



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca




Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# *The OpenGADGET code*

Milena Valentini



UNIVERSITÀ  
DEGLI STUDI  
DI TRIESTE



Dipartimento di  
**Fisica**  
Dipartimento d'Eccellenza 2023-2027

**Spoke 3 Technical Workshop, Trieste October 9 / 11, 2023**

# Scientific Rationale

- Numerical cosmology
- Structure formation and evolution






Open-GADGET3

USM

LMU

UNIVERSITÀ DEGLI STUDI DI TRIESTE

OSSERVATORIO ASTRONOMICO NAZIONALE INAF ISTITUTO TRIESTE



# Technical Objectives, Methodologies and Solutions

## The Lagrangian code Open-GADGET

Code: descendant of our developer version of GADGET-3 (TreePM+SPH, originally from [Springel 2005](#)), featuring an improved SPH formulation and several advanced physical modules (e.g. chemical evolution and enrichment by L. Tornatore)

Main differences between Open-GADGET 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 + OpenMP
- Improved SPH formalism
- Several modules for sub-resolution physics: star formation, stellar feedback, BH accretion and feedback, chemical enrichment, dust evolution
- Runs on CPUs and GPUs



# Technical Objectives, Methodologies and Solutions

## The Lagrangian code Open-GADGET

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

Key differences between Open-GADGET 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
- 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 1 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-GADGET

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

- Model based on the multi-phase effective model by Springel & Hernquist (2003)

- MUPPI ([Murante+ 2010, 2015](#), [Valentini+ 2017, 2019, 2020, 2023](#), [Granato+ 2021](#)), featuring:

AGN FEEDBACK  
ENERGY COUPLED  
TO COLD PHASE



- description of a multi-phase ISM (valid for resolution particles with mass  $\sim 10^4 - 10^7 M_{\odot}$ )
- H<sub>2</sub> - based star formation
- thermal and kinetic stellar feedback
- stellar feedback from low-metallicity environments
- stellar evolution and chemical enrichment

star formation

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

BH

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

dust

Different sub-resolution models available

# Timescale, Milestones and KPIs

## Milestones

- GPU porting
- 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

- availability of new modules of the code and improvement of existing ones for simulations carried out within the collaboration
- publication of scientific papers on refereed journals
- code release

# Accomplished Work, Results

## STRESS: inSighT on daRk mattEr with coSmological Simulations

Call for Leonardo Early Access Program

CINECA



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.

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

Halo mass: FOF; universal function: Sheth-Tormen

$\Lambda$ CDM:  $a = 0.707, p = 0.3$

WDM 1 keV:  $a = 1.000, p = 0.300$ , window = smooth- $k$ ,  $c_{M(R)} = 3.3, \beta = 4.8$

WDM 3 keV:  $a = 0.707, p = 0.300$ , window = smooth- $k$ ,  $c_{M(R)} = 3.3, \beta = 4.8$



Plot by G. Parimbelli

# Accomplished Work, Results

Dear Leonardo Beta Users,

I'm sorry to inform you that the cluster has not been released yet by NVIDIA/Eviden, and two additional reservations have been defined ending tomorrow at noon. We still have not been informed if there is the need for the cluster to be involved in additional tests after tomorrow noon. We'll update you as soon as we know more.

Best regards,

Dear Users,

we are sorry to inform you that, due to unpredictable technical issues, we have to postpone the start of the preproduction phase scheduled for today, at 12 pm. We will update you as soon as the new opening date is fixed. The access from restricted IPs will keep in place.

Best regards,  
HPC User Support and Production Team

Dear All,  
in view of Leonardo forthcoming opening to production, scheduled for May 22nd, and the atos/nvidia acceptance tests being completed, we are finally approaching the final configuration of the cluster and a stop is required tomorrow, May 17, starting at 10 am. We expect to put the cluster back to production within the early afternoon. During the stop the login node will NOT be available.

Sheth-Tormen  
0.3

Dear LEAP Users,  
the maintenance operations on Leonardo are still ongoing, due to some unexpected issues. We will inform you as soon as the access will be restored.

We apologize for the inconvenience,

Best regards

Dear all,

Today @13:00 CEST the **external network equipment** will lose power. The power will be recovered after about 2 hours. In the meanwhile, all the connections to the equipments will be truncated.

Apologies for the so short notice,

HPC User Support and production Team

Dear Users,  
the works at Tecnopolo were not concluded yesterday; a new operation is now ongoing, causing again the disconnection of the cluster from the external network. We are working to re-establish the connection within the shortest delay, we apologize for the inconvenience.

Best regards,

HPC User Support @ CINECA

Dear All,

in order to finalize the last steps of the Tecnopolo building site setup, Leonardo will undergo an electrical shutdown in blocks of nodes, starting from June 5 to June 15. The operations should not affect your jobs, but keep in mind that you may experience longer queueing times.

Best regards,

HPC User Support and Production Team @ CINECA

Dear All,

due to the ongoing activity on the electrical system of Leonardo datacenter, a cooling issue is currently affecting the machine room, and several failures occurred on the storage system is making it extremely unstable. We are trying to recover the situation within the shortest possible delay, we'll update you as soon as possible. We apologize for the inconvenience,

HPC User Support and Production Team @ CINECA

Dear Users,

due to the ongoing electrical maintenance in the data center, tomorrow at 7:30 am we have to cut the power of one of the two main electrical branches off (with the chillers powered by the branch). This may cause a general overheating of the facility with the subsequent need to shut the scratch equipment down (to avoid critical failures in the storage system). Your running jobs accessing the \$SCRATCH area may be killed, apologies for the inconvenience. The operation will end on June 14<sup>th</sup> at 1 pm.

Best regards,

HPC User Support and Production Team @ CINECA

Dear LEAP Users,


the state of the cluster, of the Data Center, and of the storage system were partially recovered, and a reduced number of nodes could be put back into production. Some additional disservices may occur until June 15.

Best regards,

HPC User Support and Production

Dear All,  
the scratch storage had again to be shut down, and all the compute nodes were drained, hence at present Leonardo is not available to production. We will inform you as soon as the cluster functionality will be recovered.

Best regards,  
HPC User Support and Production Team @ CINECA





# Accomplished Work, Results

Dear Leonardo Beta Users,

I'm sorry to inform you that the cluster has not been released yet by NVIDIA/Eviden, and two additional reservations have been defined ending tomorrow at noon. We still have not been informed if there is the need for the cluster to be involved in additional

Best regards,

Dear All,

in order to apply a few configurational changes, a series of operations are now ongoing on Leonardo which may affect the cluster availability and functionality. In particular, today starting at 14:00 the slurm service may experience some interruptions and job failures. We will inform you as soon as all operations are complete. We also inform you that at the moment login14 is not accessible for new connections, we will update you in the next few hours.

Best regards,

Dear All,

Leonardo will remain unavailable in the next hours until completion of the operations undergoing on the filesystems. The end of the maintenance is expected for tomorrow in the early afternoon.

Best regards,

Dear Users,

the works at Tecnopolo w disconnection of the clust shortest delay, we apolog

Best reg

HPC Us

Dear

due to the ongoing activity on the electrical system of Leonardo datacenter, a cooling issue is currently affecting the machine room, and several failures occurred on the storage system is making it extremely unstable. We are trying to recover the situation within the shortest possible delay, we'll update you as soon as possible. We apologize for the inconvenience,

HPC U

04/10/2023

Dear Users,

we are sorry to inform you that, due to a malfunctioning procedure managing the state of the nodes and their availability to production, at the moment it is not possible to submit jobs. The current running jobs will not be affected by the issue, and the pending jobs will remain in the queue. We will inform you as soon as the problem will be solved.

We apologize for the inconvenience,

Best regards,

HPC User Support and Production Team @ CINECA

Dear All,

in view of Leonardo forthcoming opening being completed, we are finally approach 17, starting at 10 am. We expect to put t During the stop the login node will NOT

Dear Users,

a total blackout at Leonardo building occurred last night and at present the cluster is not reachable from outside. The compute racks went in safety shutdown for cooling issues and the storage system is under investigation. We will inform you as soon as the cluster goes back to production.

Best regards,  
HPC User Support and Production

Dear Users,

in 30 minutes time we will put the /leonardo filesystem in read-only mode to perform the operations needed for the formatting of the home. We didn't find any running job working/writing in home, hence we expect no job failures. As for your activity, you will not be able to edit/write files, compile etc. in your homes, but you can keep reading files (such as input, or loading binaries for your jobs). Please, close all your open files and move your interactive activity to \$SCRATCH until further notice.

We apologize for the inconvenience,

HPC User Support and production Team

Dear Users,

this is to inform you that we are still experiencing some issues regarding the operations that are necessary to put the cluster back to production. We will promptly inform you as soon as we have updates. We apologize for the inconvenience.

Dear Users,

due to an unexpected technical issue to Leonardo cooling system, today's maintenance is still ongoing in order to properly recover the cluster's full functionality. We will inform you as soon as Leonardo will be back in production. We apologize for the inconvenience,

Dear Users,

this is to inform you that, as a result of the ongoing work to put the general purpose partition in production, you may experience interruptions in the InfiniBand network, resulting in job hangs or errors. We apologize for the inconveniences.

Dear LEAP Users,

Dear Users,

this is to inform you that Leonardo will be stopped next Tuesday, September 26th, due to scheduled maintenance. The stop will begin at 8:00 am and the cluster will be back to production within the afternoon of the same day. During the maintenance, login nodes will not be available.

Best regards,  
HPC User Support @CINECA

Dear Users,

and of the storage system were partially n. Some additional disservices may occur

All,

cratch storage had again to be shut dow t available to production. We will inform y regards,

HPC User Support and Production Team @ CINECA

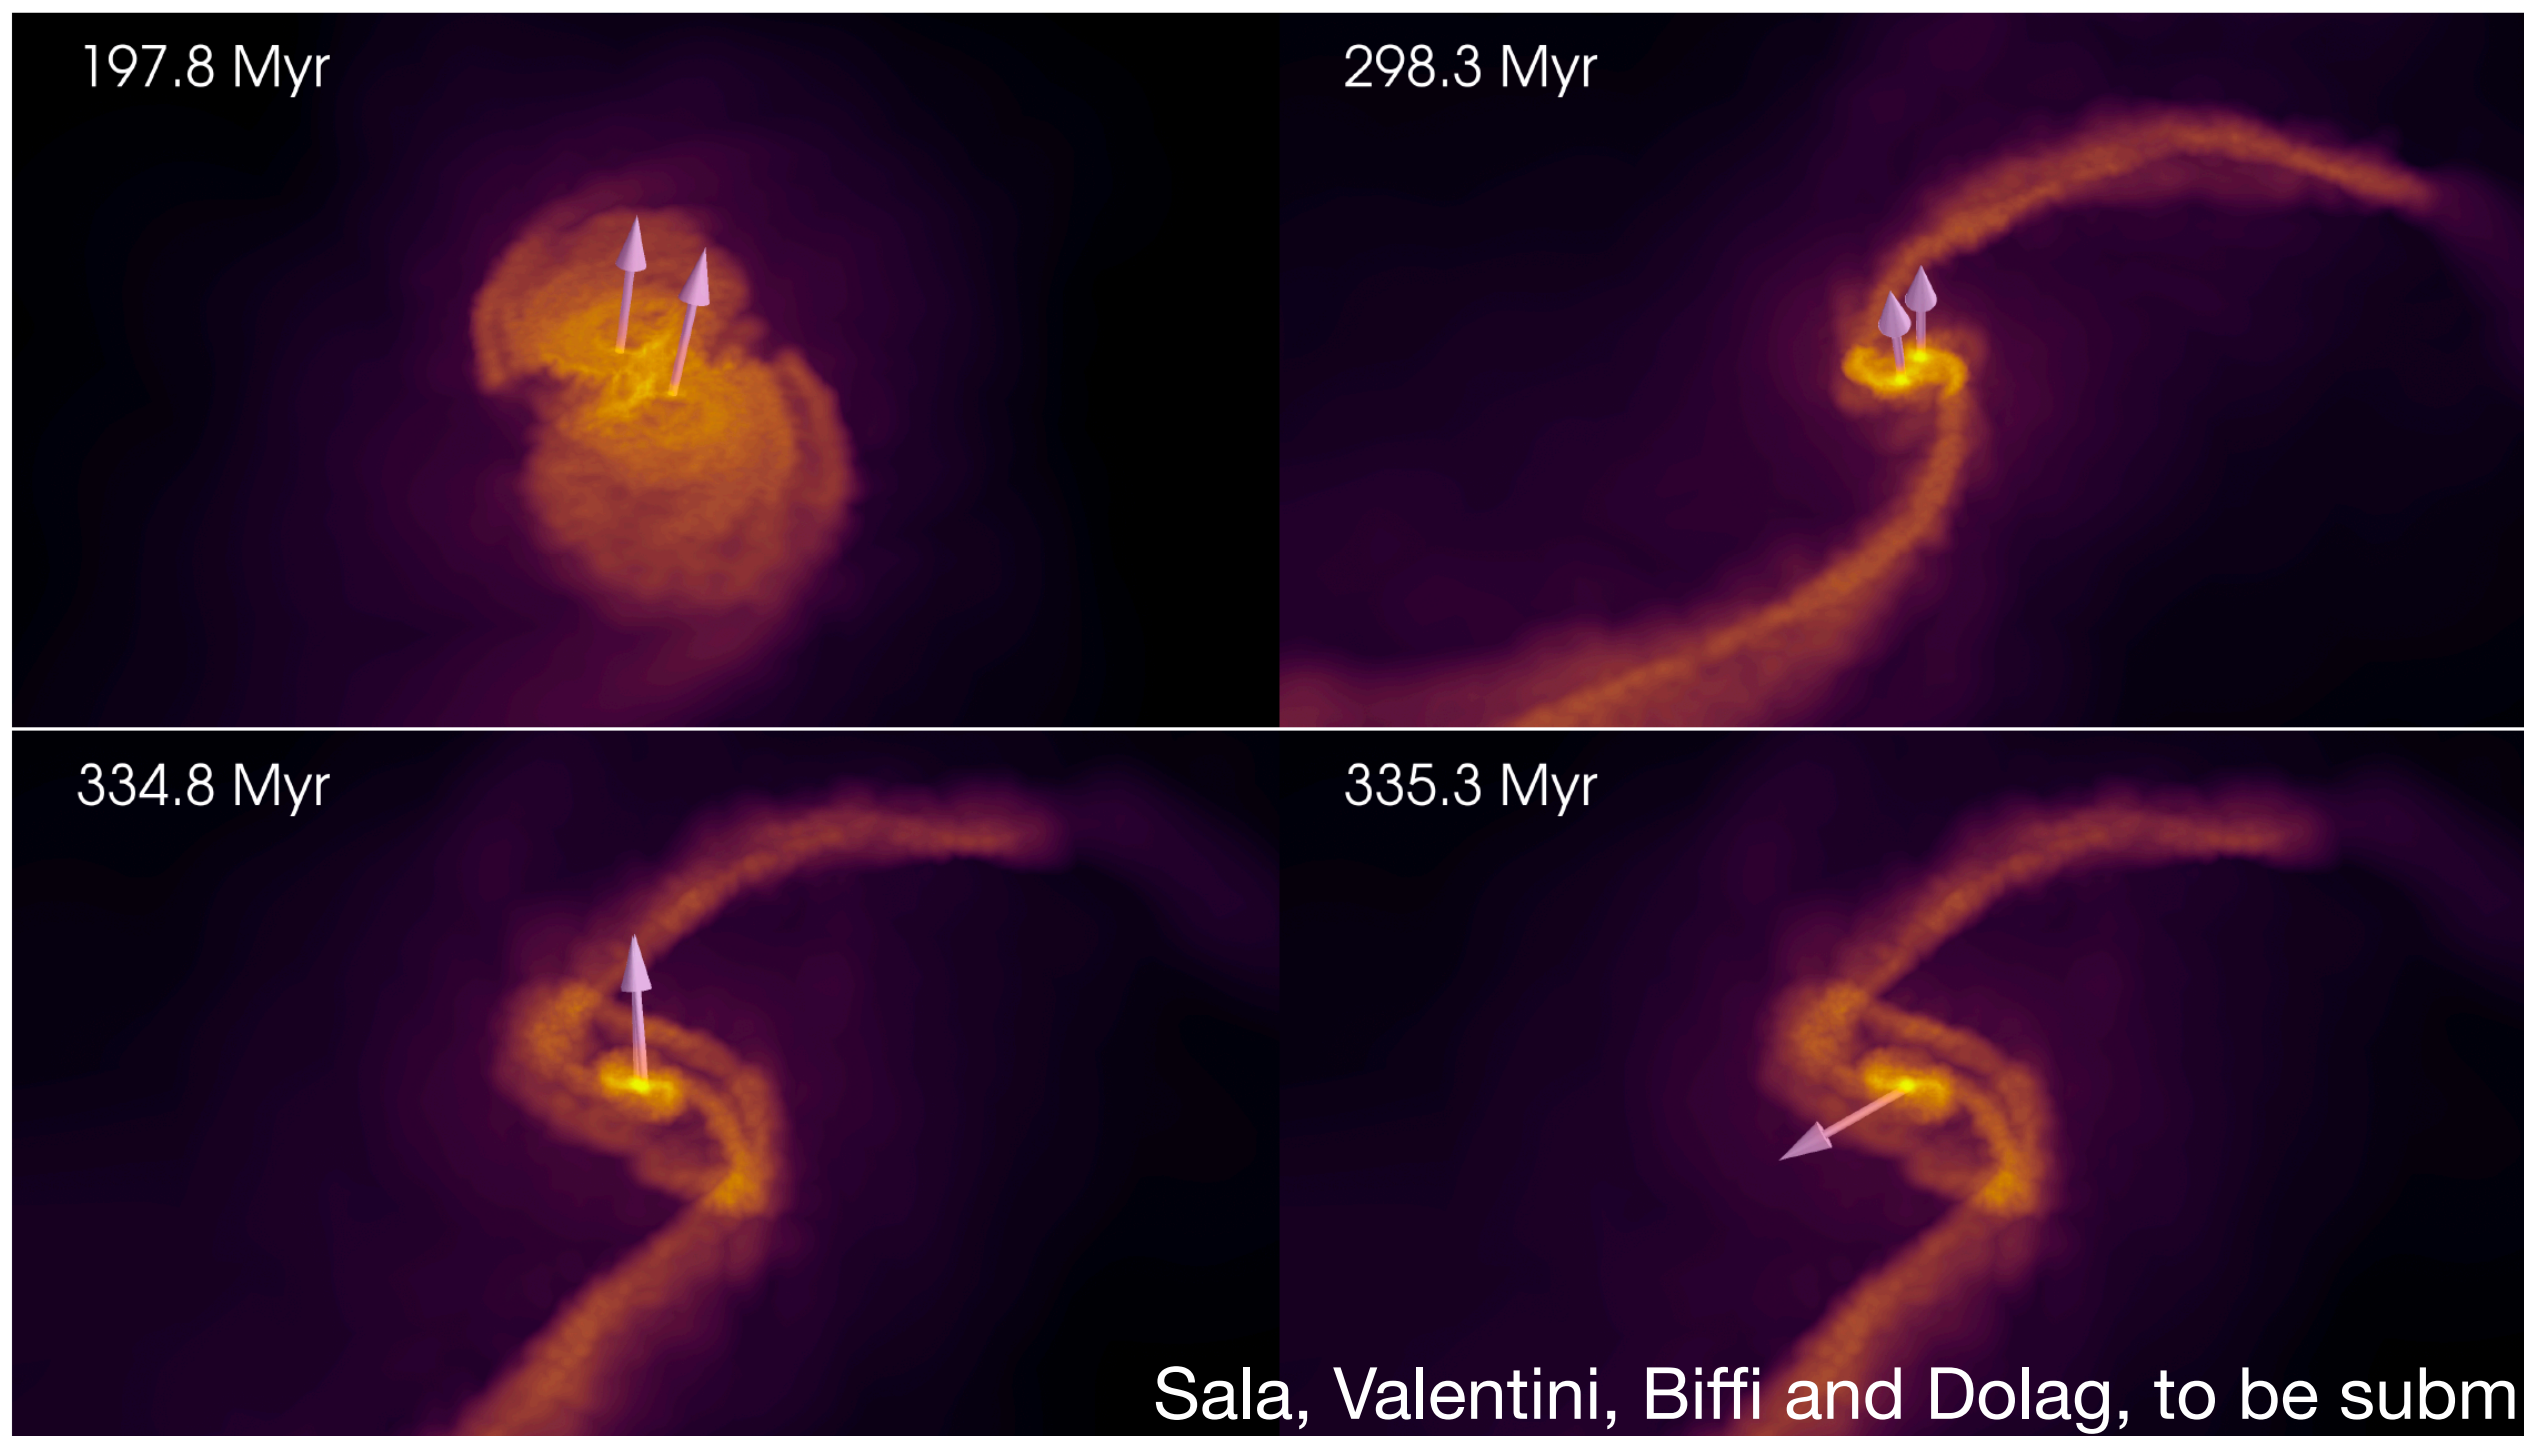

# Accomplished Work, Results

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

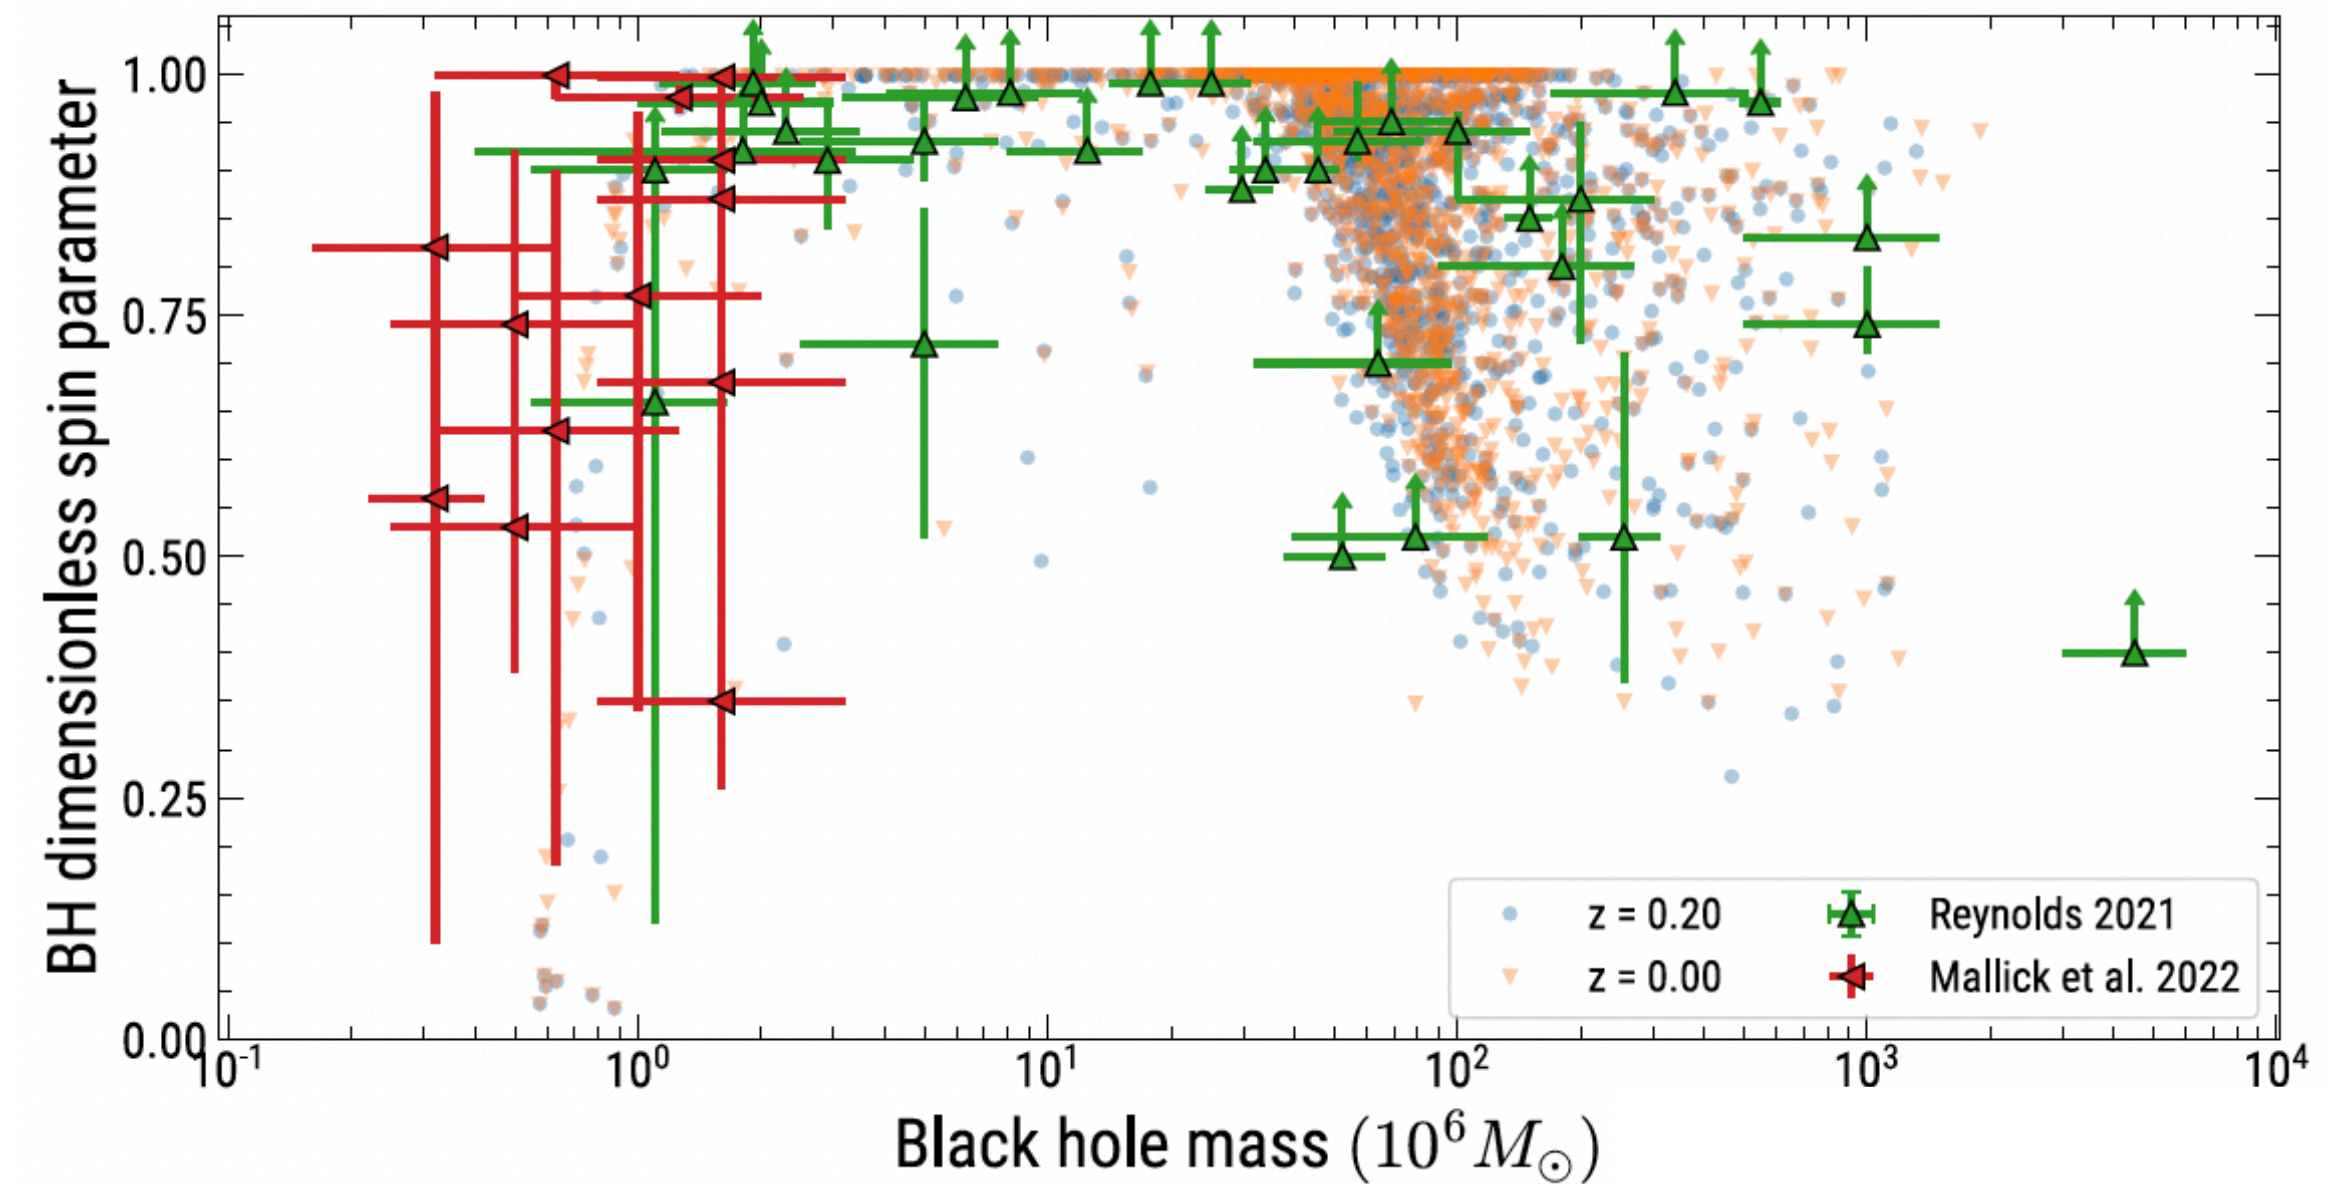
BH accretion rate  $\rightarrow$  disc mass  $\rightarrow$  accretion episodes

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

Given initial BH spin and mass  $\rightarrow$  post-merger BH spin vector

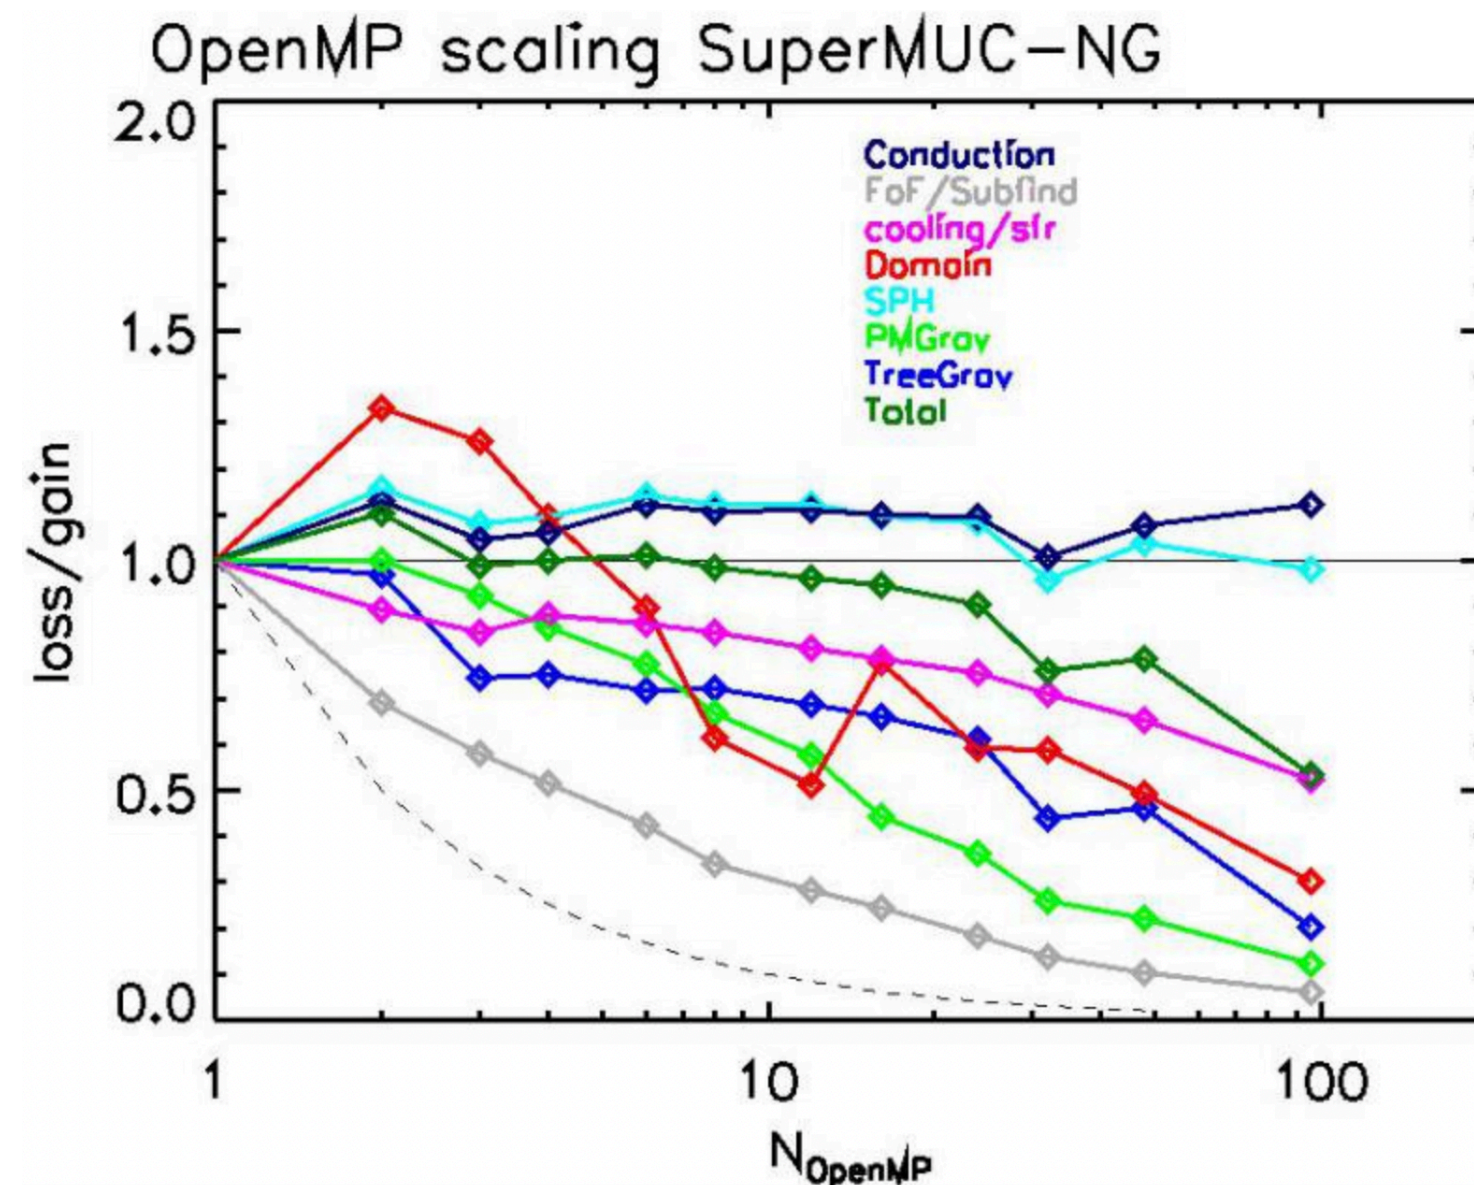
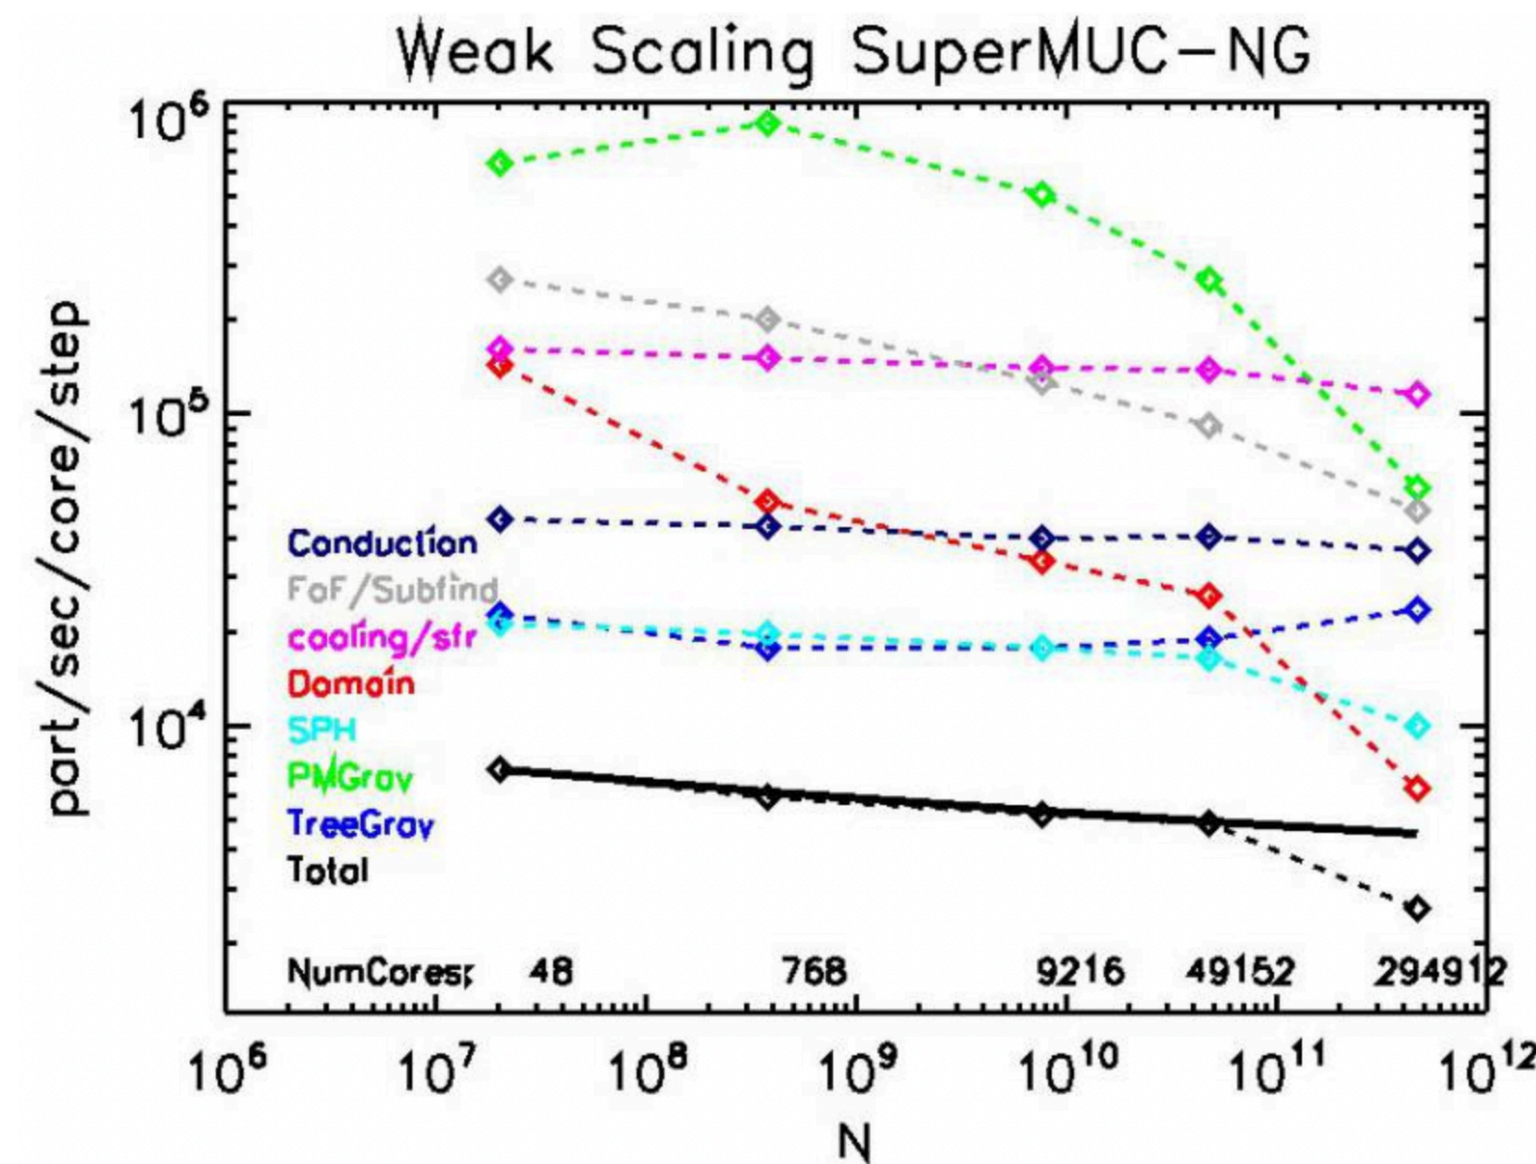
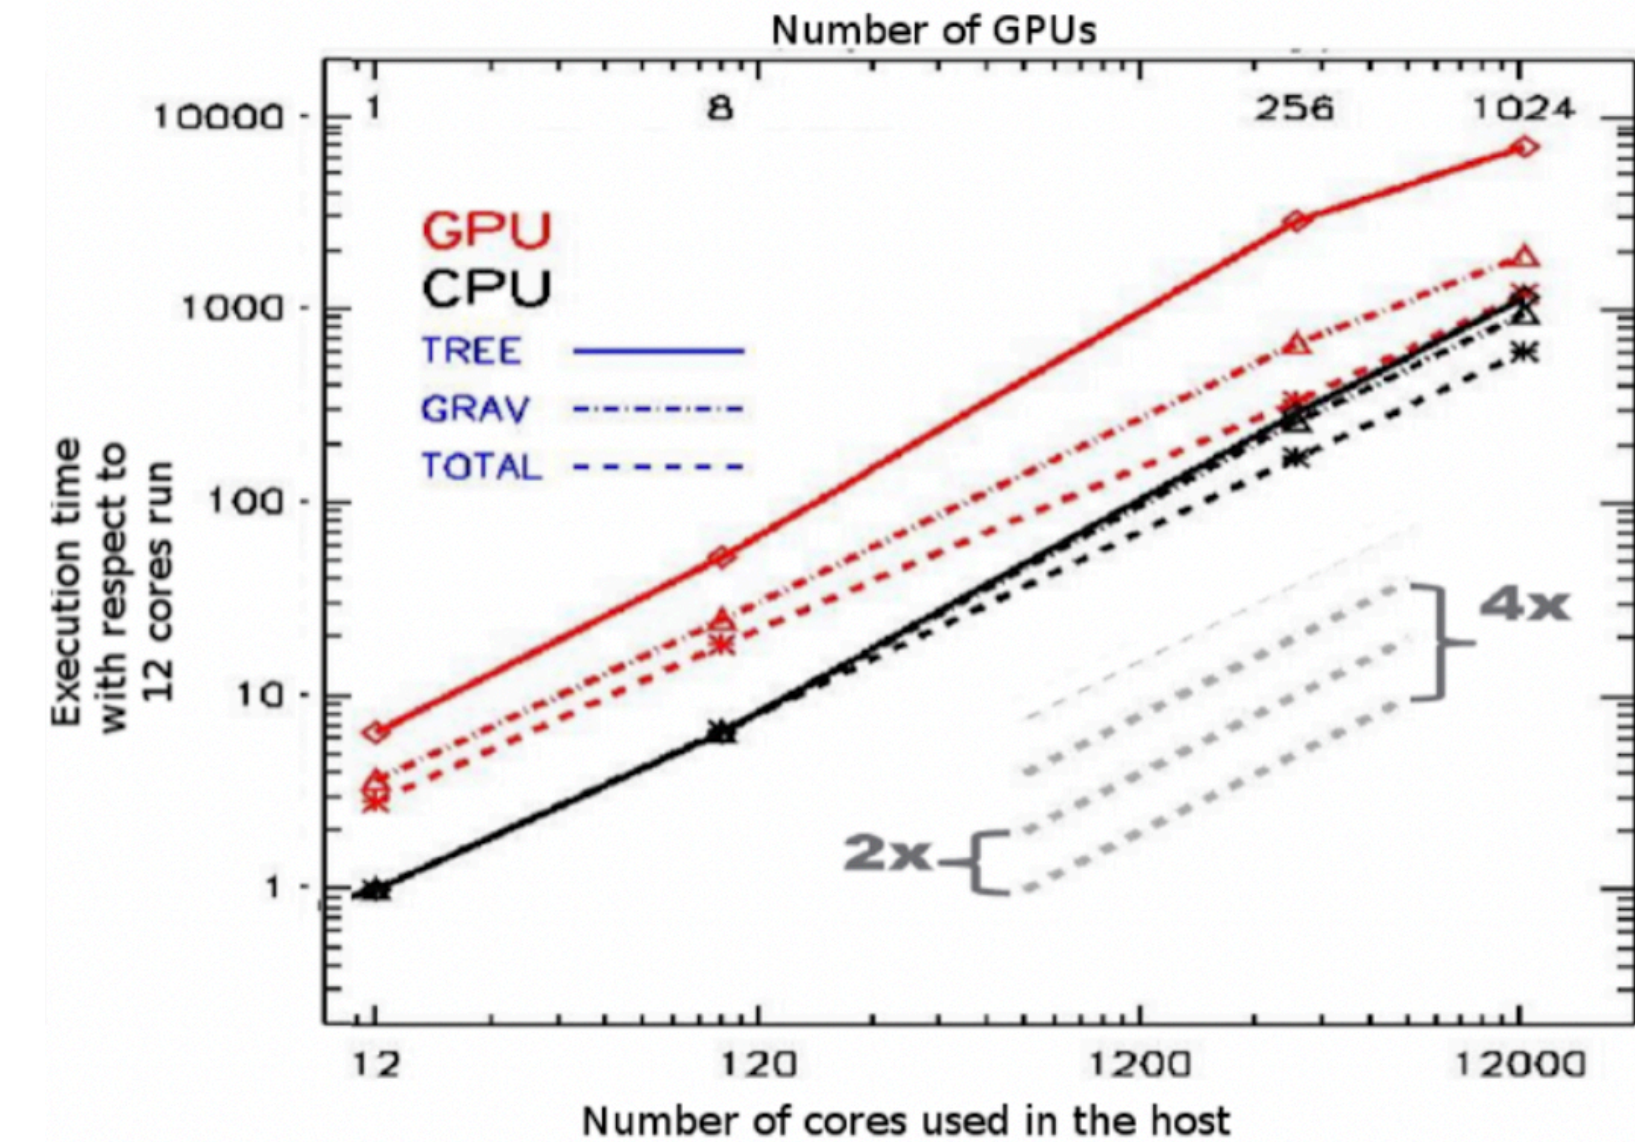


BH radiative efficiency dependent on BH spin



# Next Steps and Expected Results (by next checkpoint: April 2024)

- The MPI/OpenMP implementation of Open-GADGET3 has OpenACC support for GPUs (Ragagnin, Dolag+ 2020; GPU porting via OpenACC directives)
- Exploiting modern GPUs → total speedup by up to a factor of ~2-4 (for cosmological sims with  $> 10^7$  particles)
- Open-GADGET3 performance on modern GPUs keeps its speedup over different architectures (e.g., V100+NVLink or P100+PCI Express) and a number of devices



- Optimisation of the code performance (collaborations with Cineca, SPACE...)
- Development of the porting of OpenGADGET on GPUs: downstripping of the code to its backbone (tree build and walk, domain decomposition).

## Next Steps and Expected Results (by next checkpoint: April 2024)

---

- Development and use of modules for the description of sub-resolution astrophysical processes.
  - Completion, testing and validation of the module to describe the dynamics of Black Hole particles.
  - Completion of the testing phase of the "MUlti-Phase Particle Integrator" (MUPPI) model for star formation and stellar + AGN feedback. We will tune the MUPPI model so as to reproduce the correct observational properties of galaxies of different mass and morphology, from typical late-type disk galaxies to massive early-type galaxies.
- Include new modules and extend/improve existing ones for sub-grid physics

## Next Steps and Expected Results (by next checkpoint: April 2024)

---

- The OpenGADGET 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.
- Re-structuring of the code (modularity)
  - Cleaning the code and documenting its status

# Next Steps and Expected Results (by next checkpoint: April 2024)

---

## 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

Talk on Wednesday by Milena Valentini

### 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

Talk on Wednesday by Tiago Castro