Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA

ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

The OpenGADGET3 code for cosmological simulations

- An update in preparation of the Key Science Projects -

Milena Valentini

1924 2024

UNIVERSITÀ DEGLI STUDI DI TRIESTE

Dipartimento di Fisica
Dipartimento d'Eccellenza 2023-2027

Spoke 3 General Meeting, Elba 5-9 / 05, 2014

Scientific Rationale

- Numerical cosmology
- Structure formation and evolution

Open-GADGET3

USM LMU

UNIVERSITÀ DEGLI STUDI DI TRIESTE

INAF OSSERVATORIO ASTRONOMICO DI TRIESTE

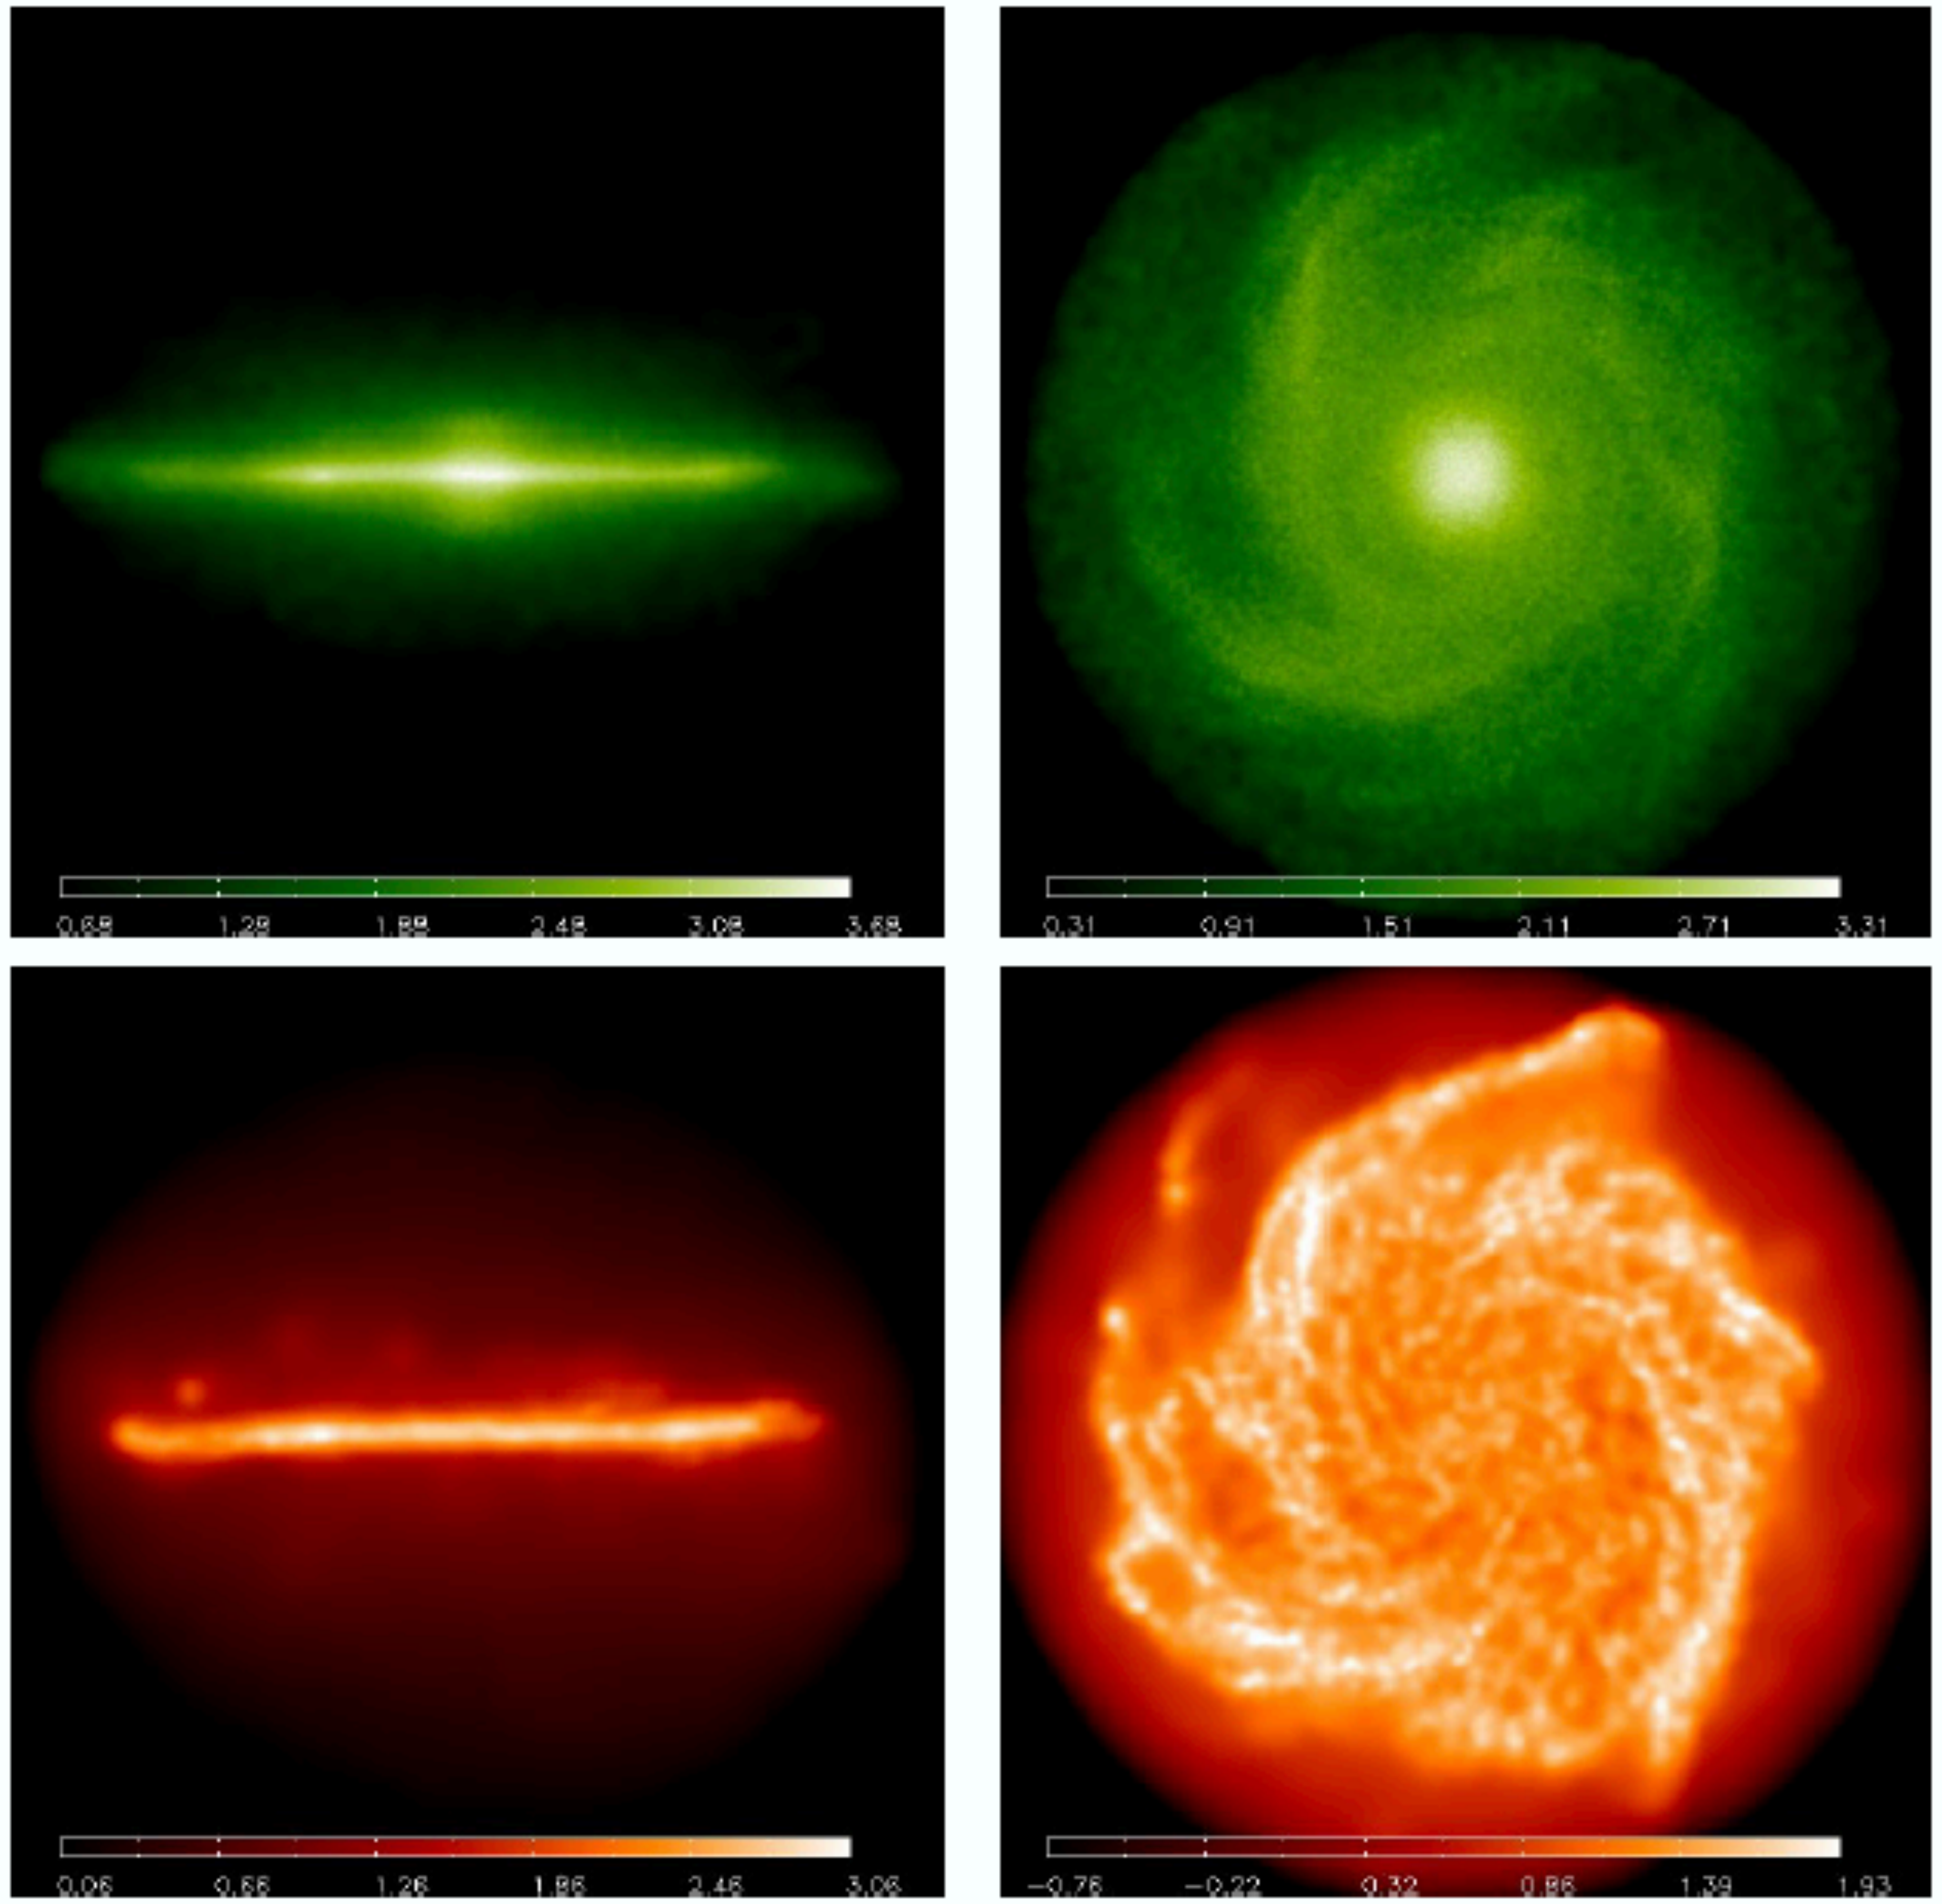

Technical Objectives, Methodologies and Solutions

The Lagrangian code Open-GADGET3

Code: descendant of our developer version of GADGET-3 (TreePM+SPH, originally from Springel 2005), featuring different hydro solvers (SPH, newSPH, MFM) and several advanced physical modules (e.g. chemical evolution & enrichment by L. Tornatore)

Main differences between Open-GADGET3 and its predecessor include: restructuring of calls to functions, tasks split in several individual functions, enhanced modularity and OpenMP parallelization…

The code

- **TreePM+SPH code**
- **Descendant of a non-public evolution of GADGET-3 code**
- **State-of-the-art code for cosmological hydrodynamical simulations**
- **Highly optimised code:** MPI parallelised (multi-segmented Peano-Hilbert curve for domain decomposition) + OpenMP
- **Improved SPH formalism (or MFM)**
- **Several modules for sub-resolution physics: star formation, stellar feedback, BH accretion and feedback, chemical enrichment, dust evolution**
- **Runs on CPUs and GPUs**

Open-GADGET3

USM

LMU

UNIVERSITÀ DEGLI STUDI DI TRIESTE

Technical Objectives, Methodologies and Solutions

The Lagrangian code Open-GADGET3

Code: descendant of our developer version of GADGET-3 (TreePM+SPH, originally from Springel 2005)

Key differences between Open-GADGET3 and its predecessor GADGET-3

The code

- **TreePM+SPH code**
- **Descendant of a non-public evolution of GADGET-3 code**
- **State-of-the-art code for cosmological hydrodynamical simulations**
- **Highly optimised code: MPI parallelised + OpenMP**
- **Improved SPH formalism (or MFM)**
- **Several modules for sub-resolution physics: star formation, stellar feedback, BH accretion and feedback, chemical enrichment, dust evolution**
- **Runs on CPUs and GPUs**

Main tasks within the WP 2 of Spoke 3

Develop Open-GADGET further:

- **including additional physics**
- **extending existing modules**
- **improving code performance**

Core team in Trieste: S. Borgani, L. Tornatore, G. Murante, M. Valentini, T. Castro, P. Monaco, G. Taffoni, A. Damiano, G. Granato, D. Goz, P. Barai, M. Parente, A. Saro, M. Viel

and collaboration in Munich led by K. Dolag

Technical Objectives, Methodologies and Solutions

The Lagrangian code Open-GADGET3

Code: descendant of our developer version of GADGET-3 (TreePM+SPH, originally from Springel 2005)

Different sub-resolution models available

- Model based on the multi-phase effective model by Springel & Hernquist (2003)
- MUPPI (Murante+ 2010, 2015, Valentini+ 2017, 2019, 2020, 2023, Granato+ 2021), featuring:

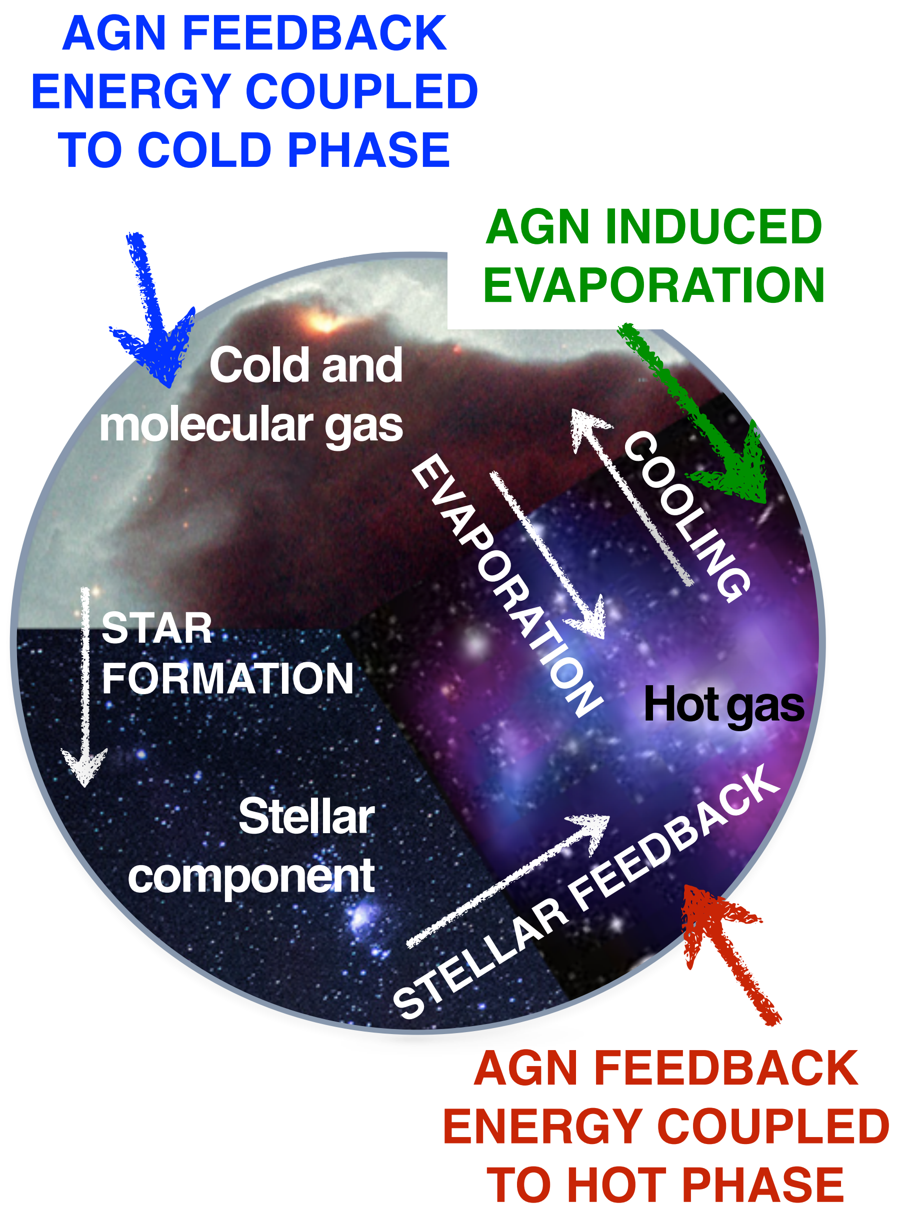

star formation
- description of a multi-phase ISM (valid for resolution particles with mass $\sim 10^4 - 10^7\ M_\odot$)
- H_2 - based star formation
- thermal and kinetic stellar feedback
- stellar feedback from low-metallicity environments
- stellar evolution and chemical enrichment

BH
- angular-momentum-dependent gas accretion
- isotropic, thermal AGN feedback

dust
- formation and evolution of dust, and dust-assisted cooling

Timescale, Milestones and KPIs

Main goals for Milestones 8, 9 and 10

- GPU optimisation
- Re-structuring of the code to enhance its modularity
- New/updated (sub-grid) physical modules for cosmological hydrodynamical simulations
- Improving OpenMP optimisation of the code and extending it to all modules

Timescale

- end of 2024: GPU porting and extension of OpenMP optimisation
- mid 2025: Code re-structuring and new sub-grid modules inclusion

Key Performance Indicators

- several talks and presentations delivered at conferences or internal/technical meetings
- publication/submission of two scientific papers on refereed journals
- code development on GitLab, in view of the upcoming public release (planned mid 2025 for code skeleton + basic modules)

Accomplished Work, Results

- We moved our code to GitLab
- We defined a more accurate working strategy
- Quite large (> 30 people from different institutes) user community

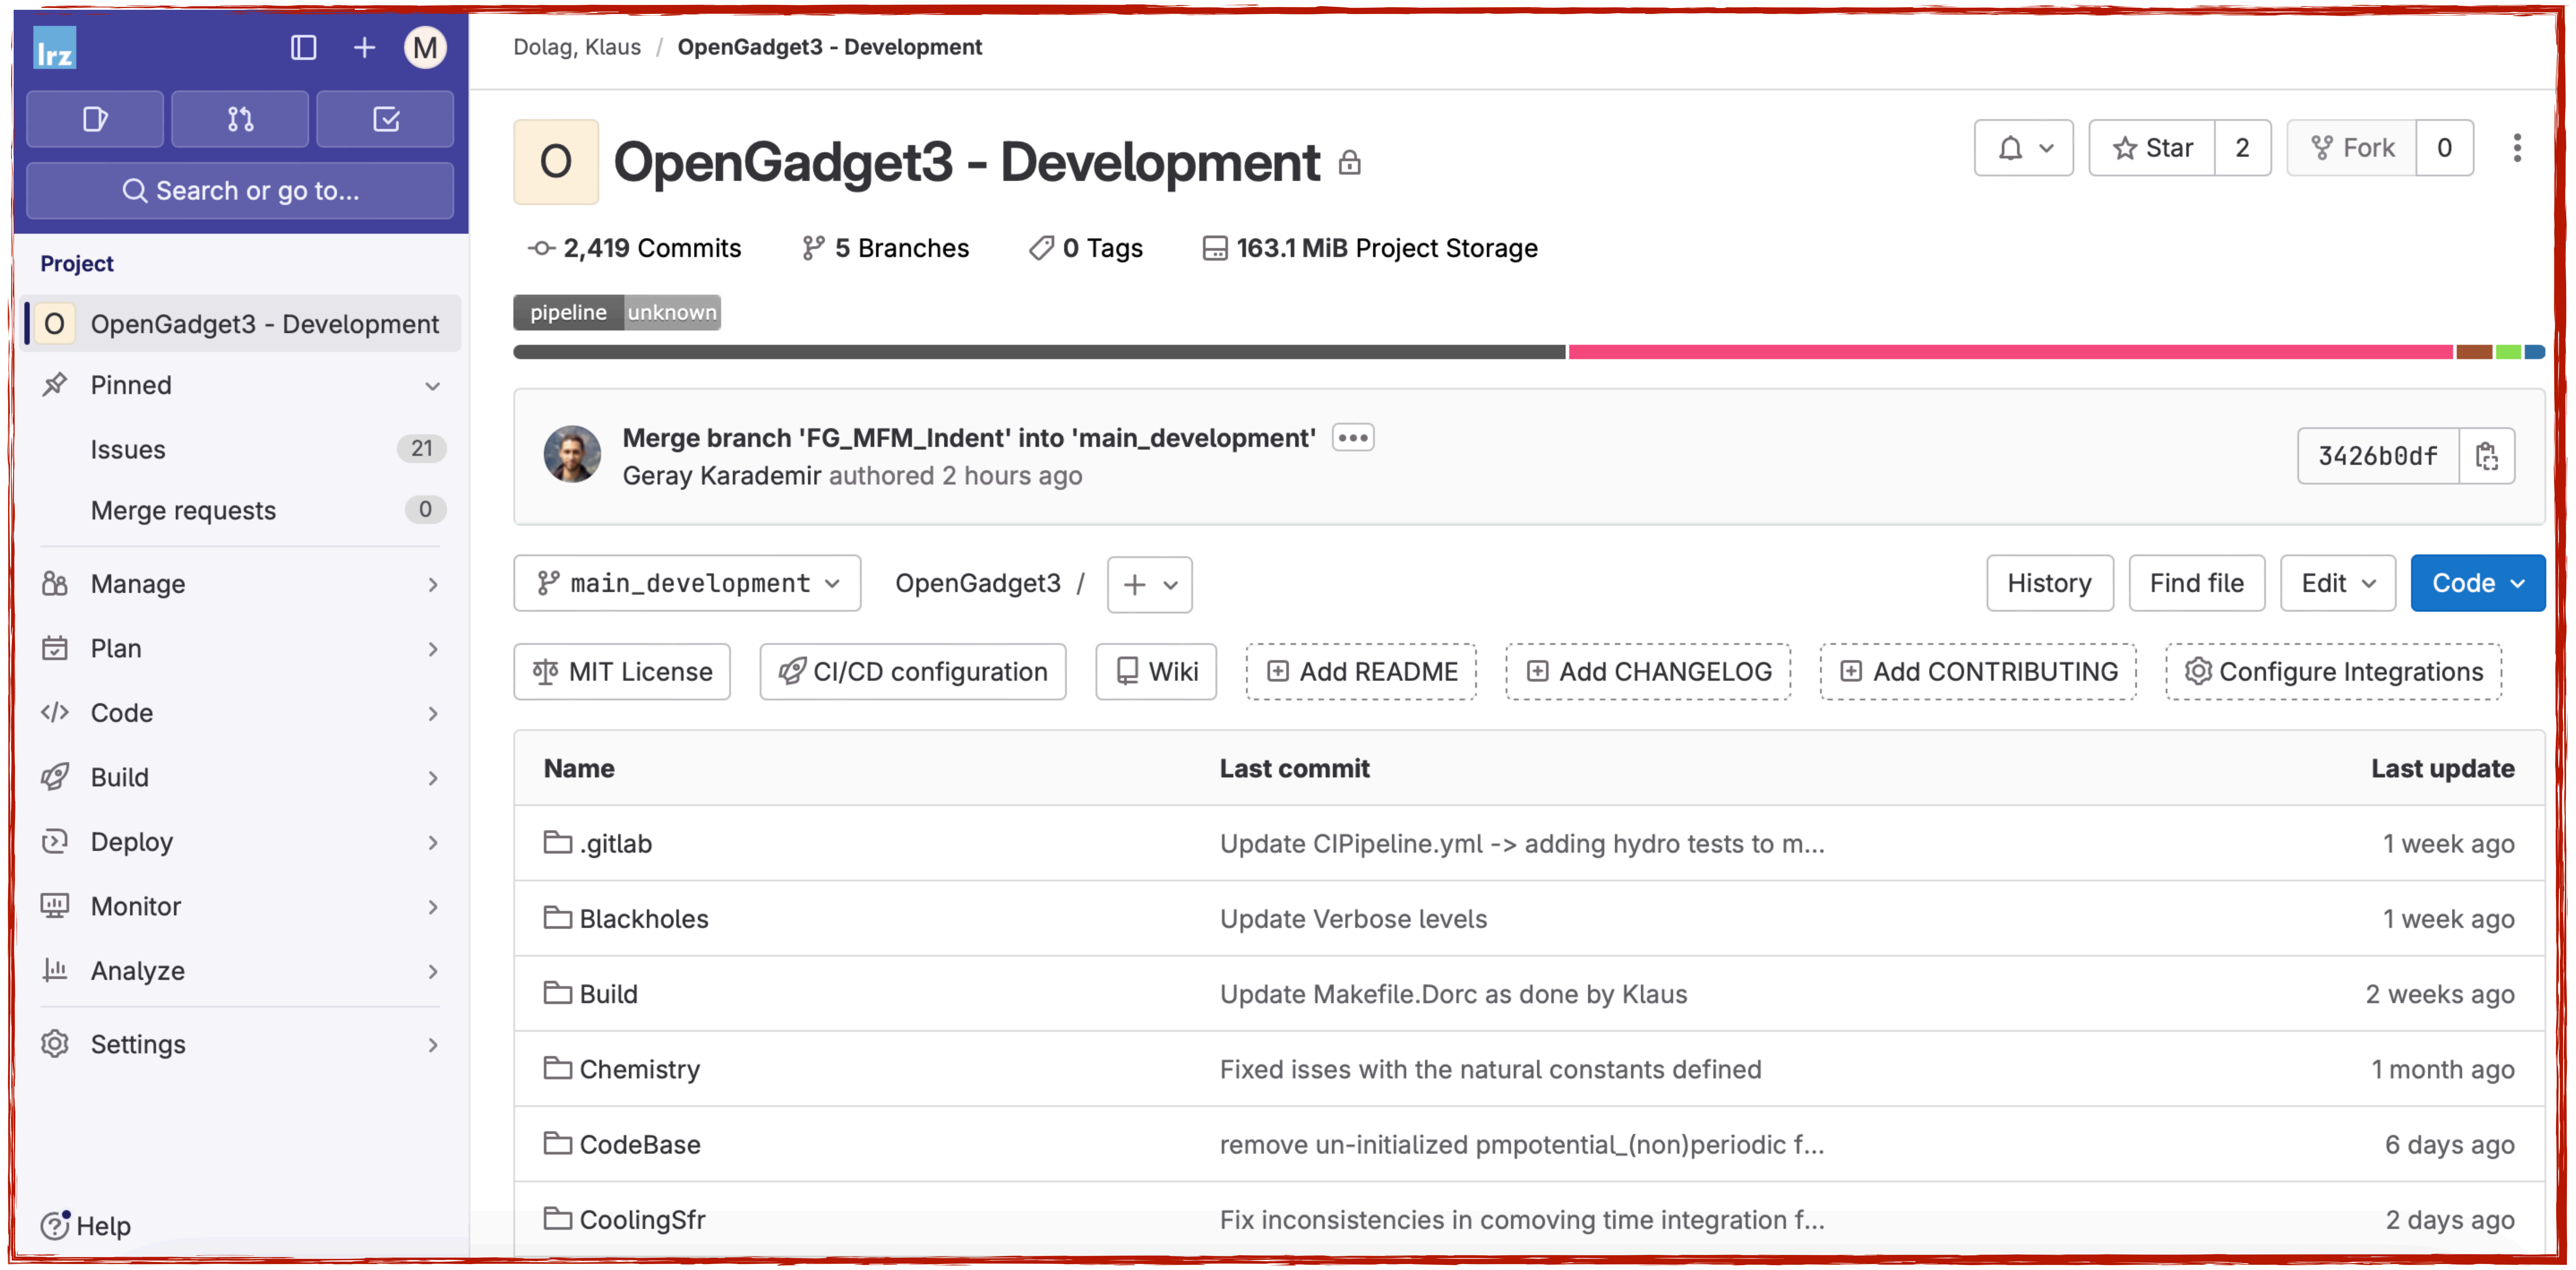

Accomplished Work, Results

- The OpenGADGET3 project aims at making the use of the many complex physics modules more user friendly.
- Substantial effort in **cleaning** and making **more transparent** the definition of the **code configurations** and of the files setting the many parameters.
- Construction of a **reference structure for the files which configure** several **reference production runs and files of parameters** for the OpenGADGET code.

Name	Last modified	Size	Description
Parent Directory		-	
Config_MV_Muppi.sh	2024-03-28 17:27	9.9K	
Config_MV_SH03.sh	2024-03-28 17:28	7.5K	
Files_aux/	2024-02-13 13:08	-	
OpenGadget3.tar	2024-02-02 15:24	14M	
OpenGadget3/	2024-02-02 15:23	-	
dfrogin_25x_Muppi/	2024-02-12 12:08	-	
dfrogin_25x_SH03/	2024-02-12 12:08	-	
dfrogin_250x_Muppi/	2024-02-12 12:08	-	
dfrogin_250x_SH03/	2024-02-12 12:08	-	
dianoga_g15_25x_Muppi/	2024-02-12 12:09	-	
dianoga_g15_25x_SH03/	2024-02-12 12:10	-	
magneticum_box4_25x_Muppi/	2024-02-12 11:57	-	
magneticum_box4_25x_SH03/	2024-02-12 12:09	-	

Accomplished Work, Results

- The OpenGADGET3 project aims at making the use of the many complex physics modules more user friendly.
- Substantial effort in **cleaning** and making **more transparent** the definition of the **code configurations** and of the files setting the many parameters.
- Construction of a **reference structure for the files which configure** several **reference production runs and files of parameters** for the OpenGADGET code.
- **Bug fixing** and tackling subtleties of the sub-grid modelling.

- **Re-structuring** of the code (**modularity**)
- **Cleaning** the code **and documenting** its status

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Accomplished Work, Results

Sub-resolution accretion disc links BH accretion on resolved scales to BH spin evolution

BH accretion rate → disc mass → accretion episodes

BH spin direction and magnitude defined by angular momentum of accreting gas

Given initial BH spin and mass → post-merger BH spin vector

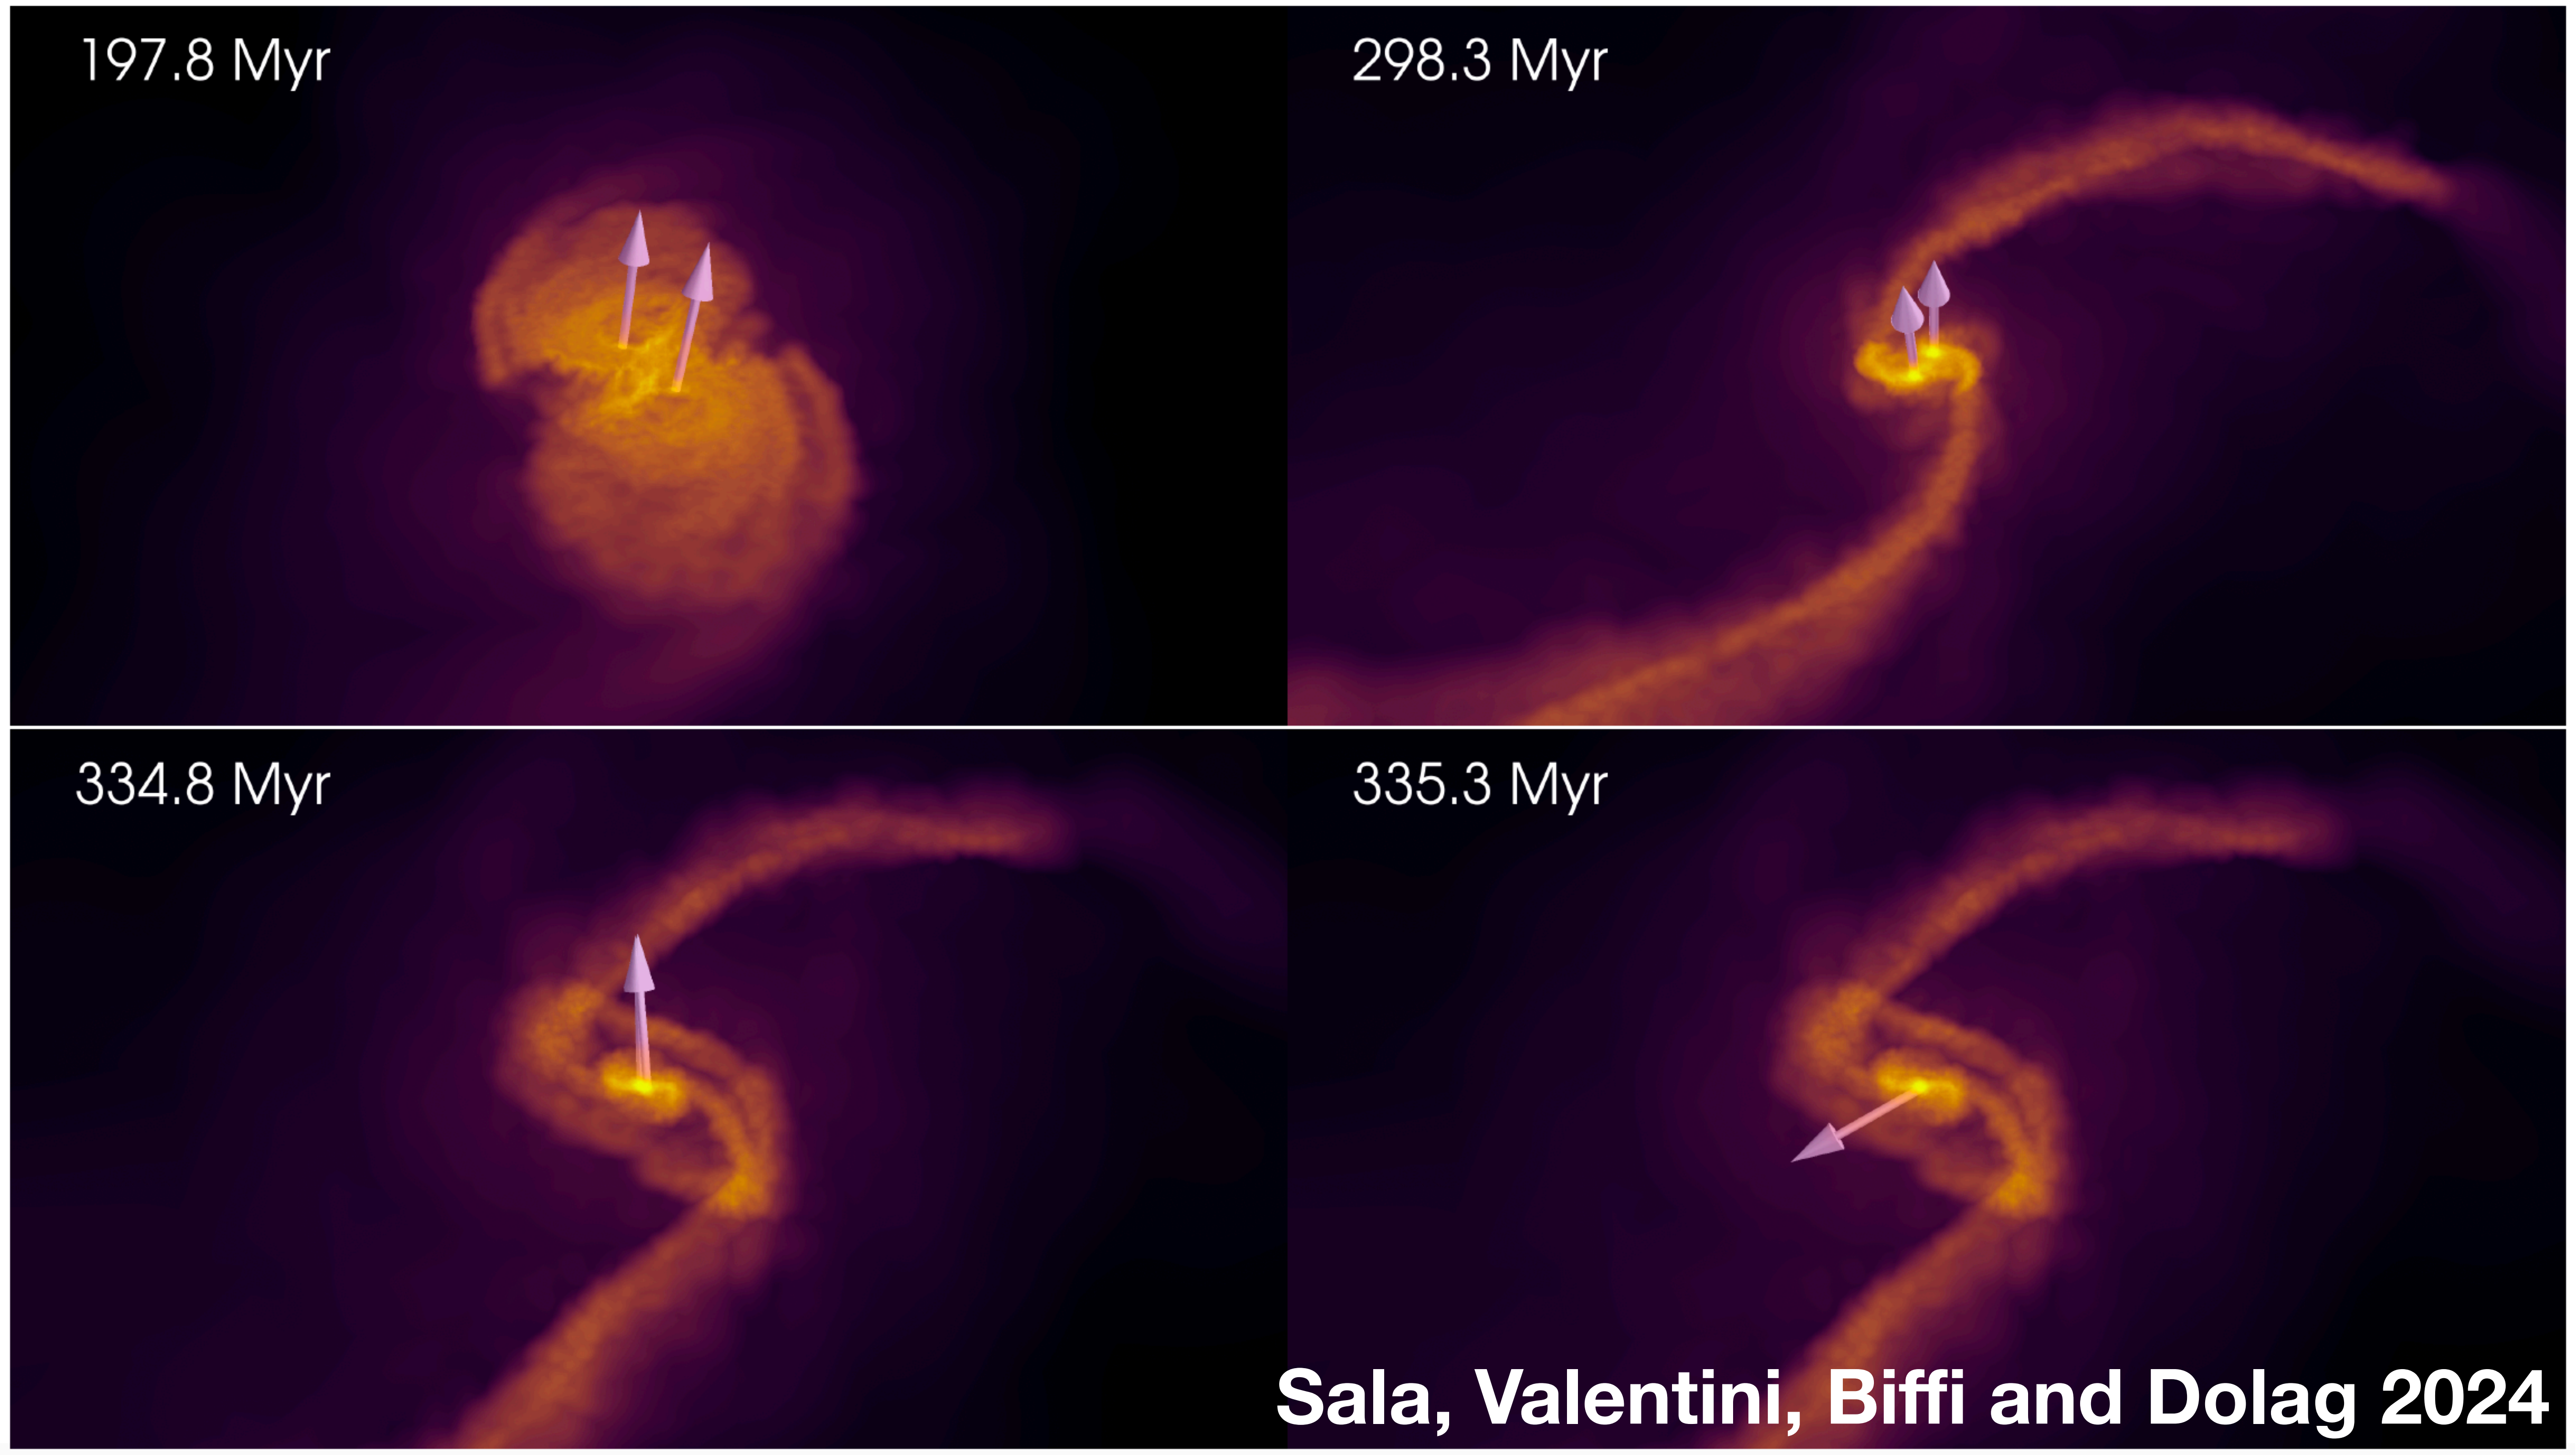

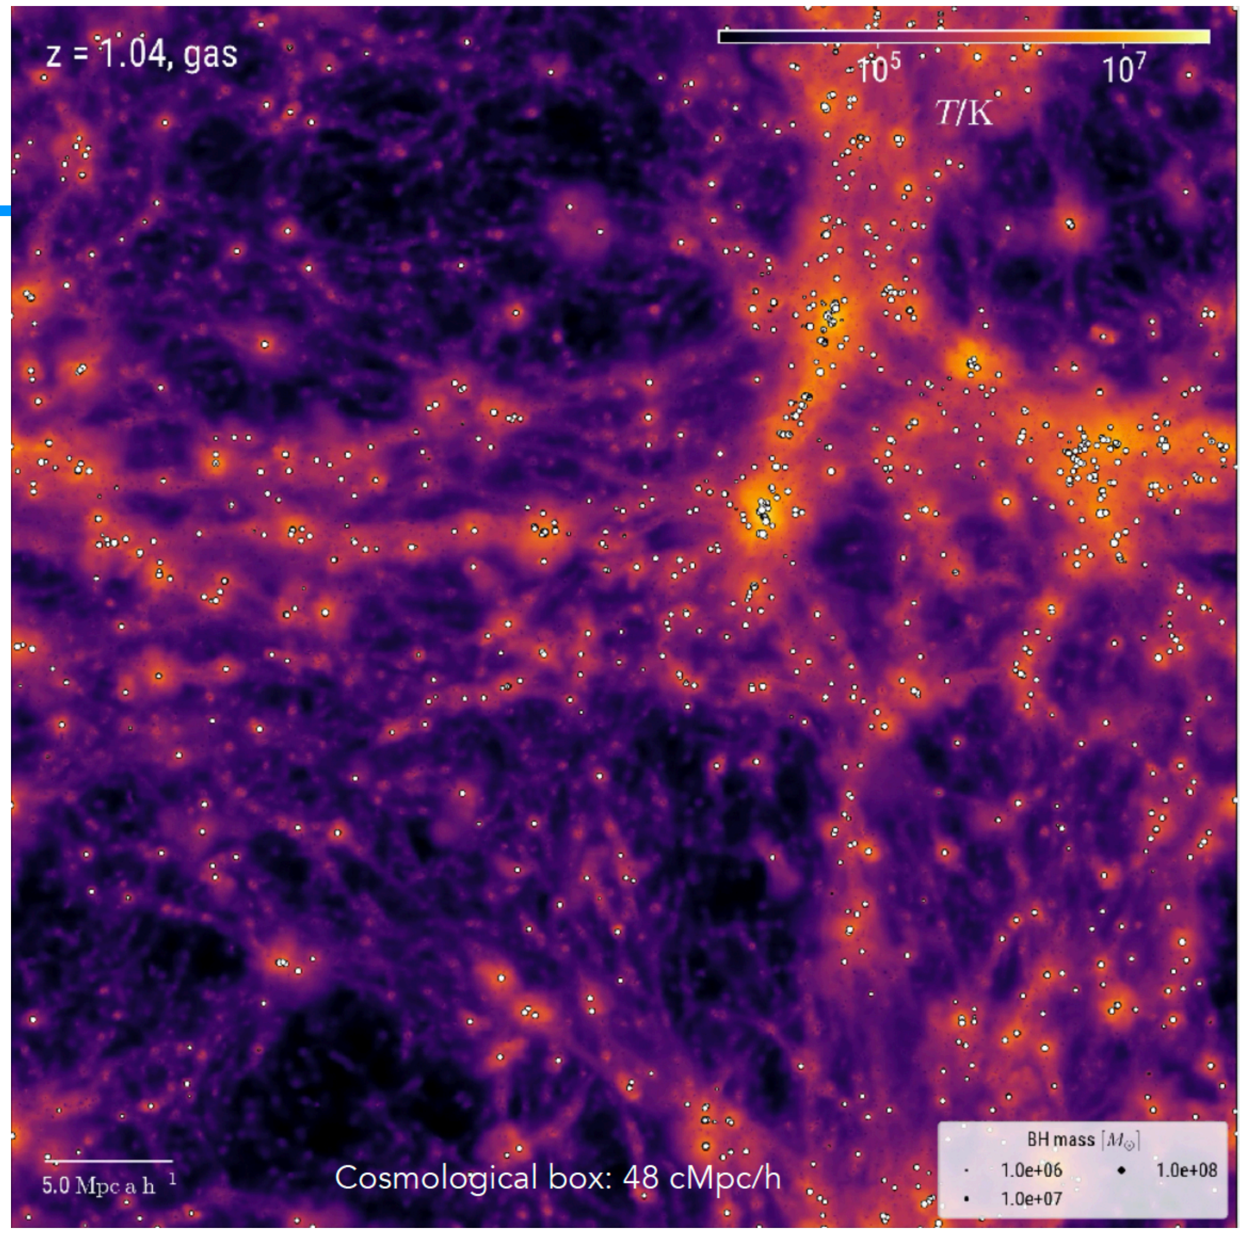

BH radiative efficiency dependent on BH spin

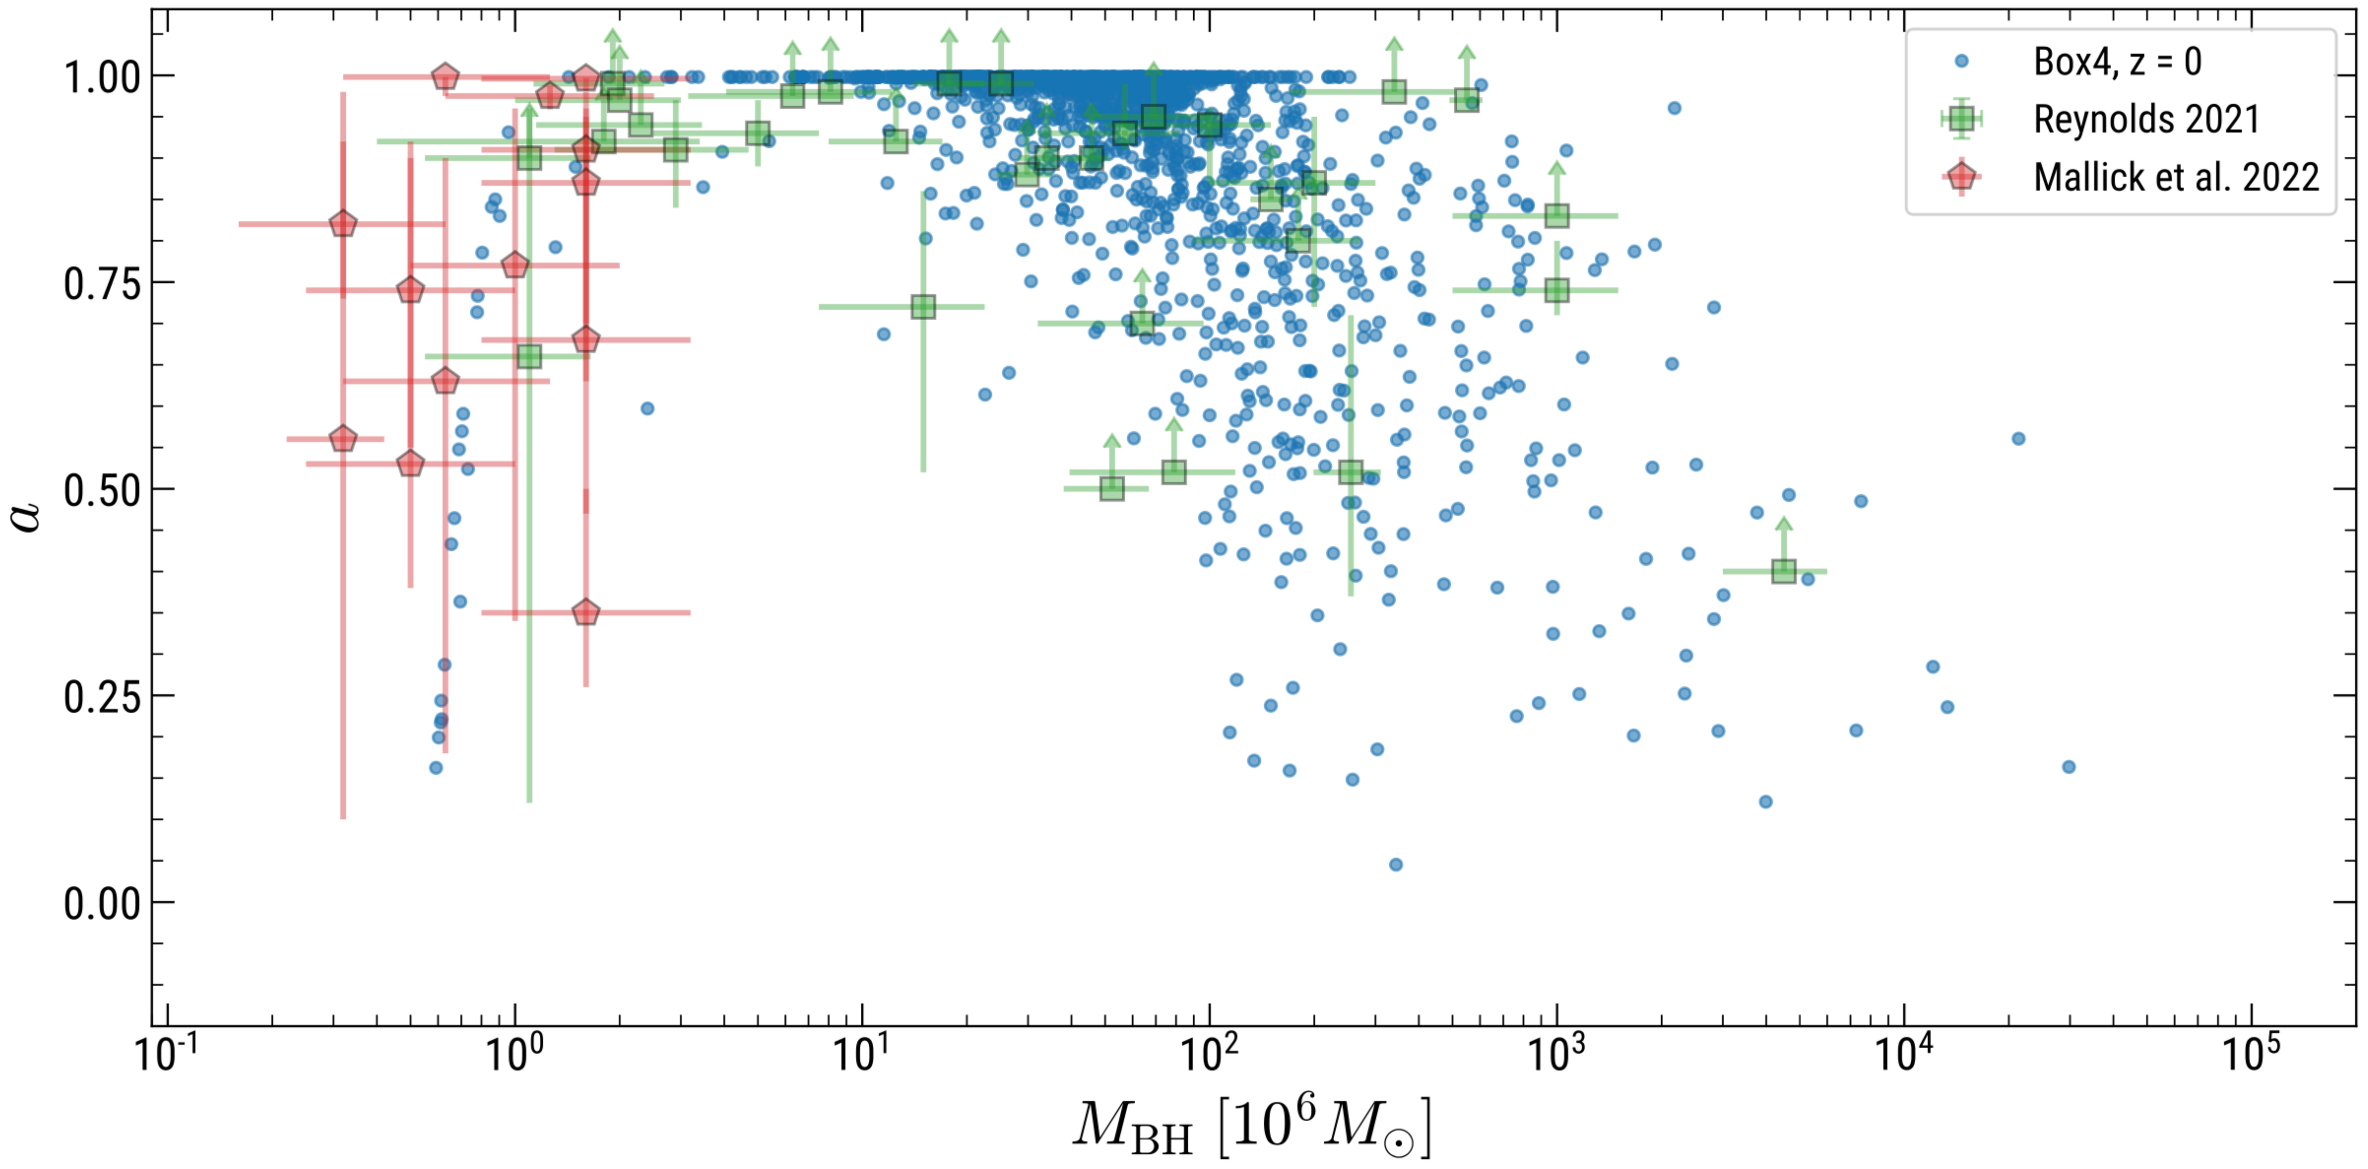

Accomplished Work, Results

Damiano, Valentini, Borgani, Tornatore+ 2024

Adopting different numerical prescriptions
for BH re-positioning has an impact
on **BH dynamics**, AGN feedback, BH-BH mergers

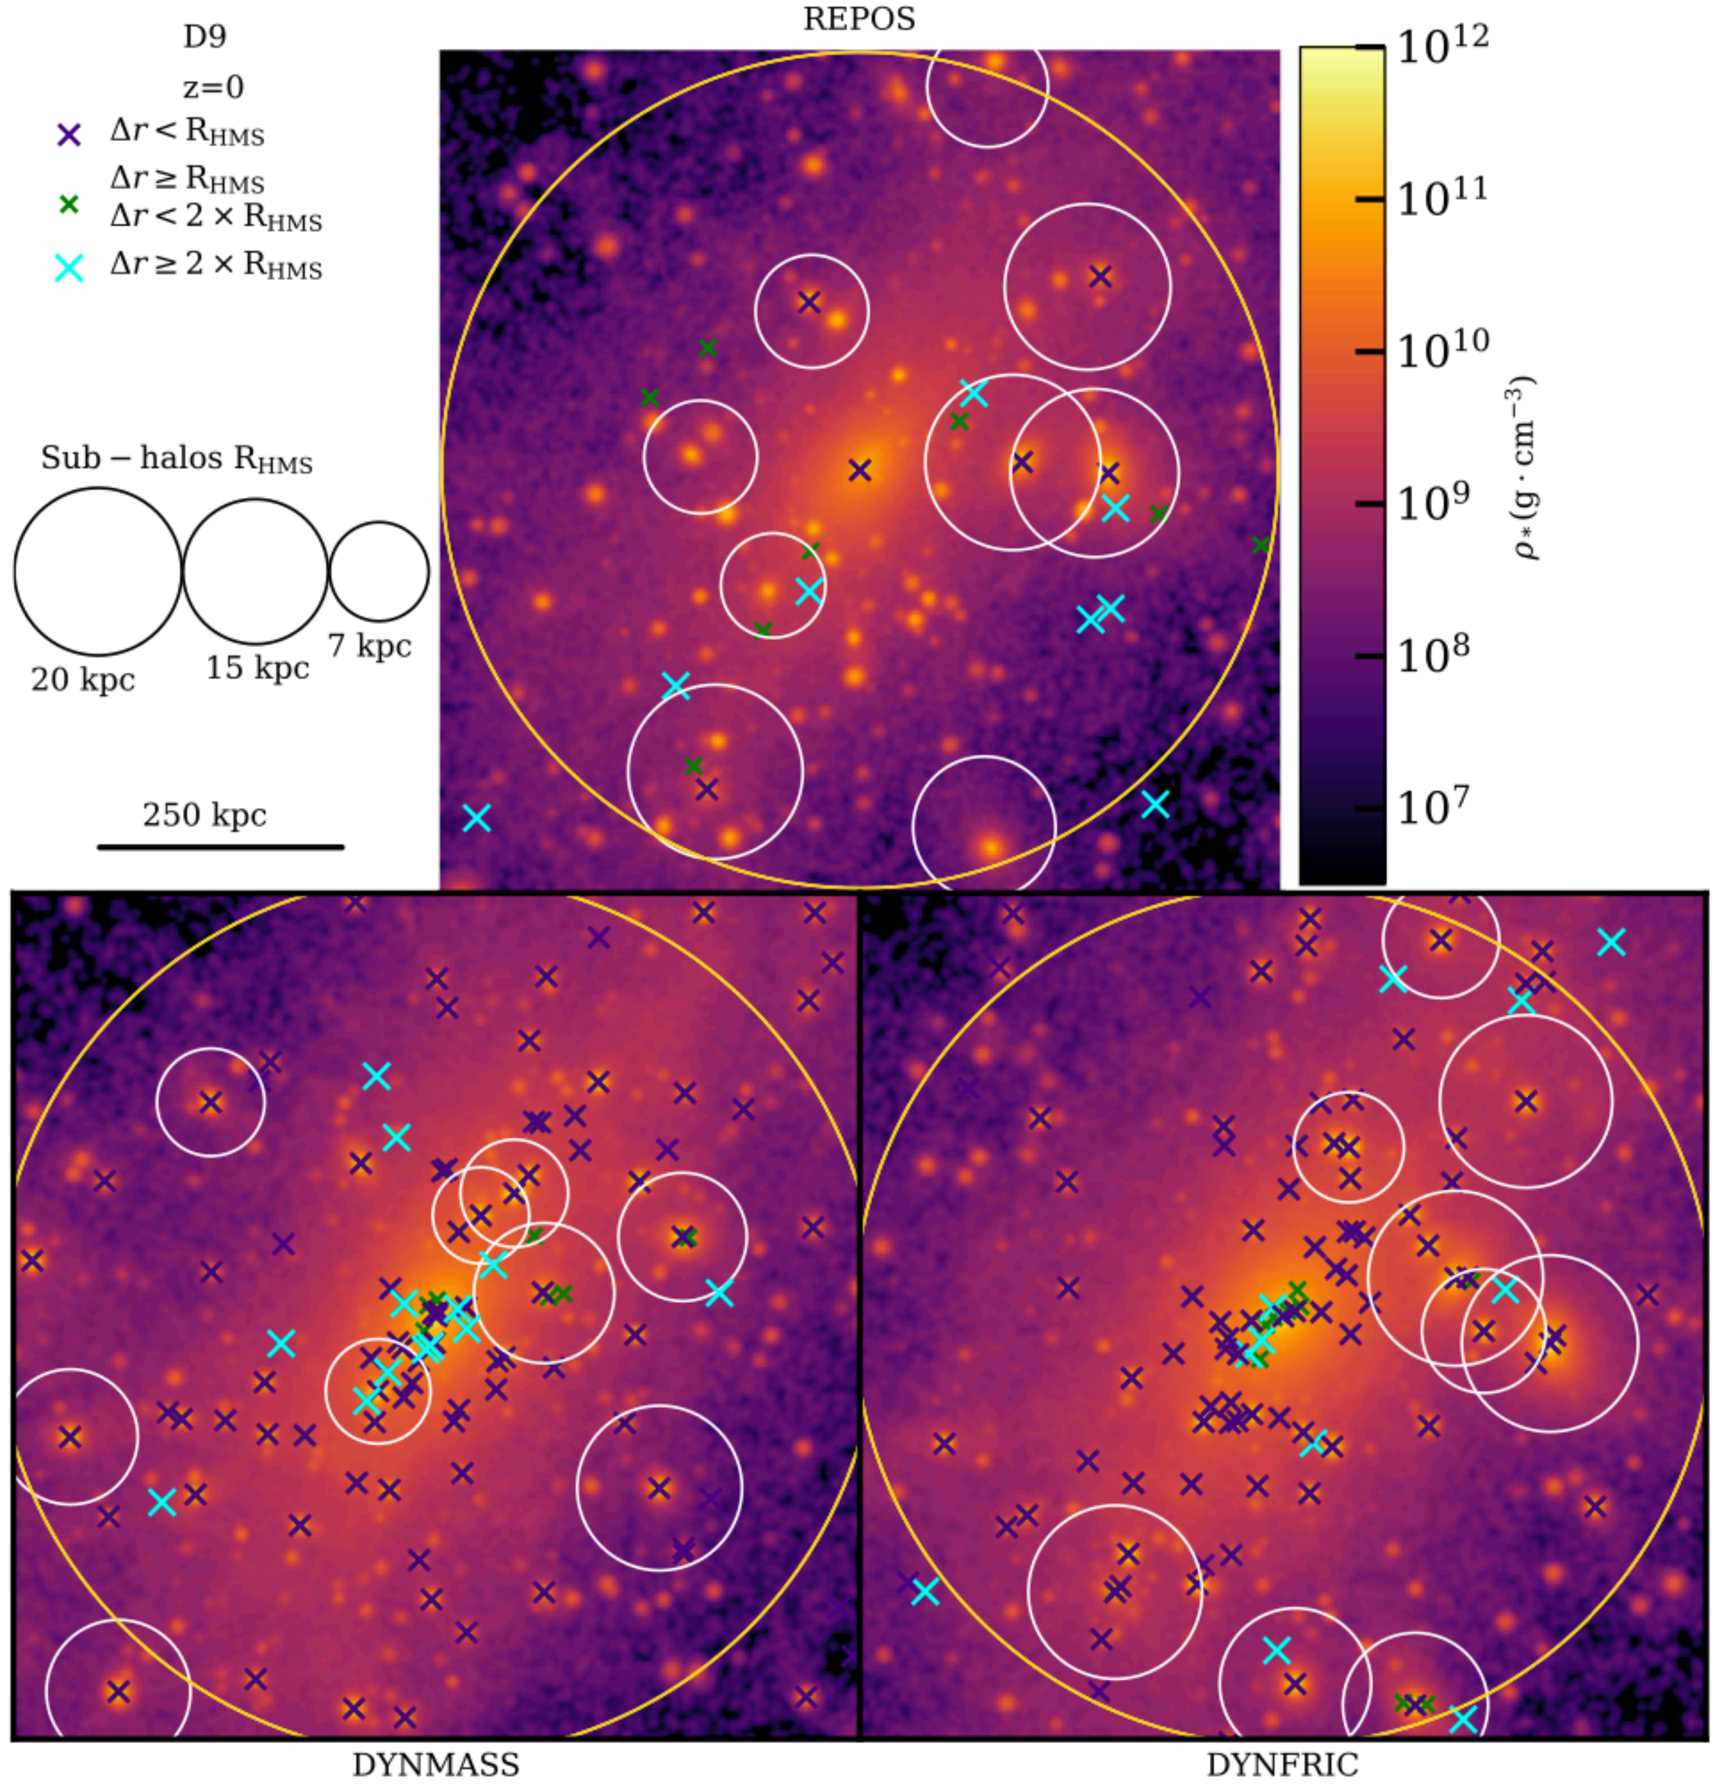

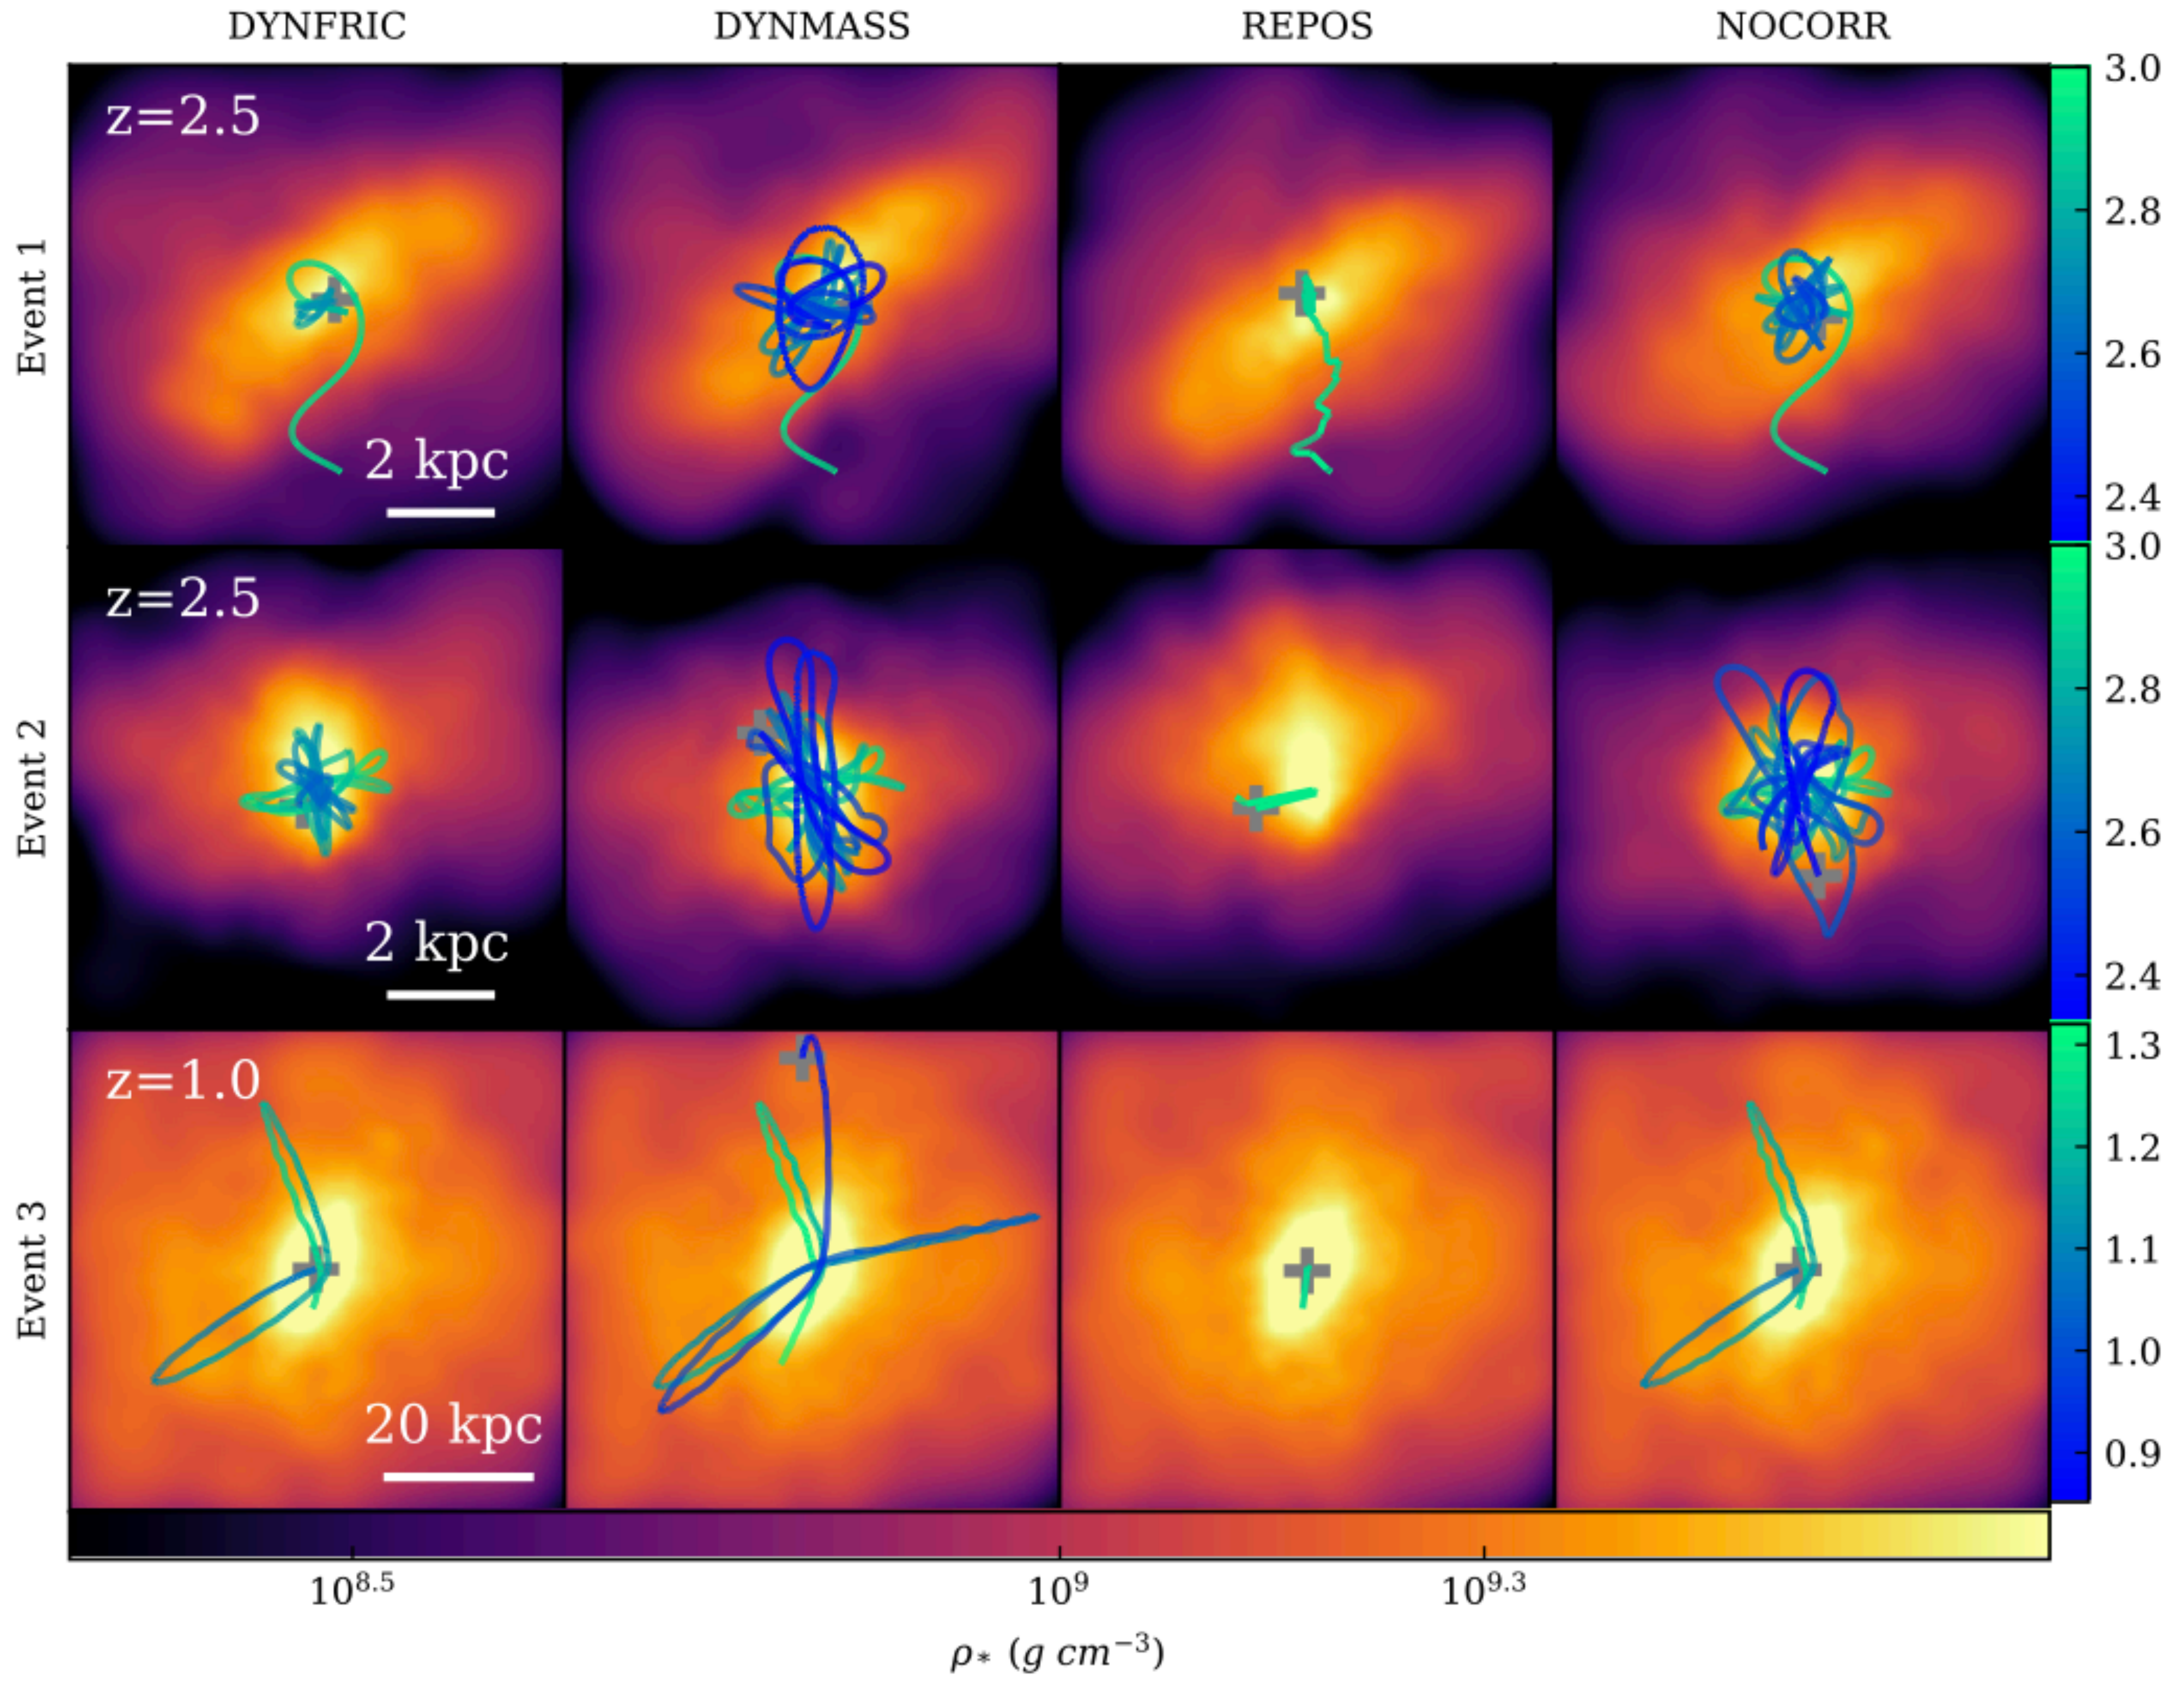

Ongoing: Assessing scalability, targeting performance issues

1. **GPU scalability**

   OpenGadget has most of the modules running on GPUs (thanks to A. Ragagnin).

   We are assessing in detail the scalability of this implementation in order to highlight the blocking factors, mitigate their impact or turn to new strategies with greater parallelism

2. **Performance issues**

   Detailed profiling with the assistance of POP and SPACE Centers of Excellence

Coordinator of the work: L. Tornatore

Ongoing: Assessing scalability, targeting performance issues

1. GPU scalability

OpenGadget has most of the modules running on GPU (thanks to A. Ragagnin).

Running a suite of tests, we are assessing in detail the scalability, from 4 nodes up to the entire Leonardo, of this implementation in order to highlight the blocking factors, mitigate their impact or turn to new strategies with greater parallelism.

size (Mpc)	resolution	physics	min Nodes	max Nodes
Cosmological Boxes at z~50				
120	$2{\times}1024^3$	gravity, grav + hydro	4	128
240	$2{\times}2048^3$	gravity, grav + hydro	32	512
480	$2{\times}4096^3$	gravity, grav + hydro	128	3400
Cosmological Boxes at z~1 (magneticum Box3, courtesy of K. Dolag)				
128	$2{\times}576^3$	gravity, grav + hydro	4	128
128	$2{\times}1536^3$	gravity, grav + hydro	32	512

Ongoing: Assessing scalability, targeting performance issues

1) GPU scalability: Speed-Up

2×1024^3, 120 Mpc, up to 512 GPUs

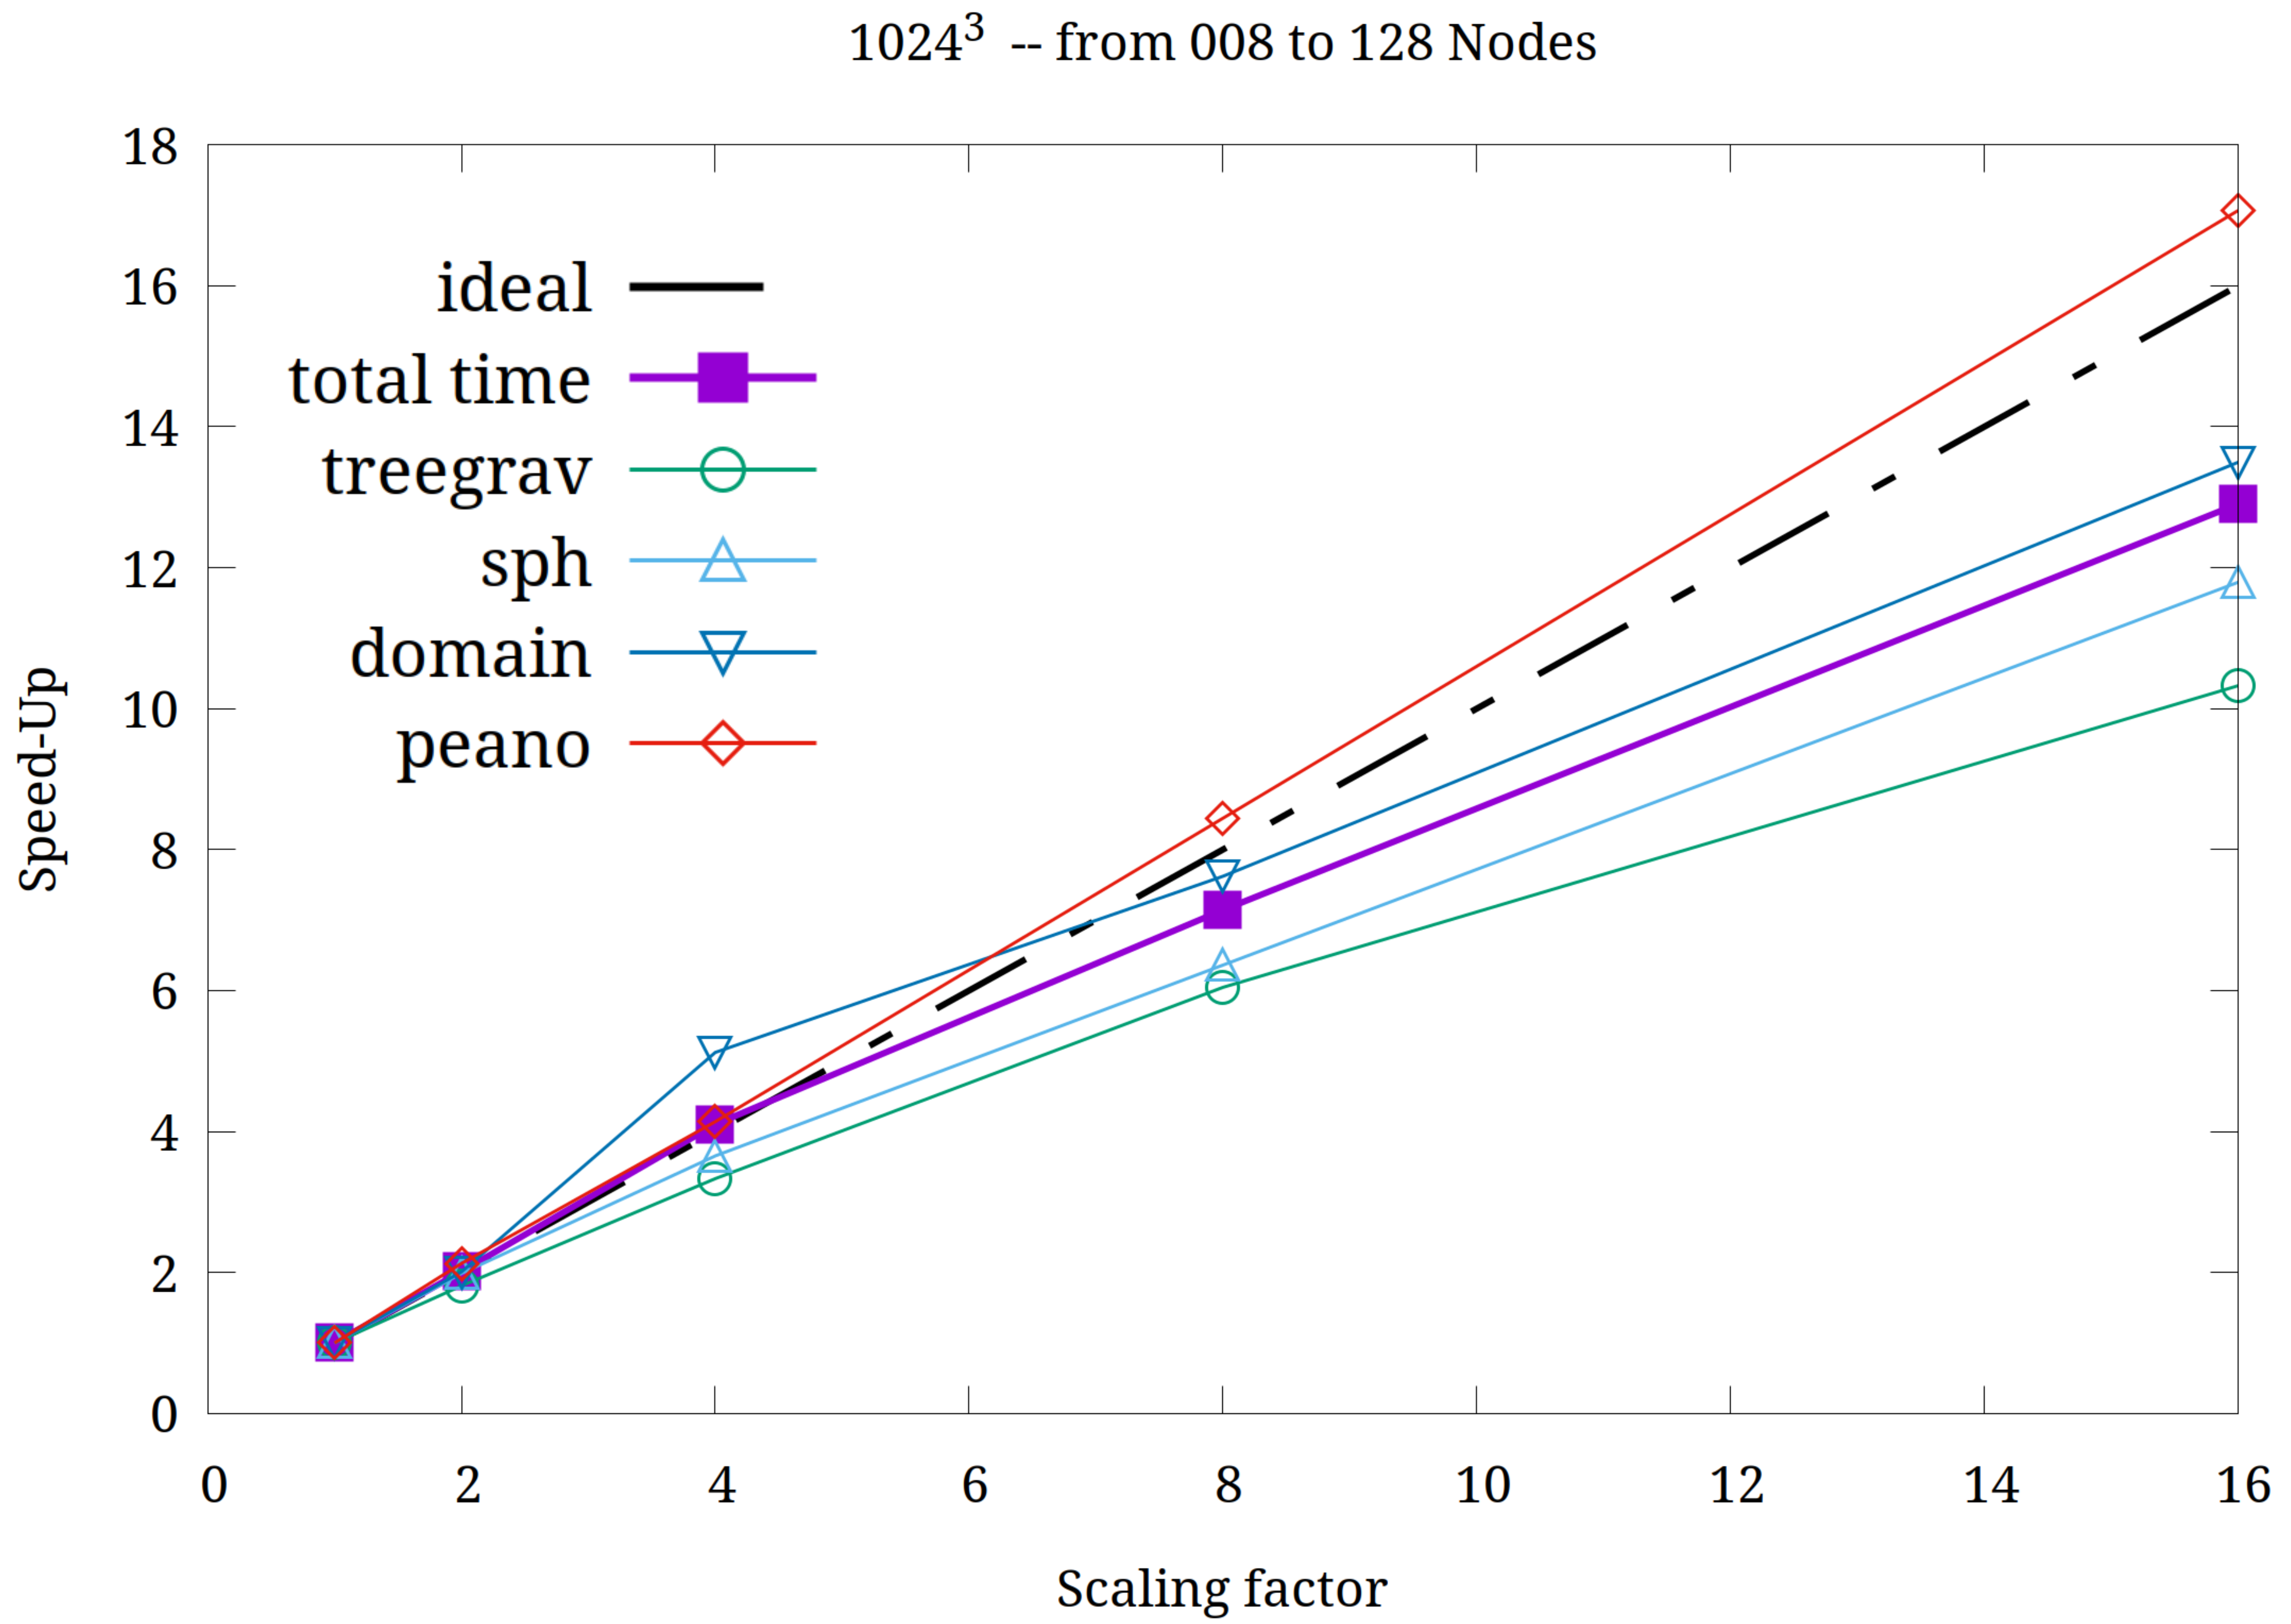

2×2048^3, 240 Mpc, up to 1024 GPUs

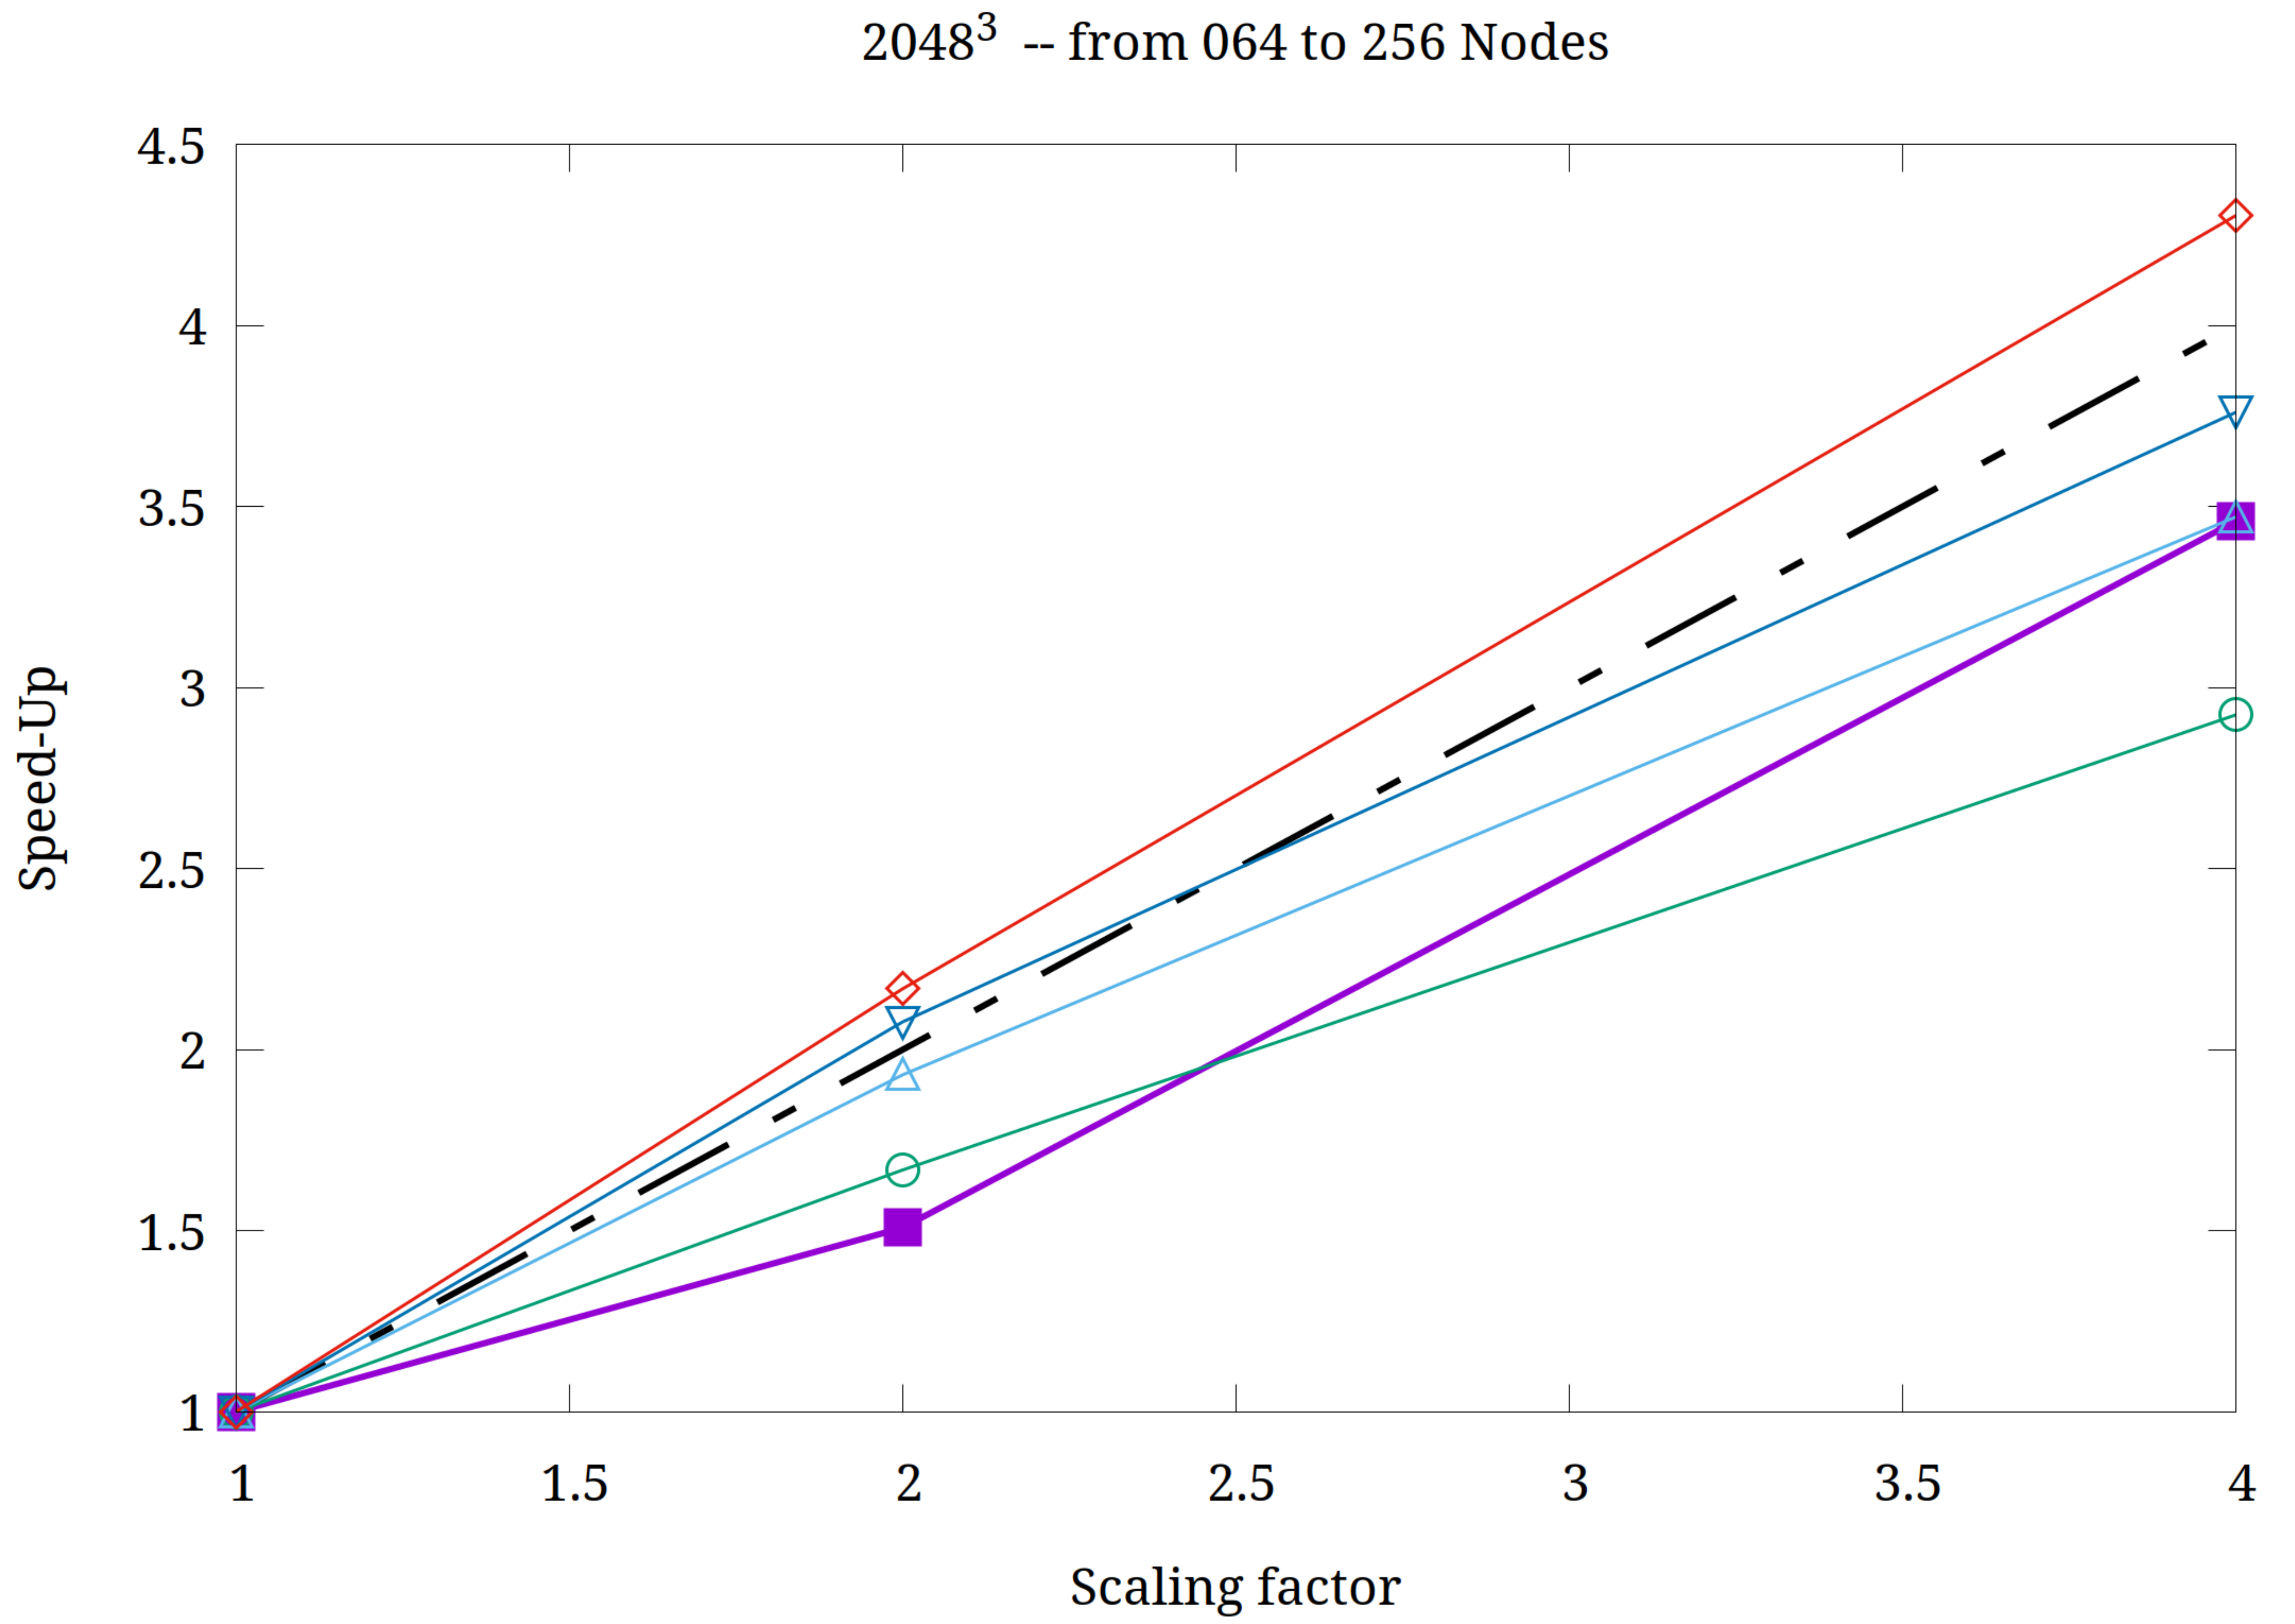

Ongoing: Assessing scalability, targeting performance issues

1) GPU scalability: Efficiency

2×1024³, 120 Mpc, up to **512 GPUs**

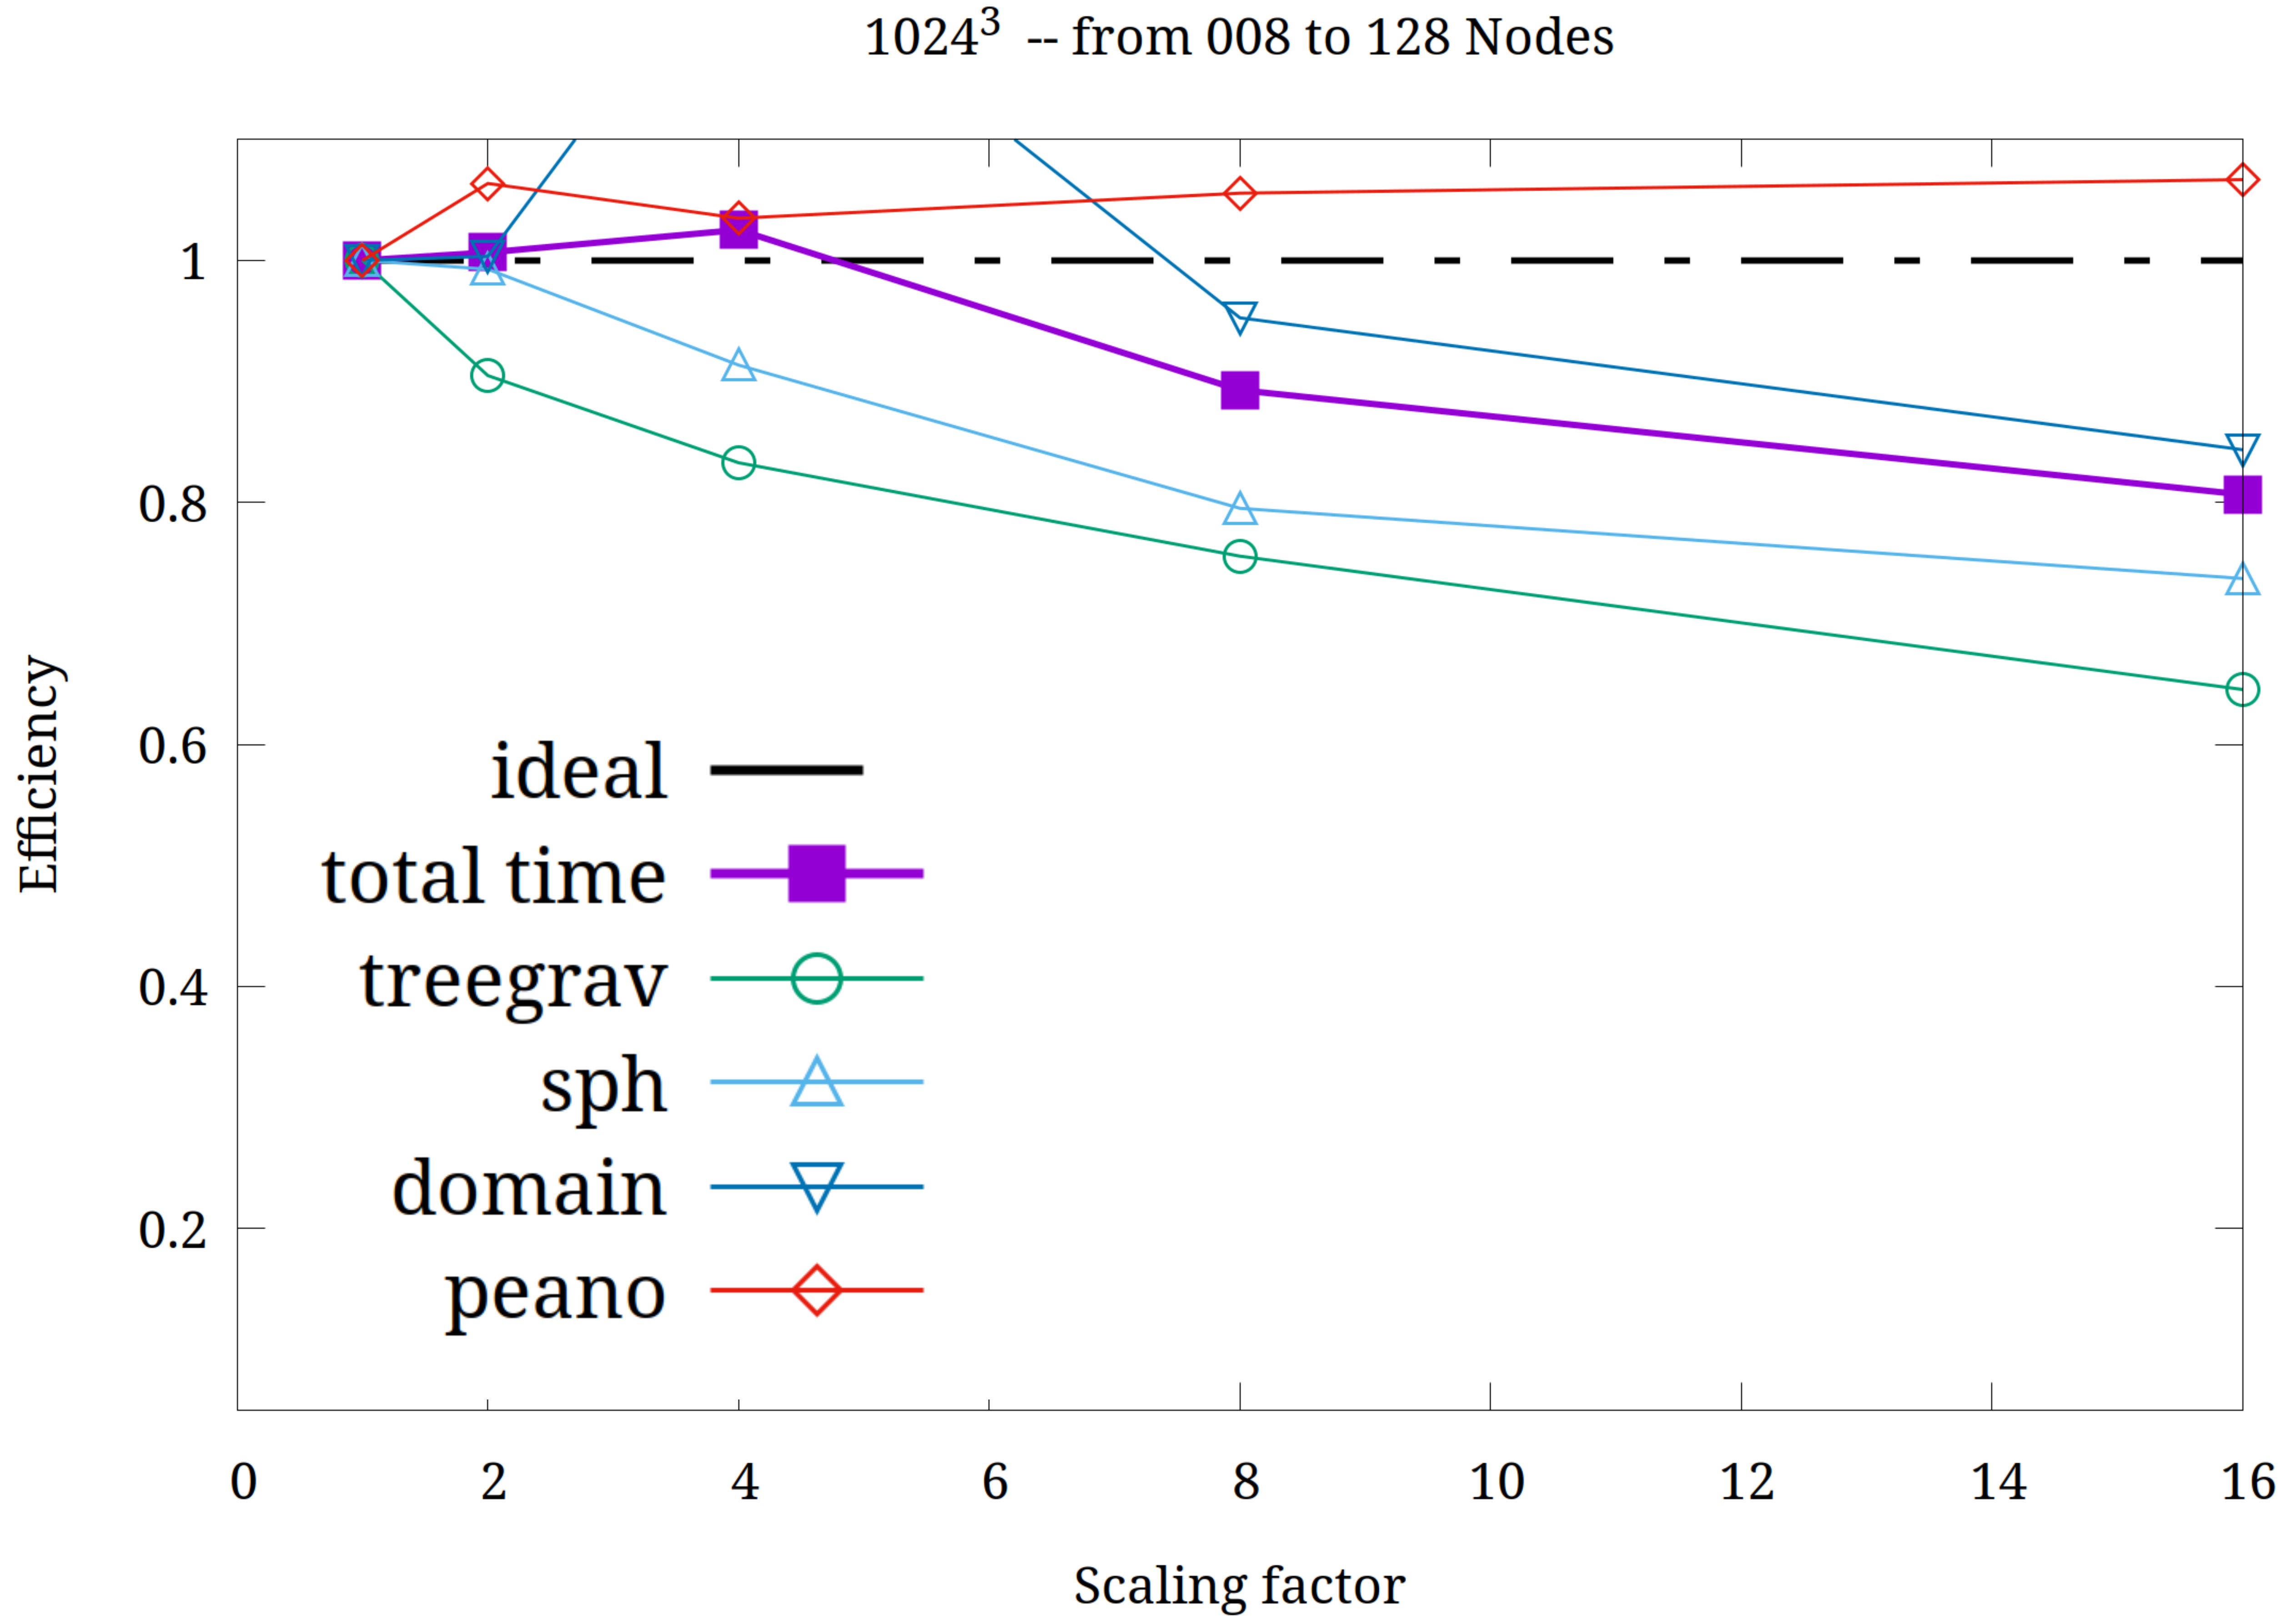

2×2048³, 240 Mpc, up to **1024 GPUs**

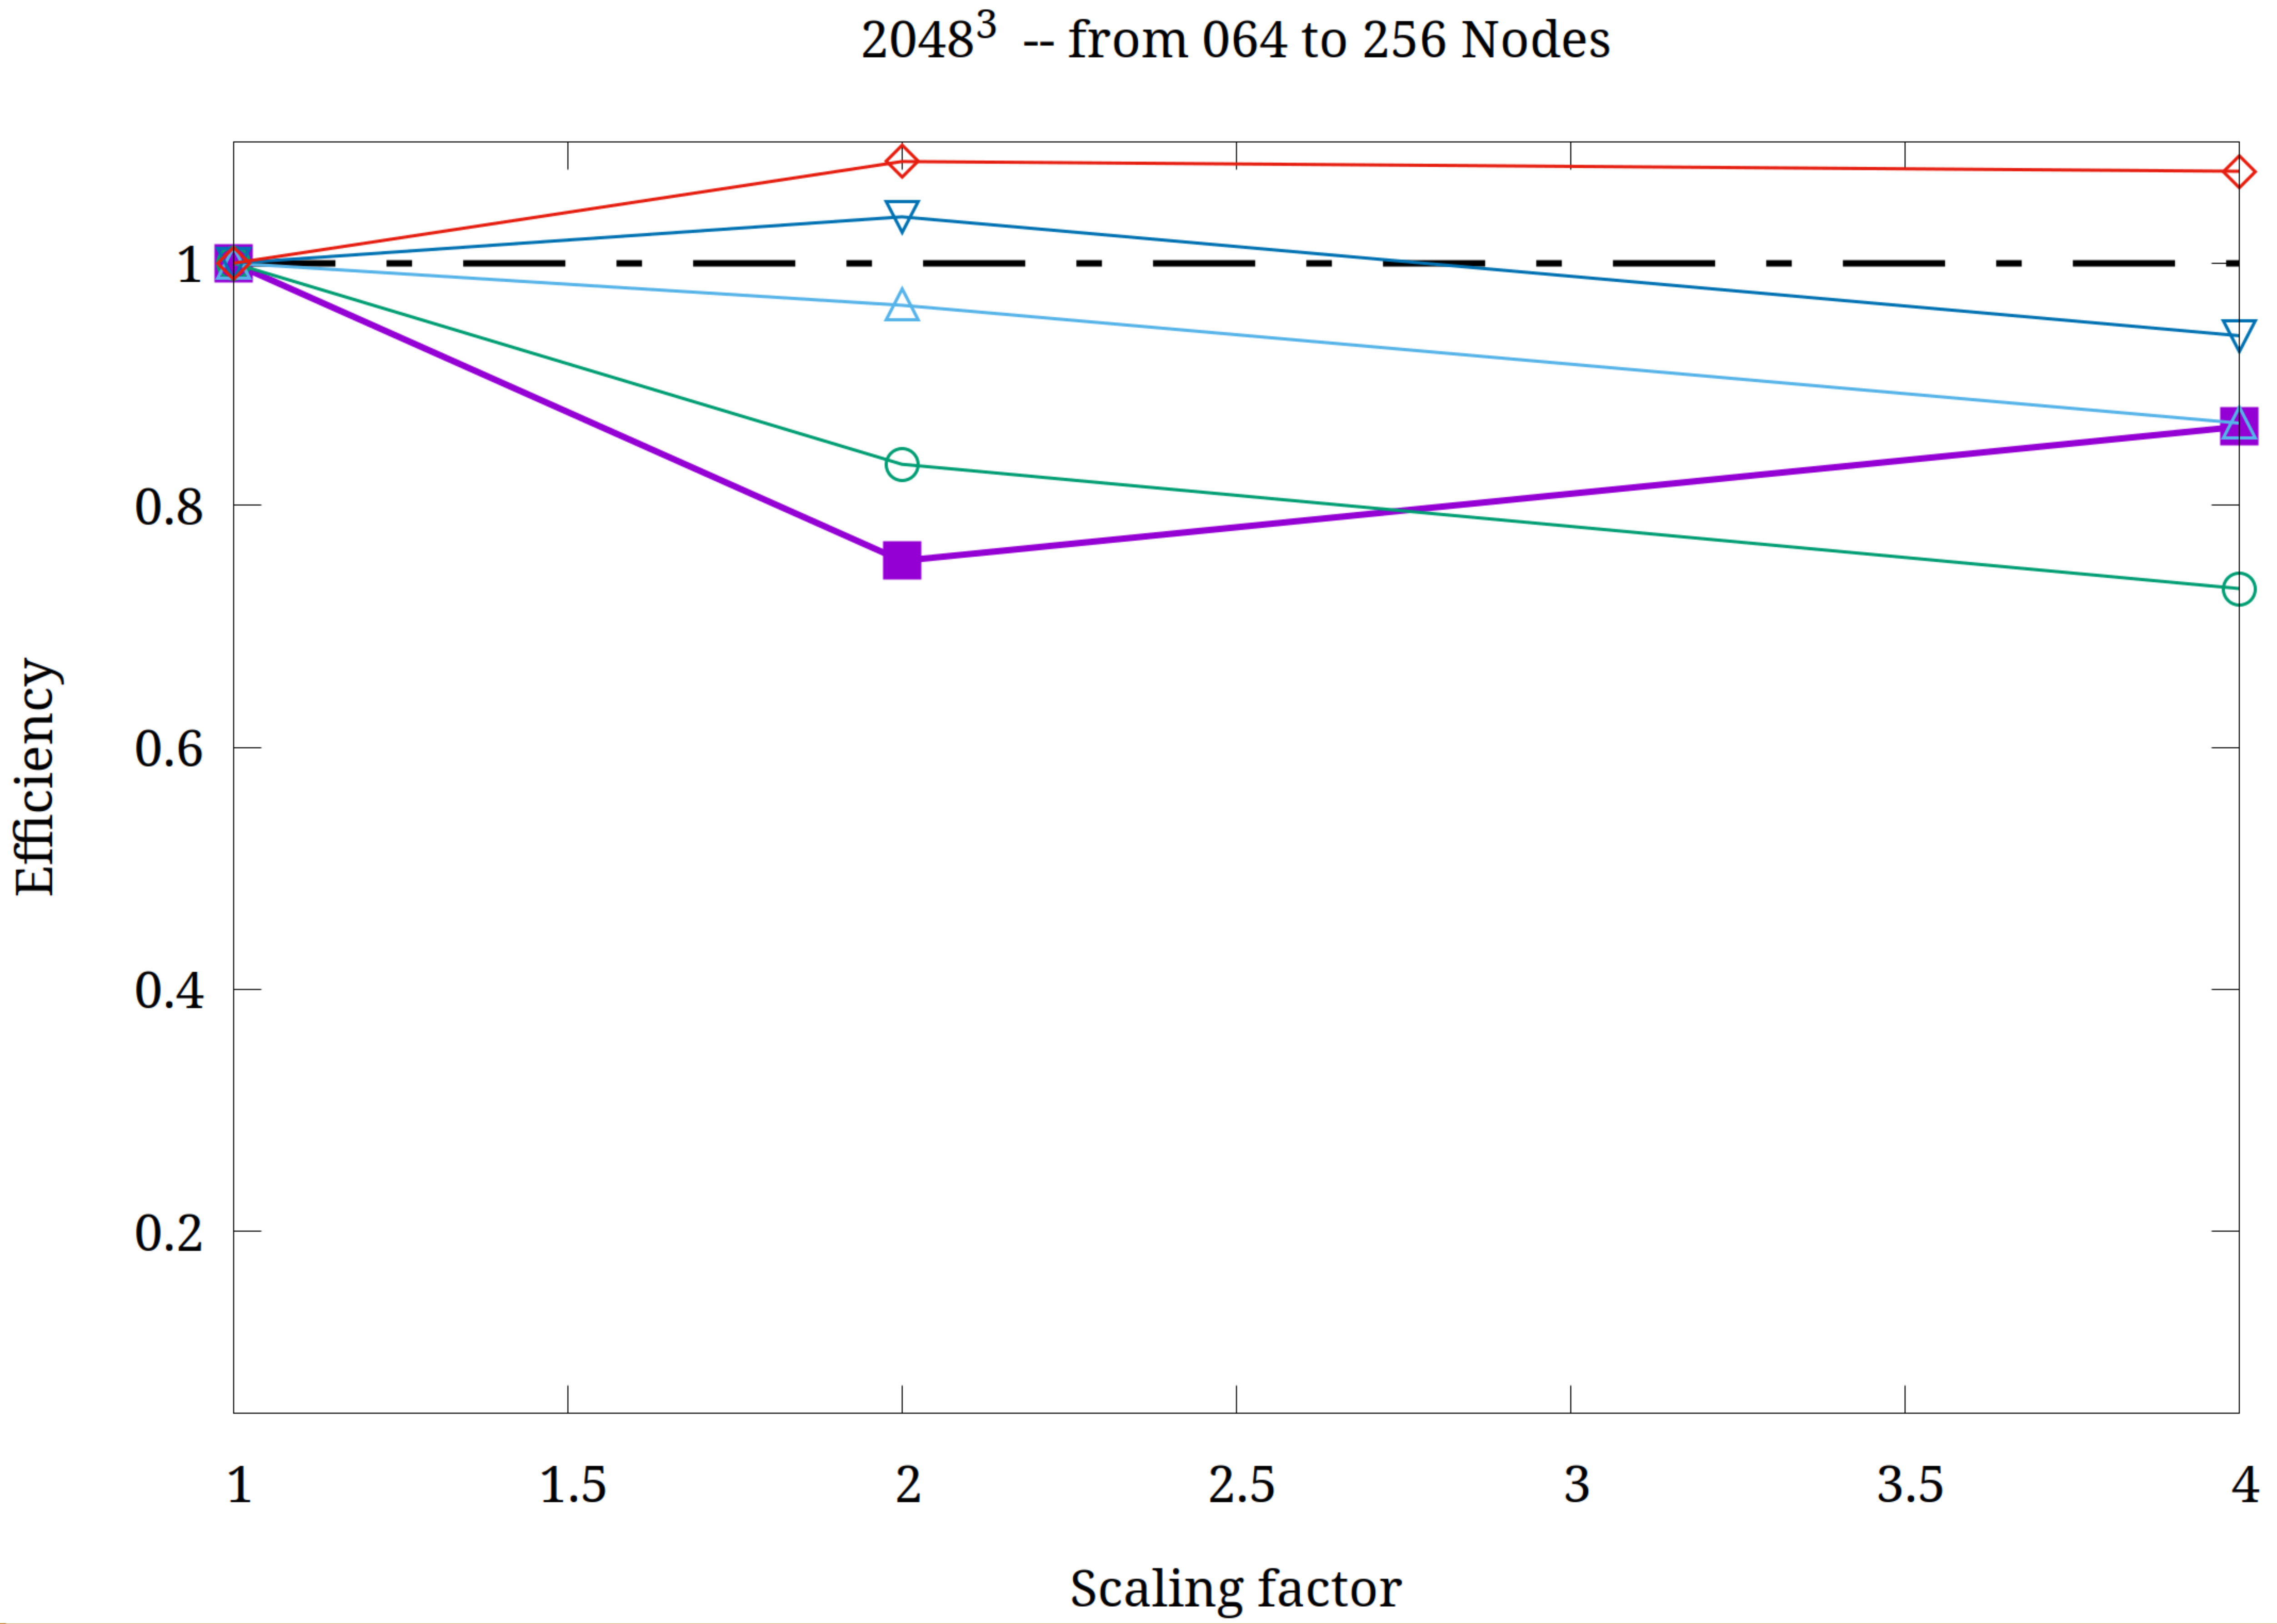

Ongoing: Assessing scalability, targeting performance issues

1) GPU scalability: more in detail

2×1024³, 120 Mpc, up to **512 GPUs**

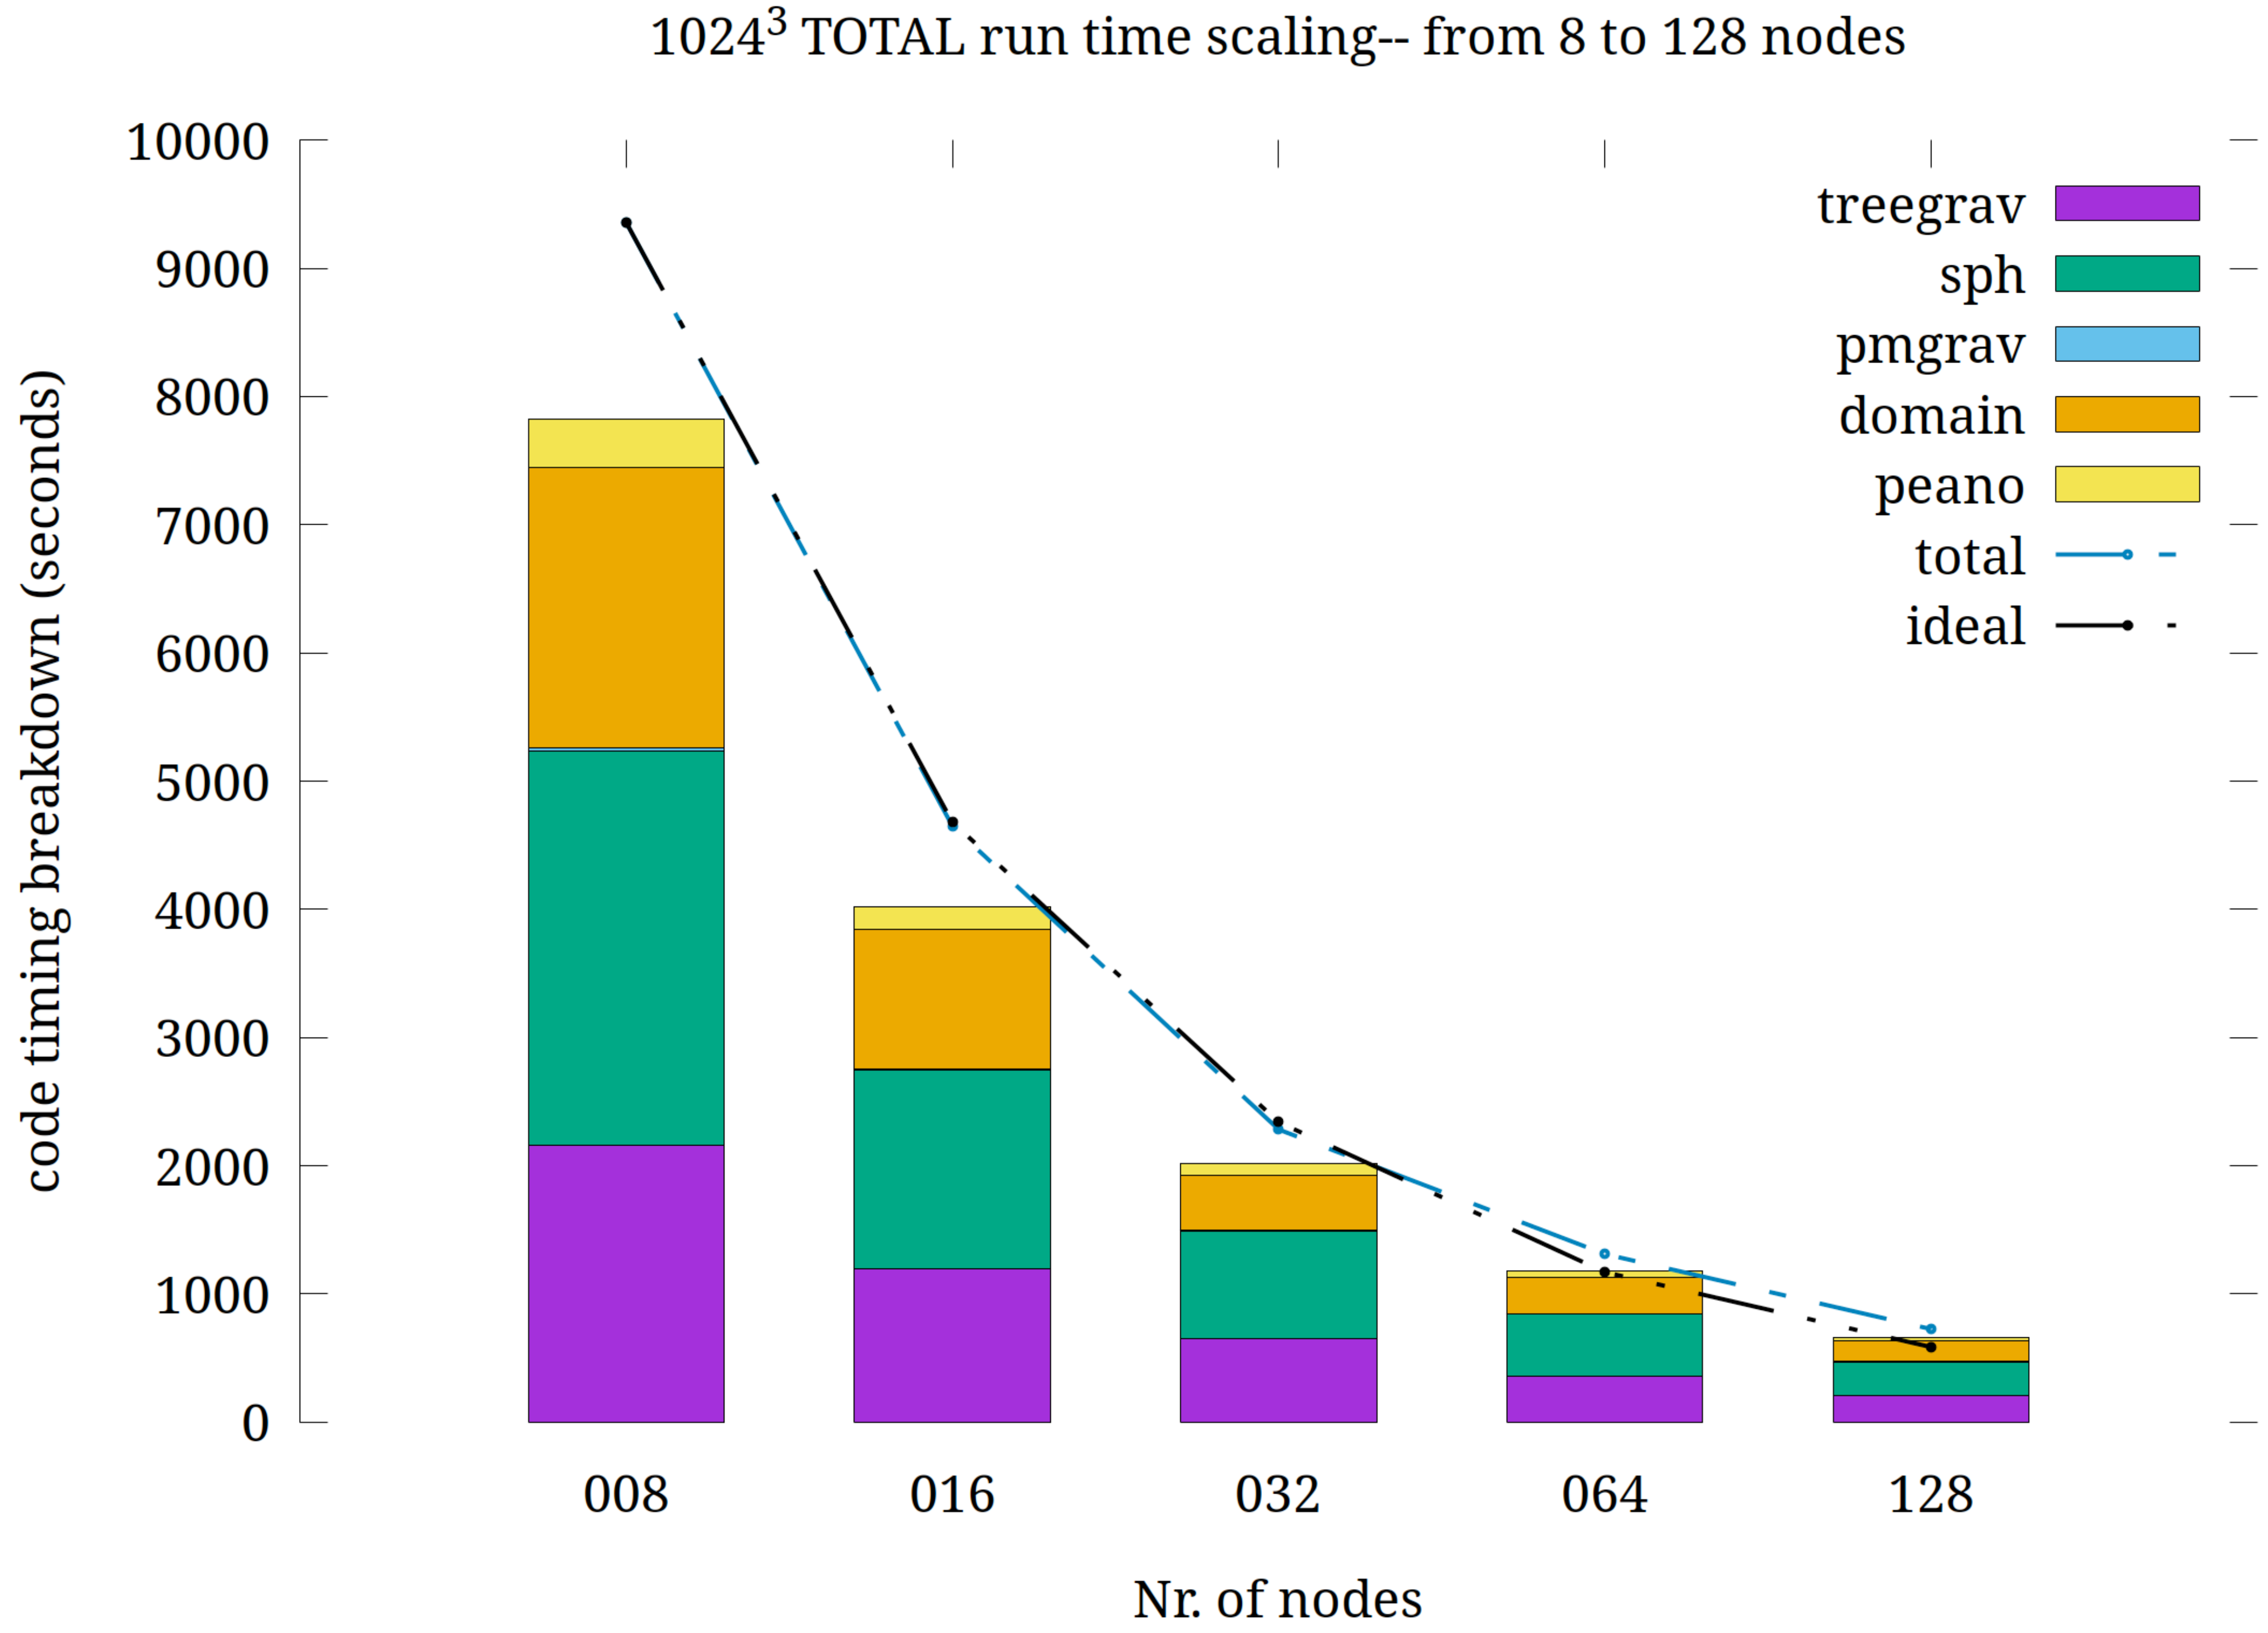

Ongoing: Assessing scalability, targeting performance issues

1) GPU scalability: more in detail

2×1024^3, 120 Mpc, up to **512 GPUs**

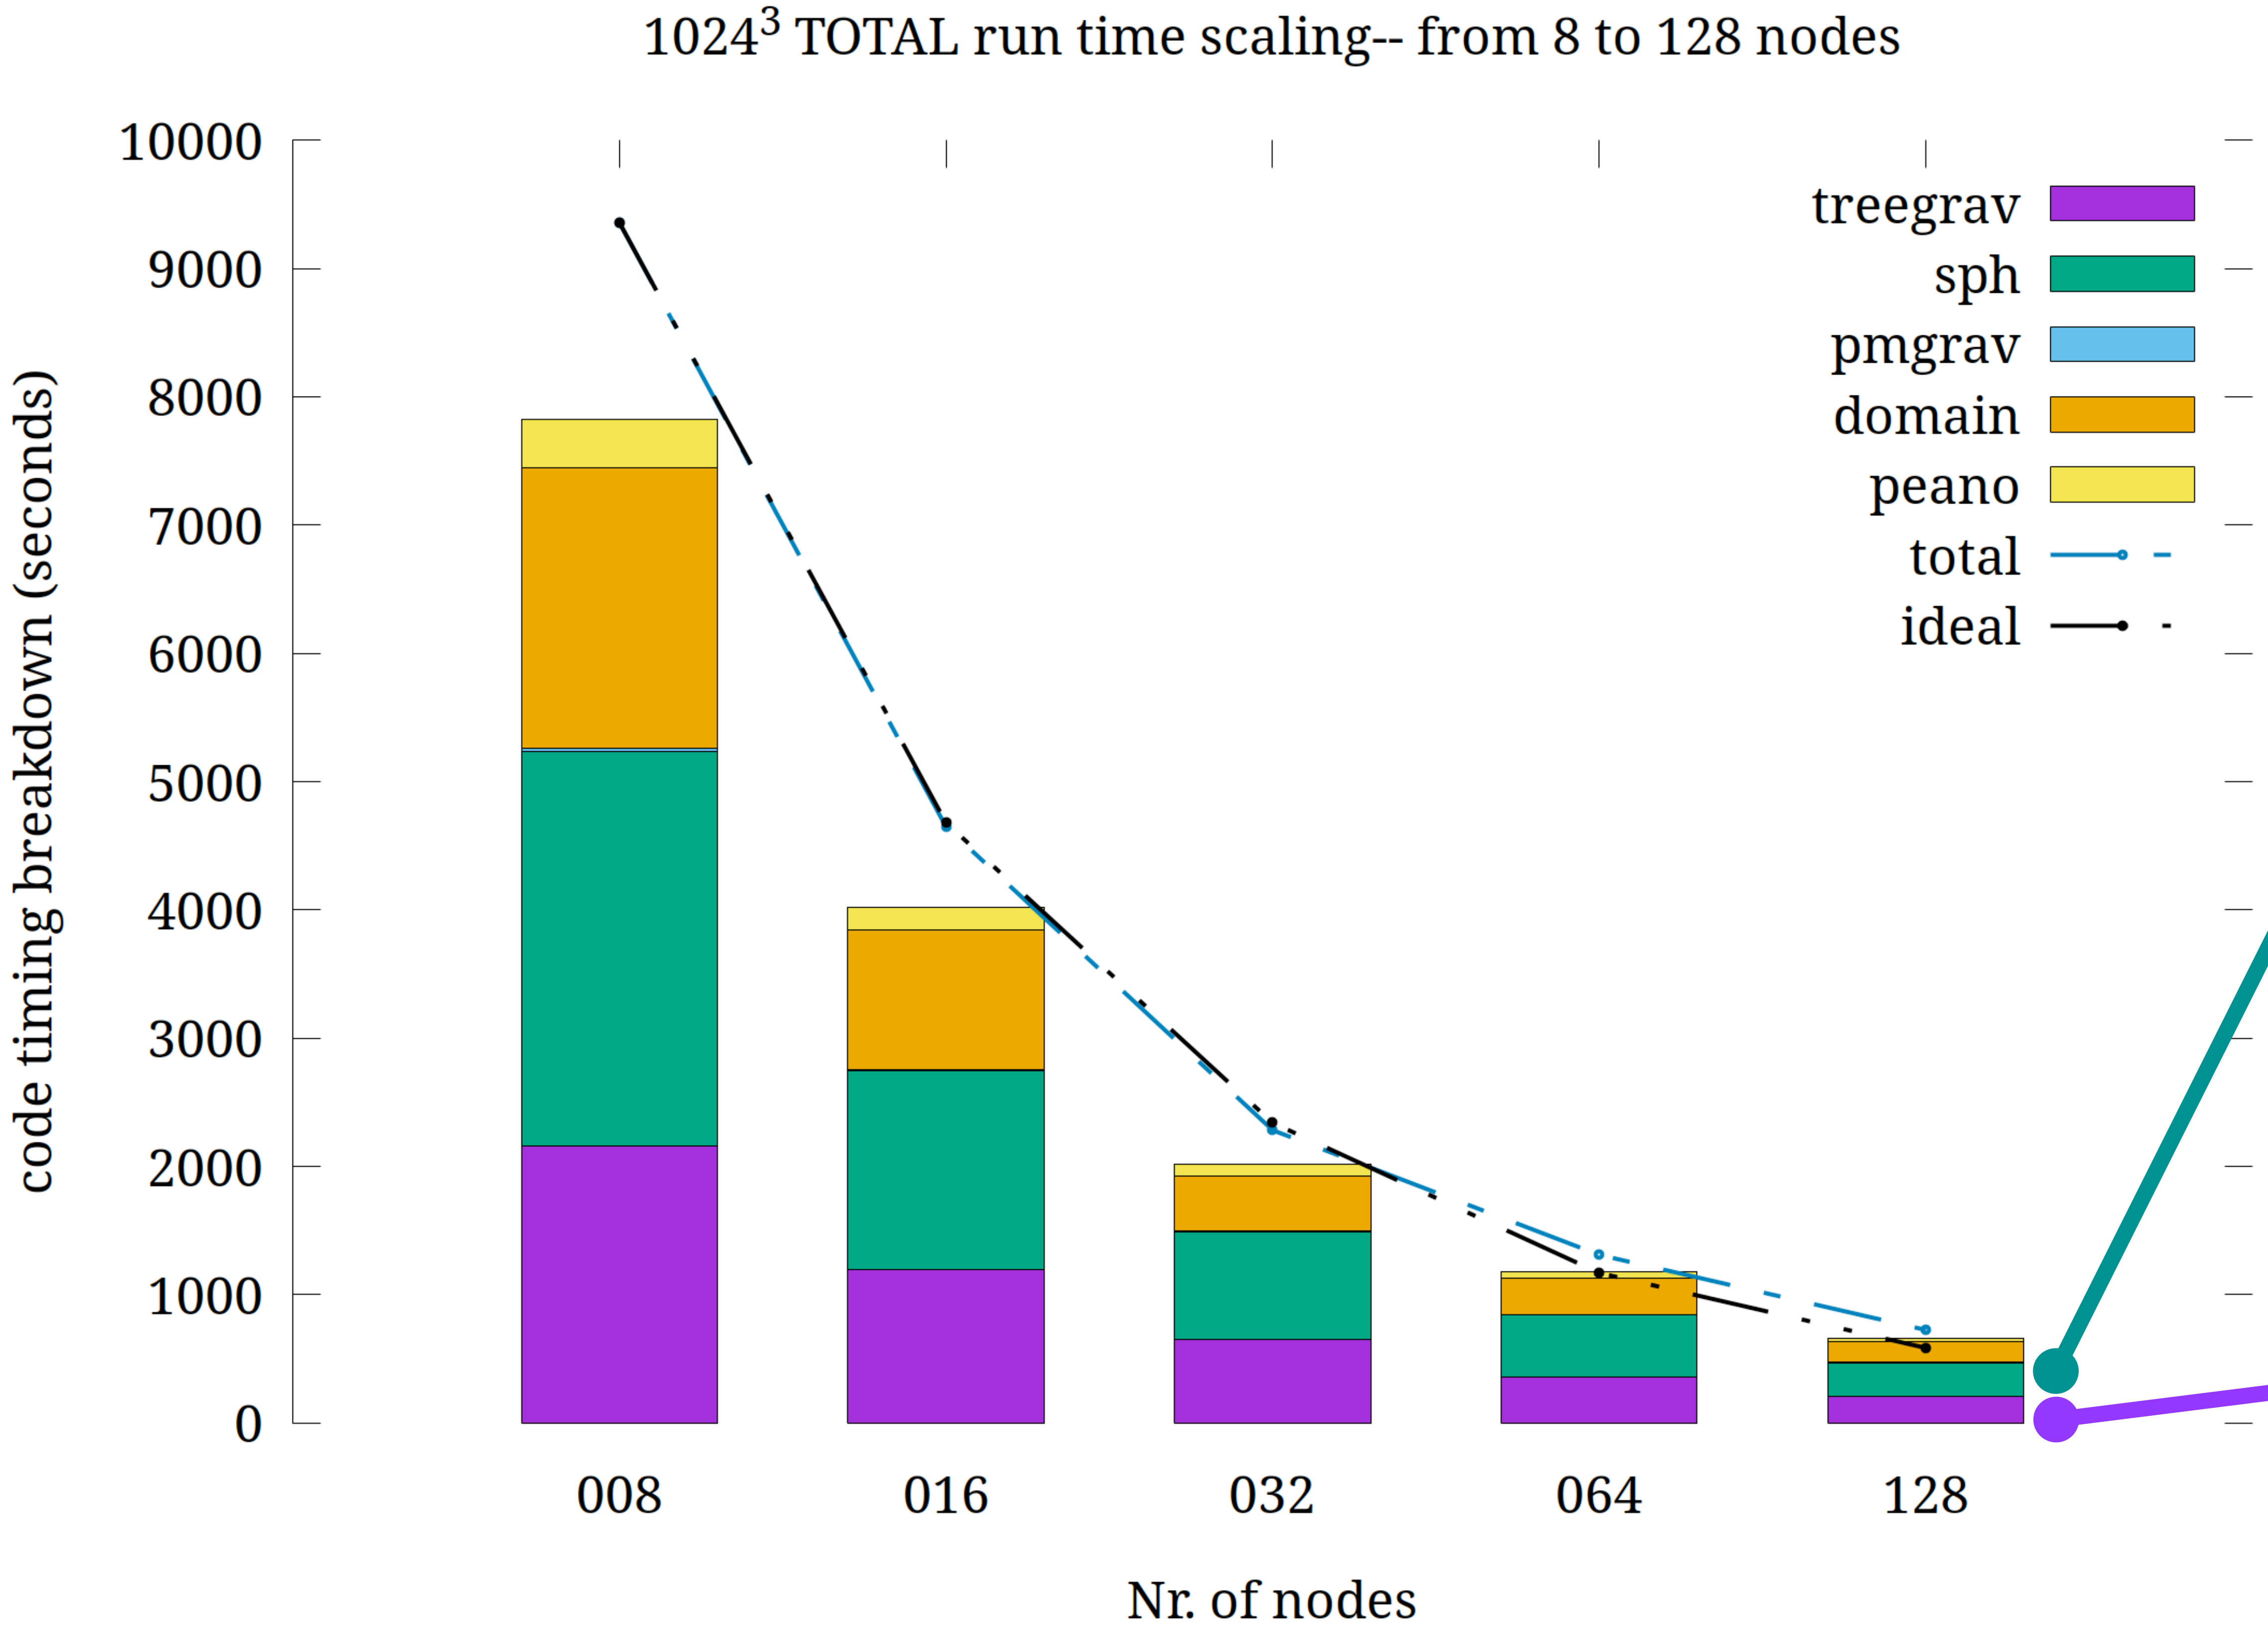

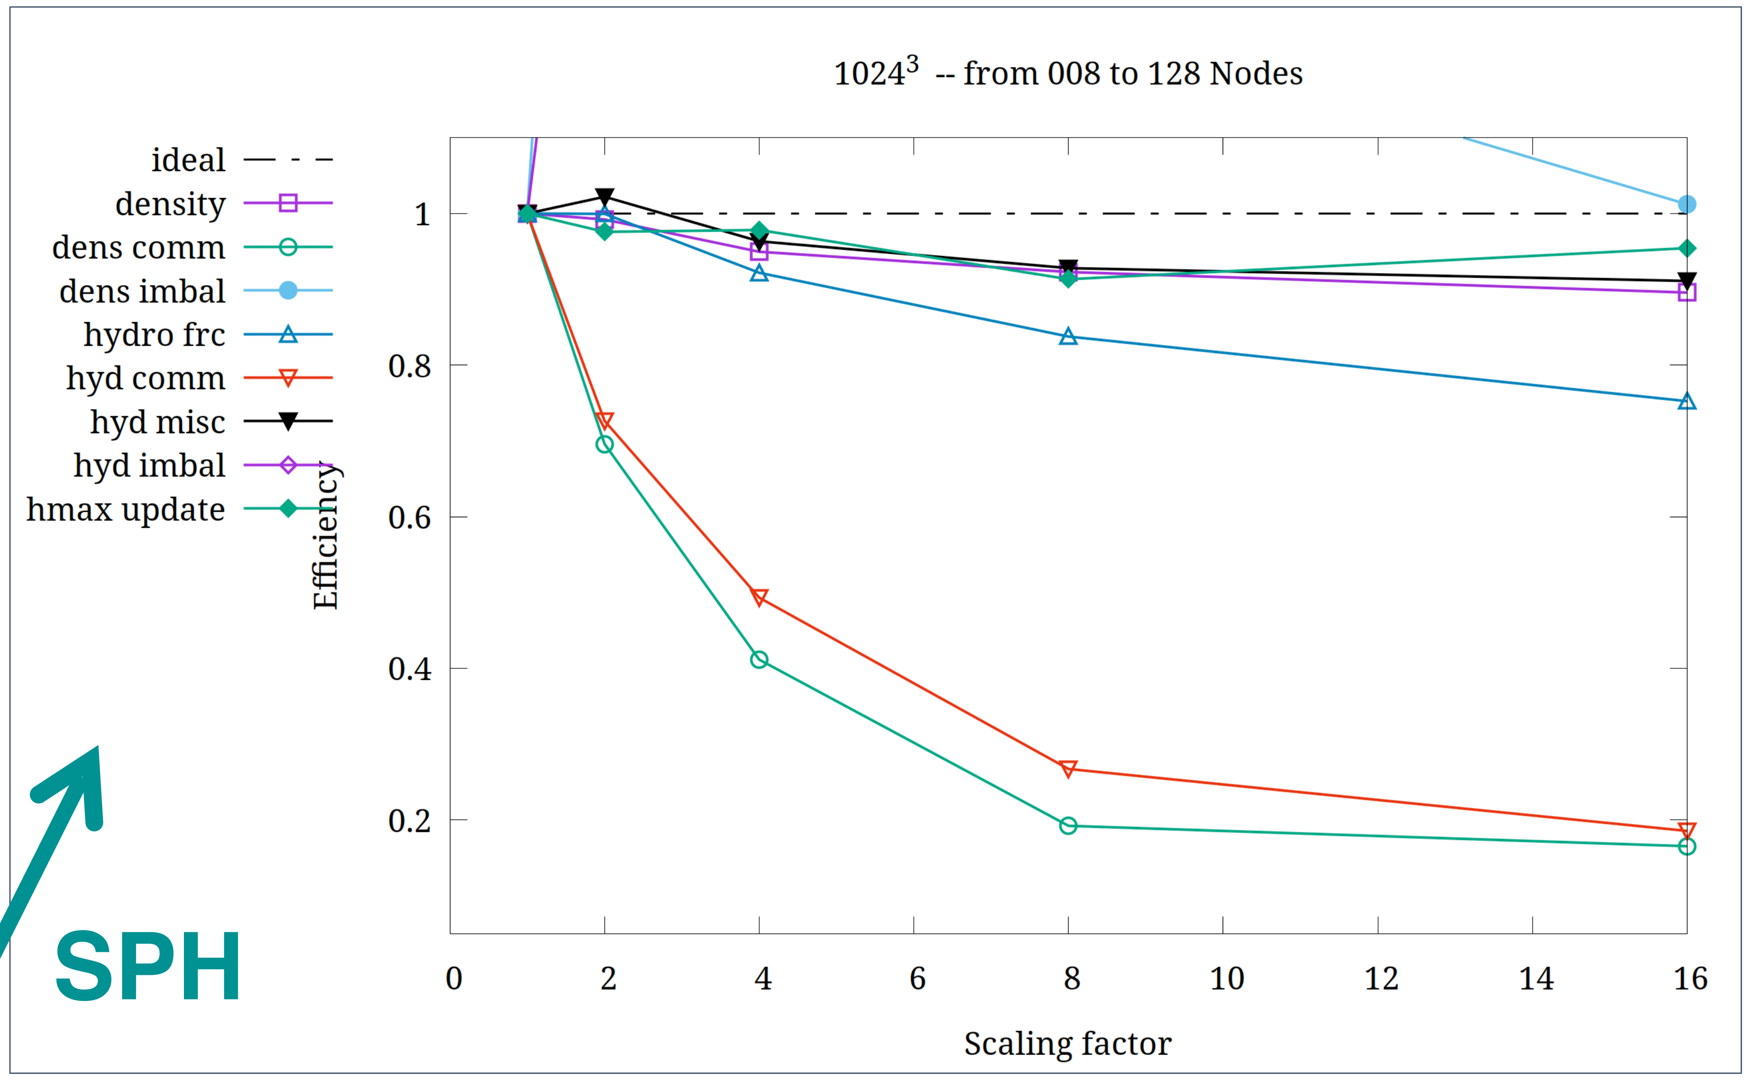

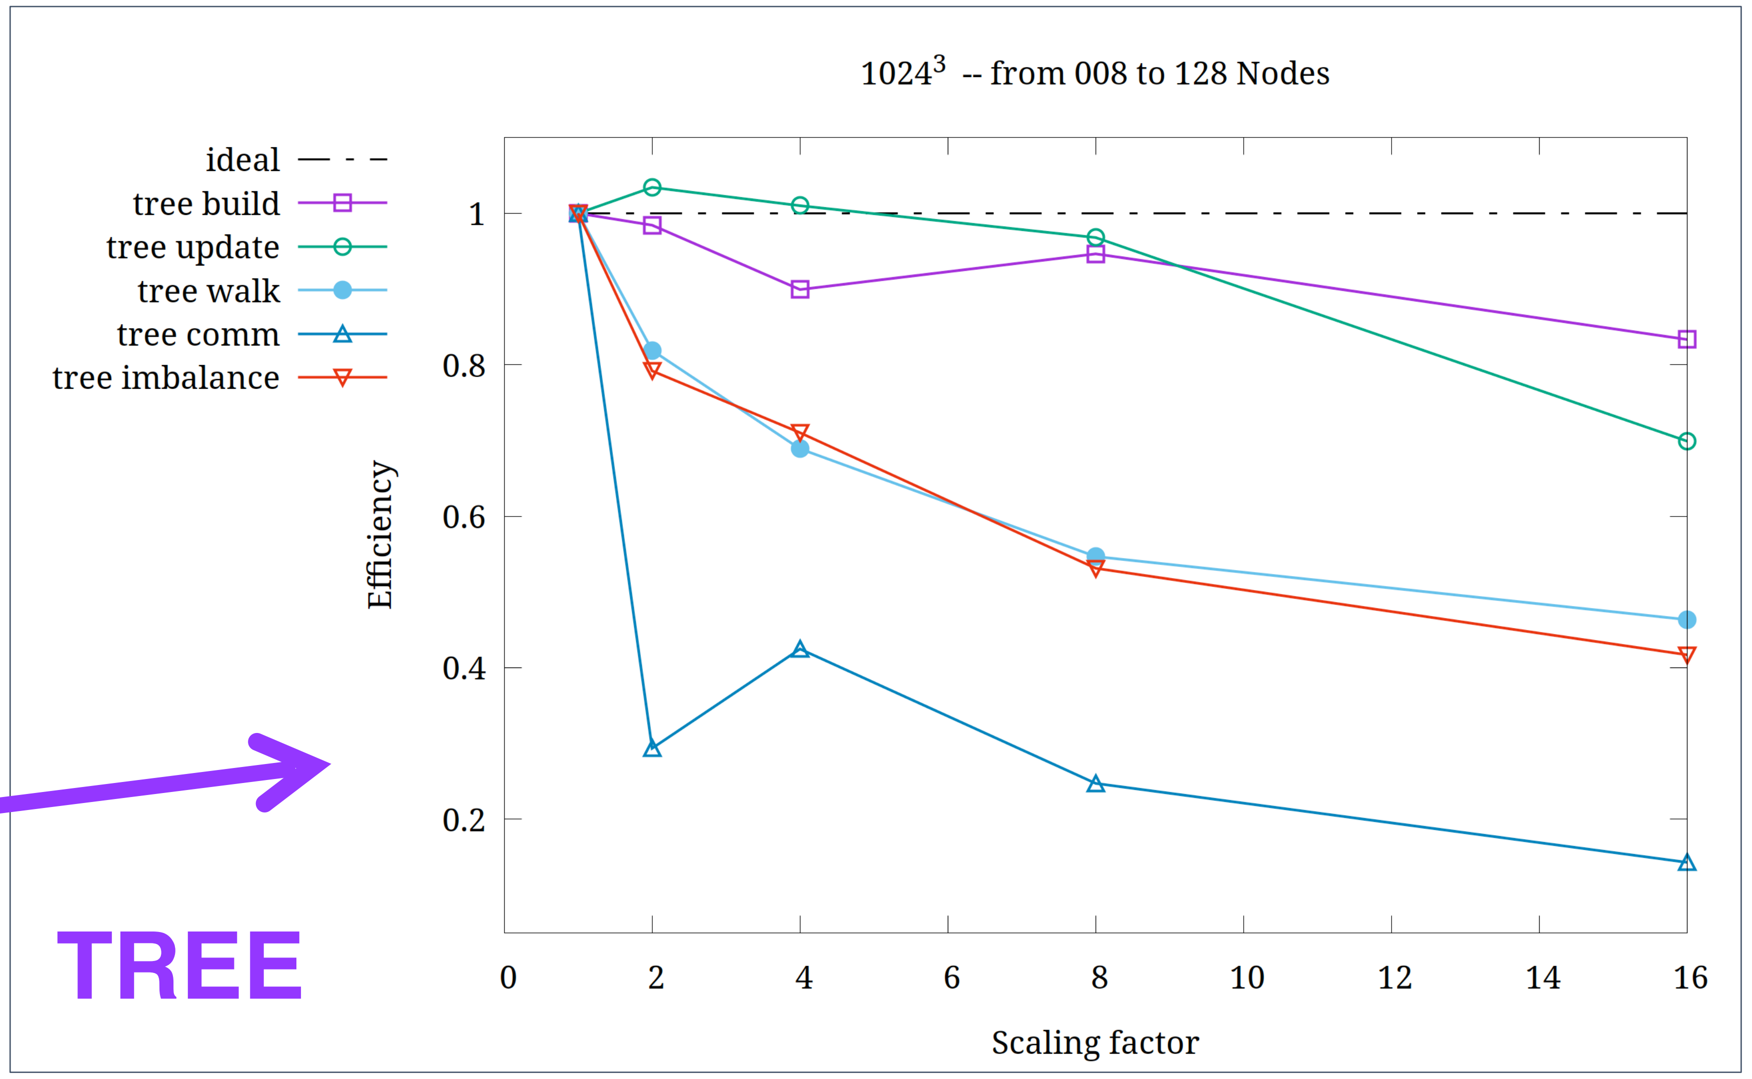

Ongoing: Assessing scalability, targeting performance issues

1) GPU scalability: more in detail

The gravity tree has some noteworthy performance issues, mostly in

- Tree Walk → Barnes&Hut is not GPU-friendly
- Communication

Communication and Nodes update has scalability issues also in the SPH part.

We are trying different Barnes&Hut implementations. Possibly, we may opt to pass from TreePM to P³M when running on GPUs.

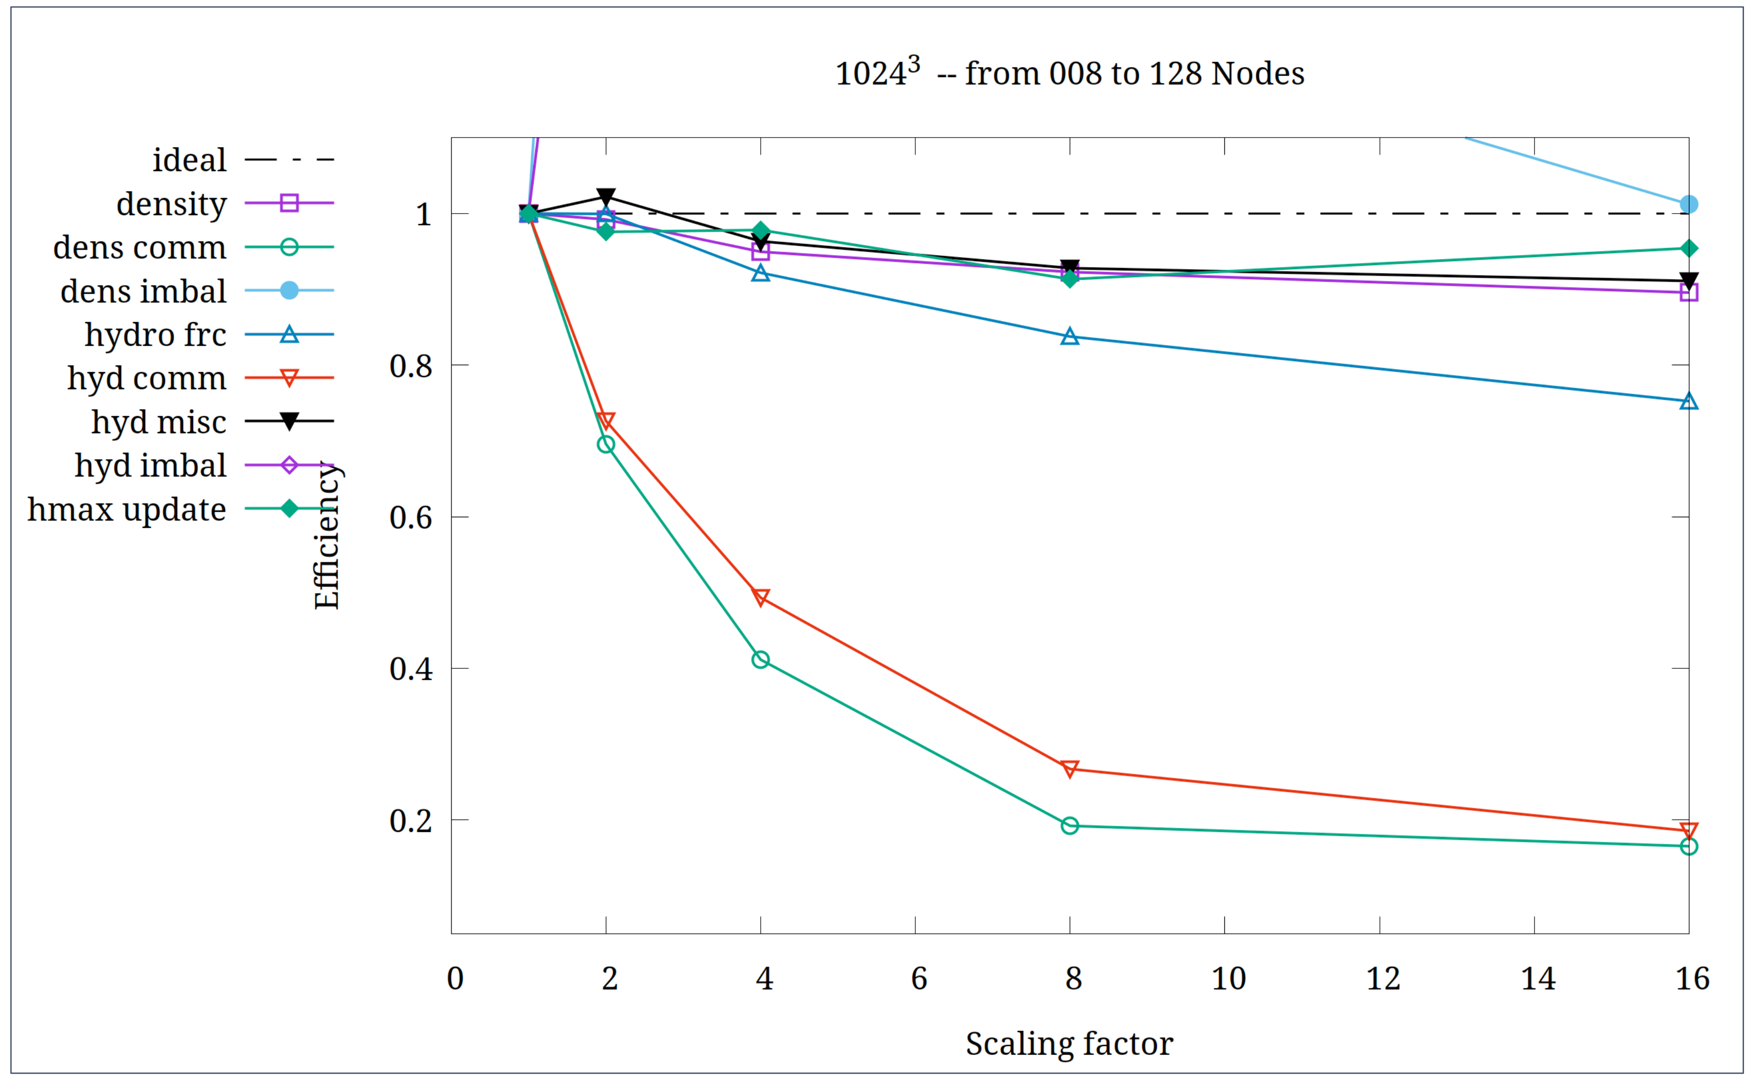

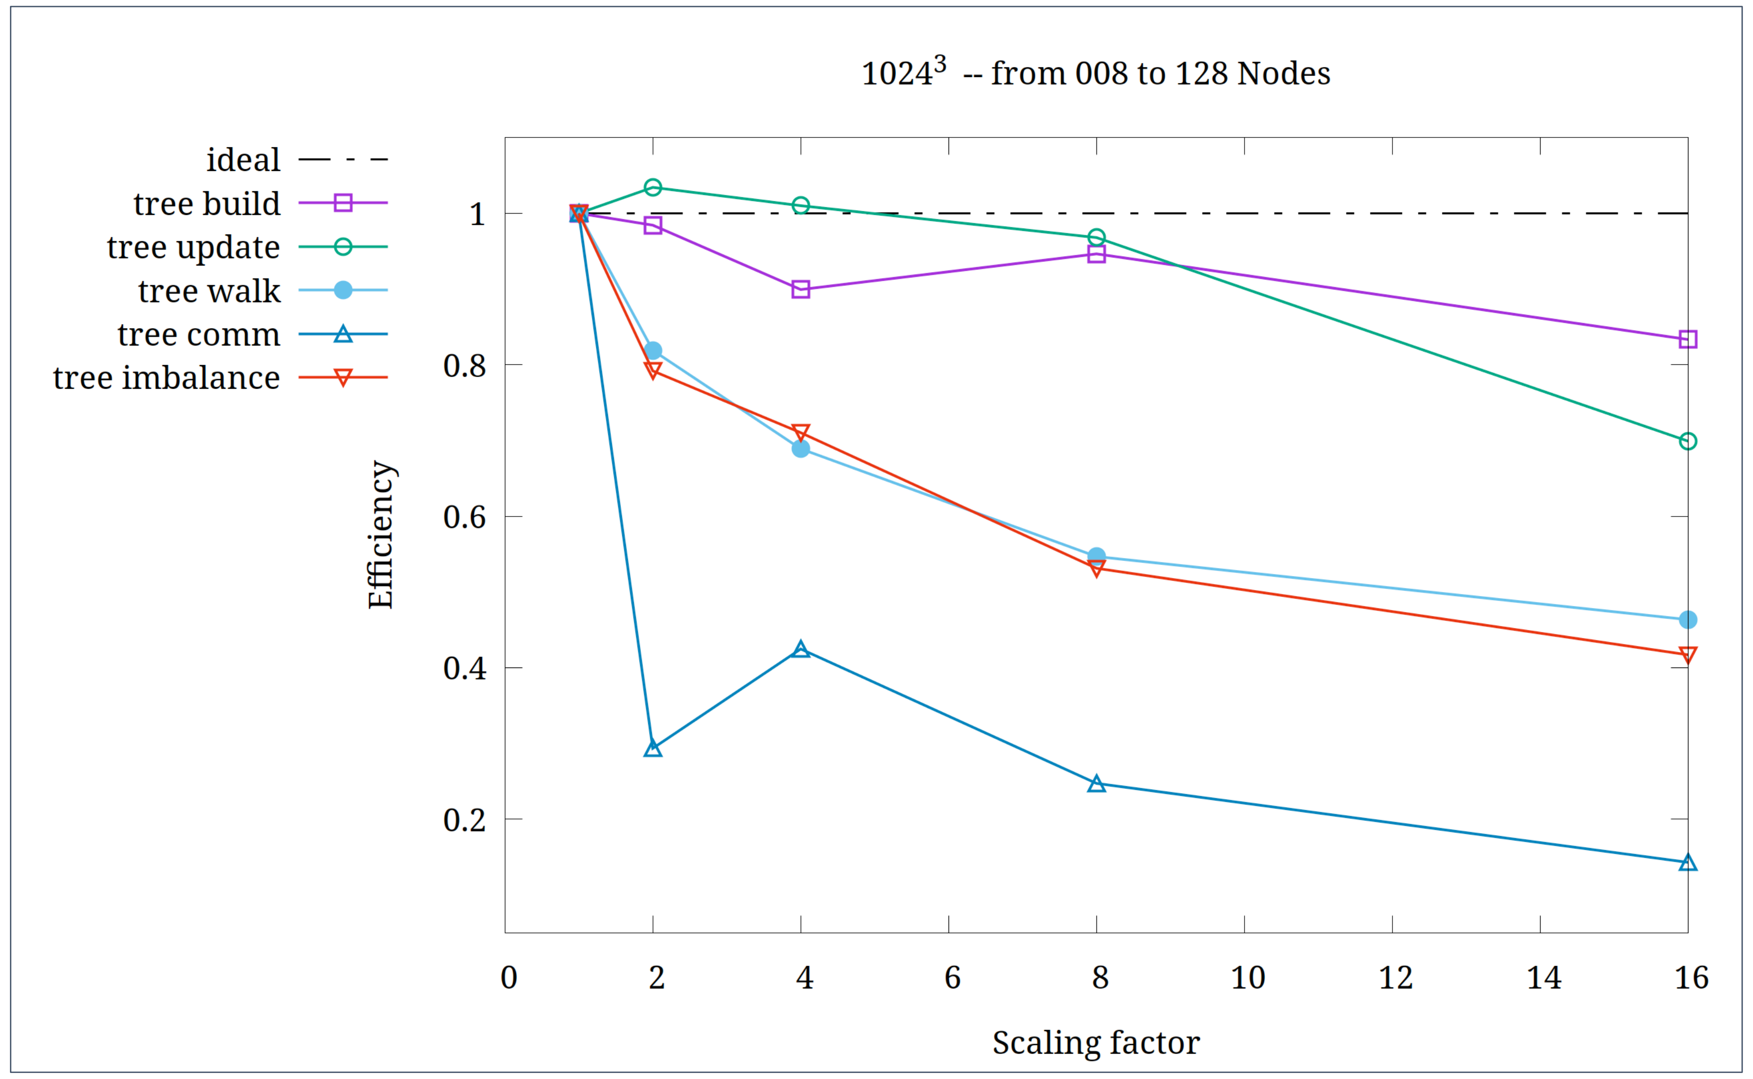

Ongoing: Assessing scalability, targeting performance issues

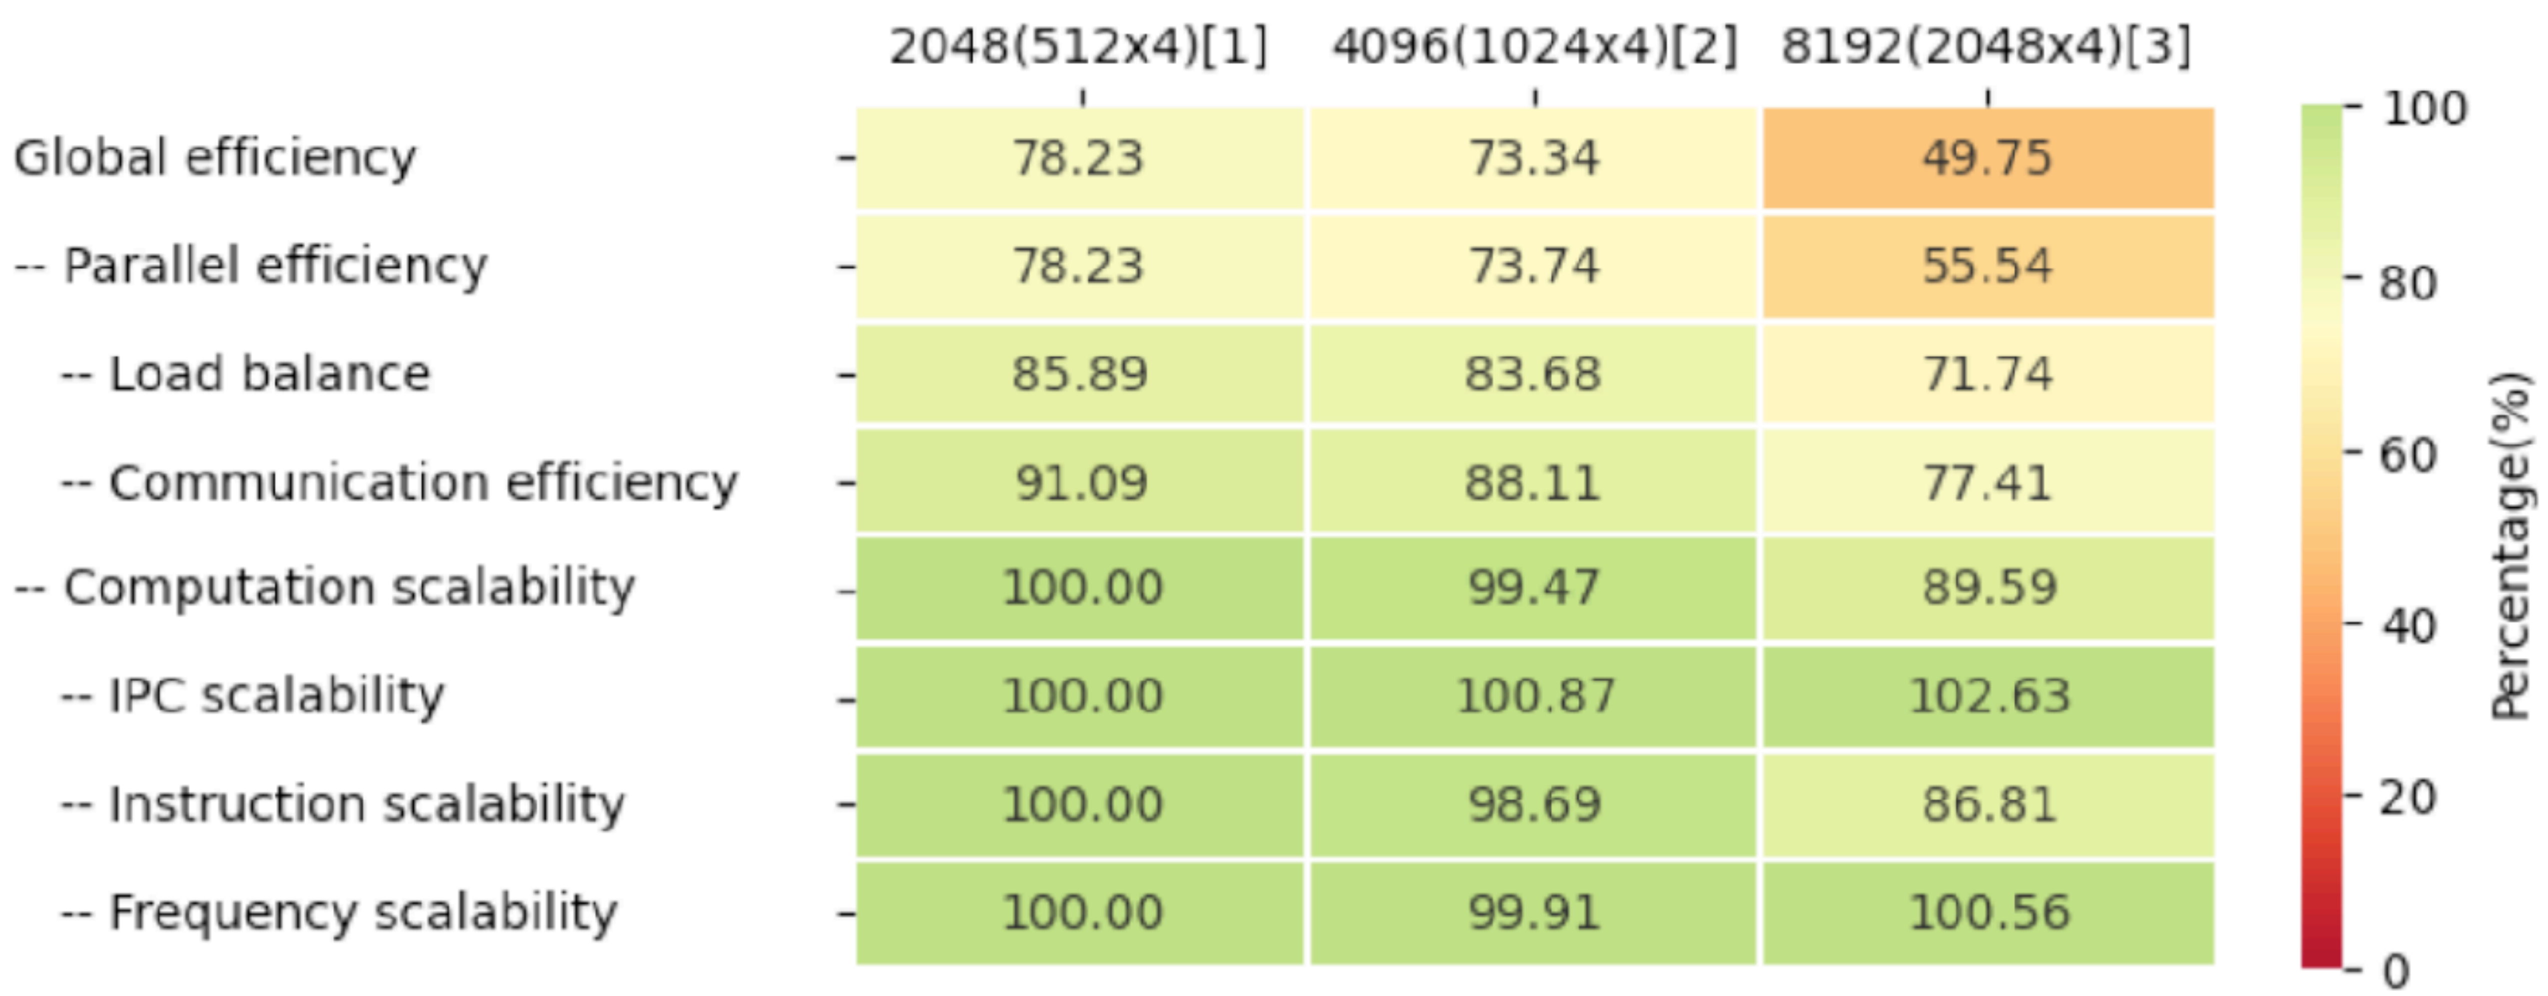

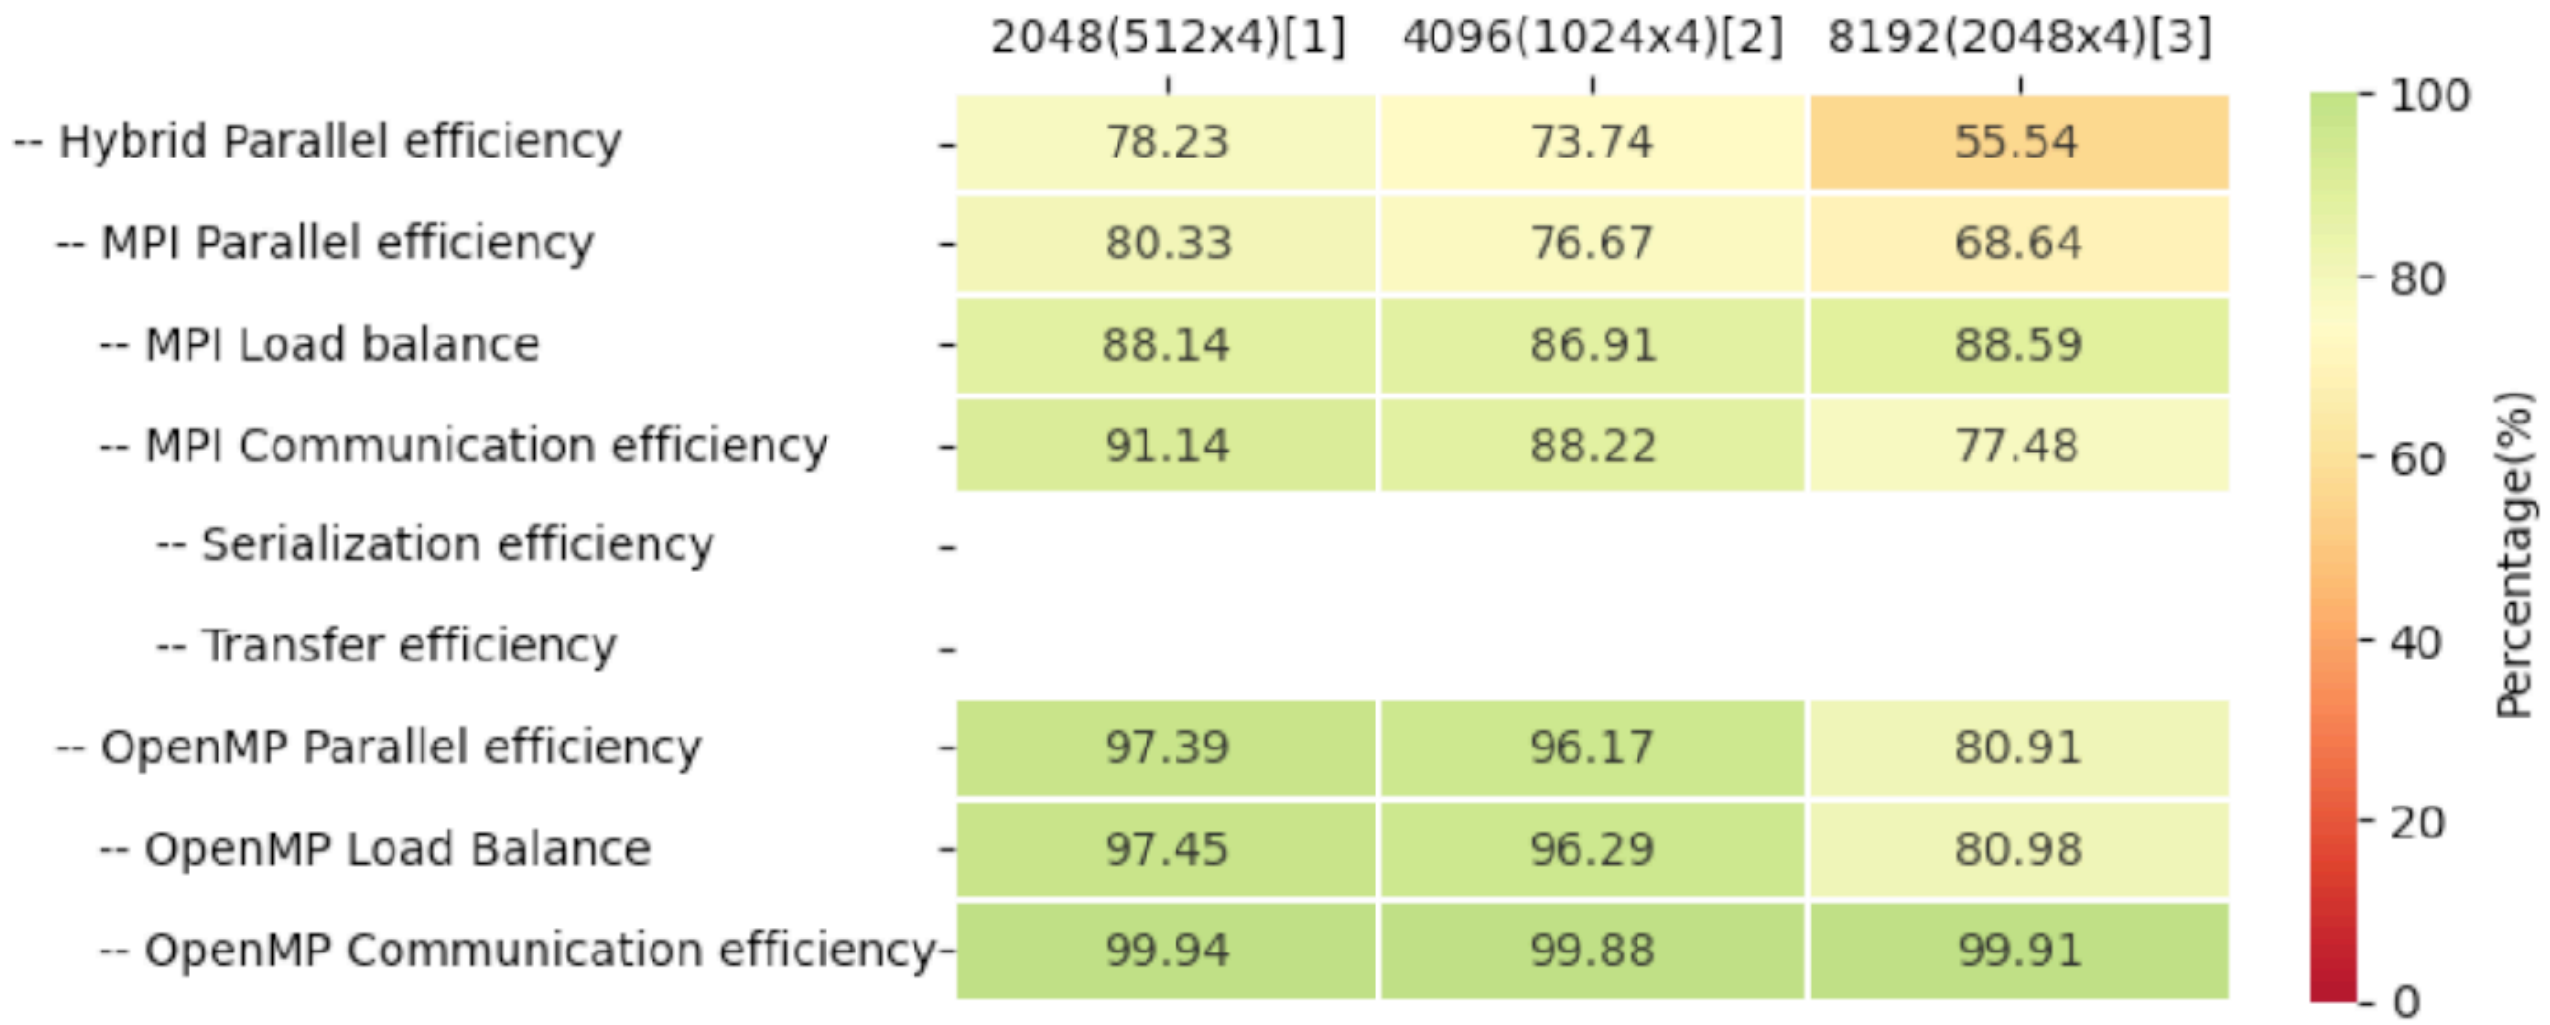

2) Performance issues: vectorization

With the assistance of the POP CoE, and within the SPACE CoE, we are profiling in details the code's behaviour.

The results are summarized in tables, as sketched in the figure on the left
(here the example is for the gravity-tree; rows are different metrics, columns refer to the total number of threads)

from which some key indicators can be collected

Number of processes	2048	4096	8192
Elapsed time (sec)	47.714394	25.446344	18.755917
Efficiency	1.0	0.937549	0.635991
Speedup	1.0	1.875098	2.543965
Average IPC	0.961925	0.970340	0.987195
Average frequency (GHz)	3.190112	3.187294	3.207830

Ongoing: Assessing scalability, targeting performance issues

	2048(512x4)[1]	4096(1024x4)[2]	8192(2048x4)[3]
Global efficiency	78.23	73.34	49.75
-- Parallel efficiency	78.23	73.74	55.54
-- Load balance	85.89	83.68	71.74
-- Communication efficiency	91.09	88.11	77.41
-- Computation scalability	100.00	99.47	89.59
-- IPC scalability	100.00	100.87	102.63
-- Instruction scalability	100.00	98.69	86.81
-- Frequency scalability	100.00	99.91	100.56

(Colour scale: Percentage(%), 0–100)

	2048(512x4)[1]	4096(1024x4)[2]	8192(2048x4)[3]
-- Hybrid Parallel efficiency	78.23	73.74	55.54
-- MPI Parallel efficiency	80.33	76.67	68.64
-- MPI Load balance	88.14	86.91	88.59
-- MPI Communication efficiency	91.14	88.22	77.48
-- Serialization efficiency			
-- Transfer efficiency			
-- OpenMP Parallel efficiency	97.39	96.17	80.91
-- OpenMP Load Balance	97.45	96.29	80.98
-- OpenMP Communication efficiency	99.94	99.88	99.91

(Colour scale: Percentage(%), 0–100)

2) Performance issues: vectorization

The low IPC (Instructions Per Cycle), although constant with decreasing workload, indicates that the computational efficiency is not high.
Further inspection returned that in particular the **vectorization ratio is very small** (~10%) and limited to 128bits registers

→ **the main target is to re-formulate the data structures that now consists in Arrays of (large)Structures**

Number of processes	2048	4096	8192
Elapsed time (sec)	47.714394	25.446344	18.755917
Efficiency	1.0	0.937549	0.635991
Speedup	1.0	1.875098	2.543965
Average IPC	0.961925	0.970340	0.987195
Average frequency (GHz)	3.190112	3.187294	3.207830

Next Steps and Expected Results

Key Science Projects

1. → **EAGER: Evolution of gAlaxies and Galaxy clustErs in high-Resolution cosmological simulations**

Stefano Borgani, Milena Valentini, Luca Tornatore, Alice Damiano, Alex Saro, Giuliano Taffoni, Tiago Castro

2. → **SLOTH: Shedding Light On dark matter wiTH cosmological simulations**

Milena Valentini, Stefano Borgani, Tiago Castro, Luca Tornatore, Matteo Viel, Alice Damiano, Pierluigi Monaco, Giuliano Taffoni

Next Steps and Expected Results

1. → EAGER: Evolution of gAlaxies and Galaxy clustErs in high-Resolution cosmological simulations

Stefano Borgani, Milena Valentini, Luca Tornatore, Alice Damiano, Alex Saro, Giuliano Taffoni, Tiago Castro

Main **plans** of the project:

- Suite of cosmological hydrodynamical simulations of **galaxy clusters** to investigate structure formation in high-density regions and the joint evolution of galaxies and their IGM within the extreme cluster environment
- **Cosmological volume**(s) for statistical studies of the properties of evolving galaxies in field environment
- Simulated boxes containing galaxies, galaxy groups and poor clusters will allow us to bridge between the **mass scale** of massive galaxies and galaxy clusters
- Investigate the connection between **super-massive BHs and host galaxies**, and the large-scale environment

Next Steps and Expected Results

Our sub-resolution model MUPPI **successfully** describes the evolution of discs **and** massive ellipticals

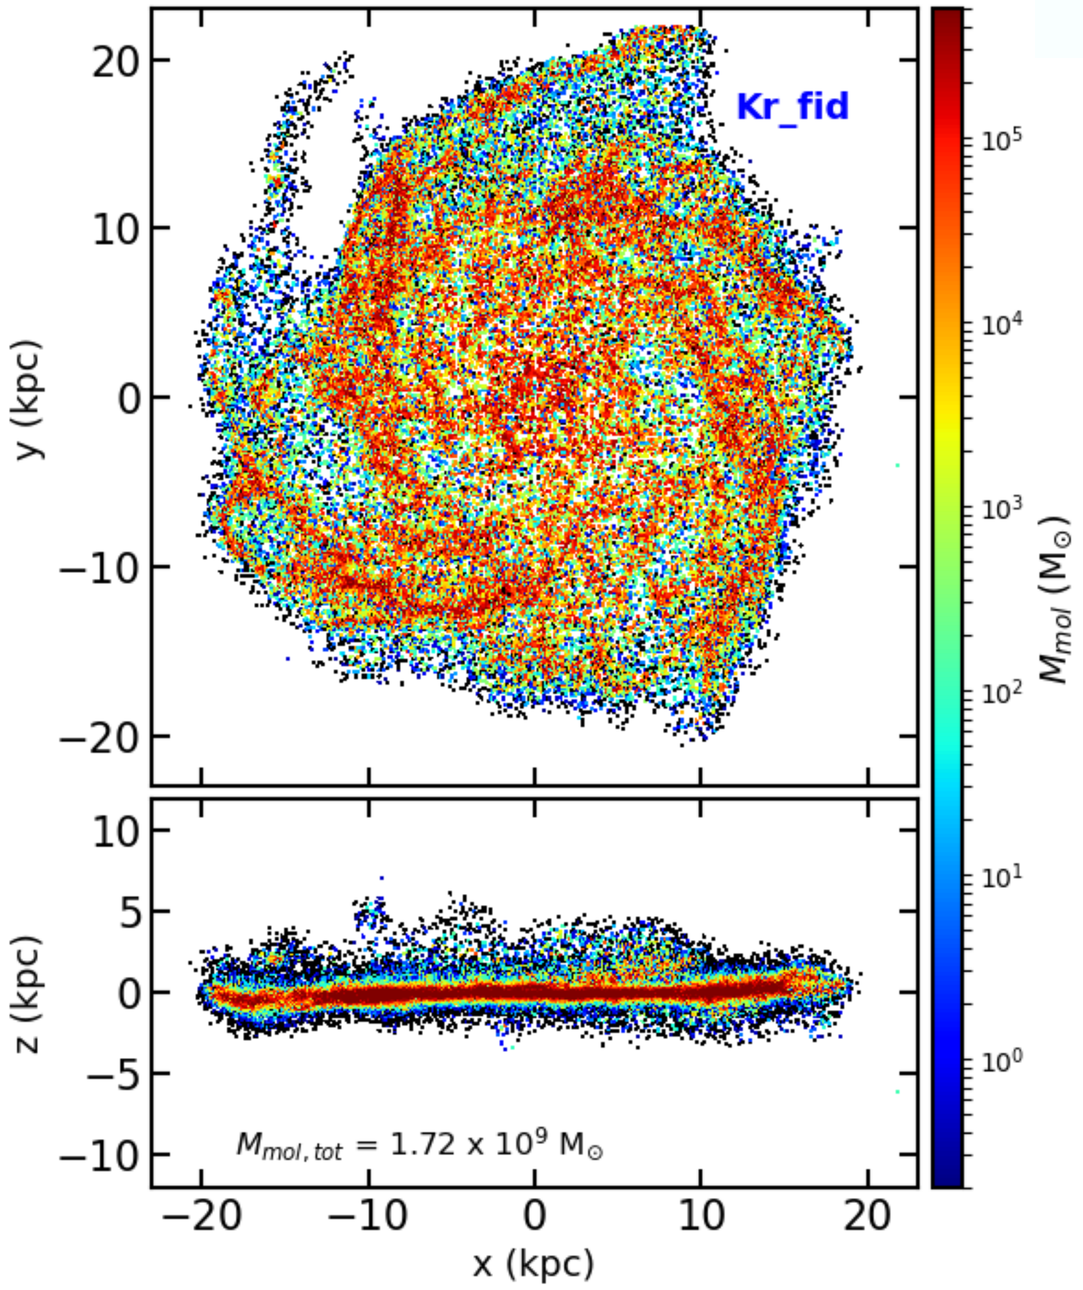

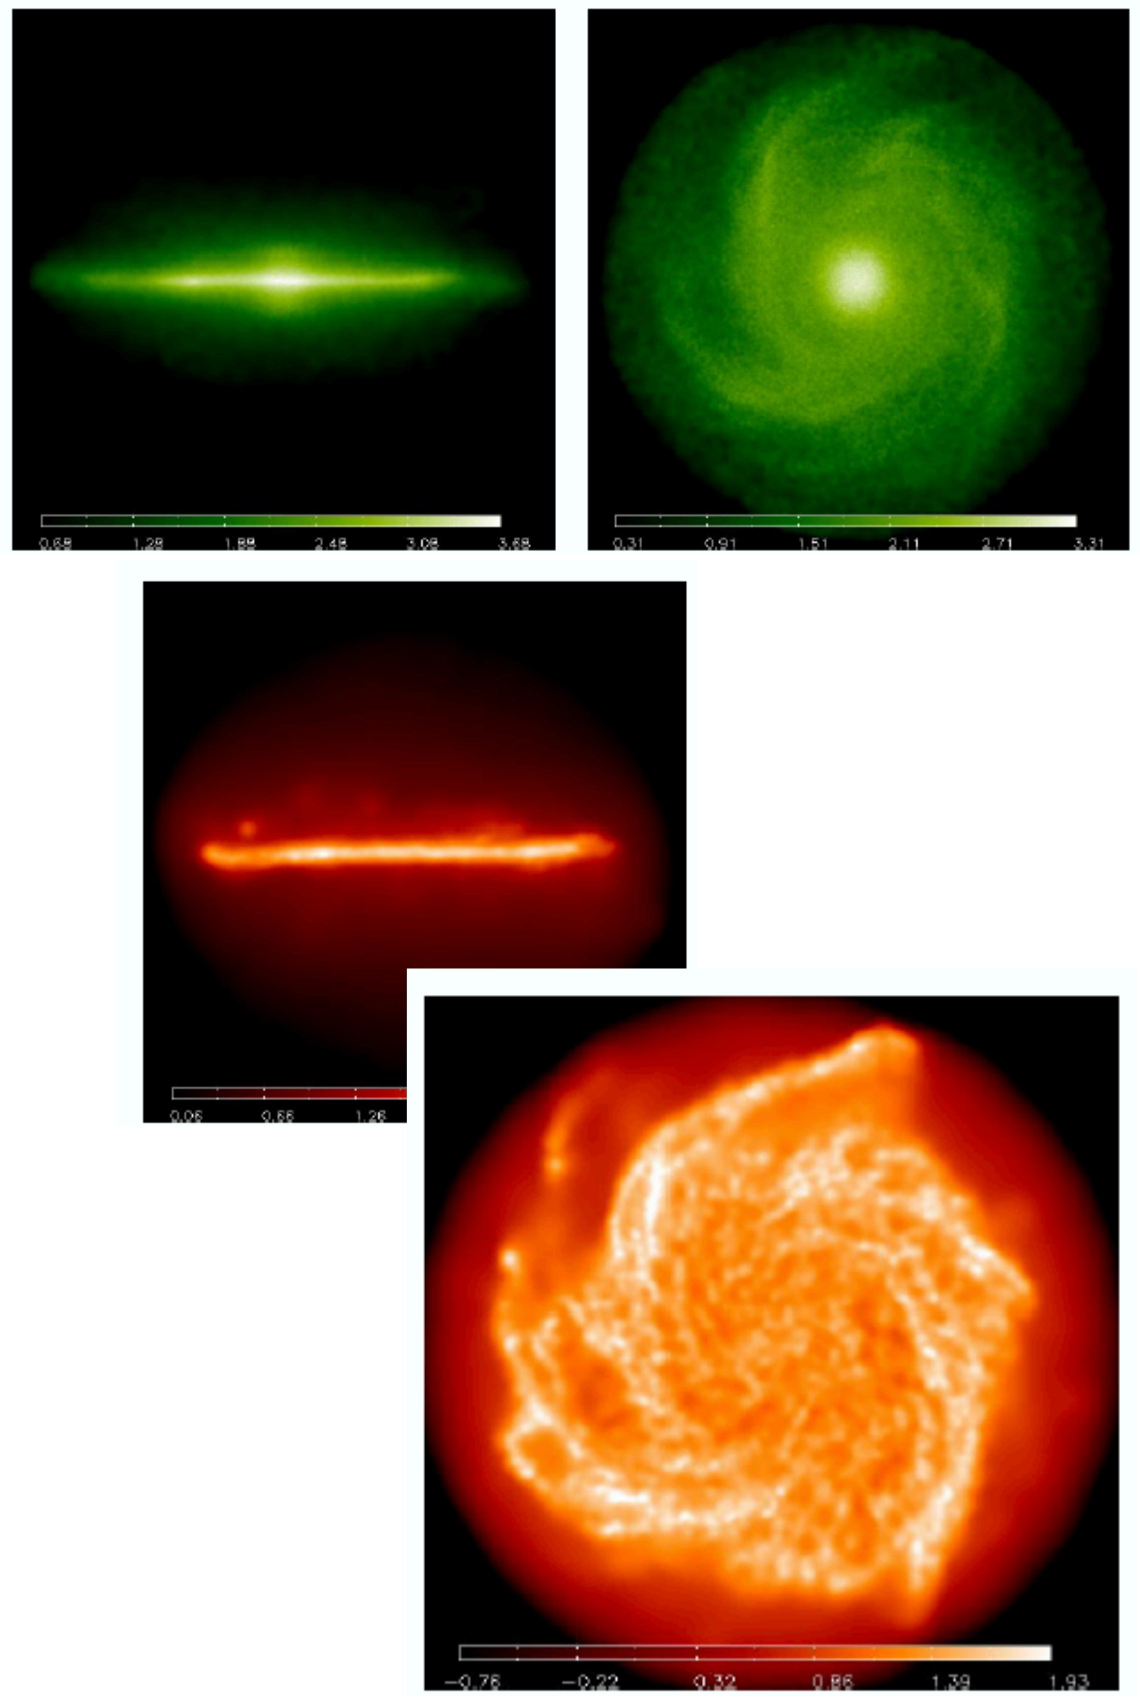

Spiral

Massive elliptical

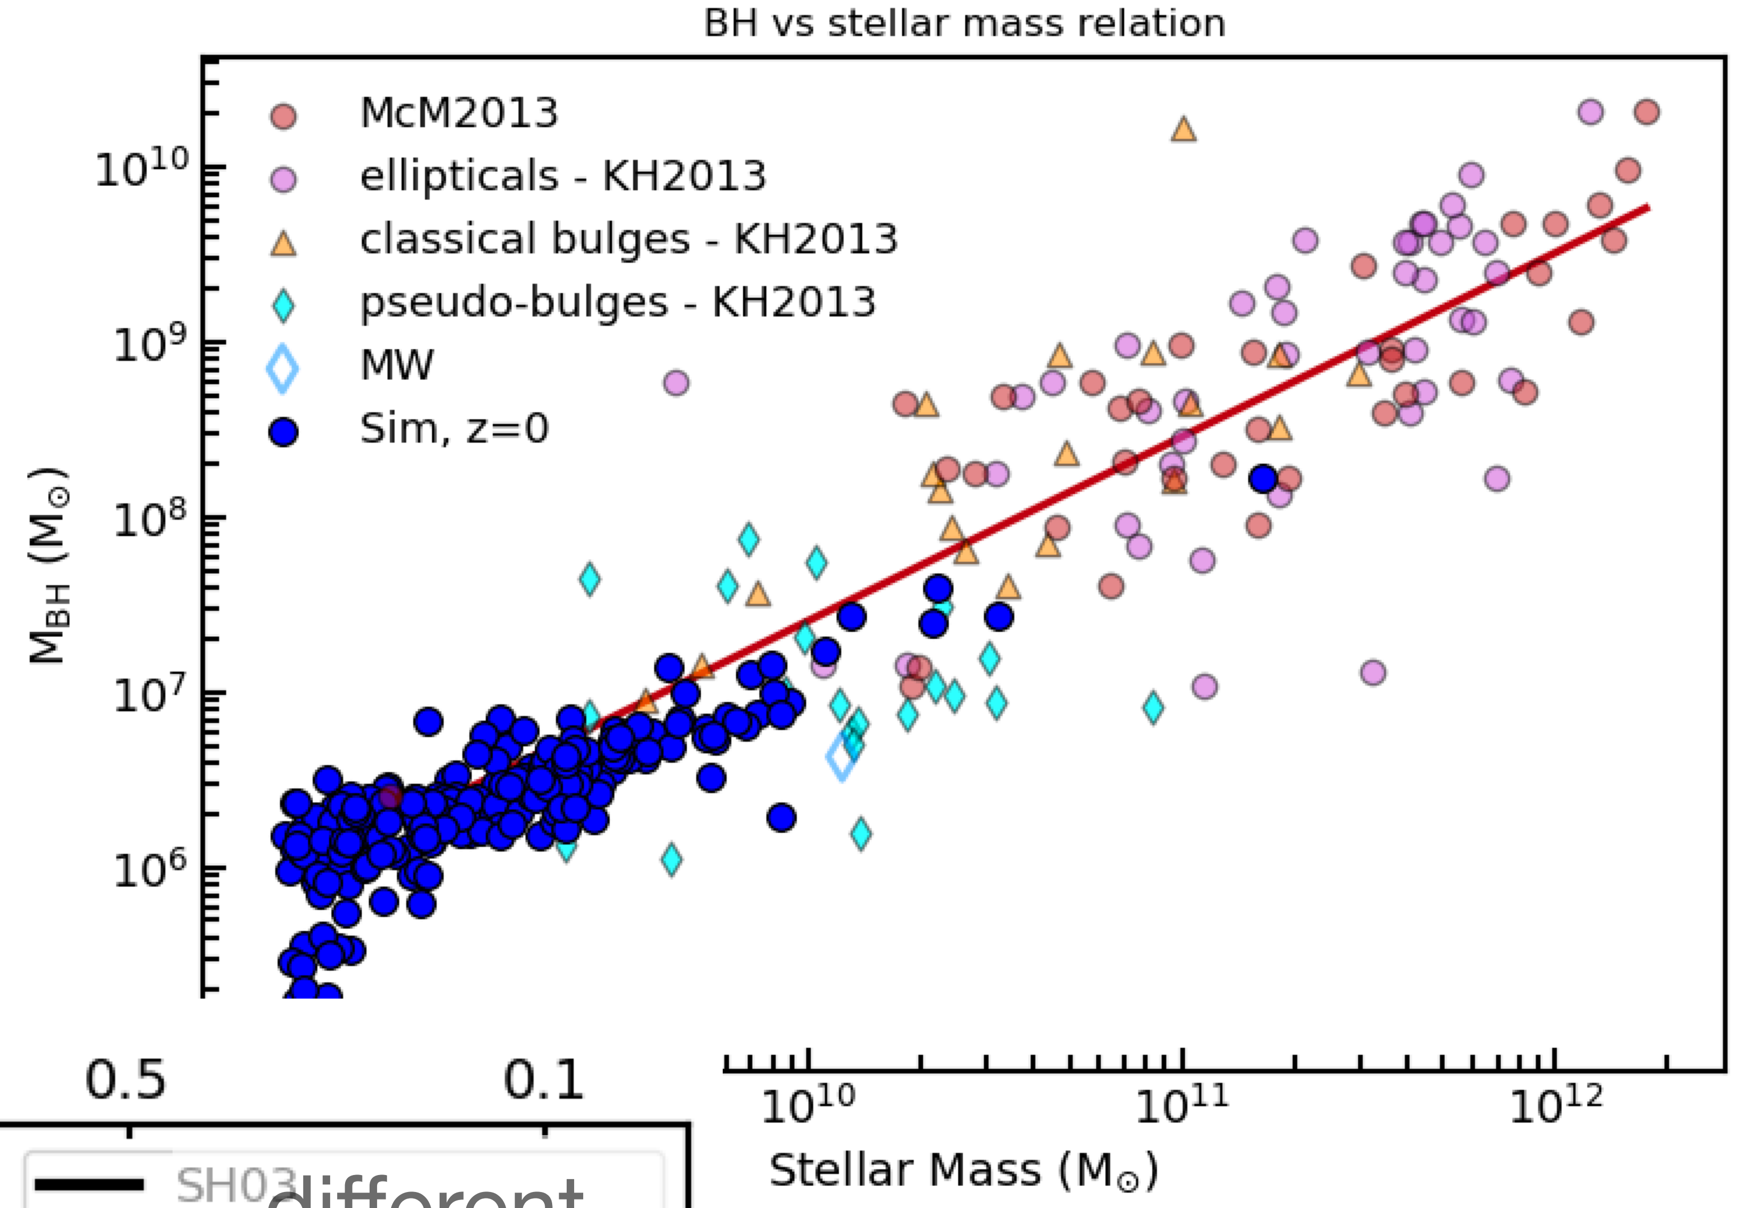

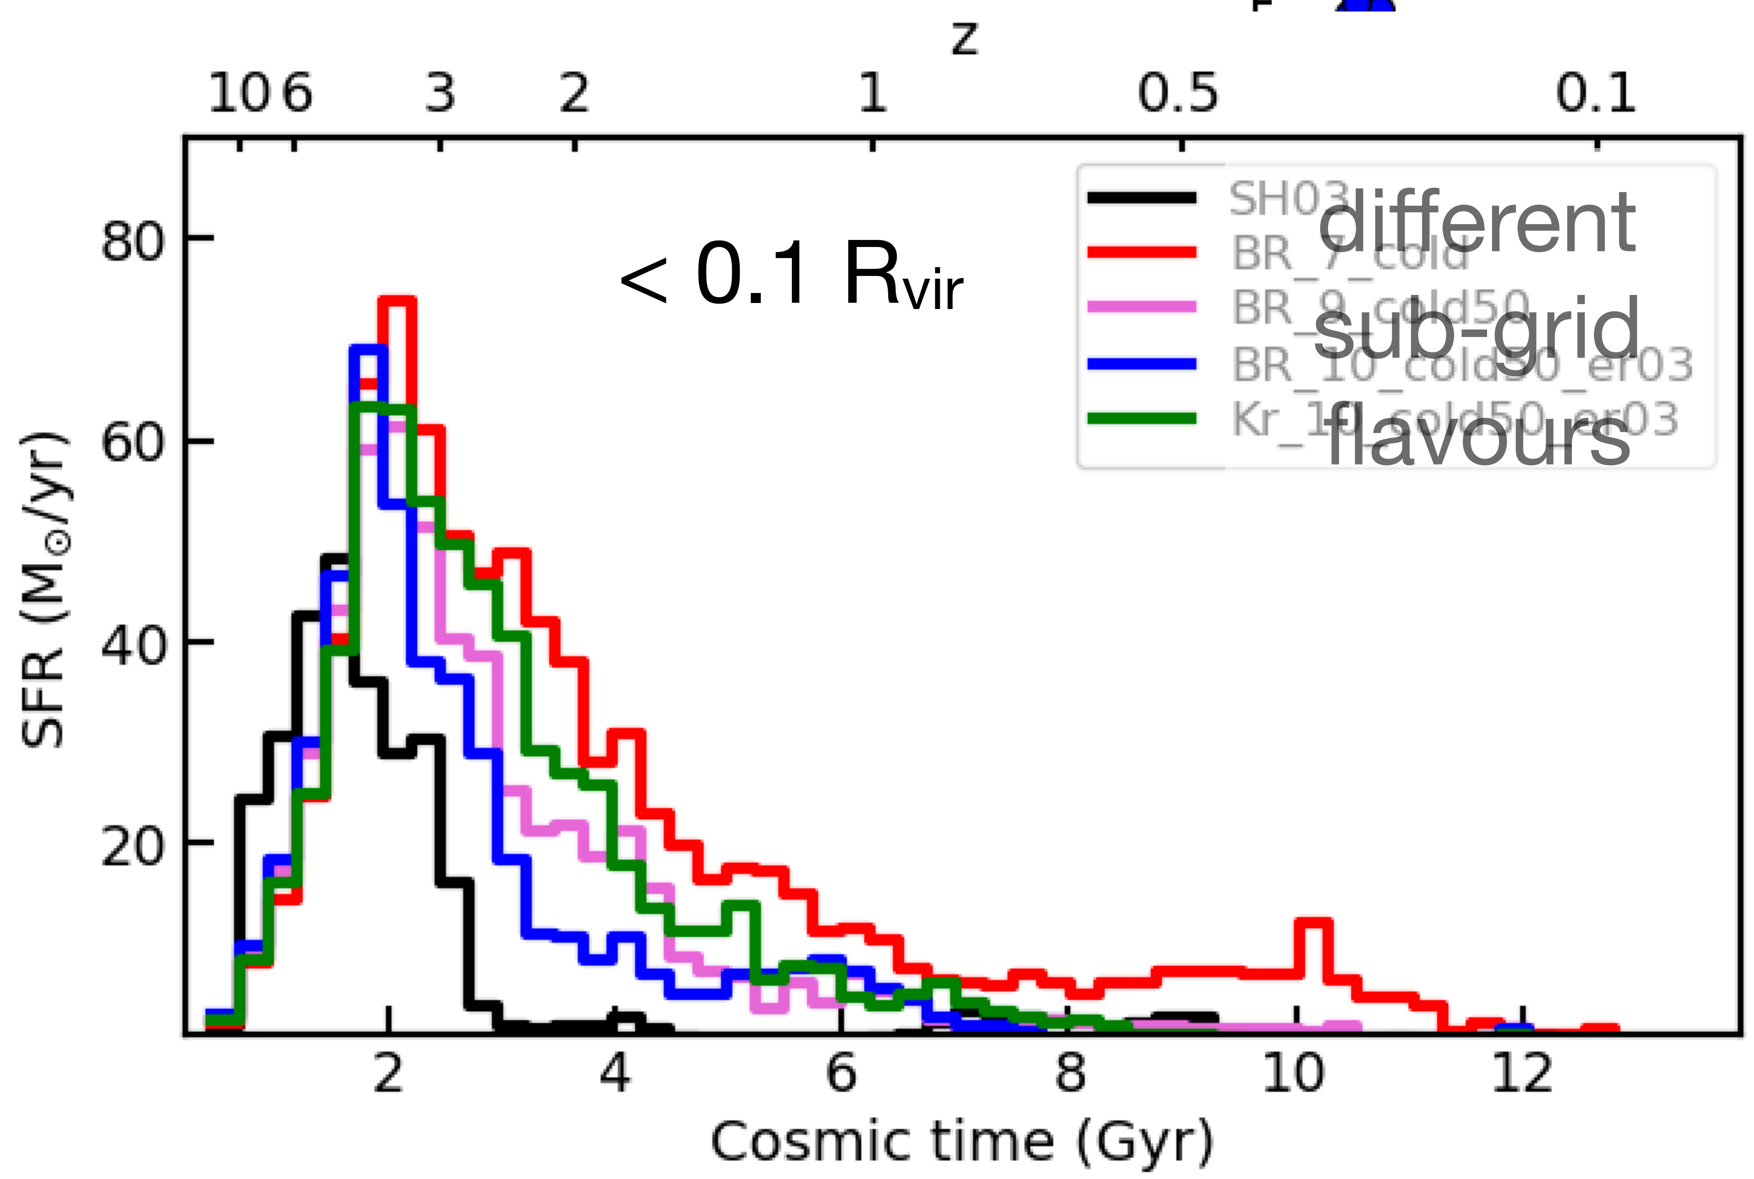

Next Steps and Expected Results

2. → SLOTH: Shedding Light On dark matter wiTH cosmological simulations

Milena Valentini, Stefano Borgani, Tiago Castro, Luca Tornatore, Matteo Viel, Alice Damiano, Pierluigi Monaco, Giuliano Taffoni

Main **scientific goals** of the project:

- theoretical understanding of primordial structure formation
- characterisation of the nature of dark matter
- exploitation and reliability of galaxy clusters as cosmological probes

Next Steps and Expected Results

2. → SLOTH: Shedding Light On dark matter wiTH cosmological simulations

Milena Valentini, Stefano Borgani, Tiago Castro, Luca Tornatore, Matteo Viel, Alice Damiano, Pierluigi Monaco, Giuliano Taffoni

Main **scientific goals** of the project:

- theoretical understanding of primordial structure formation
- characterisation of the nature of dark matter
- exploitation and reliability of galaxy clusters as cosmological probes

Simulations of warm dark matter (WDM):

- first results on a (10 Mpc/h)³ box w/ 1024³ particles

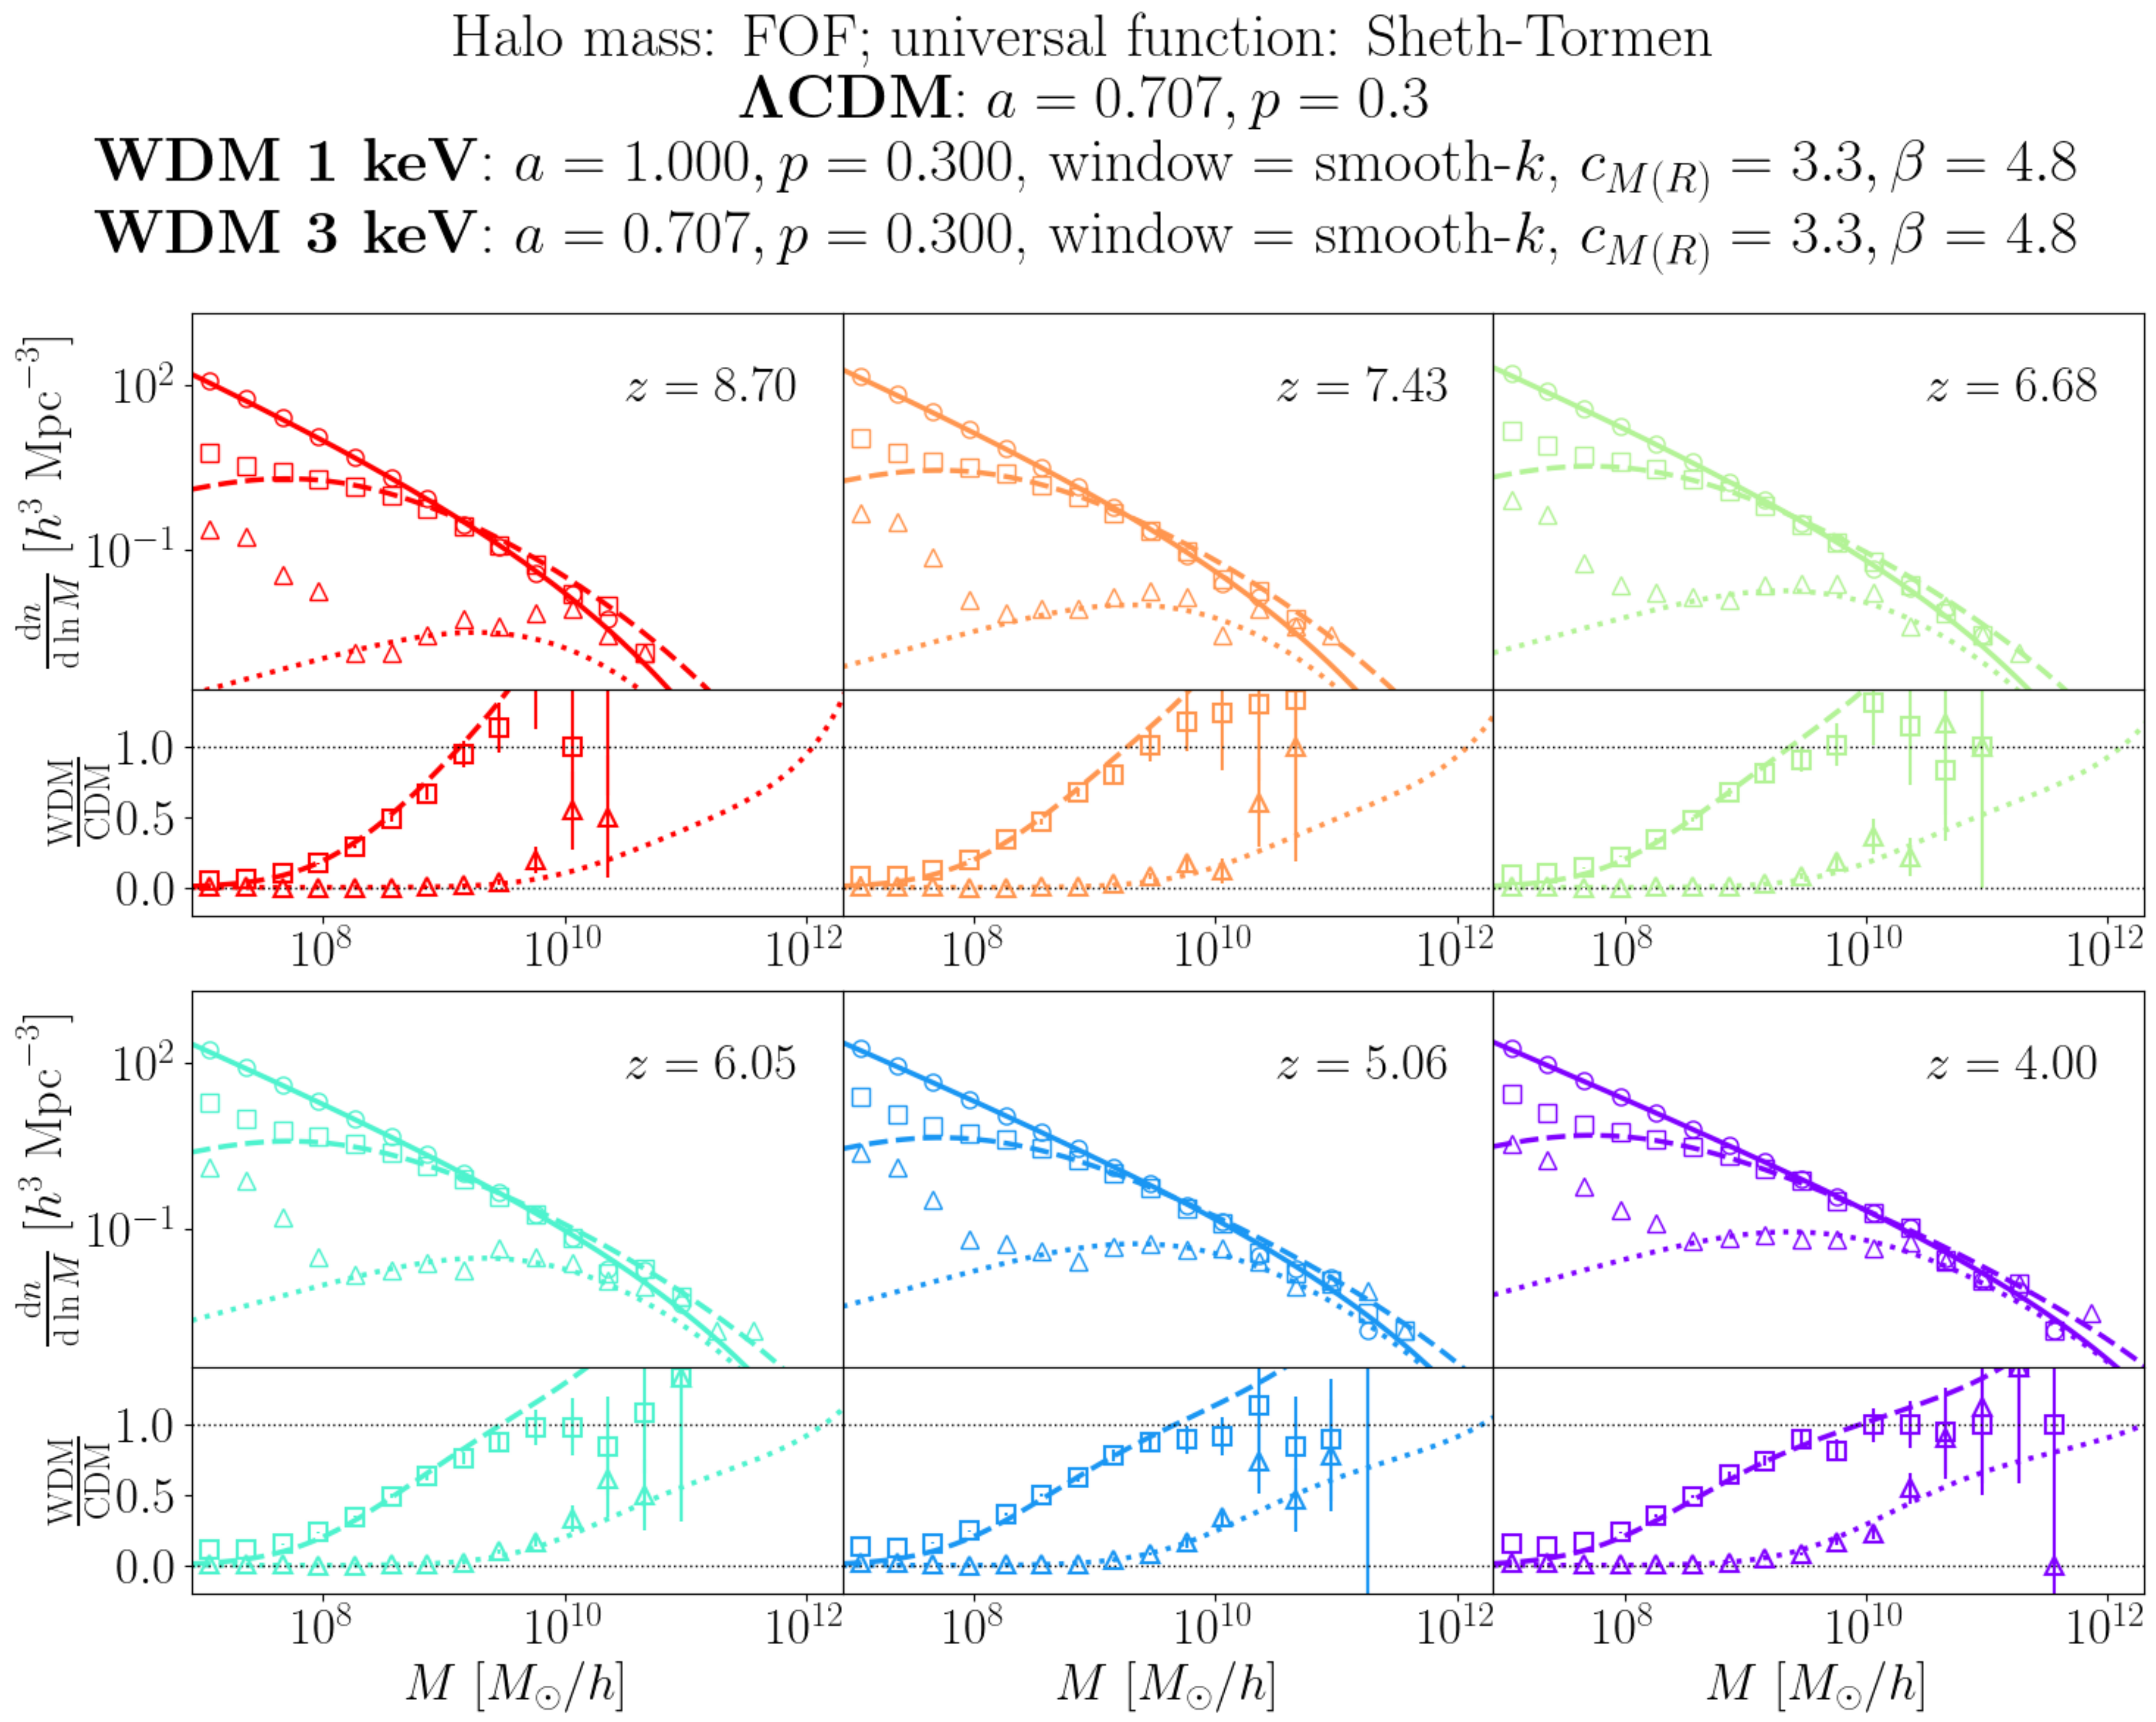

Next Steps and Expected Results

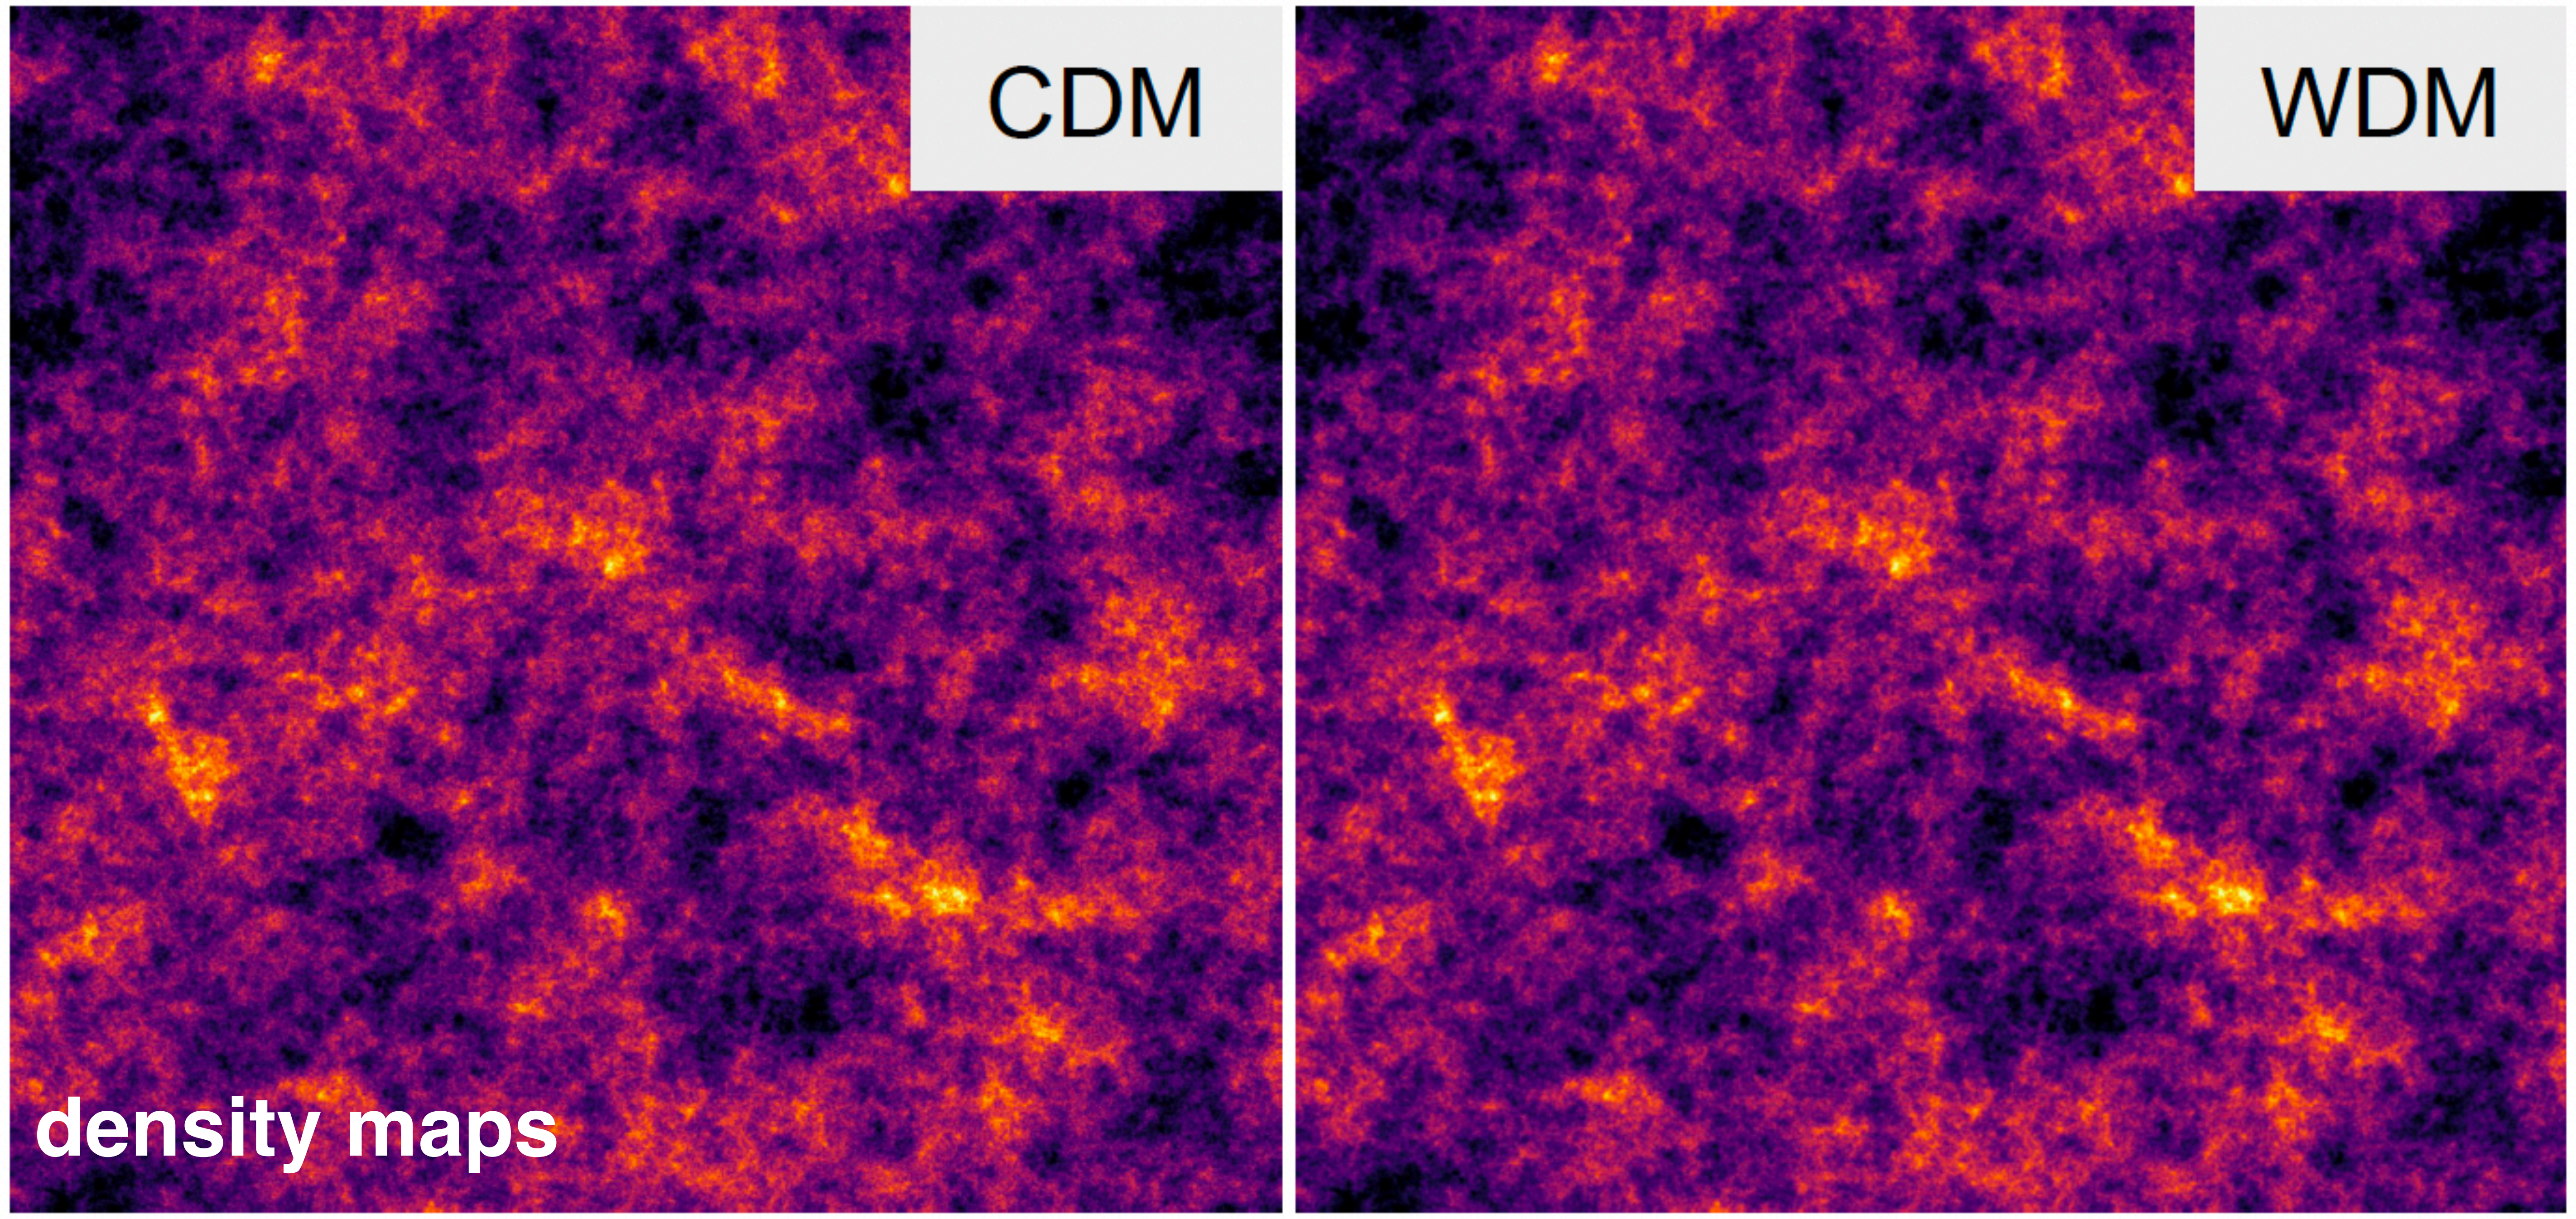

Initial conditions created for the two Flagship simulations
(6656^3 particles in a box of 65 Mpc/h on a side)

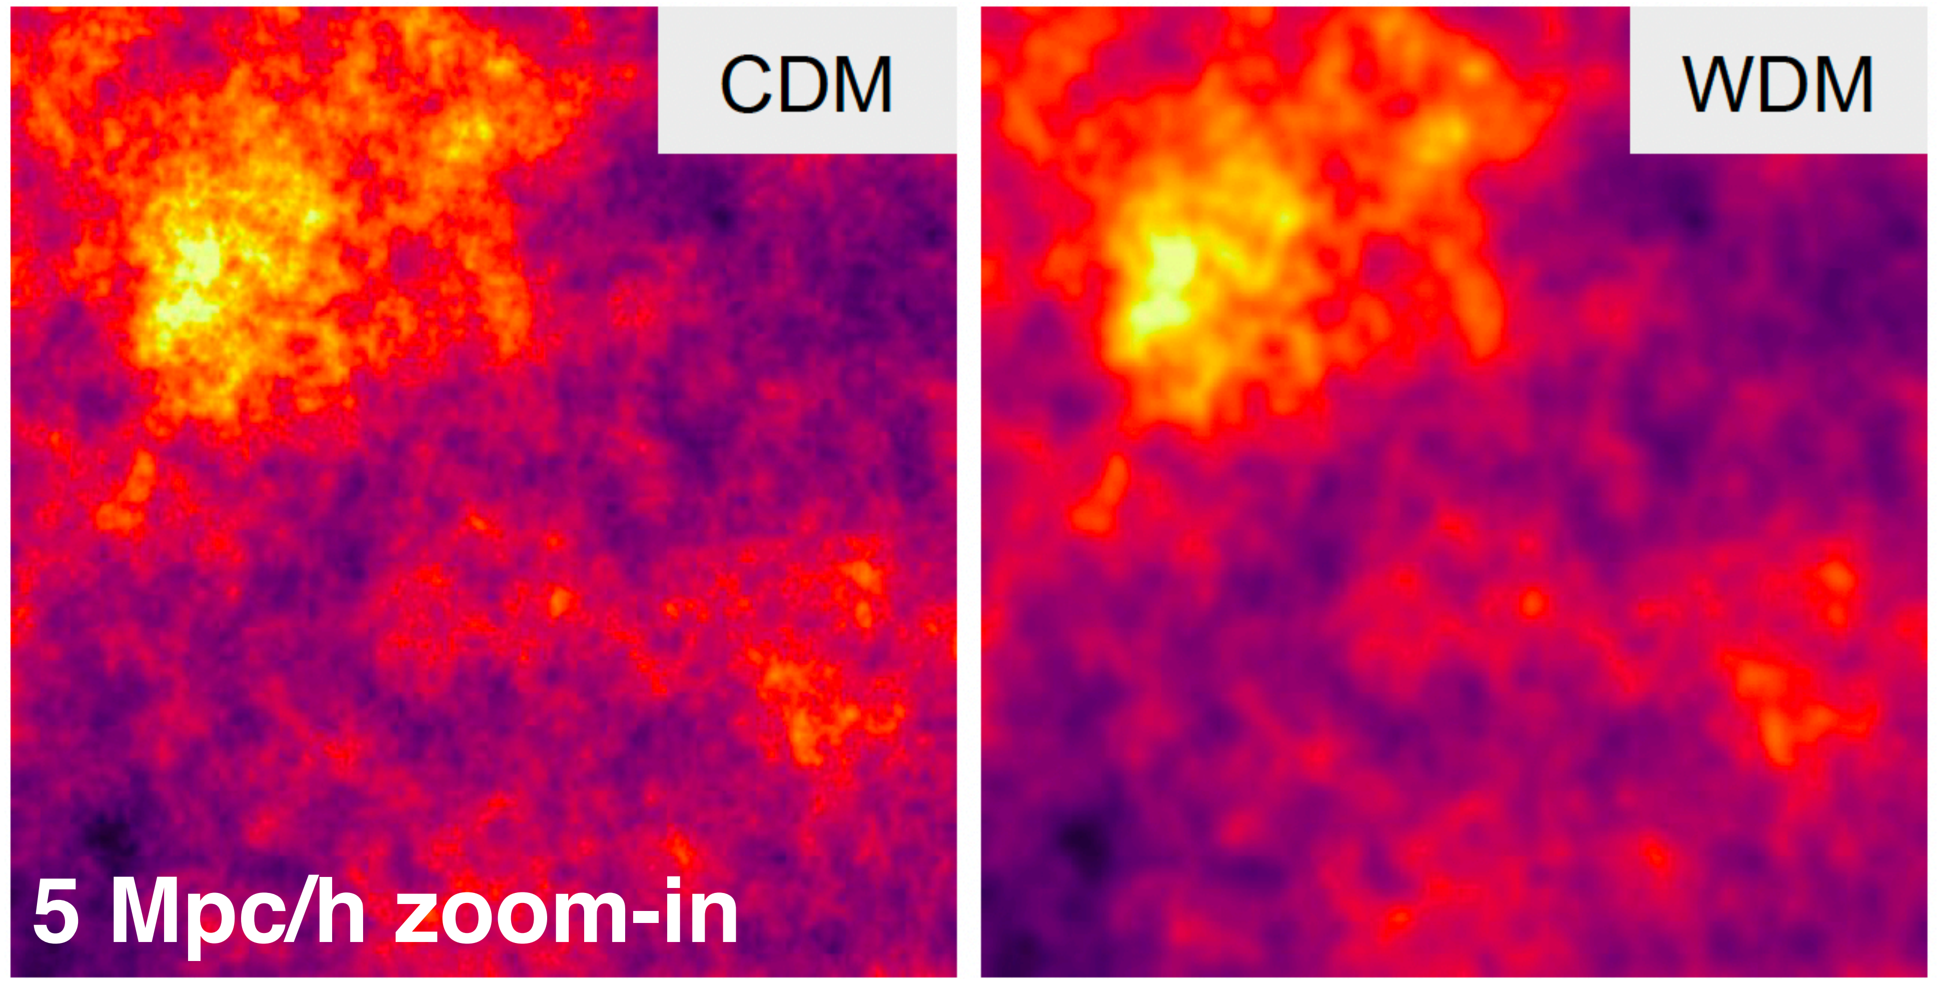

STRESS: inSighT on daRk mattEr with coSmological Simulations

Call for Leonardo Early Access Program

CINECA EuroHPC Joint Undertaking

Project Scope and Plan - Leonardo Early Access Program (LEAP)

Team: Valentini M., Castro T., Borgani S., Viel M., Tornatore L., Ragagnin A., Dolag K., Parimbelli G., Murante G., Dakin J.