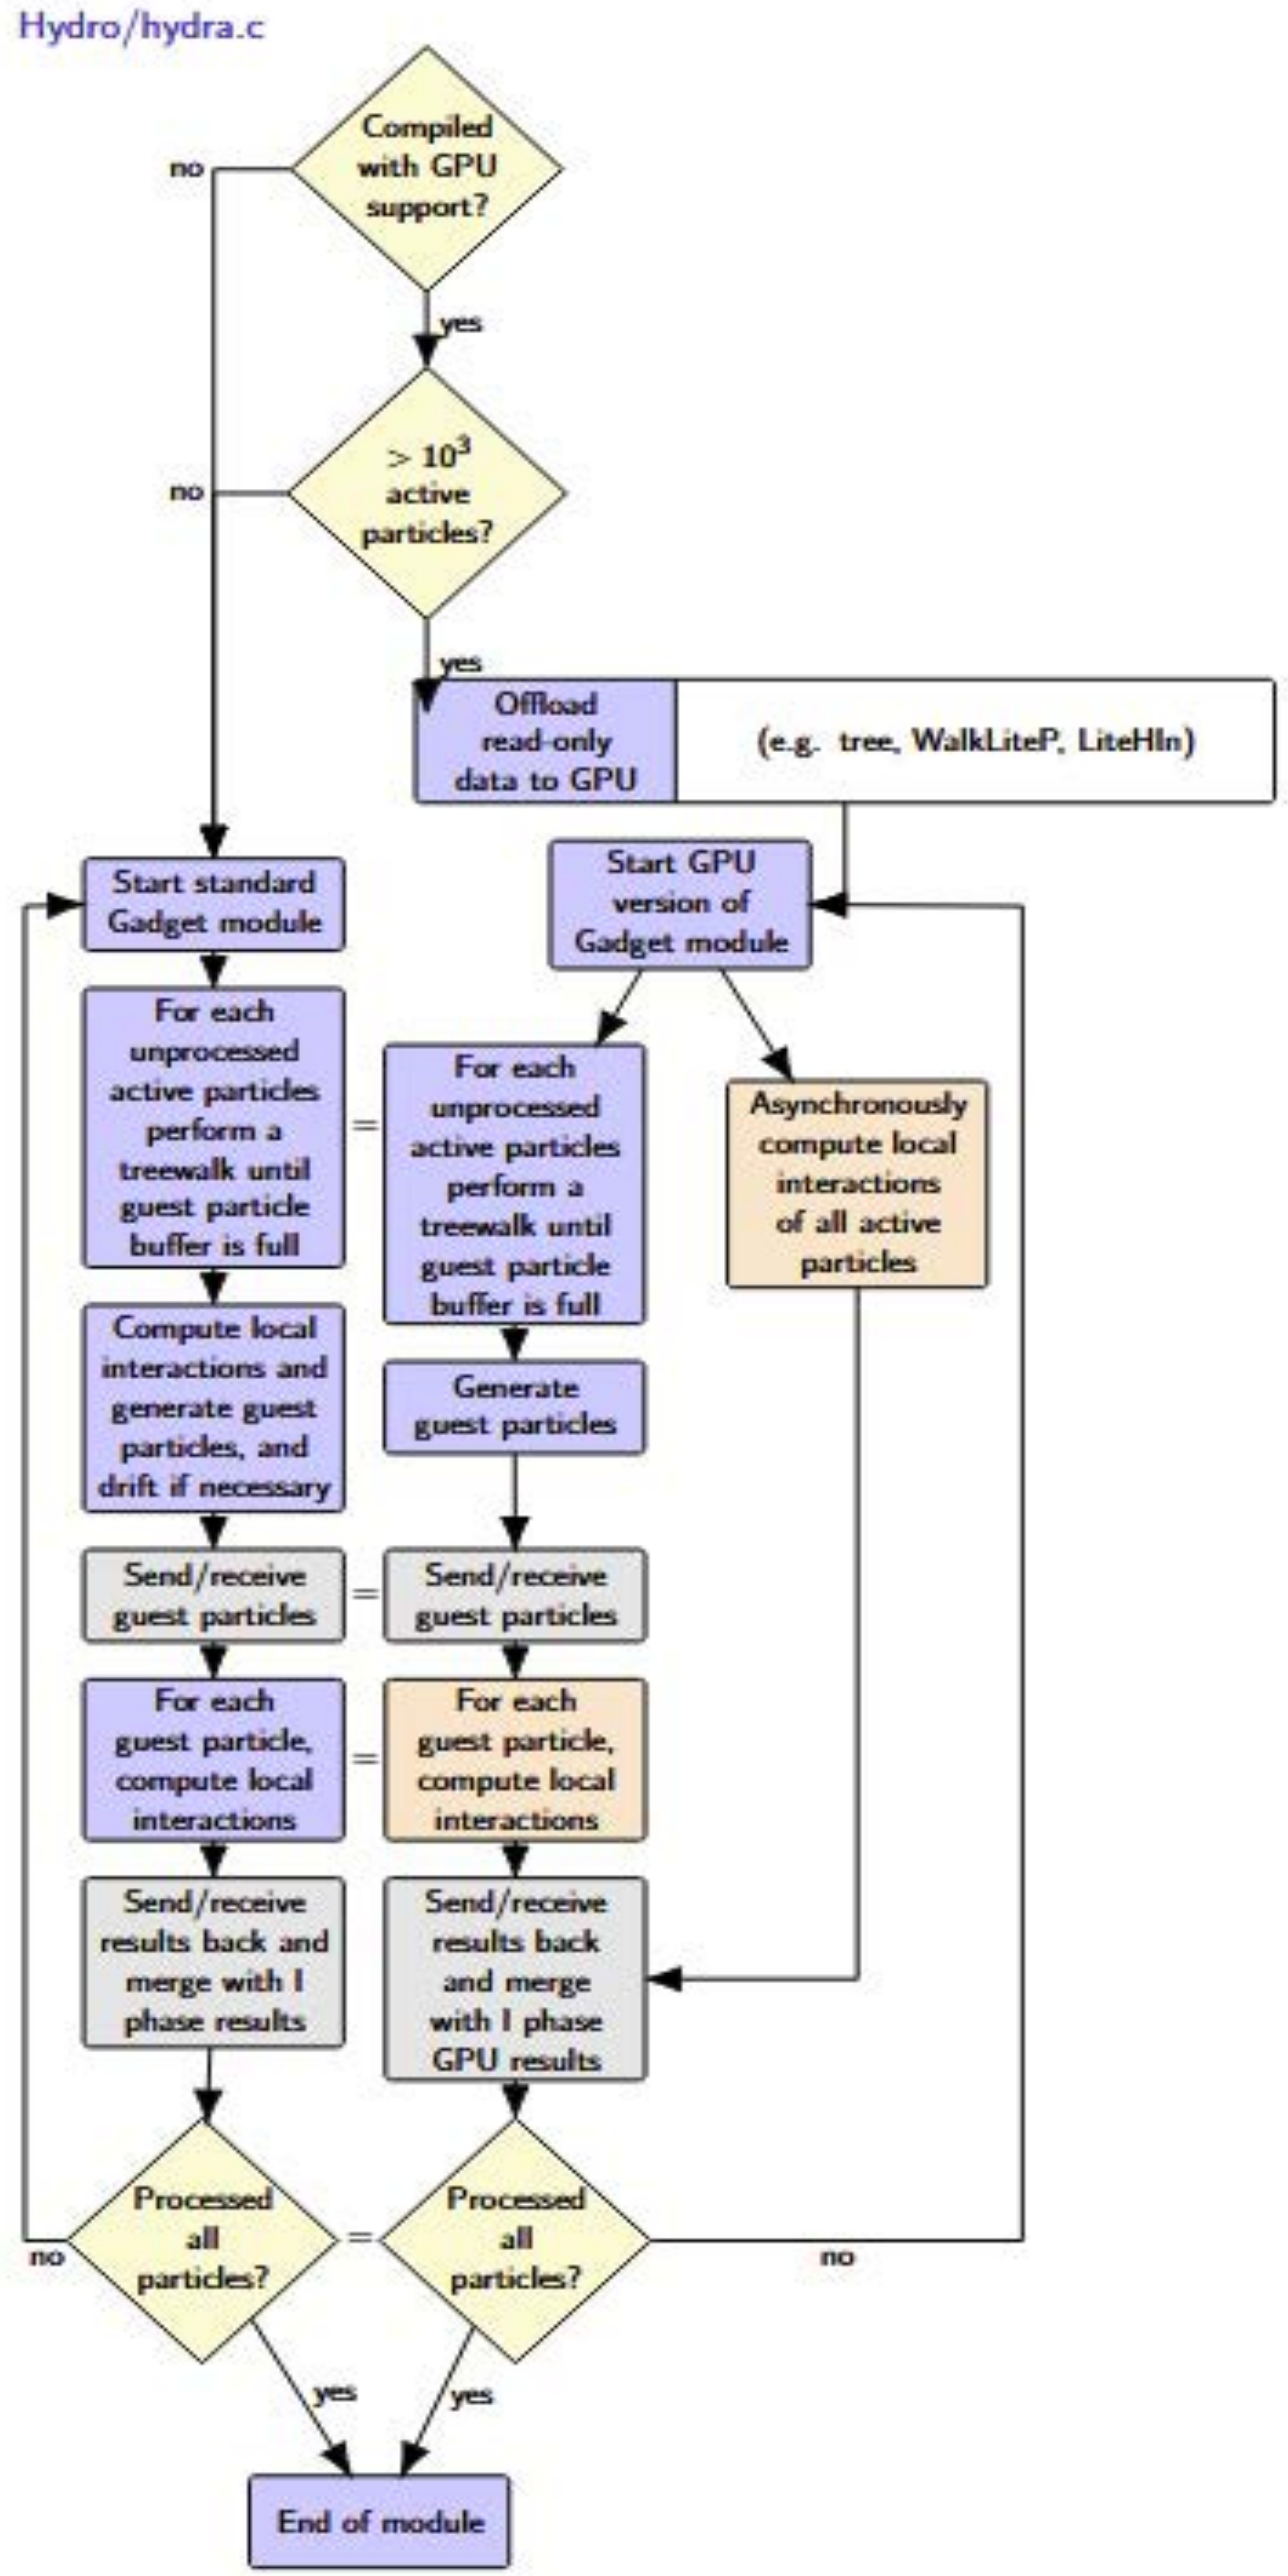

Performance Testing of OpenGadget3

WEAK scaling results obtained on LEONARDO (in joint effort with **CINECA** partners)

Scaled up to FULL Machine on Leonardo Booster ~ 8000 GPUs

Performance Testing of OpenGadget3

- **Test Systems:** Leonardo Booster (main), DCGP (comparative & full-physics tests)
- **GPU Support:** Still in progress; current results reflect limited physics setups
- **Simulation Goal:** Models cosmic evolution using dark matter & gas particles in 3D
- **Key Physics:** Gravity and hydrodynamics (SPH) dominate runtime (~40% each)
- **High-Resolution Runs: Require many particles and large volumes → intensive short-range (Barnes & Hut) calculations**
- **Redshift Dependency:**
  - Gravity and hydrodynamics (SPH) dominate runtime (~40% each)
  - **High redshift:** ~**10× speedup (homogeneous matter → simpler gravity calculations)**
  - **Low redshift (z = 0): ~ 2x speedup (dense structures → complex gravity)**
  - **Overall effective speedup:** ~ 4x-5x

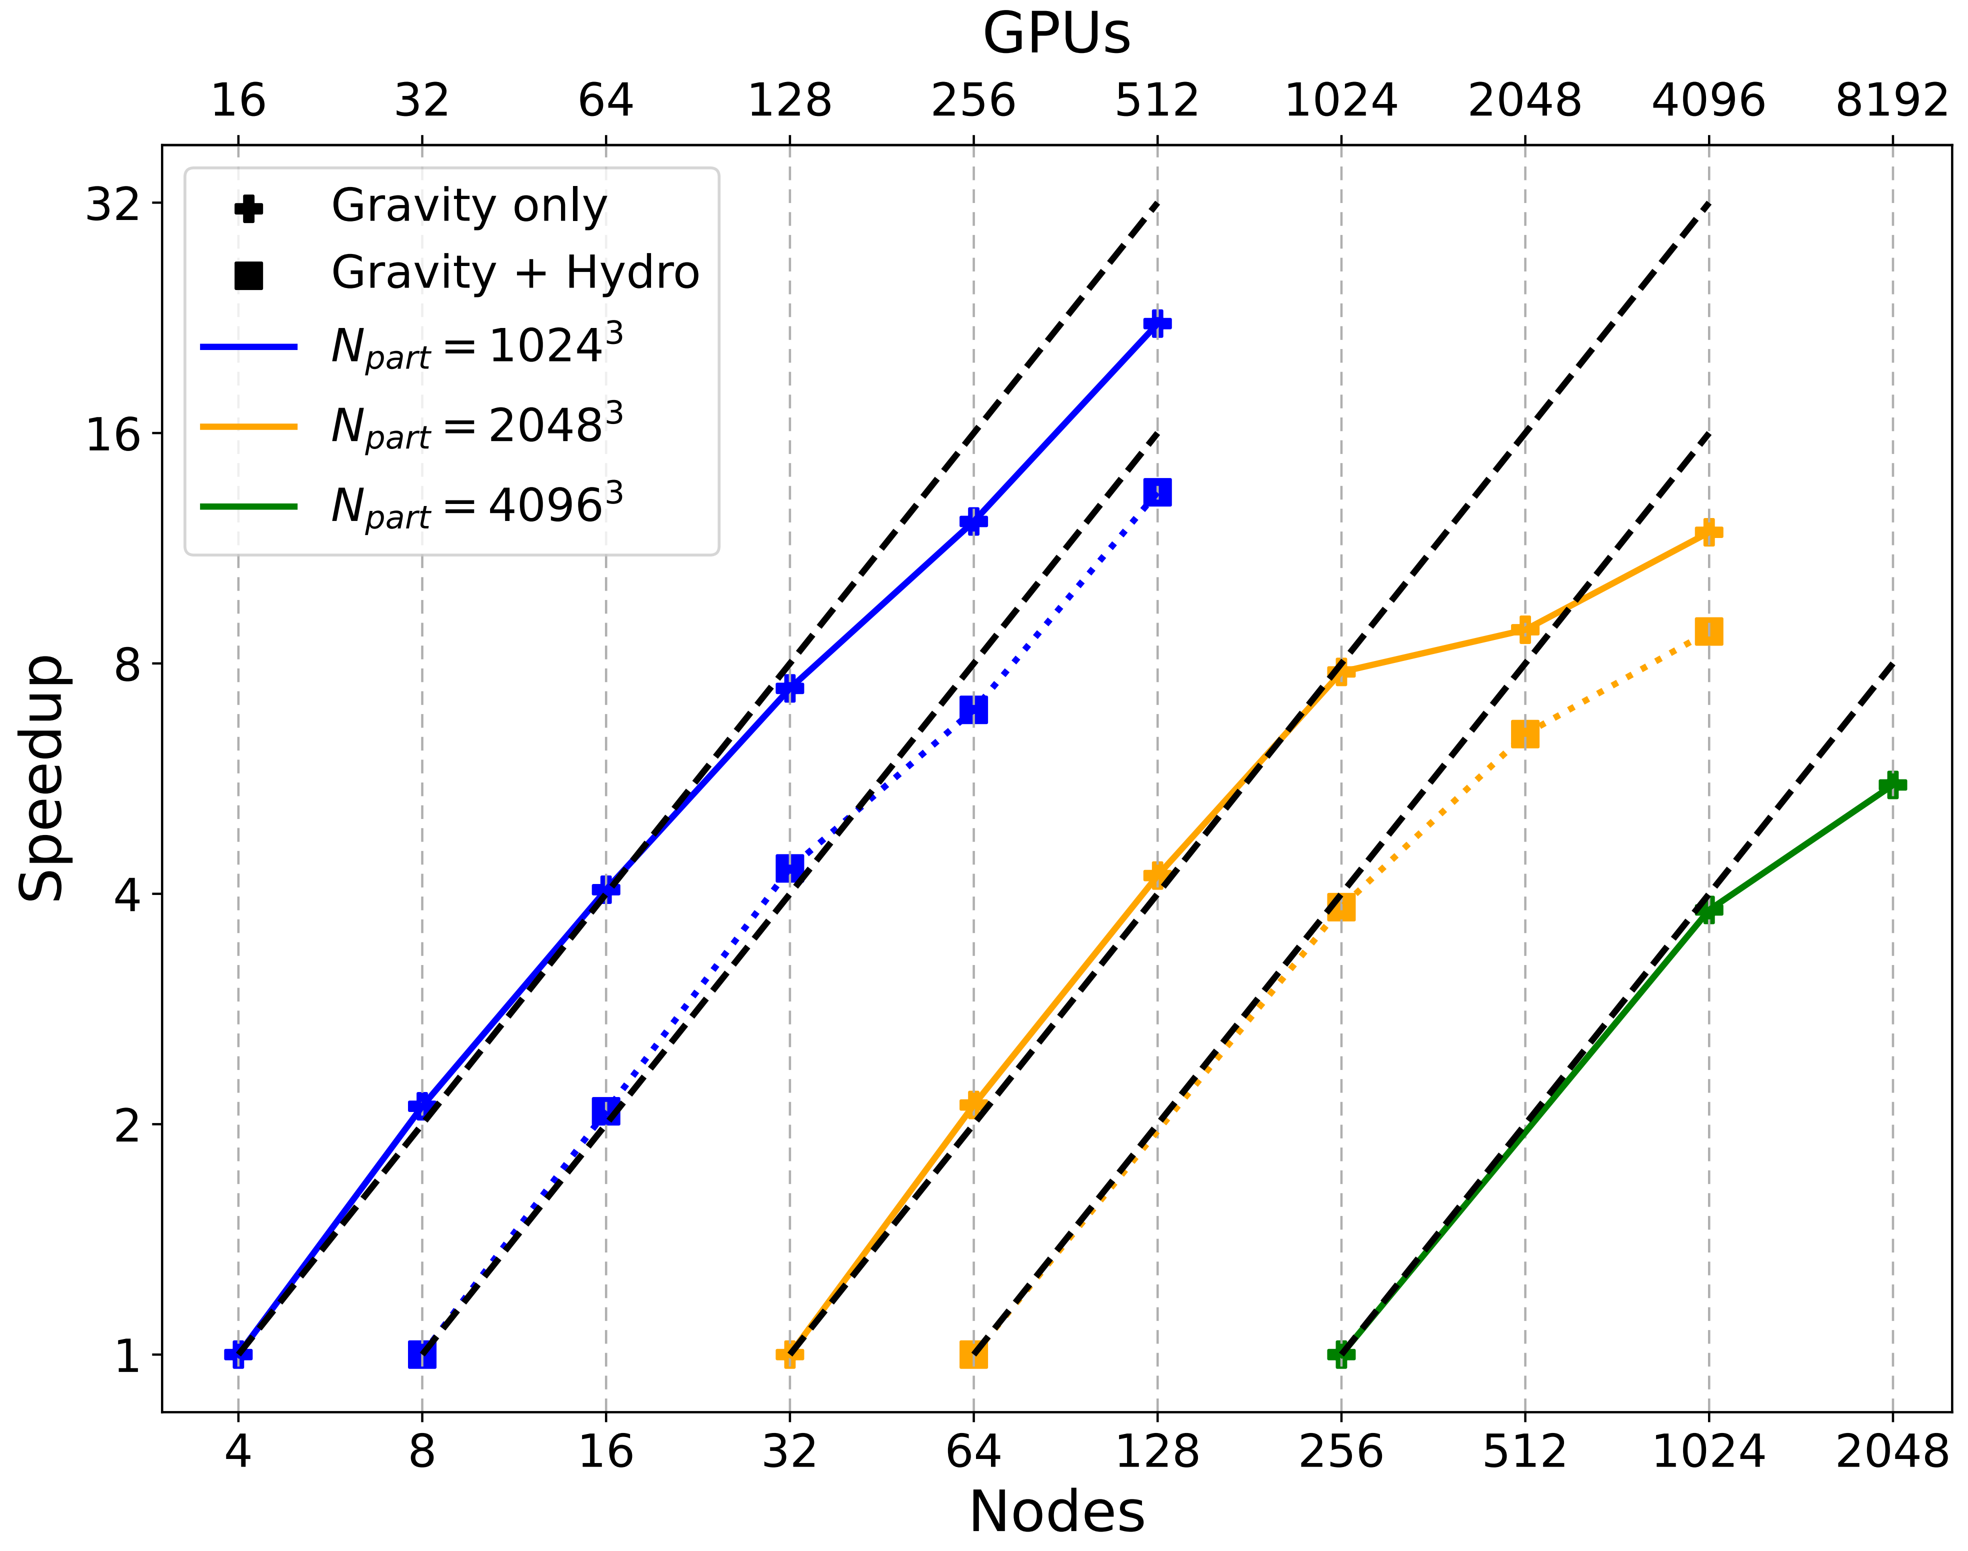

Strong scaling tests for OpenGadget3 gravity-only and gravity + hydrodynamics on Leonardo Booster. The black dashed lines indicate ideal scaling

Code owner: G. S. Karademir, K. Dolag, L. Tornatore

04

Transition iPIC3D -> iPIC3D-GPU

Utilising the capabilities of accelerators

Implicit Particle-in-Cell 3D (iPIC3D)

Simulates plasma dynamics at kinetic scales using an implicit Particle-in-Cell (PIC) method

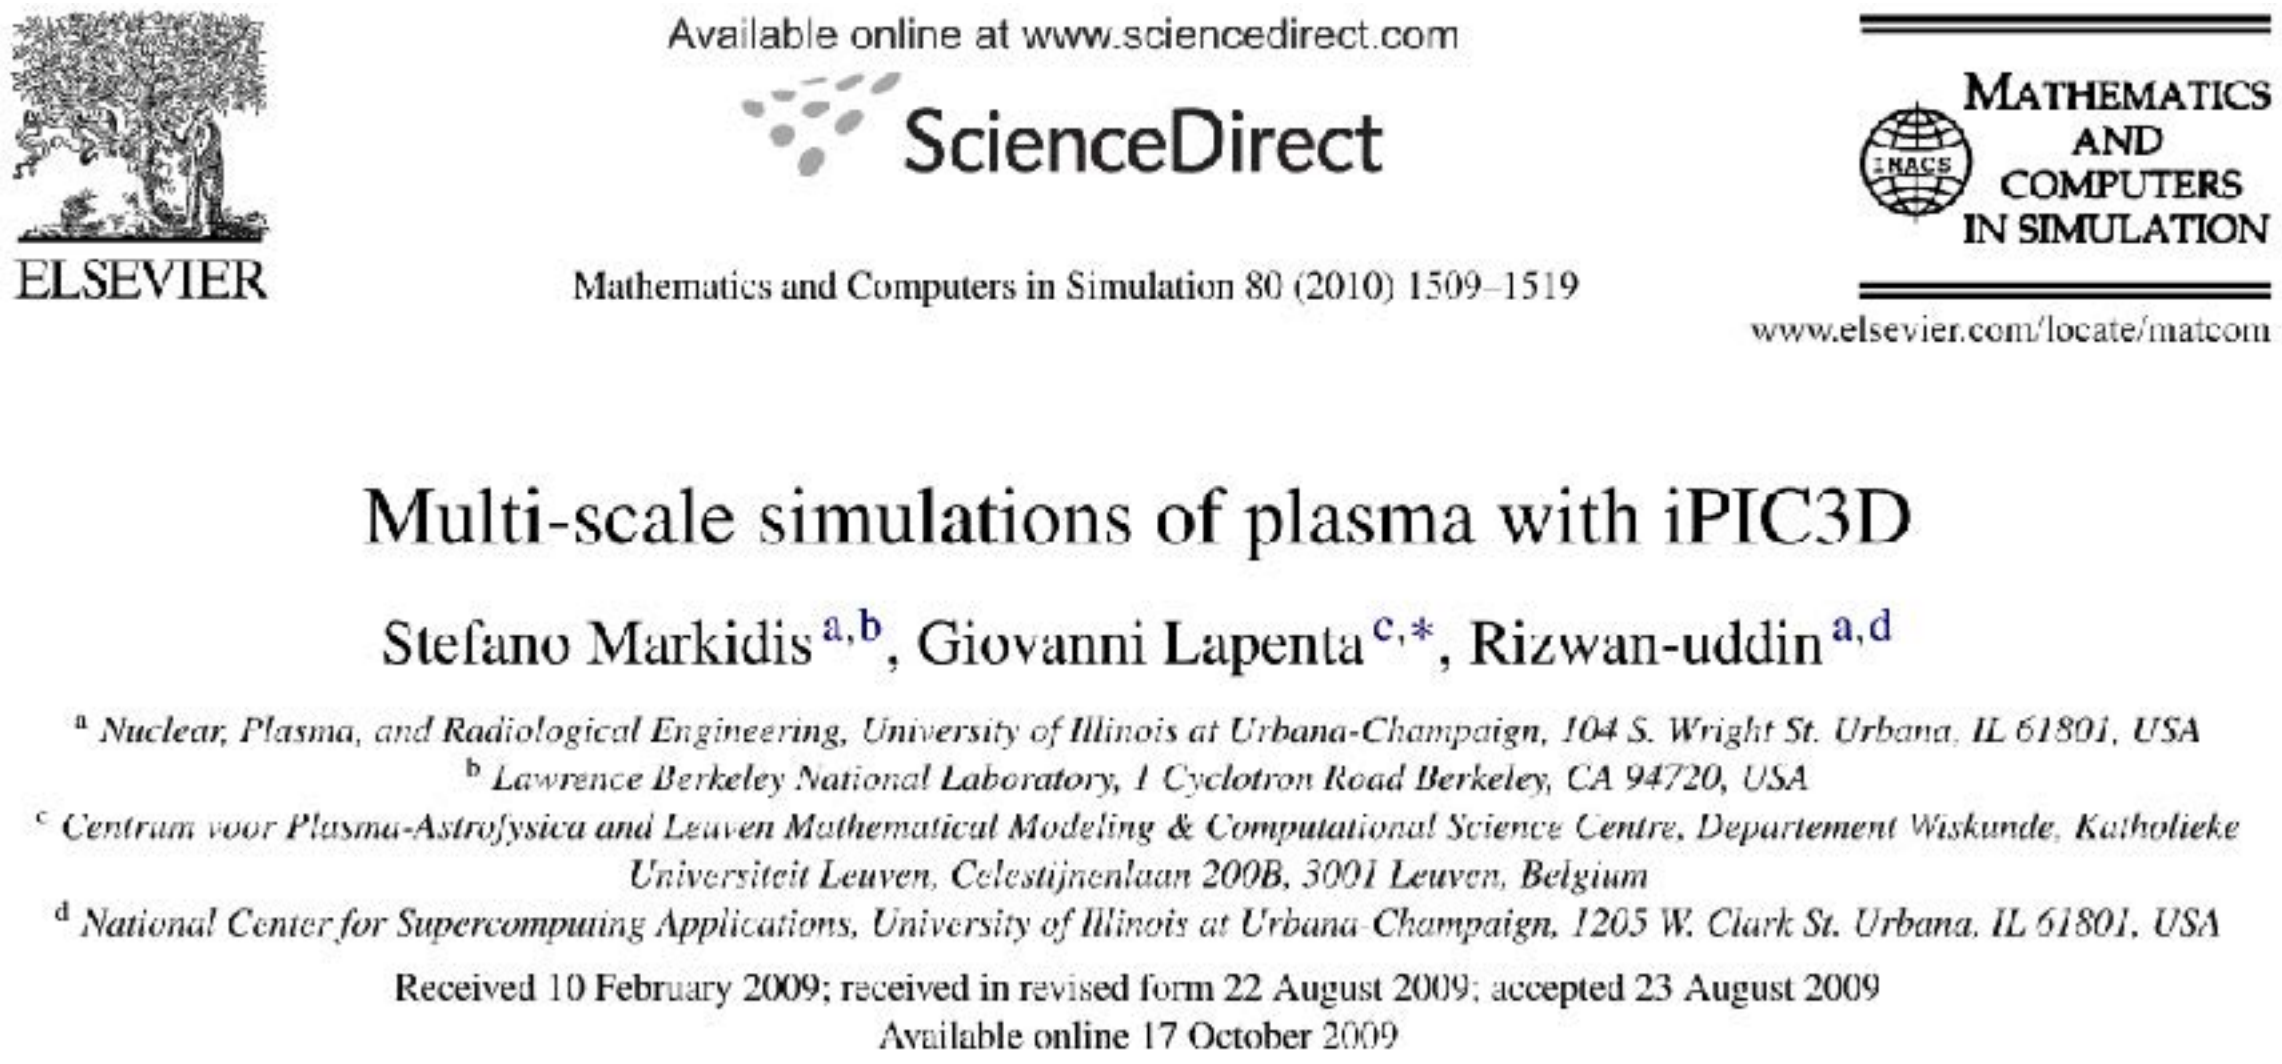
ELSEVIER

Available online at www.sciencedirect.com

ScienceDirect

Mathematics and Computers in Simulation 80 (2010) 1509–1519

MATHEMATICS AND COMPUTERS IN SIMULATION

www.elsevier.com/locate/matcom

Multi-scale simulations of plasma with iPIC3D

Stefano Markidis[a,b], Giovanni Lapenta[c,*], Rizwan-uddin[a,d]

[a] Nuclear, Plasma, and Radiological Engineering, University of Illinois at Urbana-Champaign, 104 S. Wright St. Urbana, IL 61801, USA
[b] Lawrence Berkeley National Laboratory, 1 Cyclotron Road Berkeley, CA 94720, USA
[c] Centrum voor Plasma-Astrofysica and Leuven Mathematical Modeling & Computational Science Centre, Departement Wiskunde, Katholieke Universiteit Leuven, Celestijnenlaan 200B, 3001 Leuven, Belgium
[d] National Center for Supercomputing Applications, University of Illinois at Urbana Champaign, 1205 W. Clark St. Urbana, IL 61801, USA

Received 10 February 2009; received in revised form 22 August 2009; accepted 23 August 2009
Available online 17 October 2009

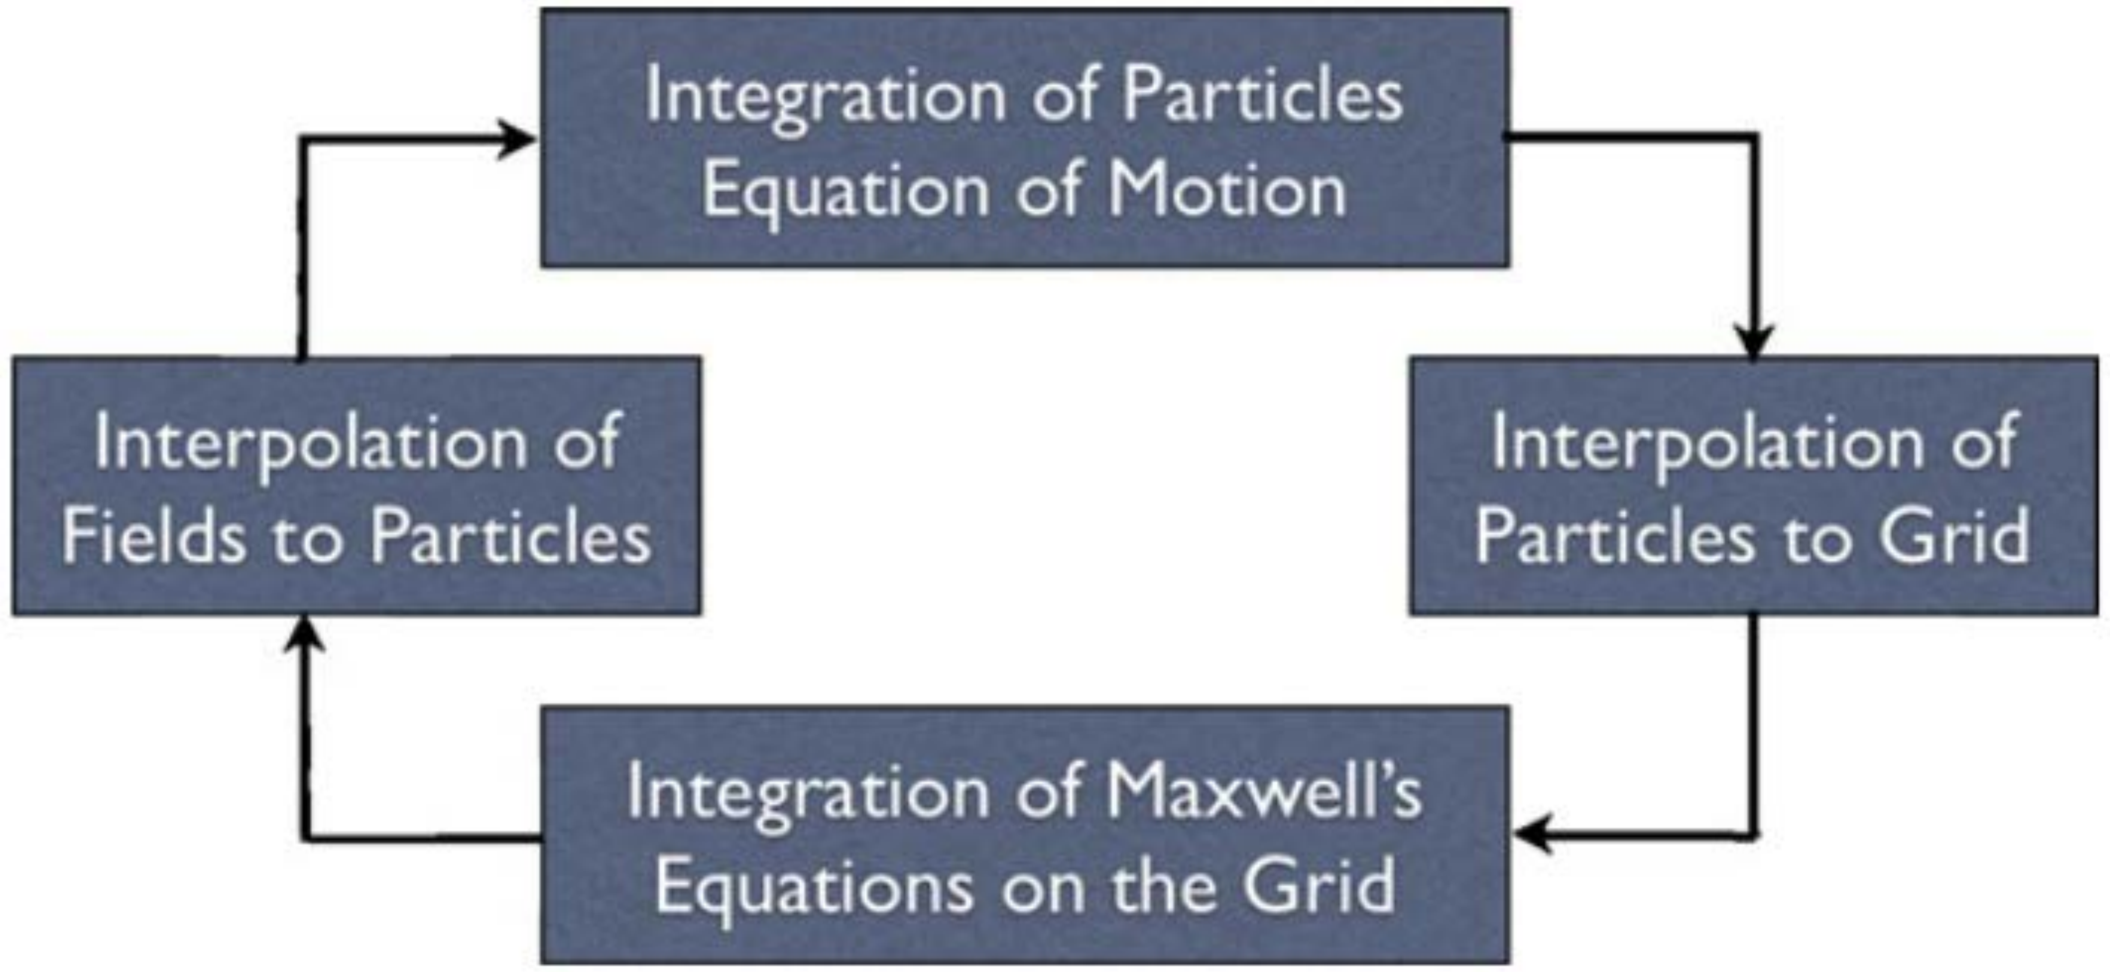

- fully kinetic (particle-in-cell) code for simulation of collisionless plasmas

- supports multidimensional plasma simulations from kinetic to fluid scales
  position & velocity – particles; E and B fields – grid

- semi-implicit time-integration (no CFL constraints)

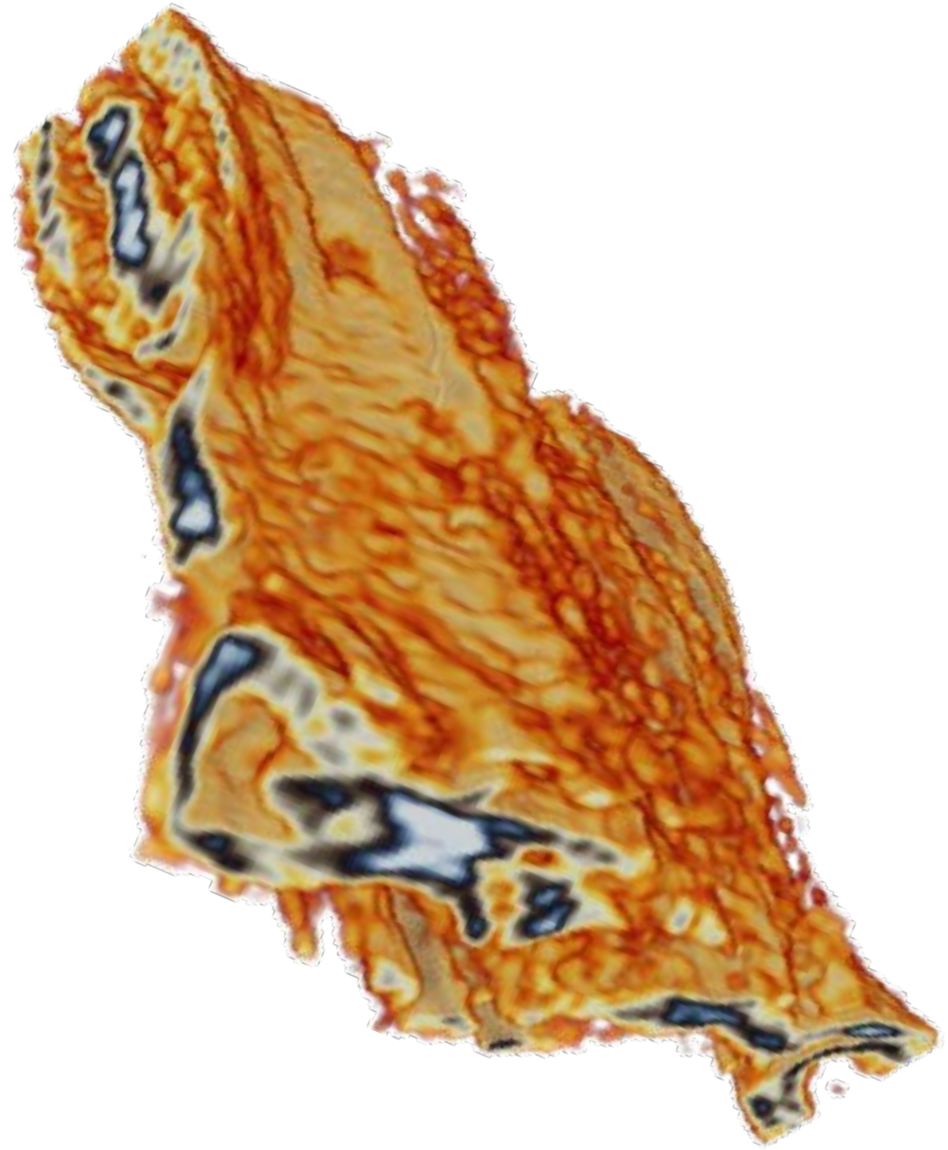

Implicit Particle-in-Cell 3D (iPIC3D)

Simulates plasma dynamics at kinetic scales using an implicit Particle-in-Cell (PIC) method

iPIC3D

- **GPU support:** CUDA (NVIDIA) & HIP (AMD)
- Three main kernels:
  - Particle Mover
  - Moment Gatherer (~80% of runtime)
  - Field Solver

iPIC3D CPU vs GPU Performance

- **Test Case:** Maxwellian distribution (2D) with 20×20×20 particles per cell, 4 species
- **Systems Used:** Leonardo DCGP (CPU), Leonardo Booster (GPU)
- **Module Speedups:**
  - **Moment Gatherer:** ~100× speedup
  - **Particle Mover:** ~40× speedup (main driver of overall speedup)
  - **Field Solver:** Runs only on CPU → no speedup
- **Weak Scaling Test (Leonardo Booster):**
  - Same 2D Maxwellian setup with 4 species, 180×180 particles/cell/species
  - Grid size scaled from 128^2 to 2048^2
  - Achieved **78% efficiency** up to **1024 GPUs** (256 nodes)

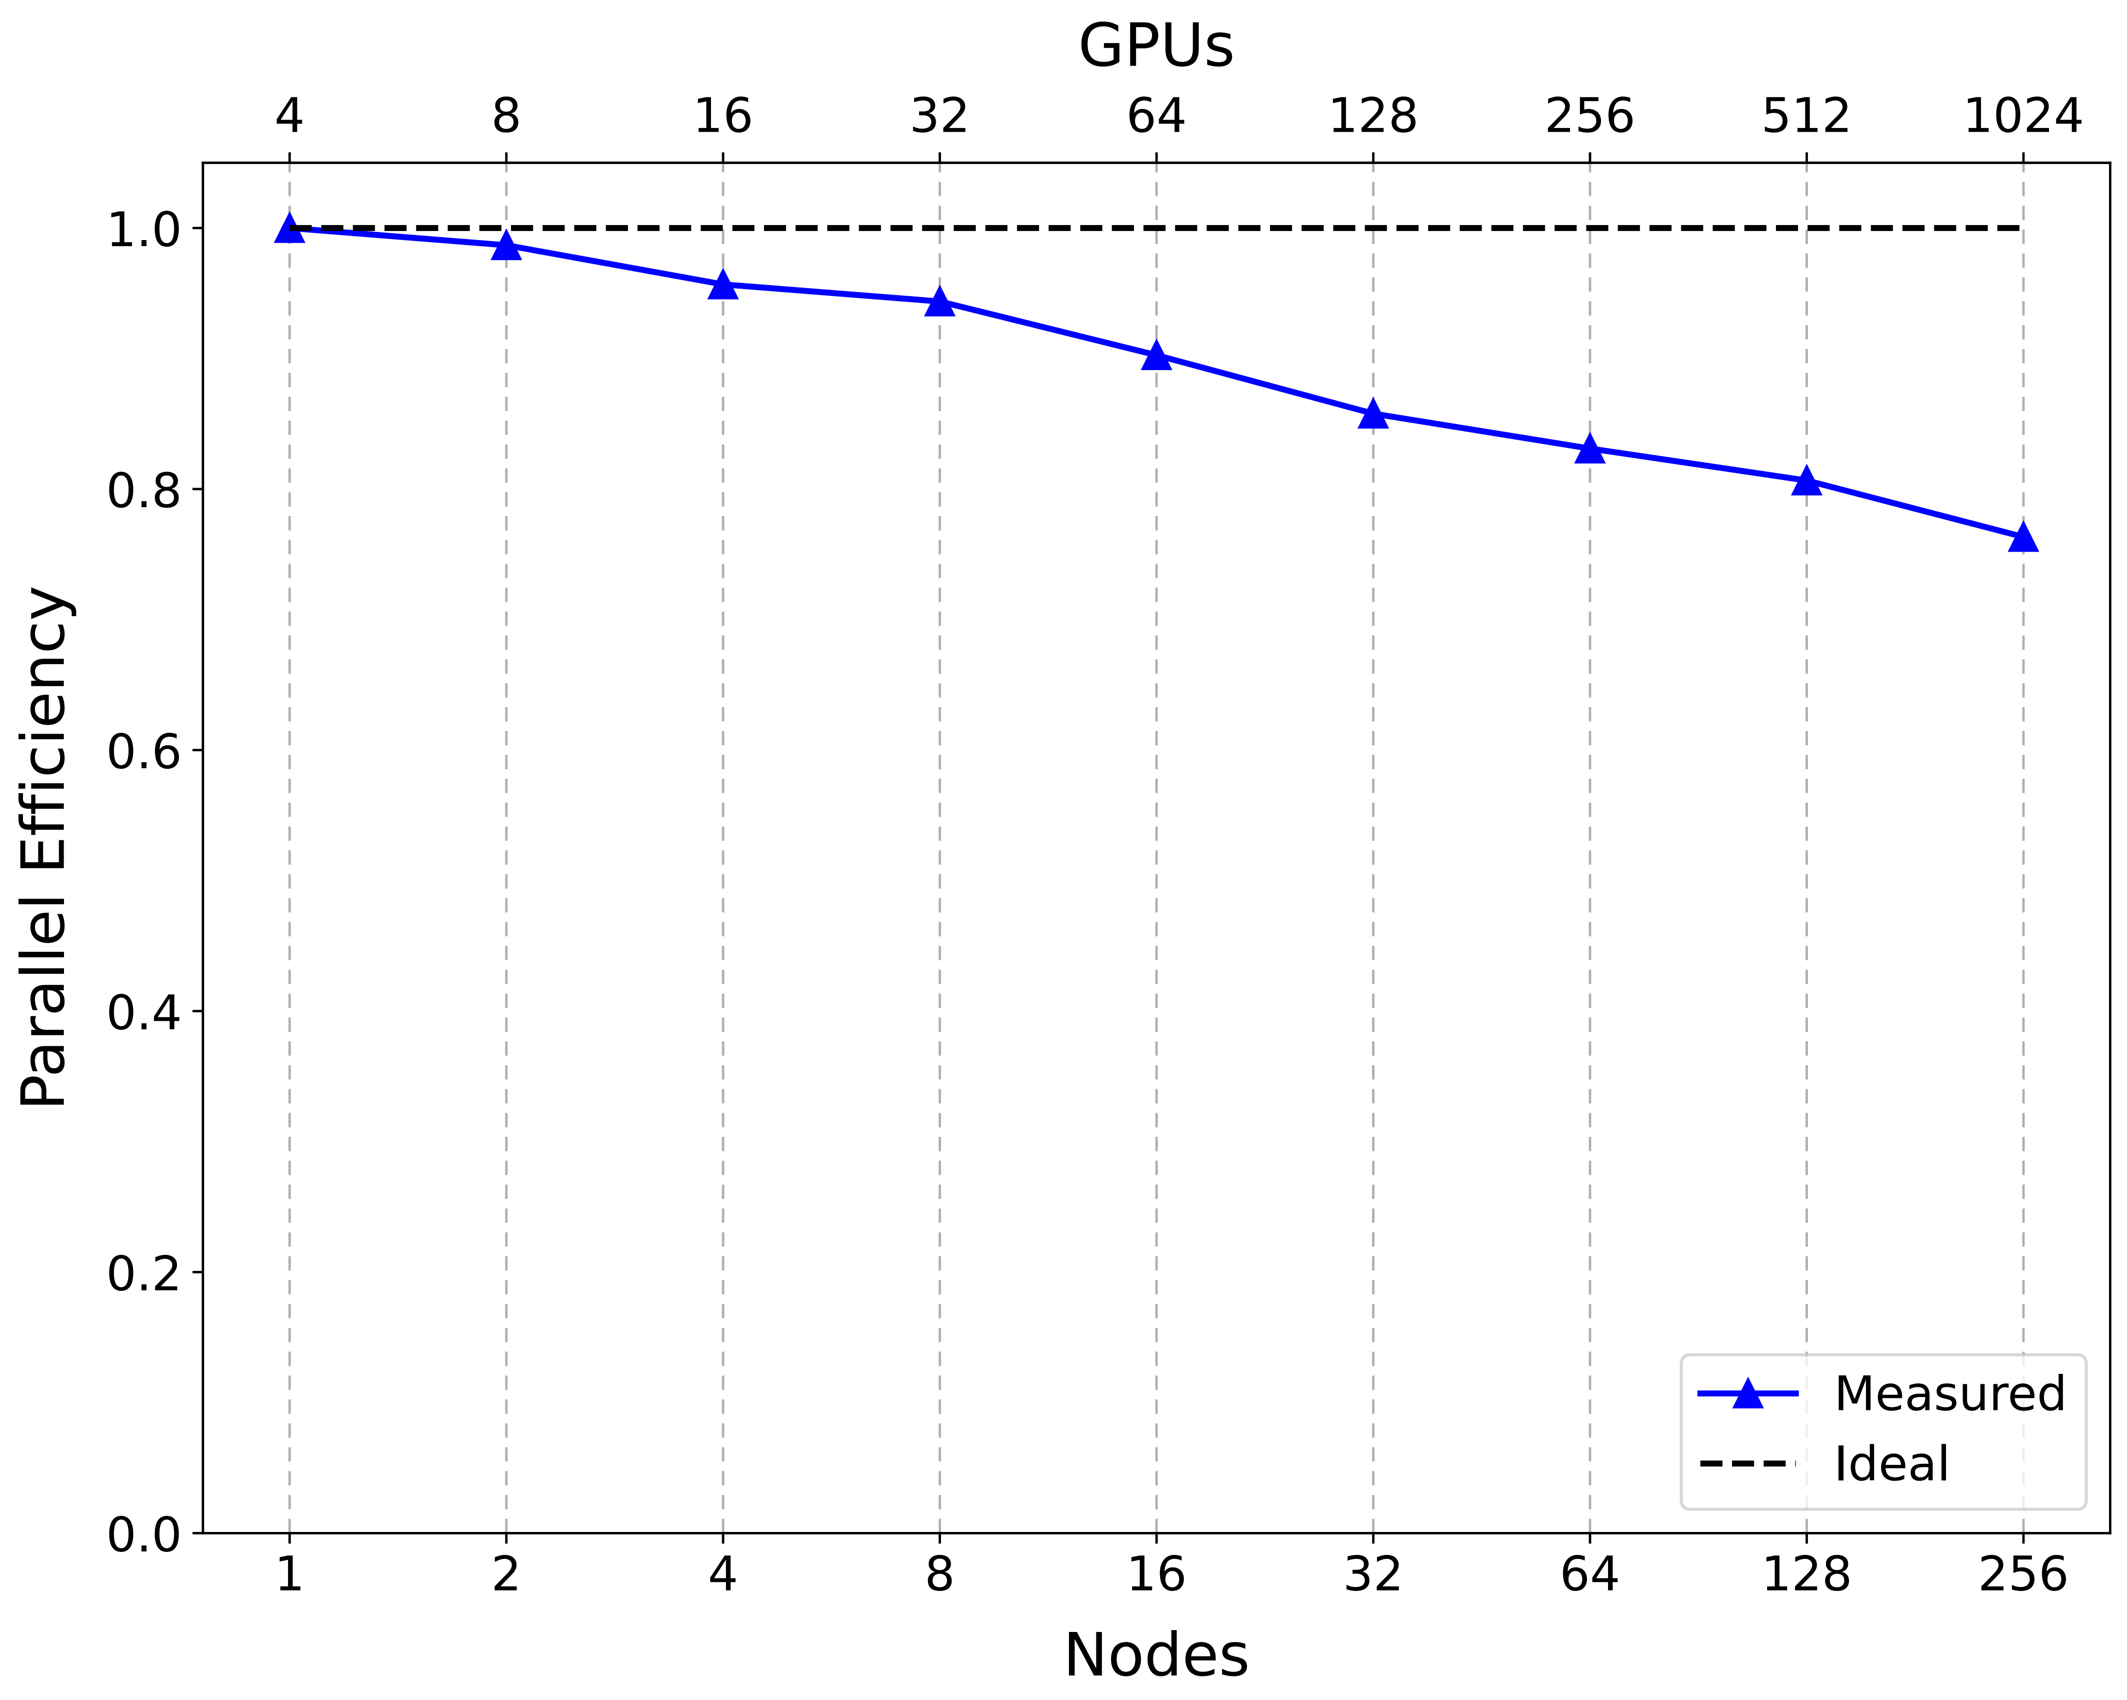

Code owner: Stefano Markidis

Summary

01 **Exascale Adaptation of Simulation Codes**: The SPACE-CoE is adapting three major astrophysical codes **gPLUTO**, **OpenGadget3**, and **iPIC3D** for GPU-based exascale computing on the **Leonardo supercomputer**

02 **GPU Offloading and Optimization**: Key components of the codes were offloaded to **GPUs** using **profiling, modularization, and kernel optimization**, enabling a transition from CPU-centric to GPU-accelerated architectures

03 **Performance and Scalability**: Early results show strong scalability on up to **1,024 GPUs** with **80–97% efficiency**, including **100× speedups** in some iPIC3D modules and excellent scaling in gPLUTO

04 Code-Specific Improvements:

- gPLUTO: Near-ideal weak and strong scaling
- OpenGadget3: Bottlenecks in the Barnes & Hut algorithm were reduced via restructuring
- iPIC3D: Achieved 78% weak scaling using selective GPU offloading and asynchronous CPU-GPU communication

Work in Collaboration With

N. Shukla, A. Romeo, C. Caravita, M. Redenti
Casalecchio di Reno, Bologna Italy

R. Vavrik, L. Riha
VSB-TU Ostrava, Ostrava, Czech Republic

A. Mignone, M. Rossazza, S. Truzzi
University of Turin, Turin, Italy

L. Tornatore, A. Ragagnin, T. Castro
INAF - Osservatorio Astronomico di Trieste, Trieste, Italy

G. S. Karademir, K. Dolag
Ludwig-Maximilians-Universitat Munchen, Munich, Germany

P. J. Deka, F. Bacchini, RP Wilhelm
KU Leuven, Belgium

ELISABETTA BOELLA
HPC Product Specialist at E4 Computer Engineering, Italy

Cosmic ray transport in anisotropic turbulence

F. Pucci[1]

In collaboration with P. Blasi[2], O. Pezzi[1], C. Evoli[2] and F. Malara[3]

1) CNR-ISTP, Bari, Italy, 2) GSSI, L'Aquila, Italy, 3) Università della Calabria, Rende, Italy

EuroHPC UserDays

Copenaghen, 1 October 2025

Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani

Cosmic rays: scientific question

- Energetic particles coming from the galaxy and beyond
- Broad energy spectrum with spectral features at distinct energies

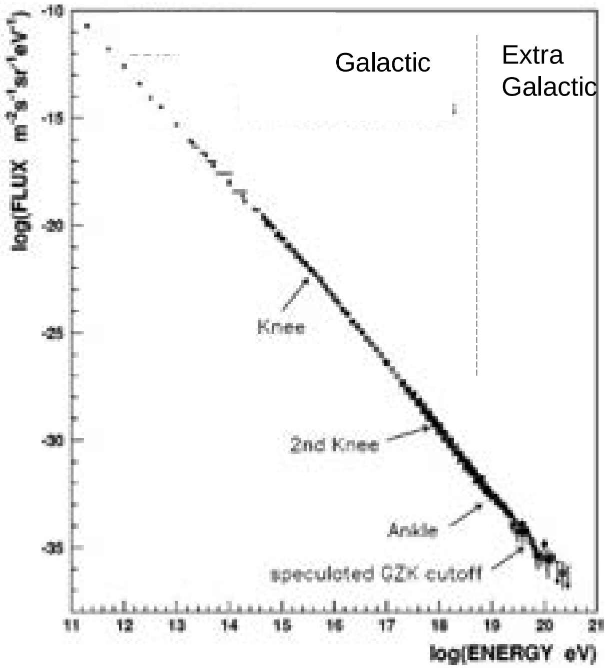

Adapted from Nagano & Watson, 2000

Cosmic rays: scientific question

- Energetic particles coming from the galaxy and beyond
- Broad energy spectrum with spectral features at distinct energies

How are they transported in the turbulent interstellar medium?

How do CR transport properties depend on their energy?

How do CR transport properties depend on the turbulence properties?

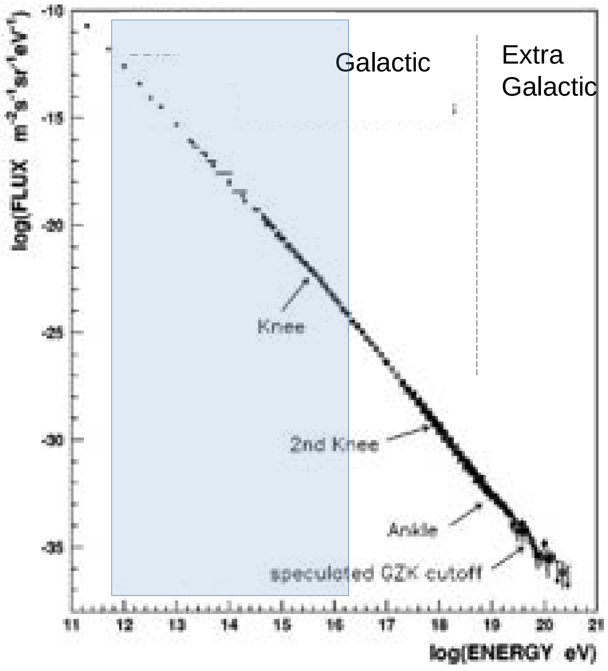

Adapted from Nagano & Watson, 2000

Turbulence: a multiscale process

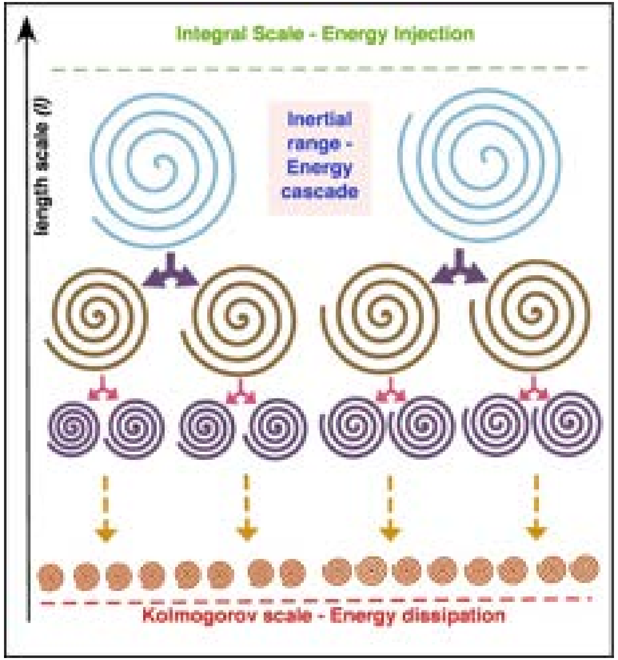

Credits: Supratik Banerjee (Research Gate)

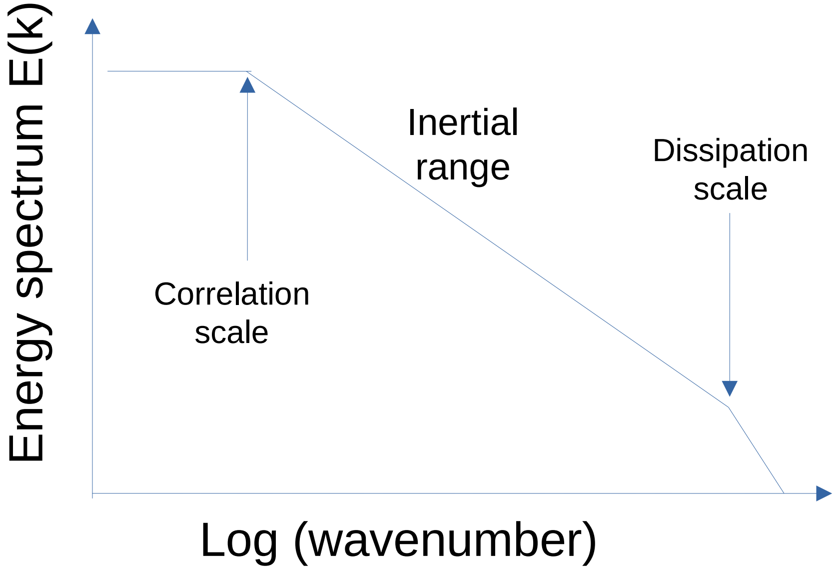

Turbulence effect on particle transport

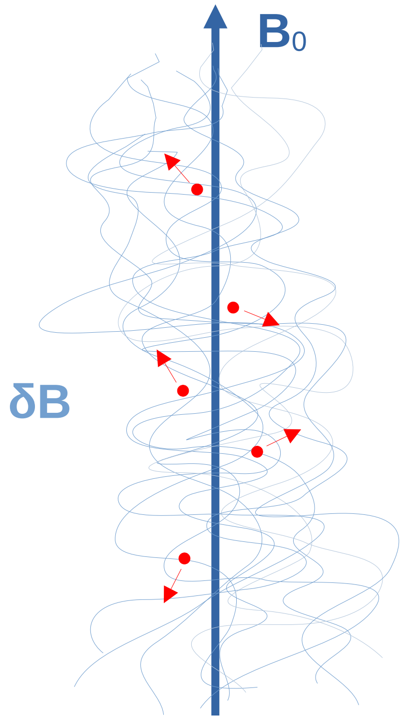

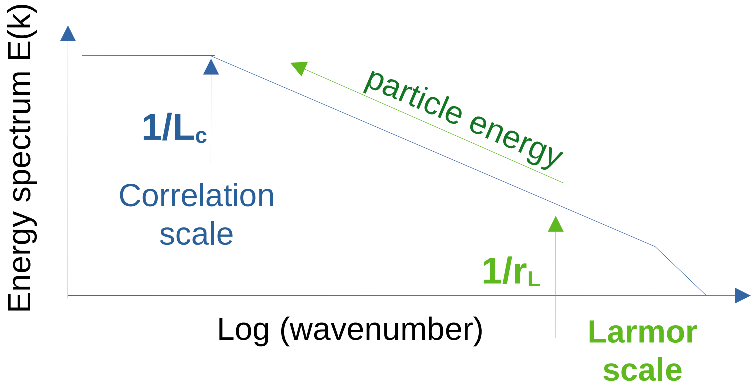

Parallel transport

Diffusive, $D_\parallel \propto (r_L/L_c)^{1/3}$ (Jokipii 1966)

Perpendicular transport

Diffusive, influenced by field line random walk (Shalchi 2024, Kuhlen et al. 2025)

$D_\parallel \gg D_\perp$

for

$(\delta B/B_0) \leq 1$

Numerical approach to the problem

1) **Model the turbulent magnetic field**

Challenge

- Turbulence multiscale nature. The number of orders of magnitude in spatial scale makes direct numerical simulations prohibitive

Solution

- Reproduce turbulence synthetic from mathematical/numerical models

2) Solve the particle motion to evaulate the diffusion coefficients

Challenge

- The single particle trajectory requires many cycles to reach a diffusive behaviour
- A large number of particle is required

Solution

- Fast computation per step
- parallel computing (**HPC resources**)

A wavelet method for synthetic turbulence I

- **The spatial domain**: a 3D periodic box

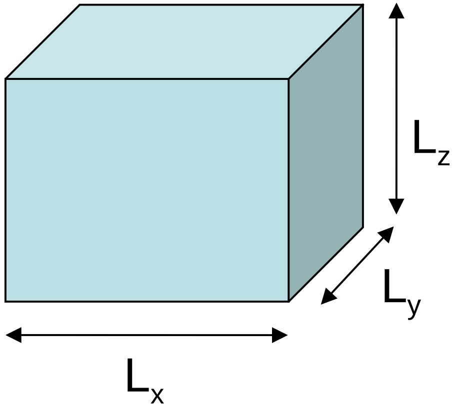

A **hierarchy of cells** at increasingly smaller scales:

i → i+1 the sides of cells are divided by 2

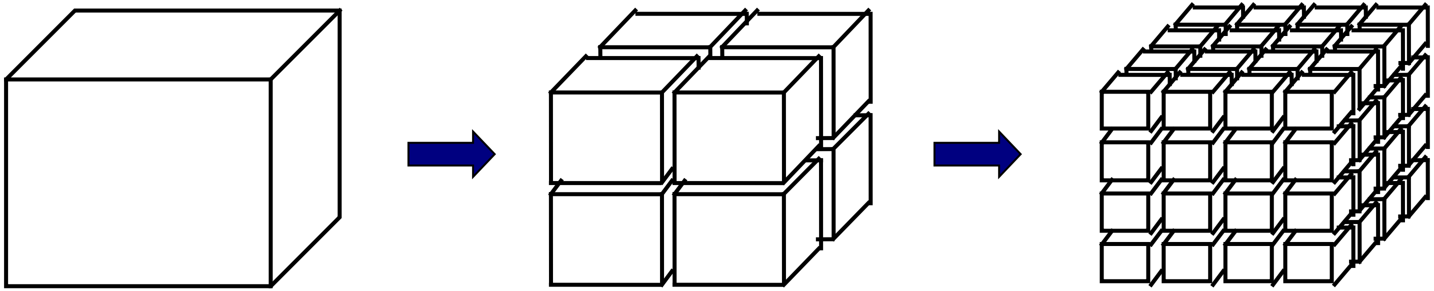

i* = *0   ***i* = *1***   ***i* = *2***

(Malara et al. 2016)

A wavelet method for synthetic turbulence II

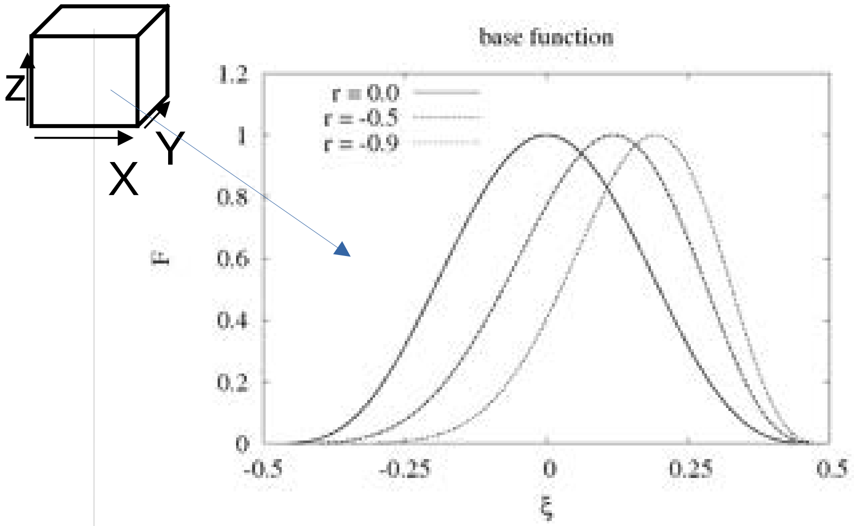

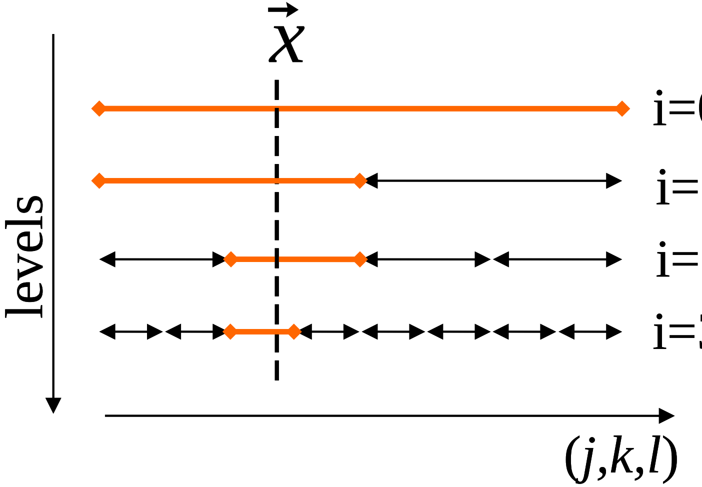

Superposition algorithm

- Each wavelet has a unique identifier
- A few random number per wavelet are used to model its shape
- The field is computed runtime (no field saved in memory)

Isotropic and anisotropic spectra

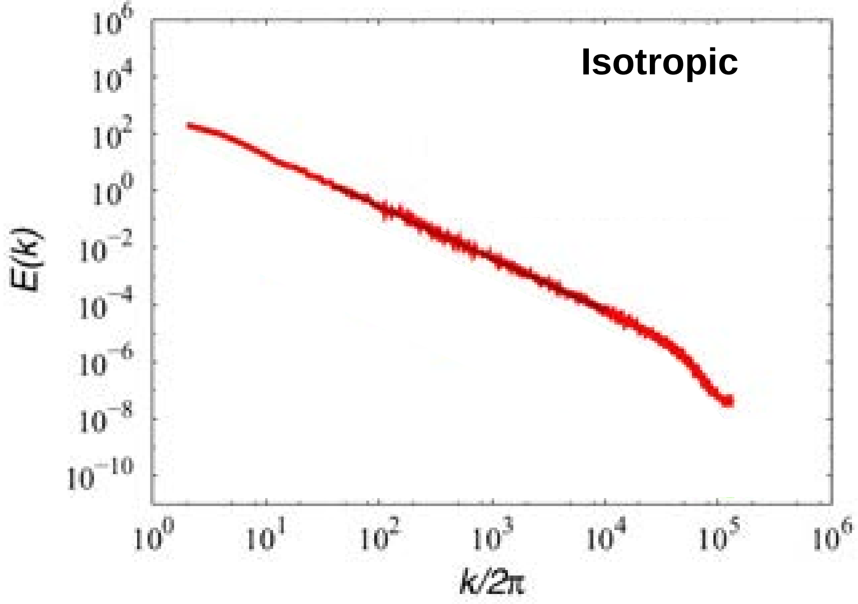

Adapted from Malara et al. 2016

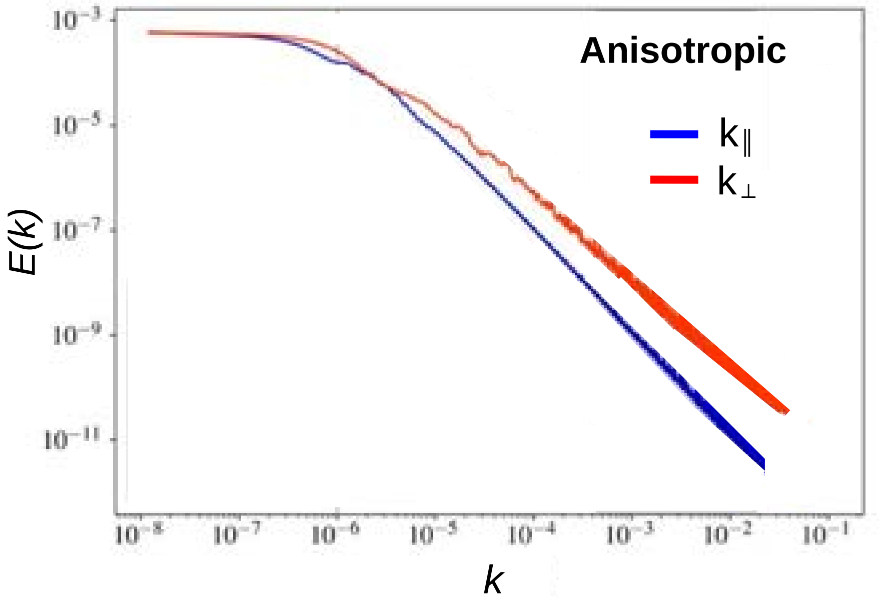

Adapted from Pecora et al. 2024

EUROHPC projects and simulations

Isotropic turbulence

- $\delta b/B_0 = 0.3, 0.5, 1.0$
- Spectral extension $L_c/L_{min} \approx 5x10^7$
- Energy range r_L/L_c [2 x 10^{-4} : 2 x 10]
- **4096 MPI processes** per simulation
- **10 mln CPU-h** on Meluxina, LUX

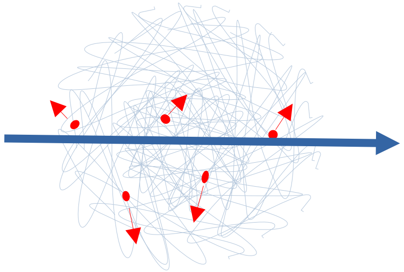

Project: "Energy dependence of cosmic ray transport in magnetic turbulence" (EHPC-REG-2022R03-217), 2023-2024

Anisotropic turbulence

- $\delta b/B_0 = 0.1, 0.3$
- Spectral extension $L_c/L_{min} \approx 5x10^7$
- Energy range r_L/L_c [3 x 10^{-3} : 2 x 10]
- **4096 MPI processes** per simulation
- **6.6mln CPU-h** on Meluxina, LUX

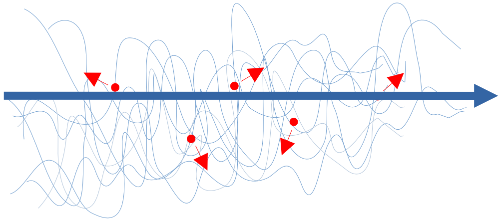

Project: "Cosmic Ray transport in anisotropic turbulence" (EHPC-REG-2023R03-197), 2024-2025

Results: isotropic turbulence

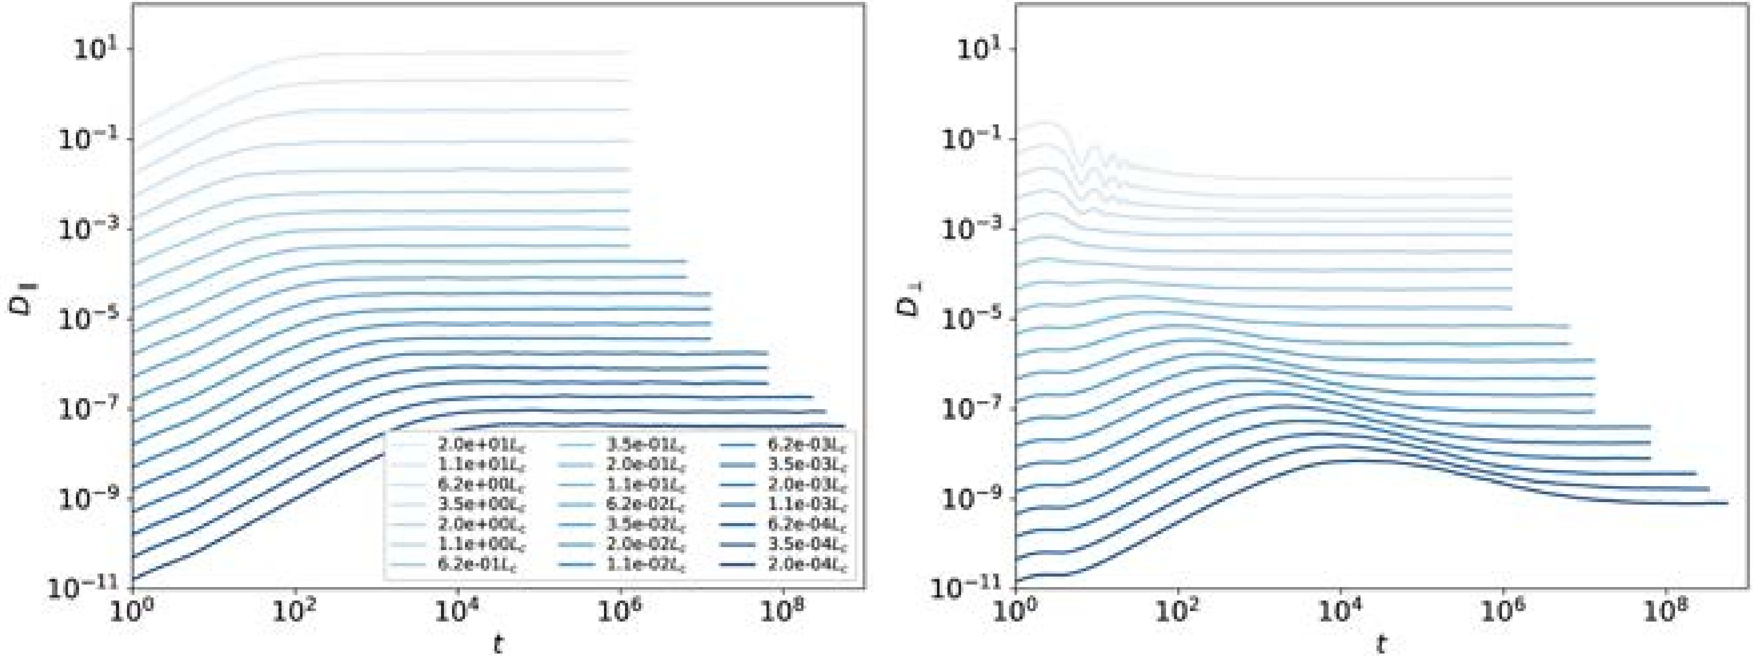

Results: isotropic turbulence

Par. and perp. diffusion coefficients as a function of r_L

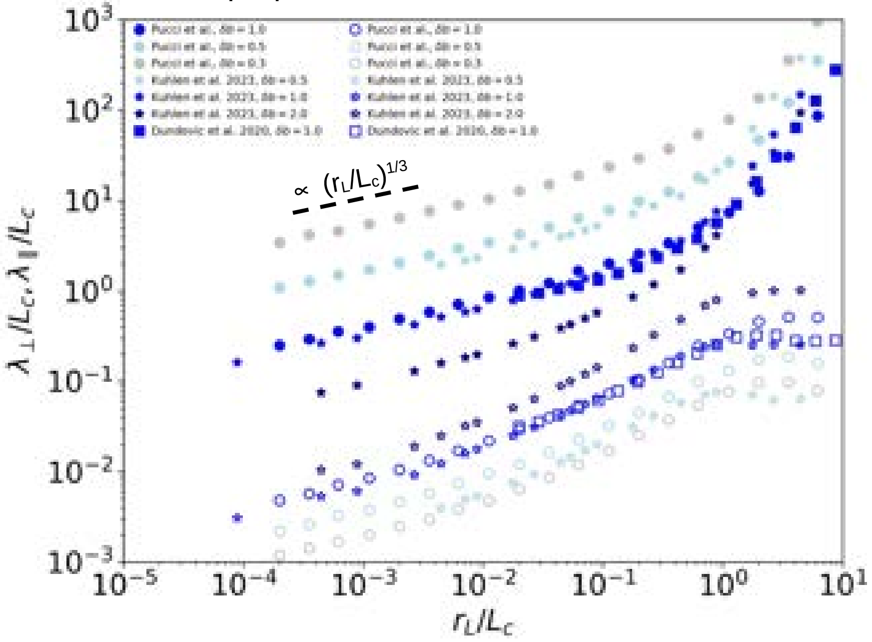

Par. and perp. diffusion coefficient ratio as a function of r_L

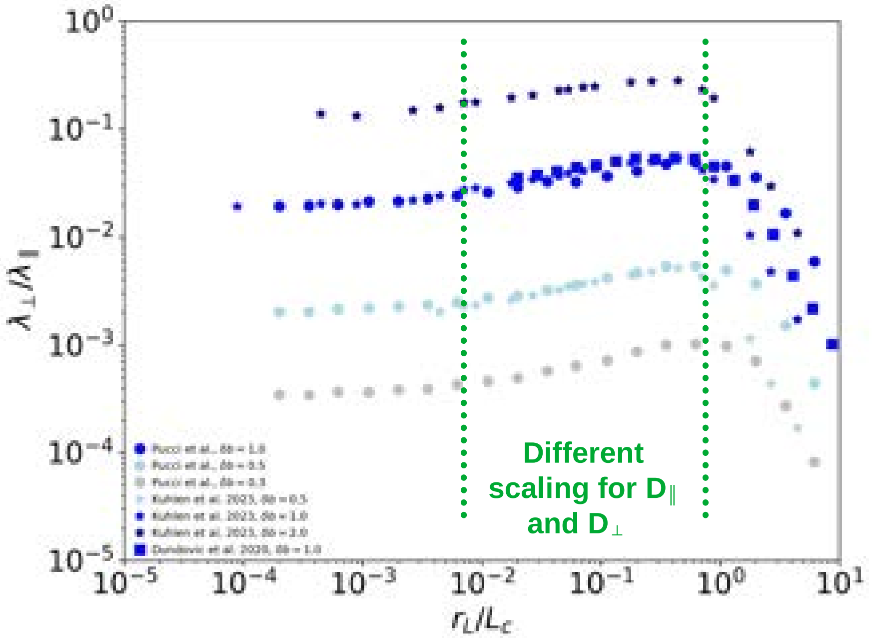

Pucci et al., in prep

Results: anisotropic turbulence

Par. and perp. diffusion coefficients as a function of r_L

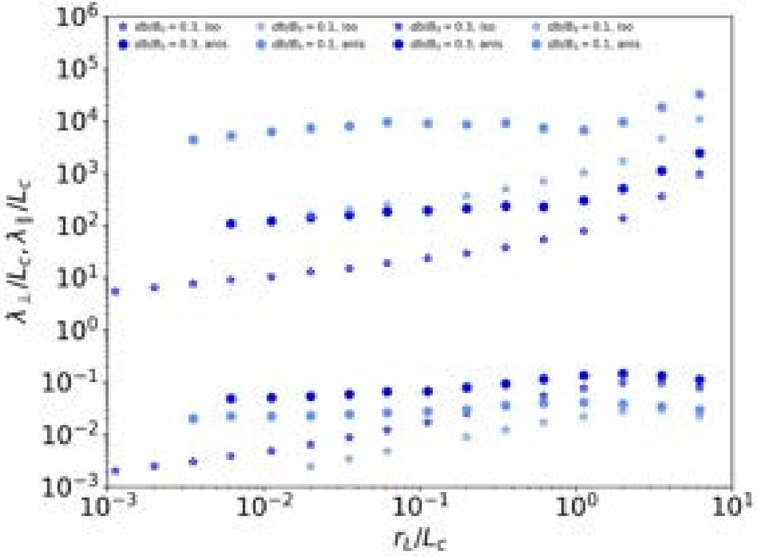

Par. and perp. diffusion coefficient ratio as a function of r_L

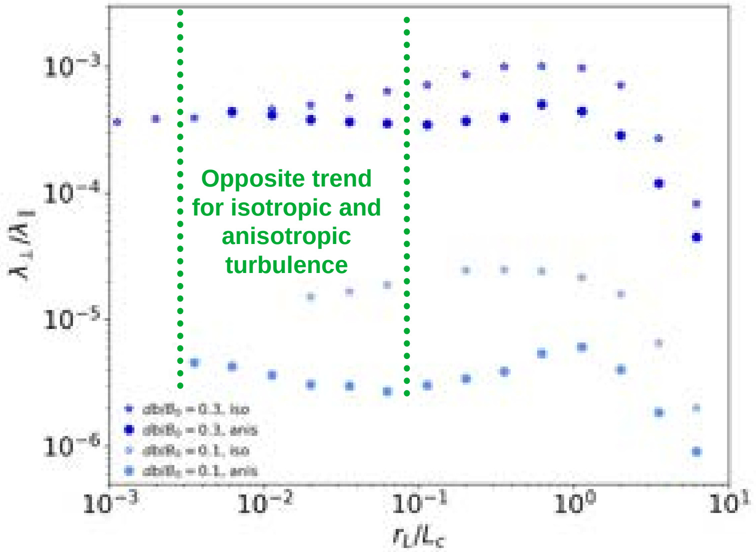

Pucci et al., in prep

Conclusions

Consiglio Nazionale delle Ricerche

- Cosmic ray transport properties are strongly influenced by the ambient magnetic turbulence
- Modelling the multi-scale interstellar magnetic turbulence and particle diffusion through it is a challenging task, which requires a numerical approach
- Using advanced numerical algorithm and HPC facilities from EUROHPC-JU, we conducted an unprecedented simulation campaign to study cosmic ray diffusion in isotropic and anisotropic turbulence
- In the isotropic case, we extended previous results to lower particle energy and magnetic field intensity, confirming an energy dependent trend of the perp to parallel diffusion coefficient ratio
- In the anisotropic case, we found an opposite trend of the diffusion coefficient ratio with respect to the isotropic case

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani

UNIVERSITY OF HELSINKI
FACULTY OF SCIENCE

European Research Council

ForEshock – Magnetosheath – Magnetosphere in 6D (FEMMA)

Minna Palmroth[1,2], Lucile Turc[1] & Vlasiator team

[1]University of Helsinki, [2]Finnish Meteorological Institute

Helsinki, Finland

ACADEMY OF FINLAND

CSC

EuroHPC Joint Undertaking

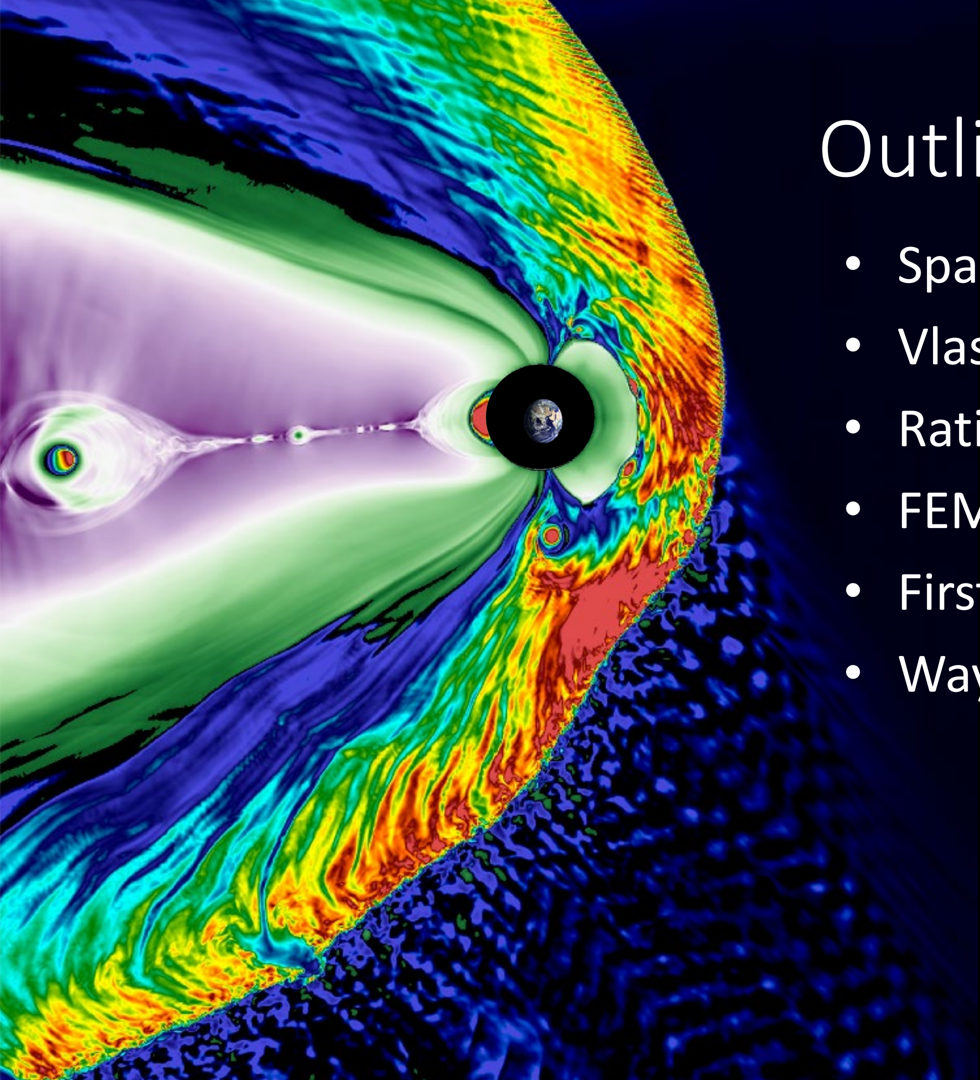

Outline

- Space weather
- Vlasiator – global ion-kinetic simulation
- Rationale for FEMMA
- FEMMA project
- First results
- Way forward

Space weather

Definition: Conditions in the near-Earth space that can affect technological reliability or human health.

Causes: Solar eruptions and dynamics in the Earth's own magnetic domain.

Strength: Varies. Small storms statistically 1/month, medium size storms a few per 11 years, megastorms 1/100 years.

Most important technological impacts:
1. Power grid
2. Satellite health
3. Signals (satellite, radar, radio, aviation, etc)

Global modelling techniques

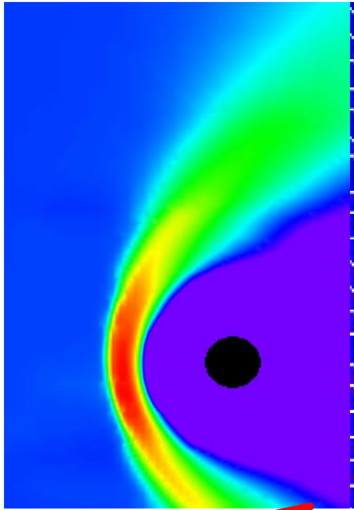

GUMICS-4, FMI

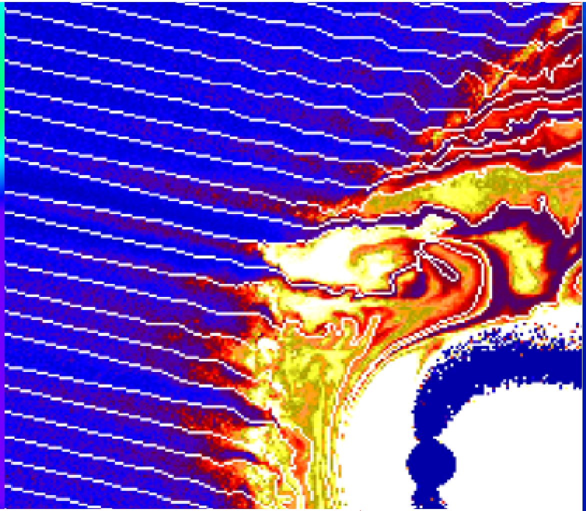

Karimabadi et al., 2014

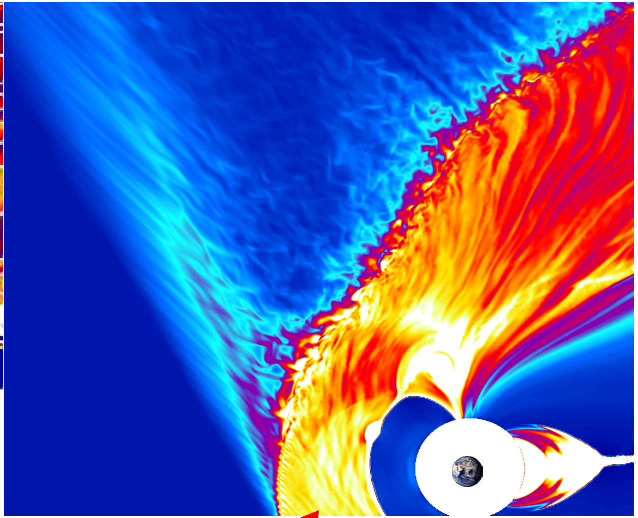

helsinki.fi/vlasiator

	Lowest order approximation: Global fluid (MHD)	More accurate, but noisy: Hybrid particle-in-cell	World's most accurate: Hybrid-Vlasov (Vlasiator)
Ions	Fluid	Particles	Distribution functions
Electrons	Fluid	Fluid	Fluid
Run time	Real-time to weeks	Weeks	Weeks – months
Scale of applications	Fluid scale (1000 km) Now: Solar system	Ion kinetic effects (10 – 1000 km) Global = Solar wind, magnetosphere (+ ionosphere)	

UNIVERSITY OF HELSINKI
FACULTY OF SCIENCE

VLASIATOR

- Global ion-kinetic plasma physics beyond MHD
- Electrons are massless charge-neutralising fluid.

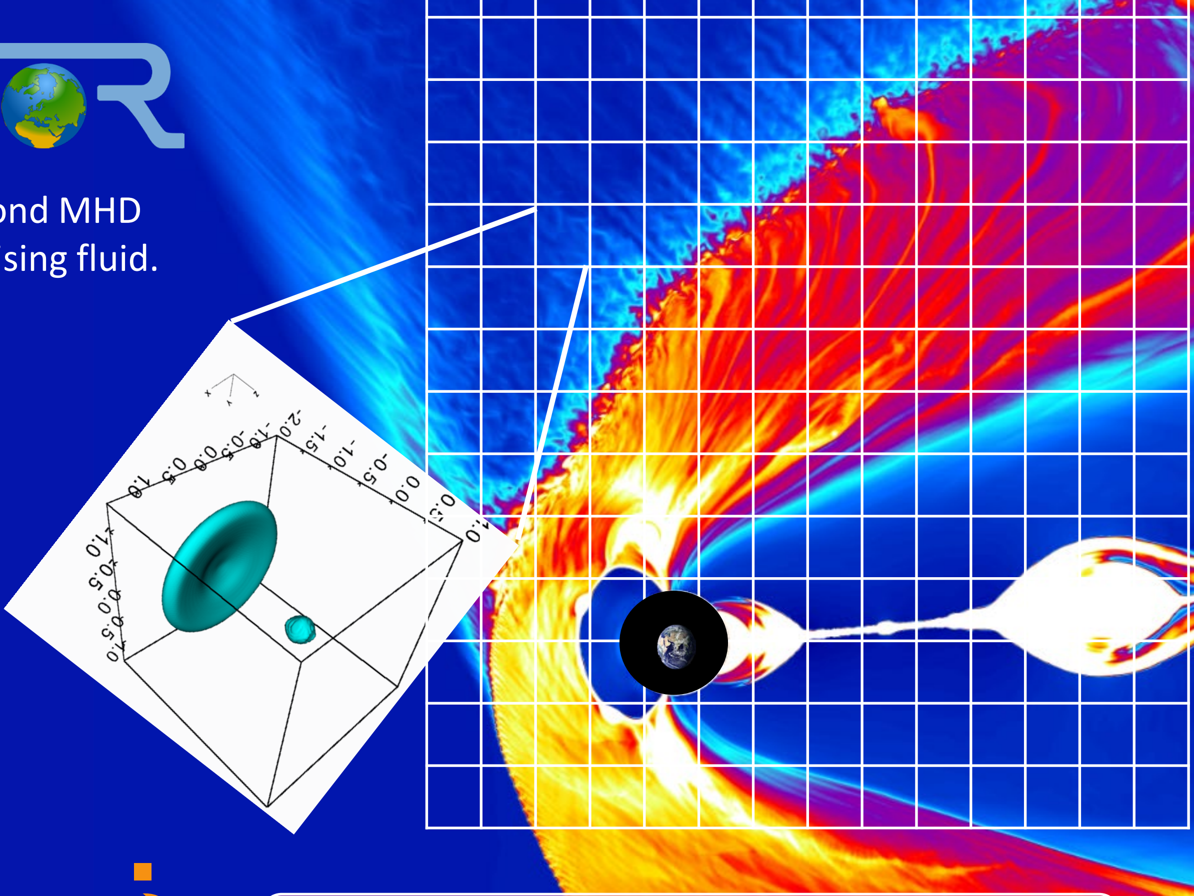

2D or 3D real space (R)

- Divide real space into grid cells
- Compute **E**, **B** fields
- Input to 3D velocity space

3D velocity space (V)

- Each R space cell contains a 3D velocity space
- Propagate and modify ion distribution function using Vlasov equation
- Couple back to ordinary space to update **E**, **B** field

Self consistent

Noise-free multi-temperature physics

More information: http://helsinki.fi/vlasiator

Contact PI: *Minna.Palmroth@helsinki.fi*

In plasma physics, everything depends on how to model the plasma distribution function

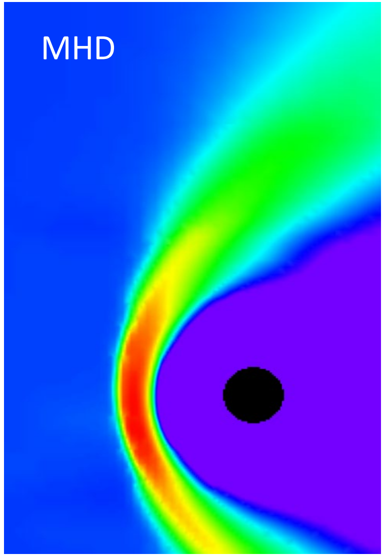

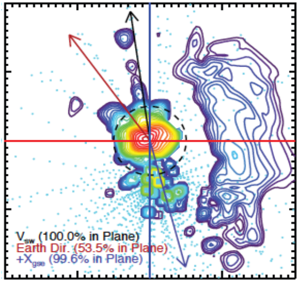

Observations (THEMIS spacecraft)

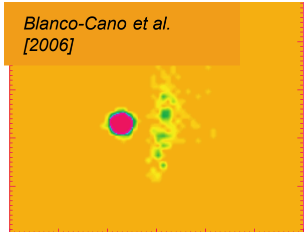

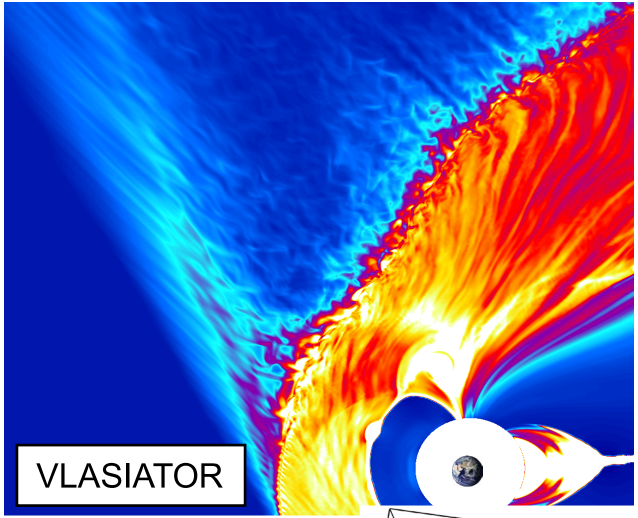

MHD: Distribution function is not modelled. Single value is used for temperature.

Particle-in-cell (PIC). Distribution is constructed from particle statistics.

Vlasiator: Distribution function is modelled perfectly. No assumptions.

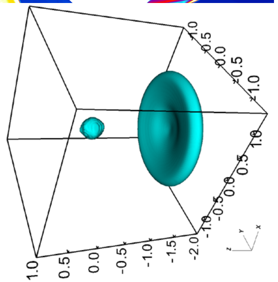

Rationale: Foreshock wave transmission

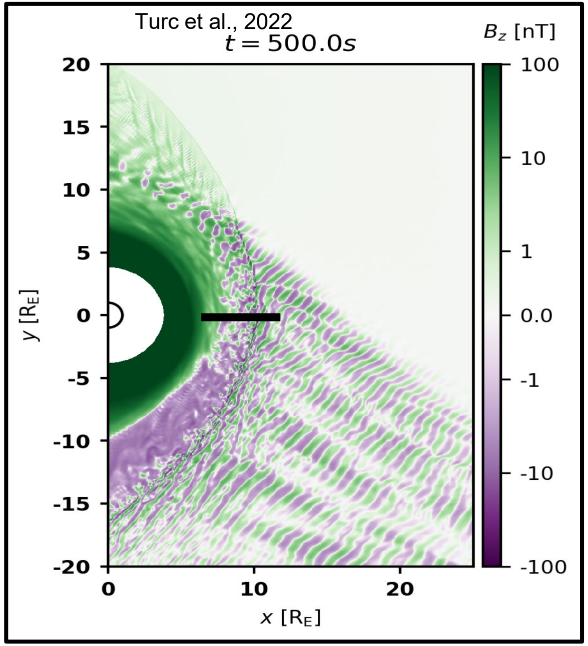

Using 2D Vlasiator simulations, we showed that foreshock waves propagate through the magnetosheath as fast-mode waves.

These waves are however not just “directly transmitted” through the shock, but are created by the foreshock waves modulating the plasma parameters just upstream of the shock

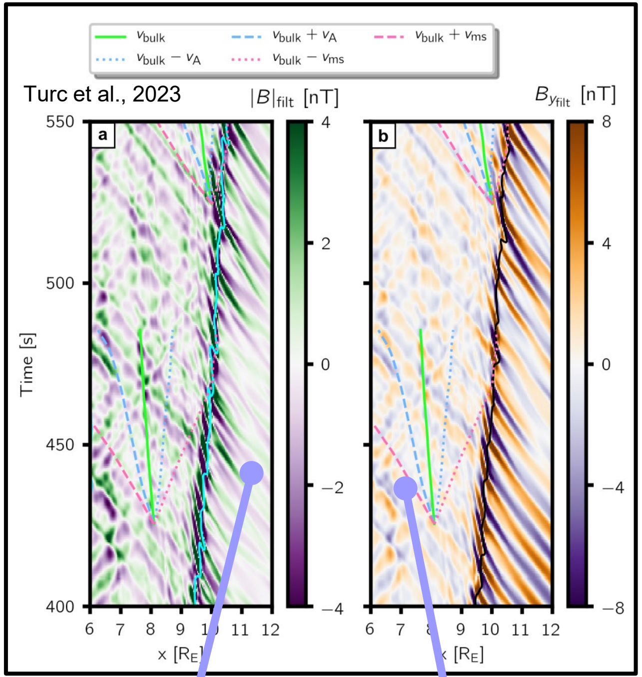

Foreshock waves

Downstream fast-mode waves

Turc et al., 2023, Nature Physics
https://www.nature.com/articles/s41567-022-01837-z

UNIVERSITY OF HELSINKI
FACULTY OF SCIENCE

FEMMA: novel data to answer key questions in shock and near-Earth space physics

FEMMA key questions:

- How does foreshock wave transmission operate in 3D?
- What is the extent of the region across which wave transmission takes place?
- How is the transmitted wave power distributed downstream of the shock?
- Where and how does the transmission into the magnetosphere take place?

IMPACTS:

- Key in transmission of electromagnetic energy from the solar wind into the magnetosphere (space weather)
- relevant for astrophysics and fusion

FEMMA campaign at LUMI

UNIVERSITY OF HELSINKI
FACULTY OF SCIENCE

- Goal: Running a 3D+3V global near-Earth simulation with a high-resolution
- Used a new pressure anisotropy-based criterion for grid refinement
- 991s simulation time, of which ~400s with for scientific investigation
- More than 3 trillion velocity space cells in the simulation
  - Maximum memory high water mark of >90 TB spread over 512 LUMI nodes
- Achieved an MPI write performance of ~250 GB/s when writing to LUMI-F flash storage (speedup of ~7.5x)
- LUMI was overall much more stable than during previous computing campaigns

Preliminary results

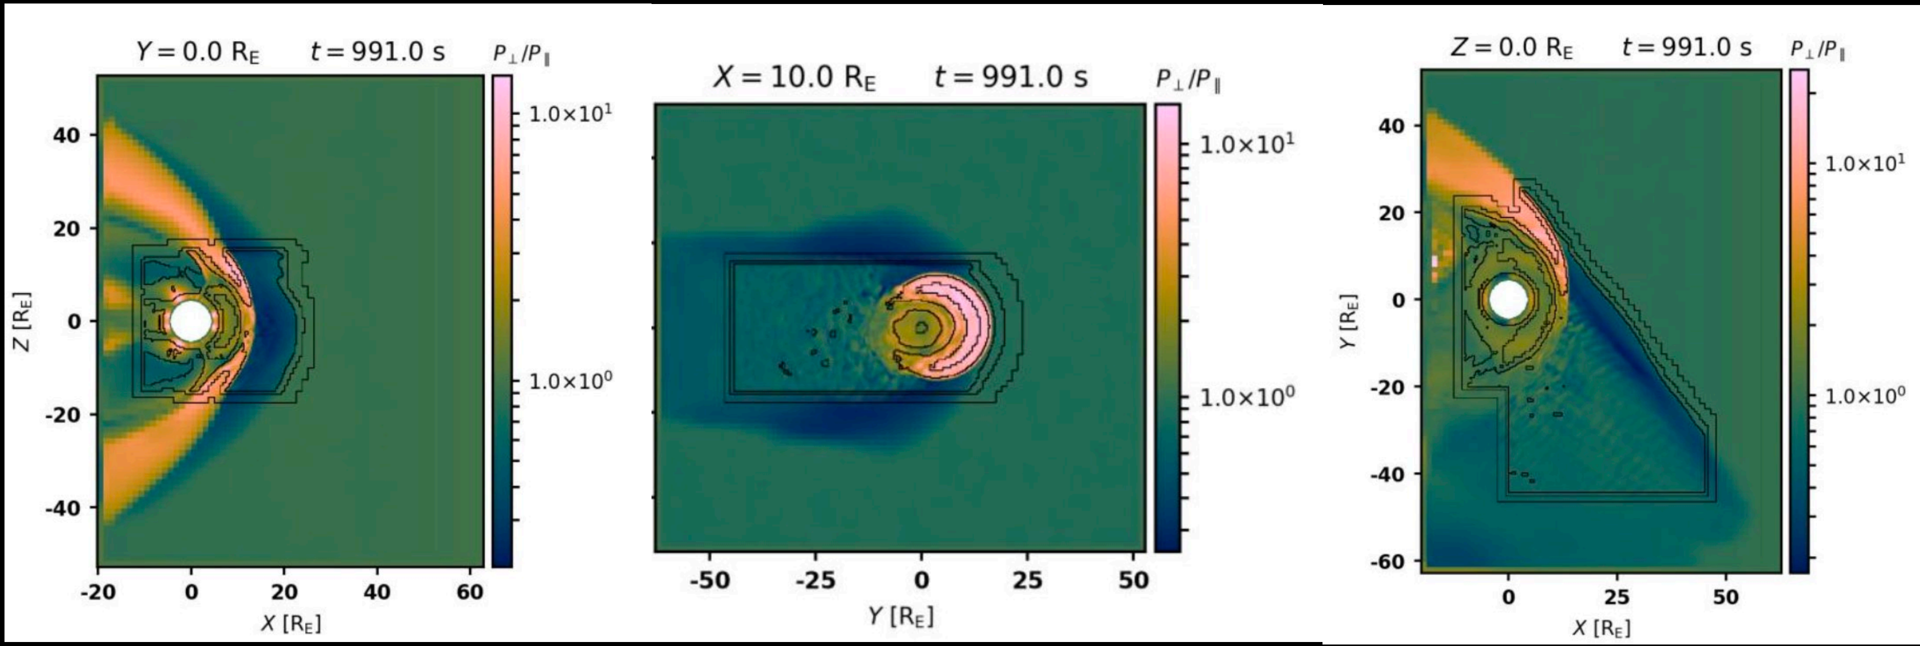

UNIVERSITY OF HELSINKI
FACULTY OF SCIENCE

Preliminary results

te new run for

New science

New methods & capabilities

ce Economy

Policy advising

etal preparedness

Outreach

EuroHPC
Joint Undertaking

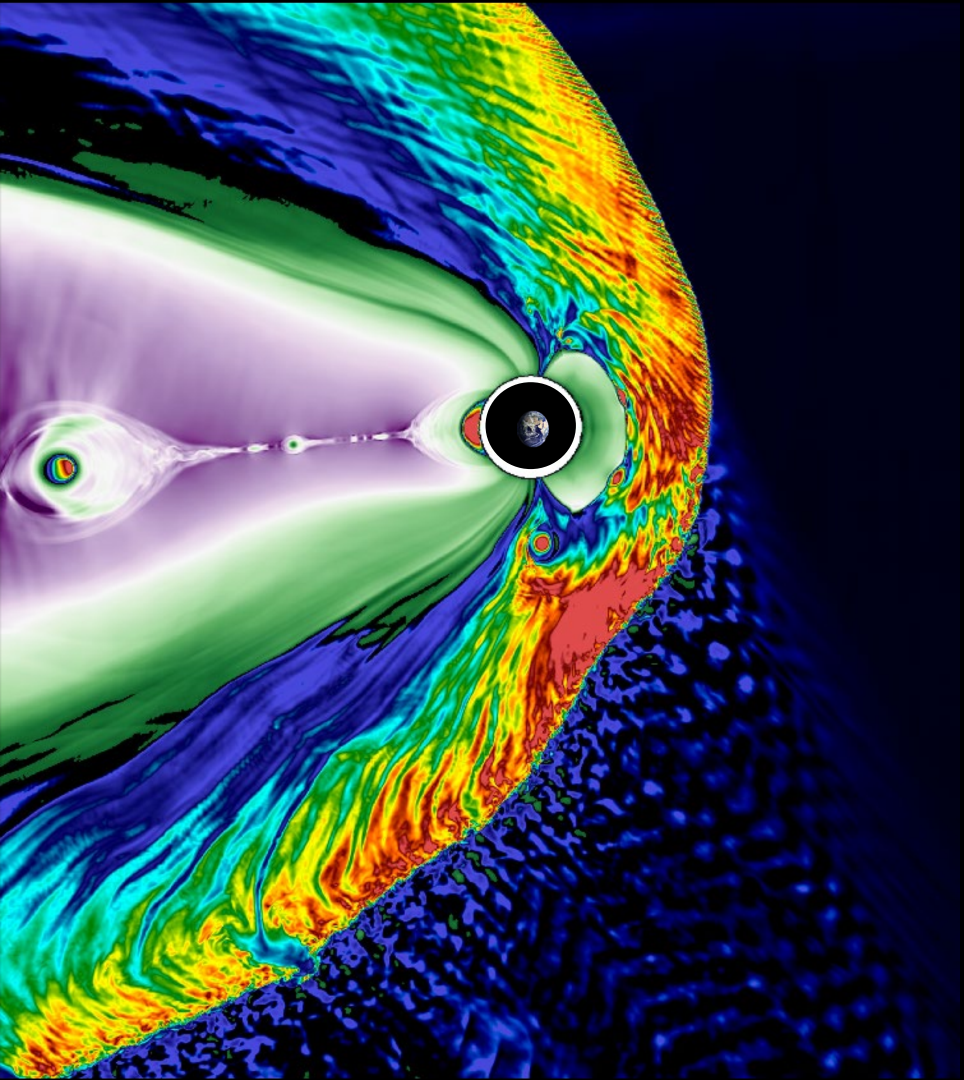

Thank you!

Contact:
minna.palmroth@helsinki.fi

Dynamos and dynamics in the solar convection zone

DYN²SOL

Maarit J. Korpi-Lagg
Associate professor (tenured),
Department of Computer Science, Aalto

Outline

- Challenge
- Solution
- Technical aspects
- Results
- Outlook

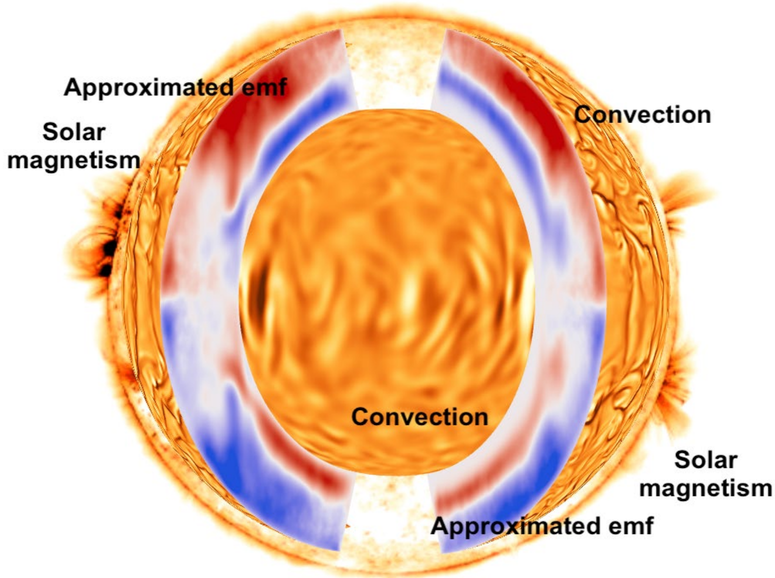

Aalto University
School of Science

European Research Council
Established by the European Commission

PRACE PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Challenge

Solar dynamo

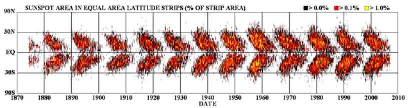

Large-scale dynamo (LSD)

Solar dynamics

Non-uniform rotation

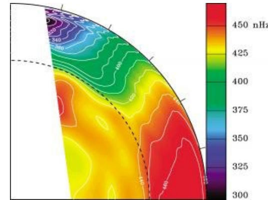

Small-scale dynamo (SSD)

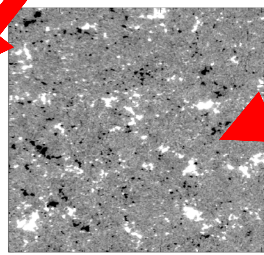

Convection

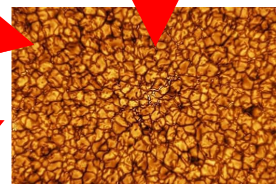

Aalto University
School of Science

European Research Council
Established by the European Commission

PRACE
PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Challenge

”Convection conundrum” – solar dynamics is not correctly modelled

Solar dynamo models cannot be trusted

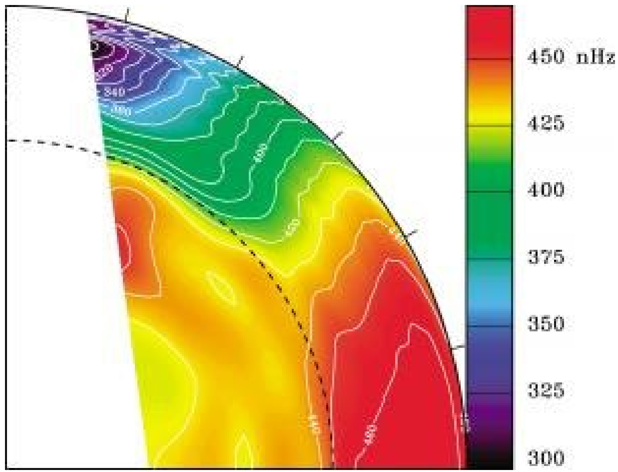

Observation

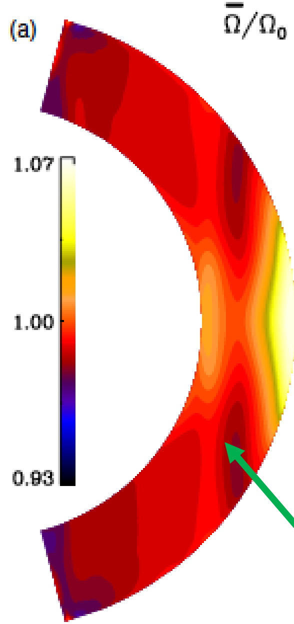

Model

“Nose”
Too cylindrical

Pencil Code

EULAG

ASH

Aalto University
School of Science

European Research Council

PRACE
PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Challenge

MHD equations

$$
\begin{aligned}
\frac{D\rho}{Dt} &= -\rho \boldsymbol{\nabla} \cdot \boldsymbol{u}, \\
\rho\frac{D\boldsymbol{u}}{Dt} &= \rho \boldsymbol{f} - \boldsymbol{\nabla} p + \boldsymbol{J} \times \boldsymbol{B} + \mu \boldsymbol{\nabla}^2 \boldsymbol{u} + (\xi + \tfrac{1}{3}\mu) \boldsymbol{\nabla}\boldsymbol{\nabla} \cdot \boldsymbol{u}, \\
\rho\frac{De}{Dt} &= -p \boldsymbol{\nabla} \cdot \boldsymbol{u} + \boldsymbol{\nabla} \cdot k \boldsymbol{\nabla} \boldsymbol{T} + \frac{\mu}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \tfrac{2}{3}\delta_{ij}\frac{\partial u_k}{\partial x_k}\right)^2 \\
&\qquad + \xi(\boldsymbol{\nabla} \cdot \boldsymbol{u})^2 + \rho\eta\mu_0 \boldsymbol{J}^2, \\
\frac{\partial \boldsymbol{B}}{\partial t} &= \boldsymbol{\nabla} \times (\boldsymbol{u} \times \boldsymbol{B} - \eta\mu_0 \boldsymbol{J}).
\end{aligned}
$$

Solution(s)

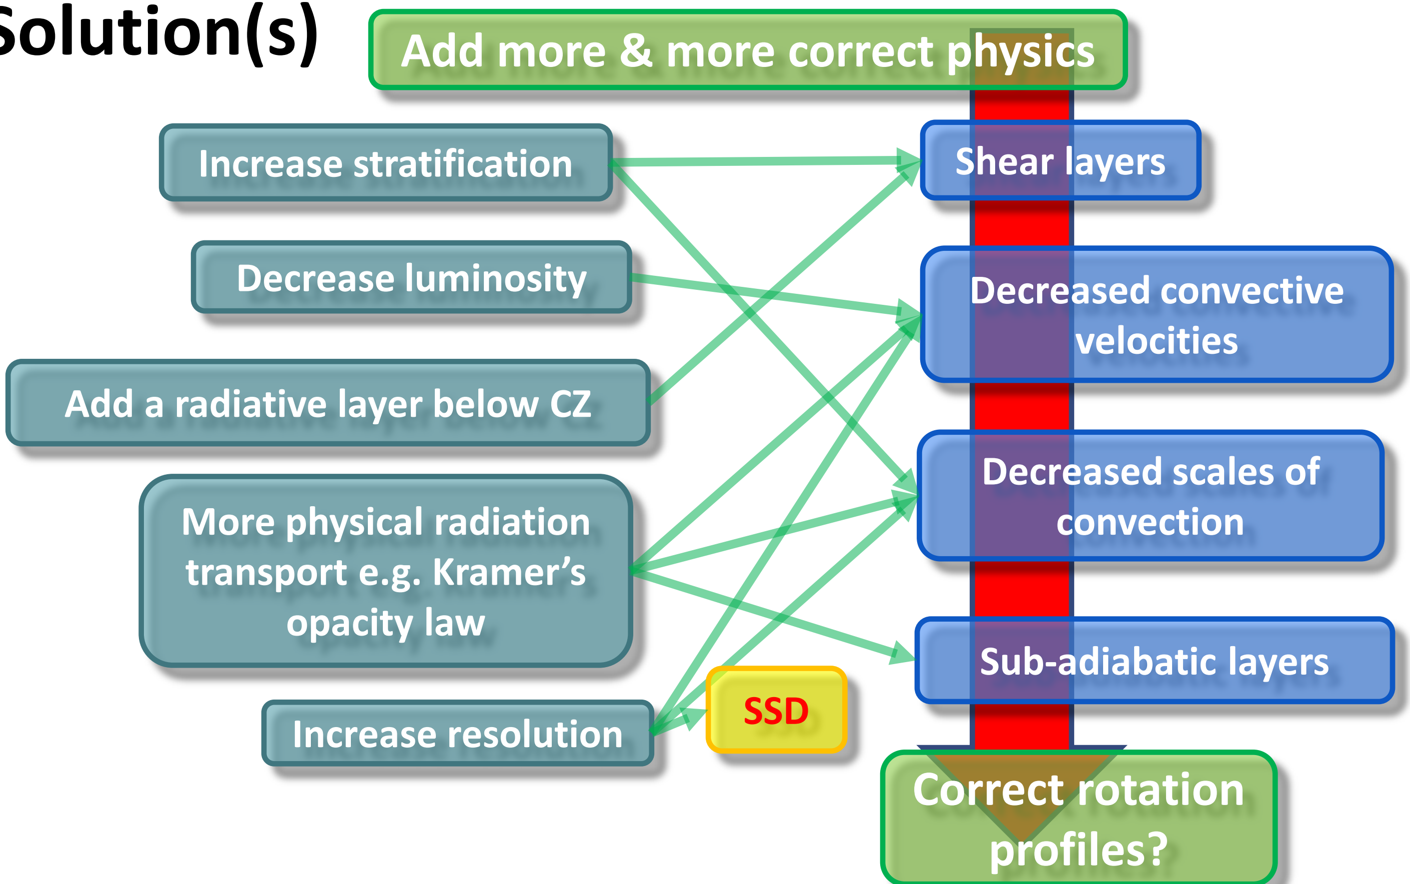

Aalto University
School of Science

European Research Council
Established by the European Commission

PRACE PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Technical aspects

High resolution

Exciting SSD requires **high resolution**; more difficult than investigating LSD alone.

Long integration times

Integration times for SSD and LSD together **are very long** (many dynamo cycles required).

Special CPU/GPU allocations required

Objective: to resolve the long-standing problem of correctly capturing the dynamics and dynamos in solar and stellar convection zones.

Aalto University
School of Science

European Research Council
Established by the European Commission

PRACE PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Technical aspects

PRACE INTERDYNS, 57 MCPUhours, 2048x4096x2048 grid simulations already made; why even higher resolution required?

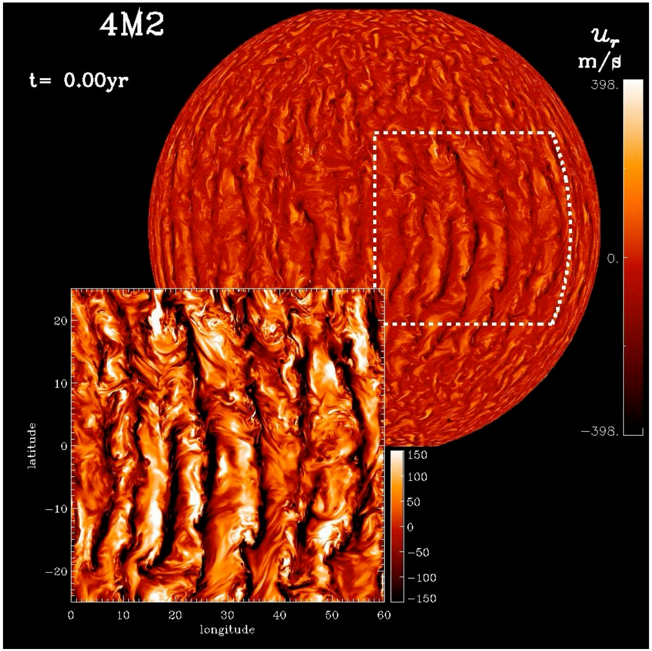

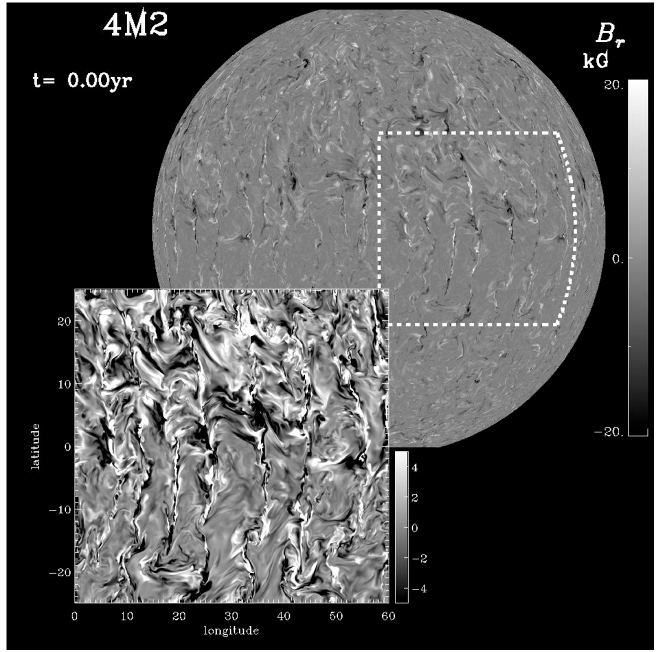

Warnecke+2025, A&A, 696, A93

Aalto University School of Science
European Research Council
Established by the European Commission
PRACE PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

No convergence!

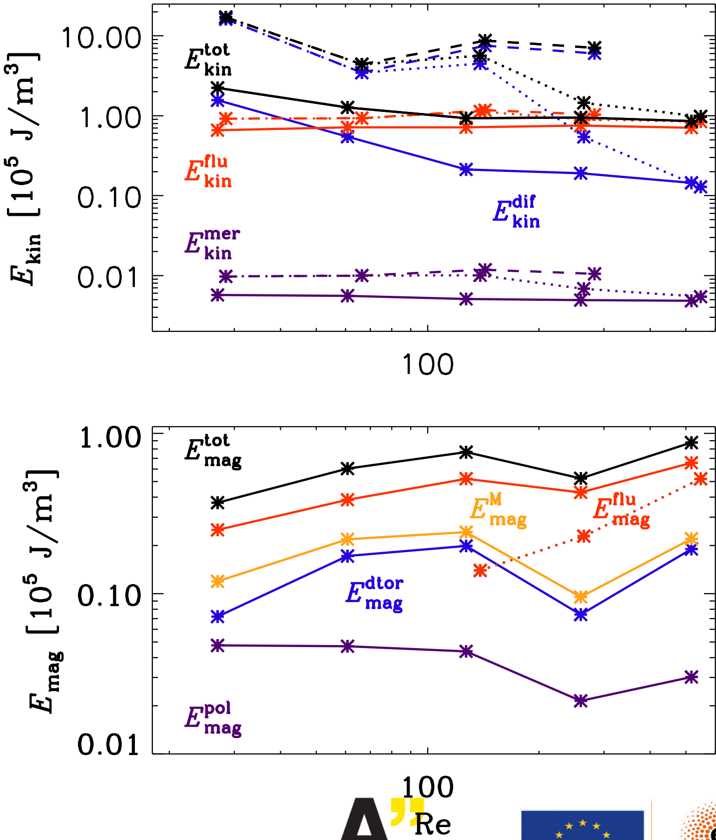

H = dashed = pure hydro

M= solid = full MHD

S = dotted = SSD only by removing the mean field at each time step

Magnetic fluctuations are still growing.

Mean field energies show a dip but then a revival.

Differential rotation is still not saturated

Aalto University School of Science

European Research Council Established by the European Commission

PRACE PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Technical aspects: how to reach even higher resolutions?

GPU-accelerated solvers

Astaroth library: order of magnitude acceleration; with 4096-8192 GPUs we can now achieve 2048x4096x512/4096x8192x1024 resolutions.

HPC ecosystem to tackle the data explosion EuroHPC Inno4scale NEOSC project

Data storage becomes a serious issue; data analysis operations on the fly

Astaroth is integrated with **multithreaded Pencil Code** to perform some of the data analysis with unused CPUs.

Oppostunistic data analysis platform (ODOP) to schedule further data analysis and movement operations to unused CPU/GPU resources.

Aalto University
School of Science
European Research Council
Established by the European Commission
PRACE PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Results

Finalizing the HPC ecosystem to enable these runs took way more humsn resources and time than we anticipated.

We had to apply for an extension; first reliable results will be obtained before Xmas

HPC ecosystem is working flawlessly and we are collecting very valuable data

Aalto University
School of Science
European Research Council
Established by the European Commission
PRACE
PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

Outlook

Computing at Exascale is an absolutely non-trivial task

Going for even higher resolutions will not be easily possible despite of the HPC ecosystem we have developed

Time to think out of the box!

Aalto University
School of Science

European Research Council
Established by the European Commission

PRACE
PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE