



SCALABLE PARALLEL ASTROPHYSICAL CODES FOR EXASCALE

Scientific Cases

Deliverable number: D1.1

Version 1.0/1.0



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

Project Information

Project Acronym: SPACE
Project Full Title: Scalable Parallel Astrophysical Codes for Exascale
Call: Horizon-EuroHPC-JU-2021-COE-01
Grant Number: 101093441
Project URL: <https://space-coe.eu>

Document Information

Editor:	Luca Tornatore (INAF)
Deliverable nature:	Report (R)
Dissemination level:	Public (PU)
Contractual Delivery Date:	31.12.2023
Actual Delivery Date	27.12.2023
Number of pages:	46
Keywords:	Astrophysics, Cosmology, Numerical Simulations, Scientific Use Cases
Authors:	Georgios Doulis – GUF Benoît Commerçon – CNRS Joakim Rosdahl – CNRS Andrea Mignone – UNITO Stefano Truzzi, UNITO Luca Tornatore – INAF Geray Karademir – LMU Robert Wissing – UiO
Peer review:	Joakim Rosdahl – CNRS Geray Karademir – LMU

History of Changes

Release	Date	Author, Organization	Description of changes
0.1	03.11.2023	Gino Perna, ENGIN-SOFT	Started the document
0.7	14.12.2023	All authors	Version for reviews
0.9	15.12.2023	Joakim Rosdahl, CNRS	internal review
0.91	18.12.2023	Luca Tornatore, INAF	corrections following reviewer 1
0.92	19.12.2023	Geray Karademir, LMU	internal review
0.93	21.12.2023	All authors	Integration as requested by reviewer 1 & 2
1.0	27.12.2023	Eva Sciacca, INAF; Luca Tornatore, INAF; Andrea Mignone, UNITO	final version assembled

Scalable Parallel Astrophysical Codes for Exascale

DISCLAIMER

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain. Neither the European Union nor the granting authority can be held responsible for them.



The space above and below the message intentionally is left blank.

Executive Summary

Astrophysics is a field of study that seeks to understand the behavior and properties of celestial objects, such as stars, galaxies, and black holes. The study of these objects requires collecting and analysing large amounts of data obtained from a large diversity of observational instruments, ground- and space-based. The data cover the whole electromagnetic spectrum, from the radio to the gamma-ray wavelengths, and include gravitational waves, cosmic rays and several other non-electromagnetic observables.

In recent years, advances in technology and computational power have enabled scientists to collect and analyze unprecedented amounts of data, leading to new discoveries and insights into the nature of the Universe.

Numerical simulations are invaluable tools for interpreting the vast amount of complex data, and unprecedented numerical laboratories to both *(i)* study the evolution of astrophysical objects and large-scale structure, and *(ii)* contrast the prediction of models with observations.

This Centre of Excellence (CoE) represent an effort to evolve representative European codes to exa-scale capability both in terms of the capability of efficiently exploiting the exa-scale platforms and in terms of handling unprecedented data volumes.

This document illustrates several scientific use cases in astrophysics that demonstrate the power of data-driven research.

The selected scientific cases have two main characteristics.

On the one hand, they represent current cutting-edge problems in the respective research fields and allow us to examine how different types of data, including observational, simulated, and experimental data, are used to address pressing questions in astrophysics.

On the other hand, they pose relevant computational challenges to the CoE codes and stress their capability to scale in several respects.

As such, the scientific cases the code owners have selected will be used as testbed throughout the project to assess the progress achieved, using several metrics (see D2.1 and D2.2).

Contents

1	Introduction	7
2	OpenGADGET	8
2.1	Case 1 - Cosmological Boxes	8
2.2	Case 2 - Zoom-in simulations	10
2.3	Summary	12
3	PLUTO	13
3.1	Case 1 - Resistive Relativistic Reconnection	13
3.2	Case 2 - Plasma Column Instability	16
3.3	Case 3 - Synthetic Spectra from 3D Relativistic Jets	19
4	BHAC	21
4.1	Case 1 - Magnetised torus around a spinning black hole	23
4.2	Summary	24
5	ChaNGa	25
5.1	Case 1 - Isolated Galaxy Mergers	25
5.2	Case 2 - Magnetized cloud collapse	27
5.3	Case 3 - Sedov-Taylor blast wave	27
5.4	Case 4 - Cosmological boxes	29
5.5	Case 5 - Zoom-in simulations	29
6	FIL	30
6.1	Case 1 - Head on Collision	31
7	iPic3D	33
7.1	Case 1 - Properties of Turbulence in Various Planetary Magnetospheres	33
7.2	Case 2 - Regions of Energy Exchange in Ions and Electrons	34
7.3	Case 3 - 3D Anti-Parallel Reconnection	35
8	RAMSES	36
8.1	Case 1 - Sedov-Taylor blast	37
8.2	Case 2 - Cosmological box	37
8.3	Case 3 - Isolated galaxy	39
9	Conclusions	41

List of Figures

1	Projection of gas (left column) and stellar (right column) distribution at $z \sim 0$ (i.e. at present time, after ~ 13 billion years of evolution) in 30 Mpc boxes. Each row shows the simulation of the same box resolved with a different number N_p of particles: 64^3 , 128^3 and 256^3 in the top, medium and bottom rows. It can be seen that the level of details increases from top to bottom with increasing N_p	9
2	Zoom-in simulation of a galaxy cluster at four different time steps from upper left to lower right. Here it can be seen how the area necessary for the simulation of the cluster (see central area) is significantly smaller than a cosmological box, allowing the simulations to be calculated with higher resolution.	11
3	Spatially-averaged transverse component of magnetic field as a function of time for different resolutions and selected numerical schemes, with $S = 10^6$ (from [1]).	14
4	From top to bottom: 2D snapshots of mass density, effective resistivity, transverse velocity, current density, and fraction of resistive electric field in a reconnecting current sheet with effective resistivity at $t = 3 \times 10^3 \omega_p$ (left column) and $t = 3.6 \times 10^3 \omega_p$ (right), where ω_p is the plasma frequency. Length are in code units (jet radius).	15
5	Plasma column at the time $t = 50$, in the high resolution cases.	18
6	Plasma column at the time $t = 75$, in the high resolution cases.	18
7	Left: in gray scale the fluid density of the jet, in scale of colours are represented the LPs, the colour is proportional to the maximum energy of the LP spectral distribution. Right: the simulated energetic spectrum of one particle at different times.	19
8	2keV emission of a jet simulated using the LP module (dimensionless units normalized to the maximum).	20
9	GRRT of 3D torus with increasing resolution. Snapshot images of 3D GRMHD simulation data with parameters chosen to mimic the emission from Sgr A*. The resolution of the simulation data is indicated in the bottom-right corner of each panel.	21
10	Evolution of the 2D magnetised torus with resolution 1024×512 for times $t/M \in \{300, 1000, 2000\}$. We show logarithmic rest-frame density (top), logarithmic plasma β (middle) and the logarithm of the magnetisation parameter $\sigma = b^2/\rho$ (bottom). Magnetic field lines are traced out in the first panel using black contour lines. One can clearly make out the development of the MRI and evacuation of a strongly magnetised funnel reaching values of $\beta < 10^{-5}$ and $\sigma = 103$	22
11	Fluid-frame density (top) and $\log_{10} b^2$ (bottom) for $t = 3000M$ on the $y = 0$ plane (left) and the $z = 0$ plane (right) in the 3D magnetised torus run with resolution $384 \times 192 \times 192$	23
12	Top: rendering of the magnetic field strength for a simulated Milky Way-type galaxy (mass $M = 10^{12} M_\odot$), with a gas mass resolution of around $m = 10^2 M_\odot$. The highest magnetic field strength is in dark blue; moderate strength is in white and the weakest magnetic field is represented as red. Bottom: shows the timeline of a major merger between two lower resolution versions of the top figure. We can see that during the merger there is significant magnetic field amplification.	26
13	Resolution study of the magnetized cloud collapse with a mass-flux ratio of $\mu = 10$ from [2]. Here we vary the resolution from left to right, in the initial cloud ($12^3, 25^3, 50^3, 100^3, 250^3$). Renderings shows the time of jet formation (around the free fall time of the cloud), which occurs due to the winding of the magnetic field during collapse, which produces a magnetic tower structure. The top row shows a rendered face-on slice ($L_{xy} = [2000AU, 2000AU]$) of the density [g/cm^3], the rest of the rows show rendered slices through the rotation axis ($L_{xz} = [2000AU, 2000AU]$), where the second shows density [g/cm^3], and the third show the absolute poloidal magnetic field [μG]. All quantities are shown in logarithmic scale. For the scientific cases in this project we will perform an even higher resolution version (500^3) and not include the lowest resolution ones ($12^3, 25^3$).	28
14	Comparison of equal mass binary neutron star merger simulations computed with the second order IllinoisGRMHD code (right) and with the fourth order FIL code (left). On the y axis is log density ρ , the image is of a snapshot of the in spiral. The Fourth order FIL is able to capture sharp features better.	30

15	Foliation of the spacetime manifold. The lapse function α measures the proper time along the normal vector n^a to the hyper-surface σ_t which is equipped with an induced metric γ_{ij} . The shift vector β^i measures the displacement, on consecutive hyper-surfaces, between the observer time lines t_a and the normal lines n_a [3].	31
16	In this visualization, we observe a turbulent region characterized by the presence of numerous reconnection sites, depicted as yellow areas. Electron flow paths weave through this tumultuous landscape, intersecting multiple reconnection points. Notably, these electron trajectories are color-coded to reflect the strength of the local electric field encountered along their journey. This electric field is instrumental in the energy exchange process between the magnetic field and the electrons, a key factor in their acceleration. The vibrant color gradient along the flow lines thus serves as a dynamic map of the electric field intensity, highlighting zones of significant energy transfer and acceleration within the turbulent region [4].	34
17	Anti-parallel reconnection depicted in the kinetic simulation [5].	36
18	Sedov-Taylor blast wave: projection of the dark matter density at different times and resolutions.	38
19	2D dark matter density projection in runs with different resolutions and at different timesteps.	39
20	Face-on and edge-on maps of the isolated disk RAMSES galaxy, in the upper and lower rows respectively, after running for a simulated 21 Myrs. The projected maps show, from left to right, maximum hydrogen density along the line of sight, maximum refinement level along the line of sight, DM surface density, and stellar surface density. Note that the maps show only a fraction of the total volume, which is 320 kpc in width.	41

List of Tables

1	The set of cosmological boxes generated for the profiling activity along the project. The boxes are referred to as BOX_ N_p -SIZE where the "size" refers to the side length of the box.	10
2	The set of zoom-in simulations used for profiling and benchmarking. Here are details about two zoom-in simulations of the same object at different resolutions and consequently different mass per particle. For comparison with the parent simulation the mass for each dark matter particle M_{dm} is displayed as well as the mass for each gas particle M_{gas}	12
3	Resolution parameters of the three initial conditions used in our simulations. Stating the number of particles (N_{tot}), initial particle masses ($m_{gas}, m_{star}, m_{dark}$) and the softening length of the gas (ϵ_{gas}). m_{star} refers to the mass of star particles that form during the simulation. In addition, there is an old stellar disk component in the ICs, which consists of star particles that have masses of around $m_{star,old} \approx 4.3 m_{star}$	27
4	Resolution and IC parameters for the cosmological benchmark tests. From left to right, the parameters are: number of particles N_{part} per dimension, box width L_{box} in co-moving Mpc, starting redshift z_{start} , matter fraction Ω_m , dark energy fraction Ω_Λ , Hubble constant H_0 , power spectrum amplitude σ_8 on an 8 Mpc scale, and seed for random number generation with MUSIC.	39
5	Fixed isolated disk setup parameters, copied from AGORA paper II [6]. The parameters are virial mass M_{200} , rotational velocity $v_{c,200}$, virial radius R_{200} , concentration parameter c , spin λ , box width L_{box} , disk stellar mass $M_{d,*}$, disk scale radius r_d , disk scale height z_d , disk gas mass $M_{d,gas}$, disk gas fraction f_{gas} , and bulge stellar mass $M_{b,*}$	40
6	Varying isolated disk parameters. The parameters are, from left to right, number of DM particles N_{DM} , number of stellar disk particles N_{SD} , number of stellar bulge particles N_{SB} , coarse refinement level ℓ_{min} , fine refinement level ℓ_{max} , coarse cell width Δx_{max} , and fine cell width Δx_{min}	40

List of Acronyms

SPACE	Scalable Parallel Astrophysical Codes for Exascale
CoE	Centre of Excellence
AGN	Active Galactic Nucleus
AMR	Adaptive Mesh Refinement
BH	Black Hole
BHOSS	Black Hole Observations in Stationary Spacetime
BHNS	Black Hole Neutron Star
BNS	Binary Neutron Star
CDI	Current Driven Instability
CGM	Circum-Galactic Medium
CRs	Cosmic Rays
CT	Constrained Transport
DSA	Diffusive Shock Acceleration
EOS	Equation Of State
EM	Electromagnetic
ET	Einstein Toolkit
FCT	Flux-interpolated Constrained Transport
FIR	Frankfurt Illinois code
GLM	Generalized Lagrangian Multiplier
GPU	Graphics Processing Unit
GRB	Gamma Ray Burst
GRMHD	General Relativistic Magnetohydrodynamics
GRRT	General Relativistic Ray Tracing
GWs	Gravitational Waves
HD	Hydrodynamics
HLLC	Harten-Lax-van Leer-Contact
HLLD	Harten-Lax-van Leer-Discontinuity
HR	High Resolution
HRSC	High Resolution Shock Capturing
ICs	Initial Conditions
IGM	Inter-Galactic Medium
IMF	Interplanetary Magnetic Field
LDR	Localized Dissipation Region

LP	Lagrangian Particle
MHD	Magneto-Hydrodynamics
MP	Monotonicity Preserving
MPI	Message Passing Interface
MRI	Magneto Rotational instability
PIC	Particle In Cell
PPM	Particle Particle Mesh
PWN	Pulsar Wind Nebula
RK	Runge Kutta
RMHD	Relativistic MHD
SPH	Smoothed Particle Hydrodynamics
SPMHD	Smoothed Particle Magneto Hydrodynamics
UHR	Ultra High Resolution

1 Introduction

Forthcoming challenges in observational astrophysics will deliver a flood of data inconceivable just years ago, both in terms of data volume and complexity. The enormous diversity of data sources and their nature will open the possibility to shed light on some of the most important open questions about the nature and evolution of the Universe. However, this comes at the price of challenges equally formidable from the computational point of view. In fact, numerical codes are unavoidable tools to process and interpret the data from observational campaigns, whose size and intricacy are far beyond human capability. Moreover, machine learning increasingly acquires a central role in data mining, data-driven discovery and data processing, taking advantage of the most recent advancements and triggering the development of even more sophisticated techniques. At the same time, digging, processing and interpreting the data requires an adequate theoretical framework from which detailed predictions on precise observables could be drawn and against which the observations could be contrasted. It is difficult to underestimate the role that numerical simulations play in astrophysics and cosmology. They allow scientists to solve complicated non-linear multi-scale multi-physics systems of differential equations that represent a huge variety of intertwined physical processes, applying them to gargantuan physical systems.

Numerical simulations can then produce massive amounts of numerical results that allow insights into the complex network of combined effects of the different physical processes and how they influence each other. That is the only way in which those physical processes and their interplay can be studied together and observable predictions can be obtained. Moreover, the observational findings convey information as it emerges from very intricate astrophysical environments; their interpretation requires the capability of individuation of the effects of single processes and mutual influence among different phenomena and environments. Hence, numerical simulations are inestimable laboratories that enable scientists to tackle that complexity.

In this framework, the SPACE Centre of Excellence (CoE) aims to evolve representative EU-based astrophysical codes to exa-scale capability and to develop new – and enhance existing – machine learning techniques and algorithms. The selected codes cover different areas and disciplines, from cosmology and galaxy formation to single compact objects and multi-system simulations.

In the CoE design, the codes evolve within an iterative collaboration between WP1, which provides the codes and individual kernels (see D1.2), and WP2, which supports the scientific partners with expertise and knowledge in HPC. The WP1 partners also provide the scientific cases that serve as testbeds for running the codes and kernels and evaluating their performances by extracting the sensitive metrics (see D2.1 and D2.2). Also, the codes will be constantly validated by comparing the results obtained by running the evolved codes and kernels on the scientific cases with the ones obtained with the initial versions of the same codes to verify the integrity of the results.

In this document, the code owners illustrate the scientific cases selected for each code, which have two main characteristics. On the one hand, they represent current cutting-edge problems in the respective research fields and allow us to examine how different types of data, including observational, simulated, and experimental data, are used to address pressing questions in astrophysics. On the other hand, they pose relevant computational challenges to the CoE codes and stress their capability to scale in several respects.

From the technical point of view, a "scientific case" consists, besides its scientific rationale, of the set of files needed to run the simulation described here. That set of files is uniquely identified and stored in our common Internal Data Archive (IDA; see D4.1).

The following chapters are organized as one per code, where each code presents its Scientific cases.

2 OpenGADGET

OpenGadget is a collisionless N-Body/Lagrangian cosmological code using the smoothed particle hydrodynamics (SPH) computational method to describe the motion of fluids in addition to the gravitational forces, which are calculated using a tree structure. The code allows simulations in a full cosmological context, i.e. accounting for an expanding background and the presence of matter, both "dark" and baryonic (ordinary matter), and dark energy. Although the full cosmological context is often the default choice, having a non-expanding background and a setup with only dark or baryonic matter is equally possible. In addition to the gravitational problem, it also simulates the evolution of the physical properties of the baryonic matter subject to hydrodynamics and other physical effects such as radiative cooling, star formation, energy feedback, radiative transfer, magnetic fields and others that we collectively refer to as "extra physics". We redirect the reader to the deliverable D1.2 - "Kernels and mini-applications" for a more detailed code description.

For OpenGadget there are two main scientific targets: (i) volumes that are representative of the Universe, that we call "cosmological boxes" (see Sec. 2.1), and (ii) volumes that contain single collapsed objects like clusters of galaxies or single galaxies (see Sec. 2.2). The boxes are used to study the evolution of large structures and how they are shaped by several astrophysical properties, such as the clustering of matter and the impact of baryonic processes in the Universe (such as the phase-space distribution, the history of stellar feedback and chemical enrichment, the growth of black holes and their relationship with the environment and hosting objects, and many others). These single-object simulations focus on the detailed evolution of gravitationally bound structures of cosmological interest, whose properties depend and impact on the fundamental observables that inform us of the nature of our Universe. Hence, the cosmological numerical simulations are very helpful in testing and discriminating cosmological theories that predict the statistical properties of the Universe and in interpreting multi-dimensional observations from many ground- and space-based experiments.

In document D2.1, we describe the fundamental actions planned for the code in more detail: enhancing the core efficiency and the scalability of both OpenMP and MPI, and re-engineering the overlapping communication and computation.

To evaluate the progress of these improvements and help us in studying and profiling the code behaviour during the optimization, we have selected two specific user cases: (i) cosmological boxes and (ii) single galaxy-cluster simulations, which we describe in the following sections.

2.1 Case 1 - Cosmological Boxes

As the denomination suggests, cosmological boxes are built as fair samples of the Universe. In these boxes, the matter distribution and clustering are substantially homogeneous throughout the box, and their statistical properties are aimed to be similar to those of the Universe (at the scale of interest). The fact that cosmological structures (e.g., galaxies and clusters of galaxies) are homogeneously distributed in the computational domain renders this problem more easily balanced for the workload. Once the cosmological parameters are set, a cosmological box is characterized by two quantities:

- i the physical size of the volume it represents, which is commonly expressed as the edge length of the box in Megaparsec (Mpc);
- ii the number N_p of particles used to sample the distribution of matter within the volume. The last number is commonly referred to as $N_p = 2 \times N_G^3$, where N_G is the grid number used to generate the particles in the initial conditions, and the factor $\times 2$ descends from the fact that we include both ordinary baryonic and dark matter particles.

The physical size of the box affects the number and size of sub-structures, such as galaxies and clusters of galaxies, which are simulated. In general, it can be said that the larger the volume, the larger the number of structures and the maximum size of those sub-structures, e.g. galaxy clusters. Meanwhile, the number of particles affects the mass resolution; in other words, the larger the number of particles used, the higher the accuracy in following the dynamics and describing the internal structures of collapsed objects.

For this study, we conduct simulations of a set of boxes whose parameters are detailed in Table 1.

Figure 1 shows a projection of the gaseous and stellar distribution, coloured by density, at redshift 0 (i.e. after ~ 13 billion years of evolution). In this simulation, the volume remains constant, while the accuracy and resolution increase with the number of particles N_p . As the number of particles grows, the computational cost of simulating the same region becomes larger, too. This increasing computation costs do not scale linearly due

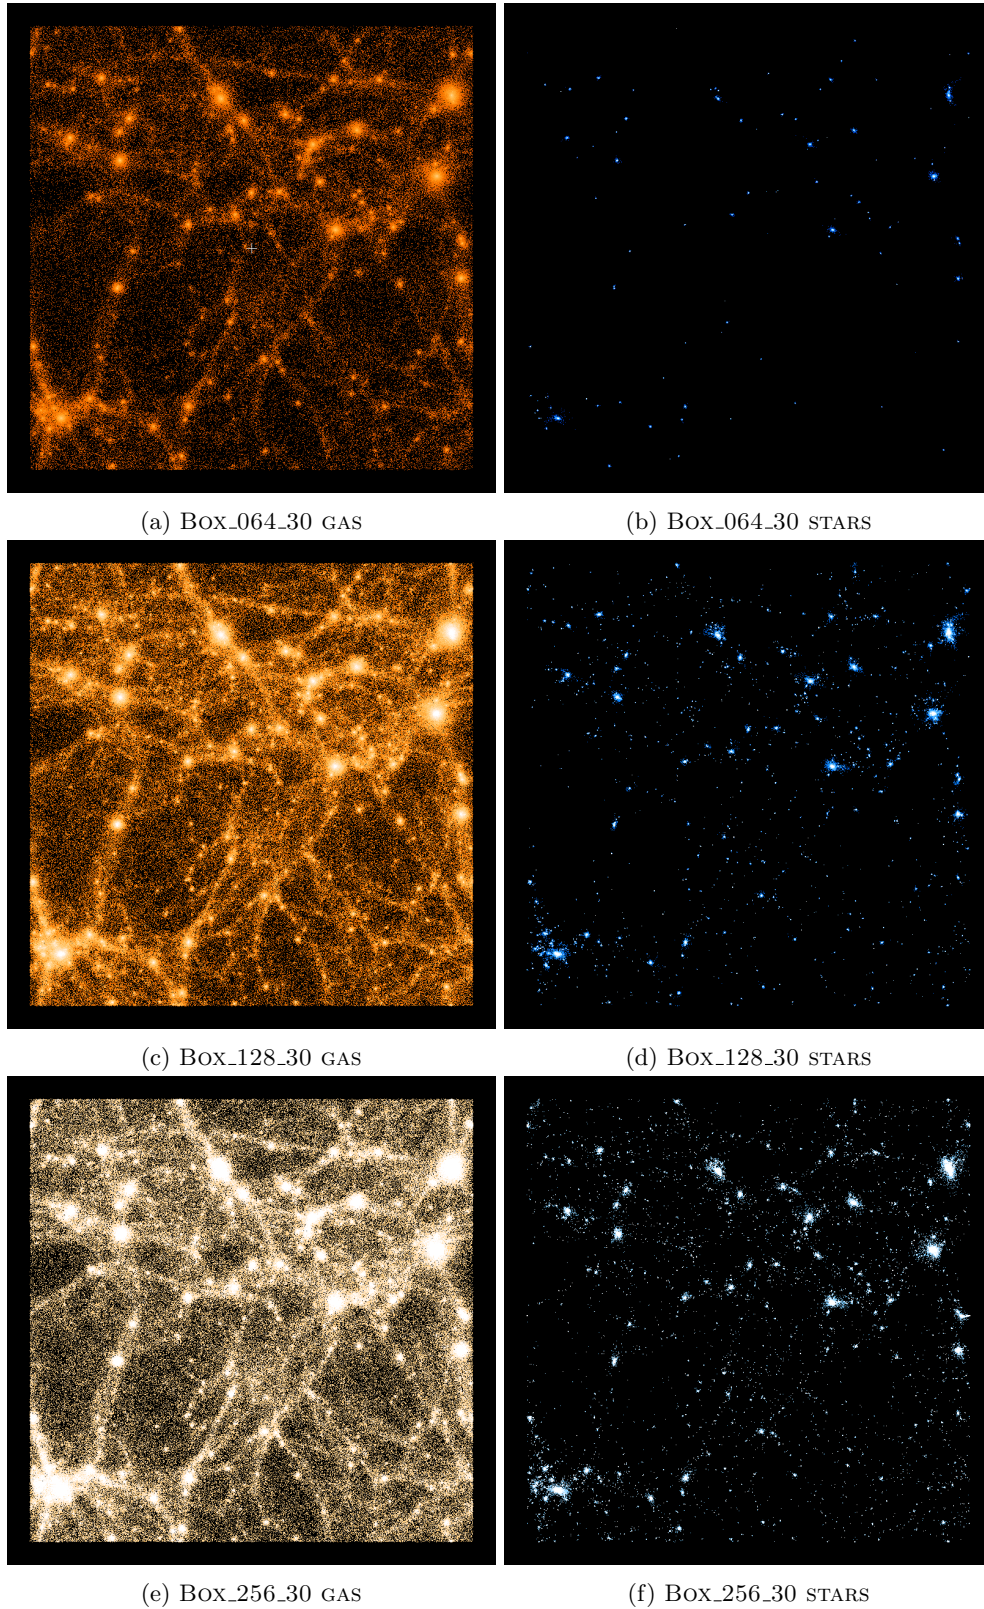


Figure 1: Projection of gas (left column) and stellar (right column) distribution at $z \sim 0$ (i.e. at present time, after ~ 13 billion years of evolution) in 30 Mpc boxes. Each row shows the simulation of the same box resolved with a different number N_p of particles: 64^3 , 128^3 and 256^3 in the top, medium and bottom rows. It can be seen that the level of details increases from top to bottom with increasing N_p .

Box size/edge length / N_p	64^3	128^3	256^3	800^3
30 Mpc	Box_064_30	Box_128_30	Box_256_30	–
60 Mpc	–	Box_128_60	–	–
120 Mpc	–	–	Box_256_120	–
375 Mpc	–	–	–	Box_800_375

Table 1: The set of cosmological boxes generated for the profiling activity along the project. The boxes are referred to as BOX_ N_p _SIZE where the "size" refers to the side length of the box.

to the costs related to calculating the gravitational tree, which scales with the number of particles as $N \log N$. Increasing the number of particles does not only affect the calculation of the gravitational tree but also the frequency of its calculation. Using more particles results in a more accurate description of forces between particles; consequently, the length of the individual time steps has to be reduced. This results in a more expensive calculation as a larger number of time steps has to be calculated to simulate the same global time. It has been found empirically – since it is challenging to draw theoretical expectations for such a complex code – that an increase in the number of particles by a factor of 2 results in a ~ 12 times longer run time. Hence, we expect that the computational cost of BOX_128_30 and BOX_256_30 are $\sim \times 12$ and $\sim \times 144$ larger than for BOX_64_30.

With the suite of boxes defined above, we can conduct the following scaling studies:

1. **Strong Scaling:** the BOX_800_375 (2×800^3 initial particles) case is suitable to run up on a significant number of nodes before incurring severe parallel overhead while keeping the run-time limited as it is required for benchmarking and profiling tasks.

The BOX_256_30 (2×256^3 initial particles) case may also be used for the same purpose, although on a smaller number of maximum nodes or, even better, to test the scaling on a single node by changing the amount of OpenMP threads per MPI task.

2. **Weak Scaling:** comparing the boxes with side 30, 60 and 120 Mpc and N_p equal to 64^3 , 128^3 and 256^3 respectively, we exploit the equal mass resolution because the volume sampled by each particle remains constant while increasing the size N_p of the problem; distributing the work on a scaled amount of nodes, we can study the weak scaling behaviour of the code.
3. **"Physical Scaling":** comparing the runs for the boxes having 30Mpc size and N_p 64^3 , 128^3 and 256^3 we incur in an increase of the computational cost due both to the larger N_p and to the more accurate dynamics, as explained above; then we can test and profile the code's behaviour under such (quite common) conditions.

It is also important to note that the same scaling comparison should be conducted at different cosmic epochs as the clustering of the matter is enormously different at different redshifts, becoming much more significant at present times compared to earlier times. That translates into two primary effects: (i) the average distance of particles in collapsed structures is smaller. This results in larger forces, which requires a more accurate integration and, therefore, smaller time steps; (ii) the gravitational tree is, on average, much more profound, and the tree-related operations become more costly.

2.2 Case 2 - Zoom-in simulations

While simulating large cosmological boxes provides large number statistics, they are limited in resolution due to the computational costs as described previously. This provides a severe restriction in the study of individual objects, e.g. galaxy clusters, for which much higher resolution is required than possible in large boxes. An alternative is to run idealized simulations of these objects. While this would provide the ability to run the code with higher resolution [7] – since it is focused only on the object of interest itself – it will lack the information from the cosmological environment and are therefore only a rough proxy to their real counterparts. One solution for this problem are so-called zoom-in simulations.

For a zoom-in simulation only a small, carefully selected volume, containing the object of interest, inside a much larger parent simulation (e.g. a dark-matter-only cosmological box) is modelled at high resolution, usually including baryons. Most of the volume is filled with low-resolution boundary particles interacting only through gravity. The purpose of these particles is the creation of appropriate tidal fields and large-scale modes in the




Figure 2: Zoom-in simulation of a galaxy cluster at four different time steps from upper left to lower right. Here it can be seen how the area necessary for the simulation of the cluster (see central area) is significantly smaller than a cosmological box, allowing the simulations to be calculated with higher resolution.

high-resolution region (see, e.g. [8, 9, 10, 11, 12, 13] for further information) for the central area, filled with high-resolution DM and baryonic particles.

Zoom-in simulations therefore allow the study of physical processes at much smaller scales thus providing deeper insights into the formation and evolution of individual objects and the physical processes involved. Figure 2 shows the high-resolution dark matter distribution for a zoom-in simulation of a galaxy cluster. Here, it can be seen very clearly how the simulated zoom-in volume is limited to the particles influencing the evolution of the central cluster, by which a significant amount of computational time is saved compared to the calculation of an entire box (see Figure 1 for comparison) containing the same object.

In contrast to the cosmological boxes, which focus on the physical processes on larger scales, zoom-in simulations can profile the modules correlated to the physical processes on smaller scales. These modules, e.g. such as star formation and black hole feedback – named extra physics within OpenGadget – become more and more important with higher resolution. Thus zoom-in simulations provide the perfect complement to profile all modules and kernels of the code.

The zoom-in simulations we use are based on the DIANOGA simulation (e.g. [14]), which is a $1h^{-3}Gpc^3$ dark-matter only cosmological box containing 1024^3 particles with a mass resolution of $M_{particle} \approx 6.62 \times 10^{10} h^{-1} M_{\odot}$ for each dark matter particle.

From the parent simulation, one volume containing a galaxy cluster is chosen. We then zoom in on this volume, by not only increasing the resolution but also adding baryonic physics. We have chosen to use two different resolutions to be able to test the scalability of the zoom-in simulation at larger systems for which the low-resolution case will be too small.

The details and mass resolutions of the two specified cases are described in Table. 2.

Name/Resolution	$M_{dm}[h^{-1}M_{\odot}]$	$M_{gas}[h^{-1}M_{\odot}]$	estimated total CPUh
Parent simulation	6.62×10^{10}	–	–
HighResolution	$\sim 1 \times 10^9$	$\sim 1 \times 10^8$	60
UltraHighResolution	$\sim 4 \times 10^7$	$\sim 6 \times 10^6$	800

Table 2: The set of zoom-in simulations used for profiling and benchmarking. Here are details about two zoom-in simulations of the same object at different resolutions and consequently different mass per particle. For comparison with the parent simulation the mass for each dark matter particle M_{dm} is displayed as well as the mass for each gas particle M_{gas} .

2.3 Summary

Combined, the cosmological boxes and the zoom-in simulations provide a test suite capable of giving us insight to identify and improve potential bottlenecks within the different modules and kernels of OpenGadget, as well as testing the basic modules (e.g. gravity and hydrodynamics) and the more detailed physical processes (e.g. star formation, black holes, etc.) for their scalability and performance.

Both of the scientific cases perform the same type of calculations and hence will be jointly used to test our progress in terms of IPC and vectorization. However, we have selected these two cases instead of just one, as they stress the domain decomposition (see D1.2 for a detailed description), the workload balance and the communication patterns very differently.

The matter distribution in the cosmological boxes is intrinsically more homogeneous across the volume, with local deepening of the potential well and, hence, of the clustering. The single-object simulations, conversely, are intrinsically much more unbalanced in the matter distribution and, above all, the local clustering of the particles is much more pronounced relative to the average distribution than in cosmological boxes (at the same mass resolution). For this reason, achieving a good workload balance is much more challenging in the second case. Moreover, the communication surface is much larger when the computational domain is fragmented into a high number of segments of the Peano curve (that is required to enhance the attainable workload balance). Furthermore, deeper potential wells determine higher accelerations and then a deeper hierarchy of individual particle timesteps that further stresses the code capabilities.

The two selected cases group all the challenges for OpenGADGET on which we will focus in the optimization activity. We want to add that we will possibly borrow the galaxy encounter case adopted by the other Lagrangian code, ChaNGA (see Sec. 5.1): this case, which falls in the "single-object" class, brings a more extreme local clustering with deeper potential and then a more violent dynamic would develop that requires more frequent domain adjustment/decomposition and a more pronounced hierarchy of particles time steps, exaggerating more the traits of the galaxy clusters case.

3 PLUTO

PLUTO [15, 16], is a modular code for plasma astrophysics and it is based on a finite-volume formalism in which conservation laws are solved by means of Godunov-type methods. At the general level, PLUTO is designed to solve a general system of conservation laws, written as

$$\frac{\partial U}{\partial t} = -\nabla \cdot \mathcal{T}(U) + S(U), \quad (1)$$

where U represents a state vector of conserved quantities (typically density, momentum and energy), $\mathcal{T}(U)$ is the flux tensor (rank 2) and $S(U)$ is a source term.

The explicit definition of \mathcal{T} , S and U depends on the particular physics module that the user selects. PLUTO comes with 4 different modules, HD (Euler equations of gasdynamics), Magneto-Hydrodynamics (MHD) (ideal and resistive magnetohydrodynamics) and their corresponding special relativistic extensions (see, e.g., [17, 18, 19]). Each module is equipped with a number of different algorithms allowing the user to choose between a number of options, ranging from schemes with more numerical diffusion (such as the Lax-Friedrichs Riemann solver with linear reconstruction) to less diffusive schemes (such as Roe / Harten-Lax-van Leer-Contact (HLLC) or Harten-Lax-van Leer-Discontinuity (HLLD) solvers with PPM or WENO reconstruction). The code can run on both static and adaptive grids, the latter functionality being provided by the CHOMBO library, see [16]. While the latter is certainly a valuable addition, the current GPU implementation plans will not include the CHOMBO library in the code as there are joint efforts - outside of this project - to develop an independent, built-in library for PLUTO.

The revised upcoming version of the code will incorporate both 2nd- and 4th-order methods using limited reconstruction - linear, Weighted Essentially Non-Oscillatory (WENO) and Monotonicity Preserving (MP) - while Runge Kutta (RK) schemes are employed to advance the solution vector in time. The divergence-free condition of magnetic fields is enforced using either a hyperbolic divergence cleaning approach (the Generalized Lagrangian Multiplier, GLM) or the Constrained Transport (CT) method, in which the magnetic field retains a staggered representation.

Within the SPACE project we foresee the employment of the GPU version of the code to make significant progresses in the following selected scientific cases:

1. Resistive Relativistic Reconnection. By means of 2D high resolutions simulations using both constant and effective resistivity models, we propose to improve our knowledge of the energy dissipation process and isolate the role of spurious numerical diffusion. This context is particularly relevant in the field of relativistic particle acceleration.
2. Plasma Column Instability. MHD instabilities are thought to play a central role in the process of energy dissipation in relativistic jets. Particle acceleration may occur, in fact, in localized dissipation regions (current sheets, turbulence, shocks) formed at smaller scales as byproduct of the instability.
3. Active Galactic Nucleus (AGN) Jets. Large-scale magnetized flows in the form of jets are commonly observed in AGN, Pulsar Wind Nebula (PWN), Gamma Ray Burst (GRB), and microquasars. The emitted radiation is the only imprint for our understanding. Numerical simulations must be endowed with advanced tools allowing computed models to be compared with observations. To this purpose we plan to incorporate the Lagrangian Particle module of the PLUTO code [20] in order to model the non-thermal radiation from relativistic electrons.

3.1 Case 1 - Resistive Relativistic Reconnection

Scientific rationale - Understanding the dynamics of relativistic plasmas is of paramount importance for the interpretation of high-energy astrophysical sources such as AGNs, GRBs, PWN, and magnetars, just to cite a few. In these systems a significant fraction of the bulk kinetic energy is dissipated by means of plasma instabilities, thus becoming available at microscopic scales to accelerate particles to relativistic velocities in localized regions of space, typically much smaller than the overall system scale. At small scales, particle acceleration occurs either at shocks, magnetic reconnection sites or via turbulence originated as byproducts of plasma instabilities. Among these acceleration mechanisms, much attention has been devoted to relativistic magnetic reconnection, which occurs in regions with thin and strong electric currents where a finite resistivity η dominates the dynamics of the plasma. This causes magnetic fields of opposite polarities to rapidly annihilate

in correspondence of so-called X-points, so that part of the magnetic energy is dissipated through heat, creating strong electric fields that can efficiently accelerate particles [21].

High- S regime - For a current sheet of length L , the typical diffusive time scale L^2/η is generally much larger than the Alfvénic crossing time L/v_A , which translates in very high values for the system’s Lundquist number $S = Lv_A/\eta$ (e.g. $S \sim 10^{12}$ for the plasma in the Solar corona) due to the generally low values of magnetic resistivity η . One of the challenges of the numerical modeling of relativistic reconnection within the MHD framework is to ensure that the numerical dissipation introduced by the discretization errors of the numerical scheme is much smaller than the physical resistivity. This is a necessary condition to obtain converged results that are independent of the grid resolution and code employed to produce the models (see Fig. 3), which determines the current limitations in simulating resistive Relativistic MHD (RMHD) models with $S \gtrsim 10^6$. Thanks to the GPU-accelerated routines of the renewed module of the PLUTO code and its genuinely 4th-order methods [22], we will access values of S more than an order of magnitude higher than current RMHD models. The expected overall speed-up of $\sim 15x$ will allow to increase the grid resolution by a factor of ~ 4 without any increase in the simulation runtime, while employing higher order algorithms will further decrease the numerical dissipation, thus enhancing the ”effective” resolution of the numerical models.

As initial conditions we will use a stationary Harris-like current sheet defined over a 2D Cartesian box $[0, L] \times [-20a, 20a]$, where L is the current sheet’s length and a its thickness. The initial magnetic field is set to

$$\mathbf{B} = B_0 \left[\tanh\left(\frac{x}{a}\right) \mathbf{e}_y + \zeta \operatorname{sech}\left(\frac{x}{a}\right) \mathbf{e}_z \right], \quad (2)$$

with B_0 being the upstream magnetic field strength and $\zeta \in [0, 1]$ a free parameter that selects between a thermal pressure-balanced configuration ($\zeta = 0$) and a purely force-free one ($\zeta = 1$). Once we set the upstream values of rest mass density ρ_0 , cold magnetization $\sigma_0 = B_0^2/\rho_0$, and ratio between thermal and magnetic pressure $\beta_0 = 2p_0/B_0^2$, we compute the thermal pressure requiring the balance $p + B^2/2 = \text{const}$, i.e.

$$p(y) = \beta_0 \frac{B_0^2}{2} + \frac{(1 - \zeta^2)}{2} \operatorname{sech}^2\left(\frac{x}{a}\right). \quad (3)$$

Finally, we will use a constant density $\rho = \rho_0$ everywhere and set to zero both velocity \mathbf{v} and electric field \mathbf{E} . In order to trigger the tearing instability, we will seed a magnetic perturbation of wavenumber $\mathbf{k} = (0, 2\pi/L, 0)$ and amplitude $\epsilon = 0.01$ as

$$\delta\mathbf{B} = \epsilon B_0 \operatorname{sech}\left(\frac{x}{a}\right) \begin{pmatrix} \cos(ky) \\ \frac{\sin(ky)}{ka} \tanh\left(\frac{x}{a}\right) \\ 0 \end{pmatrix}. \quad (4)$$

We will consider a thickness of the current sheet equal to $a = LS^{-1/3}$ (which is known to trigger the growth of the tearing mode on ideal - i.e. Alfvénic - time scales; see [23]) and probe three different values of Lundquist number: $S = \{10^6, 10^7, 10^8\}$. These high values of S will not only require higher computational efficiency (due to the shortened characteristic time-scales) but also stress the algorithms used for the time integration. This is due to the fact that in the resistive RMHD framework the electric field cannot be simply computed from Ohm’s law, but has to be promoted to fully-fledged independent variable whose evolution is determined by Ampères




Figure 3: Spatially-averaged transverse component of magnetic field as a function of time for different resolutions and selected numerical schemes, with $S = 10^6$ (from [1]).

law in its complete form (namely, including the displacement current that is usually neglected in the classical resistive MHD case). The presence of stiff source terms for the electric field forces to adopt a specific strategy to mitigate the strong limitations to the time-stepping [24], which can become unstable for very low values of the magnetic resistivity. In our tests we will use high-order IMEX (IMplicit-EXplicit) Runge-Kutta schemes that allow to maintain a similar time step with respect to the ideal case, and we will test their stability in a state-of-the-art magnetic reconnection production run with very high S . Besides different values of S , we will also characterize the reconnection dynamics across different initial equilibria $\zeta = \{0, 0.5, 1\}$ and combinations of magnetizations and pressure ratios $(\sigma_0, \beta_0) = \{(1, 1), (5, 0.1), (50, 0.01)\}$. This will allow us to explore the stability of the high-order IMEX schemes in different relativistic regimes, since the corresponding upstream Alfvén speeds will be $v_{A,0} = (1/\sigma_0 + 2\beta_0 + 1)^{1/2} = \{0.5, 0.85, 0.98\}$ in units of the speed of light c . The resistivity parameter will be constant across the numerical domain and set to $\eta = Lv_{A,0}/S$, and will therefore range from $\sim 0.5 \times 10^{-8} \sim 0.98 \times 10^{-6}$. To ensure that the numerical dissipation is negligible with respect to η , we will first perform test runs at the fiducial resolution of $N_x \times N_y = 1920 \times 768$ grid points, therefore setting a uniform grid spacing of $\Delta x = \Delta y \simeq 5.2 \times 10^{-4}$. We will then assess the amount of resolution required to obtain converged results and the impact of the high-order schemes within PLUTO by halving the grid spacing until the growth rate of the tearing mode is well resolved (see Fig. 3).

Non-collisional resistivity - The dissipation of magnetic energy during a reconnection event differs, in principle, from a collisional resistivity, which is expected to be extremely low in astrophysical plasmas.

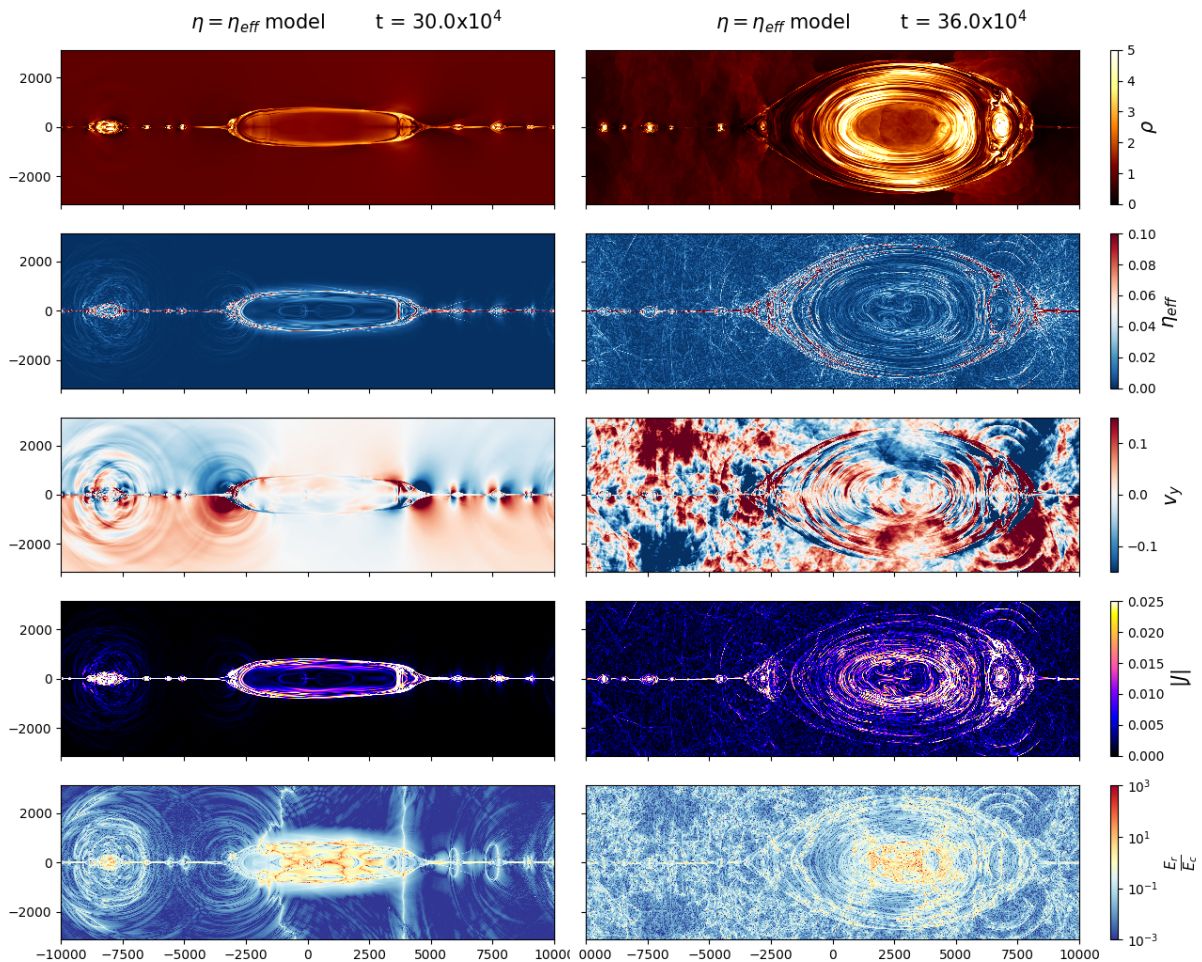


Figure 4: From top to bottom: 2D snapshots of mass density, effective resistivity, transverse velocity, current density, and fraction of resistive electric field in a reconnecting current sheet with effective resistivity at $t = 3 \times 10^5 \omega_p$ (left column) and $t = 3.6 \times 10^5 \omega_p$ (right), where ω_p is the plasma frequency. Length are in code units (jet radius).

However, resistive RMHD models are based on the over-simplifying assumption of a constant resistivity inherited from a collisional description and thus not suitable for describing collisionless plasmas typically encountered in astrophysical scenarios. Particle-In-Cell (PIC) simulations, on the other hand, provide a fully-kinetic description of the coupling between charged particles and electromagnetic fields, thus self-consistently reproducing the dissipation of magnetic fields. Recent PIC simulations [25] have shown that the dominant contribution to the magnetic dissipation within a reconnecting current sheet can be modeled as an *effective resistivity* that depends on the local properties of the plasma (such as its Lorentz factor and transverse velocity gradients). This allows to introduce non-collisional effects in the RMHD framework and model the magnetic dissipation more self-consistently, as the resistivity in the pair plasma can be expressed as

$$\eta_{eff} = \frac{m}{\rho e c} \sqrt{\left(\frac{mc}{e} \frac{\partial v_y}{\partial y}\right)^2 + E_z^2}, \quad (5)$$

where ρ is the plasma density, m is the electron mass, e is the elementary electric charge, v_y is the velocity component perpendicular to the current sheet, and E_z is the total electric field in the plasma co-moving frame. By expressing the resistivity in adimensional code units we obtain

$$\bar{\eta}_{eff} = \frac{\delta_u}{L_u} \frac{4\pi}{\bar{\rho}} \sqrt{\frac{\delta_u^2}{L_u^2} \left(\frac{\partial \bar{v}_y}{\partial \bar{y}}\right)^2 + \bar{E}_z^2}, \quad (6)$$

where δ_u is the plasma skin depth, L_u is the characteristic length unit and the barred quantities are expressed in code units. This means that Eq. 5 introduces a new characteristic length scale in the RMHD system, i.e. the plasma skin depth, whose size with respect to the system has to be specified a priori. In our case we set $\delta_u = 0.02L_u$ (which is consistent with standard PIC simulations), with the resulting η ranging from $\sim 10^{-8}$ to $\sim 10^{-1}$. Using the standard resistive RMHD module of the PLUTO code we tested such prescription and obtained models where the resistivity increases significantly during the onset of the plasmoid instability and then subsides once the saturation leads to a quenching of the velocity gradients (Fig. 4).

As for the high- S case, we will exploit the renewed resistive RMHD module of the PLUTO code to access a new dynamical regime that will introduce non-collisional effects. We will adopt the same initial conditions for the magnetic field and thermal pressure described by Eq. 2-4 but setting $\zeta = 0$ (thus a thermal pressure-balanced equilibrium), $\sigma_0 = 10$, and $\beta_0 = 0.01$. We will set, in this case, a constant temperature $\Theta = p/\rho = 1$ across the domain (rather than a constant density) and then compute ρ accordingly, which is a common choice in PIC models. The structure of current sheets is expected to become increasingly complex during the plasmoid instability, which forces one to locally identify the geometry of the current sheet in order to compute the velocity gradients in Eq. 5. Assessing the best numerical strategy to complete this stage will be crucial, since it can significantly increase the computational cost of the scheme. Moreover, the effective resistivity can drop by a few orders of magnitude during the late stages of the reconnection event, forcing to considerably increase the grid resolution in order to properly resolve the magnetic dissipative scale. We will therefore estimate what is the optimal resolution and resistivity floor value that allows properly capture the non-collisional effects introduced with the effective resistivity.

To fully explore the quantitative agreement between fully kinetic PIC models and our RMHD approach with effective resistivity we will use the Lagrangian particle module of PLUTO to estimate the non-thermal spectrum of the electrons. Our choice of magnetization and pressure ratio allows for a strongly relativistic regime (since $v_{A,0} \simeq 0.945$) but also a relatively lower cutoff to the maximum Lorentz factor that accelerated particles are expected to gain from the reconnection event (since $\Gamma_{\max} \sim 4\sigma_0$), which reduces the requirements in terms of box sizes. In fact, to properly resolve Γ_{\max} one must ensure that the size of the numerical box is sufficiently big to follow the acceleration of particles with large Larmor radius, which has been shown to translate in the condition $L/\sigma_0\rho_0 \gtrsim 40$ [26]. To capture this crucial diagnostic, after assessing the necessary resolution to properly resolve the reconnection event we will extend the numerical box described for the high- S case by a factor ~ 4 , while retaining the same grid spacing.

3.2 Case 2 - Plasma Column Instability

The aforementioned high-energy extragalactic sources are all characterized by the emission of highly collimated relativistic outflows of plasma (namely, jets) up to Mega-parsecs. Particle acceleration occurs at small spatial scales in localized dissipation regions (LDRs) forming in the jets as byproducts of plasma instabilities such as

the Current-Driven Instability (CDI). The CDI is related to the topological configuration of magnetic fields in the plasma column and it plays a key role in the conversion of the initial magnetic energy to particle’s kinetic energy. This instability is believed to occur in jets, a result already confirmed by numerical simulations using both the classical and relativistic MHD equations.

Relativistic MHD codes are more appropriate to capture the large-scale jet dynamics and explain how LDRs originate. However, the algorithm complexities and the huge resolution required to simulate such processes make them extremely computationally expensive with conventional numerical techniques relying on 2nd-order numerical schemes and traditional CPU paradigms. 4th-order methods provide more efficient strategies to perform simulations converging to the same accuracy but at much faster rates, offering cost-effective computations as the grid resolution N_{grid} increases (CPU-time gain: $\Delta t^{(4th)}/\Delta t^{(2nd)} \sim N^{-2}$).

This scientific case aims at investigating the properties of the CDI instability by means of a novel class of numerical simulations based on the renewed version of the PLUTO code that uses genuinely 4th-order-order numerical methods. These aim at providing i) better resolution of the initial plasma column equilibrium and of the linear growing modes of the instability, ii) reduced numerical dissipation introduced by the scheme and iii) quenching of spurious modes which may corrupt physical growing modes of the instability during its linear phase. In addition, the employment of 4th-order method is expected to provide dramatic efficiency gain over traditional 2nd-order computations since, for given accuracy error, lower resolutions ($\lesssim 1/2$) are needed. This translates into faster computations while retaining similar errors.

The first practical challenge chosen to assess the GPU-accelerated version of PLUTO, utilizing the directive-based programming model OpenACC, has focused on the evolution of a non-rotating stationary force-free plasma column, perturbed by a small radial velocity disturbance. It is considered the Mizuno equilibrium [27], defined by the magnetic fields (in cylindrical coordinates):

$$B_r = 0 \tag{7}$$

$$B_z = \frac{B_0}{[1 + (r/a)^2]^\alpha} \tag{8}$$

$$B_\phi = \frac{B_0}{(r/a)[1 + (r/a)^2]^\alpha} \sqrt{\frac{[1 + (r/a)^2]^{2\alpha} - 1 - 2\alpha(r/a)^2}{2\alpha - 1}}, \tag{9}$$

where:

- B_0 parametrizes the magnetic field amplitude.
- α is called *pitch profile parameter*. It determines the radial profile of the magnetic pitch $P = rB_z/B_\phi$ and provides a measure of the twist of the magnetic field lines [27]. It was set $\alpha = 1$ so that the magnetic pitch is constant.
- a is the typical radius of the plasma column and was set $a = 0.4$, in units of L_0 , an arbitrary reference length. Moreover, the choice of $\alpha = 1$ leads to $a = P$.
- $B_r = 0$ in order to respect the solenoidal condition.

The density profile can be chosen arbitrarily, as it doesn’t appear in the radial balance equation, and here is set:

$$\rho = \rho_0 \left[\eta + \frac{1 - \eta}{\cosh(r^6)} \right], \tag{10}$$

where $\rho_0 = 1$ is the on-axis density and $\eta = \rho_{r=\infty}/\rho_0$ is the density contrast. The pressure is constant everywhere and defined in accord to the latter equation:

$$p = \frac{1}{2}\beta B_0^2, \tag{11}$$

where $\beta = 8\pi p/B^2$ is the beta plasma parameter. As a consequence of the Alfvén theorem, if $\beta \ll 1$, the plasma is highly magnetized and the magnetic field dominates and it drags the matter of the plasma, while if $\beta \gg 1$, the plasma is poorly magnetized and the plasma drags the magnetic field. The same problem was studied using three different values of β : $\beta = 5.0$ (weakly magnetized), $\beta = 1.0$ (intermediate case) and $\beta = 0.2$ (strongly magnetized). For each value, two different spatial resolution were used, 192^3 and 384^3 number of points of the computational domain, for a total of six simulations. The domain is the same in all cases: $x \in [-10, 10]$,

$y \in [-10, 10]$, $z \in [0, 20]$. For each simulation, the equilibrium configuration of Eq. (7) was initialized and a small perturbation in the radial velocity was introduced:

$$\delta u_r = \epsilon u_A r e^{-r^4} f, \quad (12)$$

The algorithm employs RK3 as the time marching algorithm, parabolic reconstruction, the HLLD Riemann solver and outflow boundary conditions. The divergence cleaning method is used to control the divergence-free constraint.

Simulations, ranging from 32 to 512 processors, were conducted on CINECA's Marconi100, revealing GPU acceleration factors of approximately 7.5 compared to CPU runs. Fig. 5 and 6 show the plasma column approximately at the beginning and at the end of the linear phase (in the 384^3 points resolution case). In these figures, the helicoidal deformation, typical of the kink instability, can be seen clearly, especially in the strong magnetized case on the right. In Fig. 5 and 6 we were able to distinguish the material initially inside the column and the external ambient medium introducing the variable T , called tracer. The tracer obeys the following equation:

$$\frac{(\partial \rho T)}{\partial t} + \nabla \cdot (\rho \mathbf{u} T) = 0, \quad (13)$$

and is initially defined as:

$$\begin{cases} T = 1 & \text{if } r < a \\ T = 0 & \text{if } r > a. \end{cases} \quad (14)$$




Figure 5: Plasma column at the time $t = 50$, in the high resolution cases.




Figure 6: Plasma column at the time $t = 75$, in the high resolution cases.

In the next set of computations, we plan to incorporate the constrained transport method to enforce the solenoidal constraint of magnetic field, upgrade to 4th accuracy and explore more advanced configurations, namely, those in which the jet moves relativistically and carries different magnetic field topologies (e.g. toroidal

/ poloidal fields). The foreseen set of simulations will also consider longer computation domains and, possibly, open jets (that is, non-periodic configurations).

3.3 Case 3 - Synthetic Spectra from 3D Relativistic Jets

Scientific rationale Magnetized and relativistic large-scale flows in the form of jets are a common observational feature seen for example in active galactic nuclei (AGNs), gamma-ray bursts, and microquasars. Multi-wavelength observations covering a wide spectrum from radio wavelengths to TeV gamma-ray emission witness that the dominant emission is of non-thermal nature. The observed radiation has been interpreted as a results of the interaction between high-energy electrons and the in-situ magnetic and radiation fields. These electrons must be energised inside the jet cocoon after ejection from the central nucleus, being accelerated at strong shocks via Diffusive Shock Acceleration (DSA).

The length scales associated with these micro-physical processes are many orders of magnitude smaller than the physical jet scales, which can range up to few tens of kiloparsecs. Connecting a bridge between these scales poses a serious challenge to theoretical modeling of the emission from AGN jets. The aim of the present study case is to build a quantitative connection between such disjoint scales exploiting numerical tools that can simulate multi-dimensional flow patterns while treating small-scale processes in a sub-grid manner.

In [28] a detailed analytical and numerical framework was developed to evolve a distribution of non-thermal particles both spatially and in momentum space. This was introduced as the Lagrangian Particle (LP) module in the PLUTO code. The non-thermal relativistic electrons at a region in space are modelled as ensembles of macro-particles (also LPs) and are advected spatially following the fluid motions. Their energy spectrum is found by evolving their phase space distribution function, duly accounting for energy losses due to radiative processes such as synchrotron emission and inverse-Compton interaction with CMB photons. Moreover, when a LP crosses a DSA, the energetic spectrum is updated in function of the shock strength. In figure 7 a snapshot of a simulation made in [20] alongside the spectral evolution of one particle across the simulation are reported.

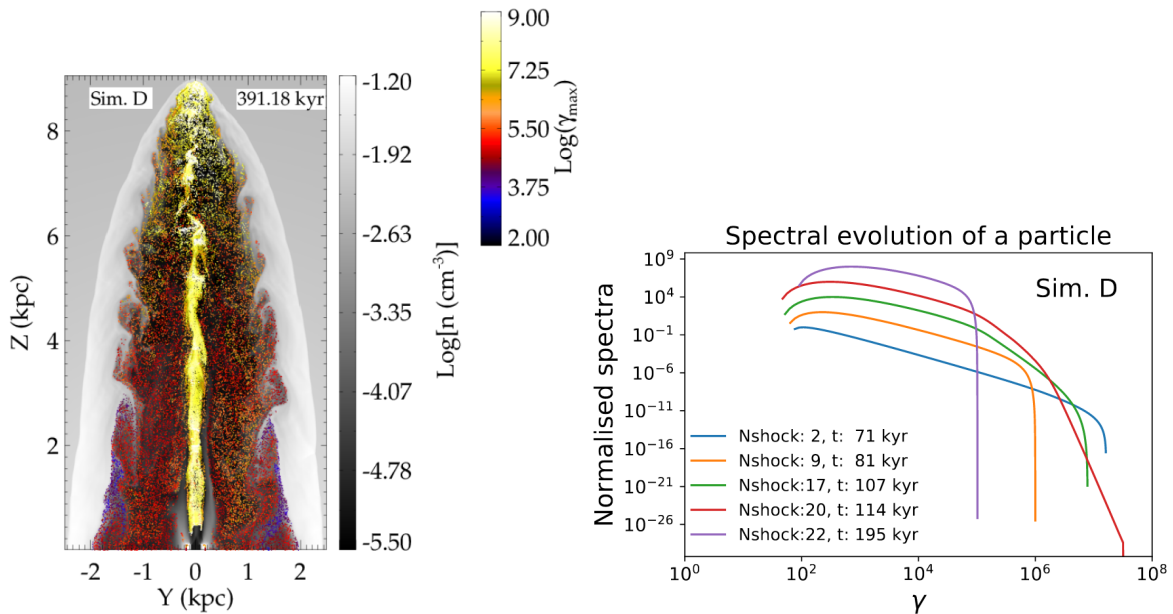


Figure 7: Left: in gray scale the fluid density of the jet, in scale of colours are represented the LPs, the colour is proportional to the maximum energy of the LP spectral distribution. Right: the simulated energetic spectrum of one particle at different times.

Expected simulation cases. Within the present study, we plan to use and improve the module to provide a reliable and efficient code to study the evolution of non-thermal electrons in relativistic jets. The project fits in the SPACE program since it aims to meet the exascale paradigm. This will be done exploiting algorithm vectorization, allowing for parallel GPU execution. From extensive internal testing, we forecast a speedup of

~ 20 times with respect to a standard CPU code. Larger and more resolved simulations can be conducted using the same amount of computational hours and energy consumption of a comparable CPU run.

To efficiently capture the physics of the expanding jet a computation domain of (approximately) $\sim 1500 \times 2500 \times 1500$ grid points is needed, covering a physical domain of $\sim 18 \times 30 \times 18$ kpc. With a resolution of 10 points on the jet radius, this amounts to a domain of size $\sim 150 \times 250 \times 150$ jet radii. The high resolution simulation, coupled with LP module, allows for the creation of high resolution non-thermal emission, polarisation and rotation maps which can be readily compared with radio observations of such systems. Figure 8 shows a synthetic emission map obtained with the LP module. We will employ the 2nd-order version of the code, using outflow boundary conditions on all sides with the exception of the jet inlet, where constant flow parameters are prescribed inside the $r < r_j$ (nozzle radius). The jet is lighter than the ambient medium, by a factor of 10^{-3} and its velocity is set to $v_j = 0.99$ (in units of c).

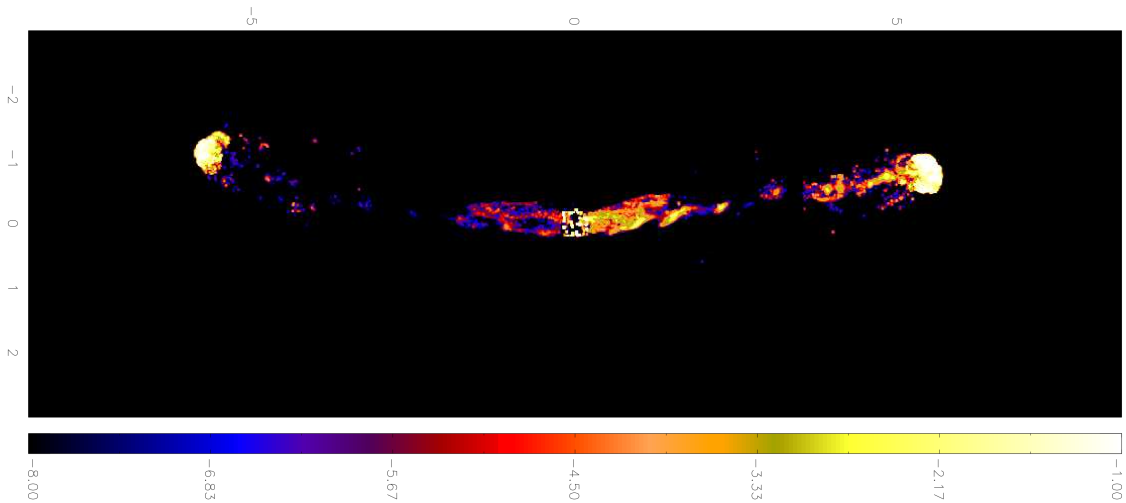


Figure 8: 2keV emission of a jet simulated using the LP module (dimensionless units normalized to the maximum).

The project will explore different jet powers ($P_j \sim 10^{44} - 10^{46} \text{ergs}^{-1}$) and magnetisations ($\sigma_B \sim 0.1 - 0.01$). The jet power P_j is the total enthalpy flux (without the rest mass energy) integrated over the injection surface, including the contribution of the magnetic field. For a flow with a total enthalpy $w_t = \rho_j h_j + B^2/(\gamma^2 4\pi) + (\mathbf{v} \cdot \mathbf{B})^2/(4\pi)$, the enthalpy flux per unit area along the z axis, excluding the rest mass energy, is [19]

$$F_z^T = (\gamma^2 \rho_j h_j - \gamma \rho_j c^2) v_z + \frac{B^2}{4\pi} v_z - (\mathbf{v} \cdot \mathbf{B}) \frac{B_z}{4\pi}. \quad (15)$$

Here ρ_j is the gas density inside the jet, h_j is the gas enthalpy and γ is the Lorentz factor of the gas injected. Different injection pressures will be also studied. Based on the previous experience of similar calculations but on smaller scales and for lower jet powers (e.g. [29, 30]), we estimate the simulations to take about 1×10^6 CPU hours, or $\sim 50 \times 10^3$ GPU hours. The memory required for fluid variables for the proposed grid size is $1500 \times 2500 \times 1500 \times 8(\text{bytes}) \times 12(\text{variables}) \times 6(\text{arrays}) = 3.02$ TB. In addition to the fluid variables, we will inject particles with the jet at the point of injection at each time step. The memory load of particles thus depends on the time evolution of the jet. However, we can compute an approximate upper limit to the memory requirement if we consider that the jet plasma to spread over the entire domain at late stages such that each computational cell has one particle. Each particle requires 128 bytes of memory. Thus the total memory in that case would be $1500 \times 2500 \times 1500 \times 8(\text{bytes}) = 720$ GB. This implies a maximum total memory requirement of 2.34 TB. A data snapshot of the simulation takes about $1500 \times 2500 \times 1500 \times 8(\text{bytes}) = 22.5$ GB and double the value for restart files. We plan to save 10 restarts and about ~ 100 snapshots which makes a total of 3 TB per case.

4 BHAC

BHAC [31, 32, 33, 34] is a multidimensional General Relativistic Magnetohydrodynamics (GRMHD) code that is mainly used to study accretion flows onto compact objects. BHAC has been designed to solve the GRMHD equations in arbitrary (stationary) space-times/coordinates and exploits adaptive mesh refinement techniques with an oct-tree block-based approach provided by the MPI-AMRVAC framework [35, 36, 37, 38]. The code is second-order accurate and uses finite volume and high-resolution shock-capturing (HRSC) methods. Various schemes for the treatment of the magnetic field update have been implemented on ordinary and staggered grids. Originally designed to study black hole (BH) accretion in ideal GRMHD, BHAC has been extended to incorporate nuclear equations of state, neutrino leakage, charged and purely geodesic test particles, and non-black hole fully numerical metrics. In addition, a non-ideal resistive GRMHD module has been developed and implemented. BHAC has been employed in a number of studies of accretion into supermassive black holes [39, 40] and other compact objects [41]. In addition, BHAC's results, after a general-relativistic ray-tracing (GRRT) post-processing, can be used to compute synthetic observable images of BH shadows and the surrounding accretion flows. These calculations are performed with the GRRT Black Hole Observations in Stationary Spacetimes (BHOSS) code [42]. The GRMHD simulation data produced by BHAC are used as an input for BHOSS to produce accretion flow and BH shadow images like the one shown in Figure 9 [31].




Figure 9: GRRT of 3D torus with increasing resolution. Snapshot images of 3D GRMHD simulation data with parameters chosen to mimic the emission from Sgr A*. The resolution of the simulation data is indicated in the bottom-right corner of each panel.

As BHAC is designed to simulate accretion flows onto compact objects, we chose to discuss the following application.

In what follows we will adopt a system where lengths are expressed in units of GM/c^2 (where M is the mass of the central object) and velocities in units of the speed of light c . Time is expressed in units of GM/c^3 . For convenience, we will also set $G = c = 1$ henceforth.




Figure 10: Evolution of the 2D magnetised torus with resolution 1024×512 for times $t/M \in \{300, 1000, 2000\}$. We show logarithmic rest-frame density (top), logarithmic plasma β (middle) and the logarithm of the magnetisation parameter $\sigma = b^2/\rho$ (bottom). Magnetic field lines are traced out in the first panel using black contour lines. One can clearly make out the development of the MRI and evacuation of a strongly magnetised funnel reaching values of $\beta < 10^{-5}$ and $\sigma = 103$.

4.1 Case 1 - Magnetised torus around a spinning black hole

The Fishbone-Moncrief equilibrium torus [43] with a weak magnetic field around a spinning BH constitutes the main GRMHD target application of BHAC. The GRRT post-processing of such kind of simulations leads to images like in Figure 9. Supermassive BHs of this type, like Sgr A* and M87, which are the primary targets of observations by the Event-Horizon-Telescope are accreting well below the Eddington accretion rate. In this regime, the accretion flow advects most of the viscously released energy into the BH rather than radiating it to infinity. Such optically thin, radiatively inefficient and geometrically thick flows can be modelled without radiation feedback. Modern BH accretion-disk theory suggests that angular momentum transport is due to MHD turbulence driven by the magnetorotational instability (MRI) within a differentially rotating disk. Recent non-radiative GRMHD simulations of BH accretion systems have resolved these processes and reveal a flow structure that can be decomposed into a disk, a corona, a disk-wind and a highly magnetized polar funnel. The simulations show complex time-dependent behaviour in the disk, corona and wind. Depending on BH spin, the polar regions of the flow contain a nearly force-free, Poynting-flux-dominated jet.

The magnetised torus around a spinning BH is a generalisation of the steady-state solution of the standard hydrodynamical equilibrium torus with constant angular momentum to MHD equilibria with a weak poloidal magnetic field. This weak single magnetic field loop is defined by the vector potential

$$A_\phi \propto \max(\rho/\rho_{max} - 0.2, 0),$$

and is added to the stationary Fishbone-Moncrief solution, where ρ_{max} is the maximum initial density. The spinning BH is a Kerr BH with dimensionless spin parameter $a = 0.9375$. Given the BH mass M and spin and the fluid equation of state, the solution is completely defined by specifying the radius of the inner edge of the torus and the radius of the maximum of the pressure. To break this equilibrium state, random perturbations are added to the pressure. This eventually triggers the MRI, allowing the plasma to accrete. It is worth mentioning that even without explicitly adding a perturbation, numerical errors produced by the discretisation would be amplified to produce a turbulent state very similar to the one obtained with the perturbation.




Figure 11: Fluid-frame density (top) and $\log_{10} b^2$ (bottom) for $t = 3000M$ on the $y = 0$ plane (left) and the $z = 0$ plane (right) in the 3D magnetised torus run with resolution $384 \times 192 \times 192$.

Specifically, we will use initial conditions describing a torus around a Kerr BH with dimensionless spin parameter $a = 0.9375$. The inner radius of the torus is set to $r_{in} = 6M$ and the density maximum is located at $r_{max} = 12M$, where radial and azimuthal positions refer to Boyer-Lindquist coordinates. With these choices,

the orbital period of the torus at the density maximum becomes $T = 247M$. We adopt an ideal gas EoS with an adiabatic index of $\gamma = 4/3$. The field strength is set such that $2p_{\max}/b_{\max}^2 = 100$, where global maxima of pressure p_{\max} and magnetic field strength b_{\max}^2 do not necessarily coincide. In order to excite the MRI inside the torus, the thermal pressure is perturbed by 4% white noise. As with any fluid code, vacuum regions must be avoided and hence we apply floor values for the rest-mass density ($\rho_{\text{fl}} = 10^{-5}r^{-3/2}$) and the gas pressure ($p_{\text{fl}} = 1/3 \times 10^{-7}r^{-5/2}$). In practice, for all cells which satisfy $\rho \leq \rho_{\text{fl}}$ we set $\rho = \rho_{\text{fl}}$, in addition if $p \leq p_{\text{fl}}$, we set $p = p_{\text{fl}}$. The simulations are performed using horizon penetrating logarithmic Kerr-Schild (KS) coordinates. In the 2D case, the simulation domain covers $r_{\text{KS}} \in [0.96r_h, 2500M]$ and $\theta \in [0, \pi]$, where $r_h = 1.35M$. In the 3D case, we slightly excise the axial region $\theta \in [0.02\pi, 0.98\pi]$ and adopt $\phi \in [0, 2\pi]$. We set the boundary conditions in the horizon at $r = 2500M$ to zero gradient in primitive variables. The θ -boundary is handled as follows: when the domain extends all the way to the poles (as in our 2D cases), we adopt ‘hard’ boundary conditions, thus setting the flux through the pole manually to zero. For the excised cone in the 3D cases, we use reflecting ‘soft’ boundary conditions on primitive variables. The time-update is performed with a two-step predictor corrector based on the total variation diminishing Lax-Friedrichs fluxes and PPM reconstruction. Furthermore, we set the CFL number to 0.4 and use the flux-interpolated constrained transport algorithm [44] as a divergence control method. Typically, the runs are stopped after an evolution for $t = 5000M$, ensuring that no signal from the outflow boundaries can disturb the inner regions.

Figure 10 illustrates the qualitative 2D time evolution of the torus by means of the rest-frame density ρ , plasma- β and the magnetisation $\sigma = b^2/\rho$, where b is the magnetic field strength. After $t \approx 300M$, the MRI-driven turbulence leads to accretion onto the central BH. The accretion rate and magnetic flux threading the BH then quickly saturate into a quasi-stationary state. The accreted magnetic flux fills the polar regions and gives rise to a strongly magnetised funnel with densities and pressures dropping to their floor values. Notice that in the deliverable [45] the 2D case is used as use-case.

Since self-sustaining dynamo activity leading to the perpetuation of the MRI cannot occur in strict azimuthal symmetry, 3D simulations are necessary to study the accretion flow in the saturated state. Therefore, we turn to the 3D runs performed with BHAC. The qualitative evolution of a high resolution run is illustrated in Figure 11 showing rest-frame density ρ and magnetic field strength b^2 in the two slices $z = 0$ and $y = 0$. Overall, the evolution progresses in a similar manner to the 2D cases: MRI-driven accretion starts at $t \approx 300M$ and enters saturation at around $t \approx 1000M$. Similar values for the magnetisation in the funnel region are also obtained. However, since the MRI cannot be sustained in axisymmetry as poloidal field cannot be re-generated via the ideal MHD induction equation, we expect to see qualitative differences between the 2D and 3D cases at late times.

4.2 Summary

With the choice of the aforescribed simulation, the performance of the basic modules of BHAC can be thoroughly tested for the main target applications of BHAC. The magnetised torus enable us to test the conservative to primitive recovery module which is the most resource demanding module of BHAC. Thus, potential shortcomings and bottlenecks of the code can be rigorously spotted and improved.

5 ChaNGa

ChaNGa [46][47][48] is an N-body and smoothed particle magneto-hydrodynamics (SPMHD) code which is used to study a wide array of astrophysical systems. While the gravity and SPMHD algorithms are based on the gasoline [49] and pkdgrav [50] codes, the unique addition to ChaNGa is its implementation of the Charm++ framework, which enables highly efficient parallel scaling. The central feature of ChaNGa is a tree-based gravity solver using a variant of the Barnes-Hut algorithm. This solver is combined with several other features which include: gravitational interaction between dark matter, stars and gas, Ewald summation to handle cosmological boundary conditions, multi-time stepping, hydrodynamics, magnetic field and several subgrid physics such as star formation, feedback from supernova, stellar winds, black holes and the radiative cooling of gas. We can summarize these into 4 major parts: gravity, Ewald/boundary, SPMHD and subgrid physics. We have five scientific cases for ChaNGa, each focusing or expanding on a particular subset of physics modules. These include:

- Case 1 **Isolated galaxy mergers:** as the name suggest here we investigate mergers between galaxies and study their evolution and impact on galactic properties. These are done in isolated environments and as such there are no cosmological boundary conditions in this case. This case includes gravity, magnetohydrodynamics, and subgrid physics. See Section 5.1 for more info.
- Case 2 **Magnetized cloud collapse:** here we study the collapse of a rotating magnetized cloud, modeling the early stage of the star formation process. This case involves a large span of physical, spatial and temporal scales and proves to be a challenging test for the numerical MHD model. It includes gravity and magnetohydrodynamics but no subgrid physics or special boundary conditions. See Section 5.2 for more info.
- Case 3 **Sedov-Taylor blast wave:** the primary phase of a supernova explosion follows what is known as the Sedov-Taylor blast wave solution. This phase can be modelled by injecting an immense amount of energy at a point in a homogeneous medium of a given density. This is a purely hydrodynamic simulation and is an excellent case to test energy conservation, gradient accuracy and scaling of the hydrodynamics. See Section 5.3 for more info.
- Case 4 **Cosmological boxes:** cosmological boxes are cubic regions of space that are used to simulate the evolution of a patch of the Universe through cosmic time. The main focus in these simulations is usually to capture the large-scale structure and statistical properties of the Universe. Cosmological periodic boundary conditions are required to model effect of the rest of the Universe on the cosmological box. The relevant physics for this kind of system can be immense, but heavily depends on what one seeks to study. For capturing the general large-scale structure of the Universe modeling gravity is enough, but for accurately modeling structures like galaxies and galaxy clusters more physics is required. This allows us to explore patches of the Universe with different physics "enabled", from pure gravity simulations to simulations with a plethora of physics modules. Computationally these simulations can be very challenging and involve a large span of physical, spatial and temporal scales. See Section 5.4 for more info.
- Case 5 **Zoom-in simulations:** while the cosmological boxes mainly focus on the large-scale, zoom-in simulations focus on resolving individual objects (galaxy, galaxy cluster, etc.) in a cosmological environment. These simulations use a parent simulation (a cosmological box), to pick out regions of interest. These regions are then re-simulated at much higher resolution together with the surrounding low resolution version of the parent simulation. Thereby the name "zoom-in" simulation. The benefit of these simulations is the modeling of individual objects in a realistic cosmic environment, compared to the isolated environment in case 1. This case poses one of the most formidable challenges for both the computational and the numerical modeling aspects of simulation codes. In addition to the vast amount of physical processes involved, it also includes the largest span of physical, spatial and temporal scales. See Section 5.5 for more info.

5.1 Case 1 - Isolated Galaxy Mergers

Galaxy mergers are very important for the evolution of galaxies and play a crucial role in shaping their structure, dynamics, and star formation activity. They also play a significant role in the hierarchical structure formation of the Universe, where small structures merge to form larger ones over time. They occur when two or more




Figure 12: **Top:** rendering of the magnetic field strength for a simulated Milky Way-type galaxy (mass $M = 10^{12}M_{\odot}$), with a gas mass resolution of around $m = 10^2M_{\odot}$. The highest magnetic field strength is in dark blue; moderate strength is in white and the weakest magnetic field is represented as red. **Bottom:** shows the timeline of a major merger between two lower resolution versions of the top figure. We can see that during the merger there is significant magnetic field amplification.

galaxies gravitationally interact and eventually collide and coalesce into a single larger galaxy. For this scientific case we are particularly interested in the evolution of magnetic fields during galaxy mergers.

The importance of magnetic fields in galaxy formation is clear from observations of the Milky Way and nearby galaxies, which reveal that the magnetic energy is in equipartition with the thermal and turbulent energies [51, 52]. This means that they are likely to have a large dynamical effect on the evolution of the galaxy, adding significant non-thermal pressure to the system. This additional non-thermal pressure can suppress star-formation rates [53] and can heavily affect the structure of the interstellar medium (ISM). In addition, it has been shown that magnetic fields have a strong impact on fluid instabilities [54, 55], which may affect how gas in the intergalactic medium (IGM) accretes onto galaxies and how gas in galactic outflows leaves (or cycles back to) galaxies. The strength and structure of magnetic fields in galaxies also determine the transport of cosmic rays (CRs), which have recently emerged as a promising candidate for driving galactic outflows because of their long cooling times [56, 57, 58].

Galaxy mergers are relevant for magnetic field amplification due to the significant increase in star formation rates and the injection of large-scale turbulence. This can enhance the small-scale magnetic fields produced

through the small-scale dynamo [59, 60, 61, 62] and affect the mean-field dynamo processes relevant to the production of large-scale fields [63, 64, 65].

For this case we intend to perform simulations of a major merger between two Milky-Way type galaxies. These galaxies are initially identical. The initial galaxy represents an evolved version of the isolated disk galaxies from the AGORA code comparison project, which ChaNGa, RAMSES, and Gadget-3 all participated in (more detail about how the ICs are generated can be seen in Section 8.3). The evolved version has been run with feedback, cooling, star formation and magnetic field turned on, and has reached a stable state, featuring both a CGM and a fully saturated magnetic field in the disk [66]. The initial disk structure can be seen at the top of Figure 12, and in the bottom we can see a timeline of the initial impact between the two galaxies (impact parameter $b = \sin\theta_{\text{impact}} = 0.1$), using our lowest resolution case. The resolution of our merger setup can be seen in Table 3. As mentioned in the introduction, this simulation includes a plethora of physics (gravity,

	N_{tot}	$m_{\text{gas}}[M_{\odot}]$	$m_{\text{star}}[M_{\odot}]$	$m_{\text{dark}}[M_{\odot}]$	$\epsilon_{\text{gas}}[kpc]$
Low	$5 \cdot 10^6$	$1.0 \cdot 10^4$	$1.0 \cdot 10^4$	$1.6 \cdot 10^6$	0.02
Medium	$4.0 \cdot 10^7$	$1.3 \cdot 10^3$	$1.3 \cdot 10^3$	$2.0 \cdot 10^5$	0.01
High	$3.2 \cdot 10^8$	$1.7 \cdot 10^2$	$1.7 \cdot 10^2$	$2.5 \cdot 10^4$	0.005

Table 3: Resolution parameters of the three initial conditions used in our simulations. Stating the number of particles (N_{tot}), initial particle masses ($m_{\text{gas}}, m_{\text{star}}, m_{\text{dark}}$) and the softening length of the gas (ϵ_{gas}). m_{star} refers to the mass of star particles that form during the simulation. In addition, there is an old stellar disk component in the ICs, which consists of star particles that have masses of around $m_{\text{star,old}} \approx 4.3 m_{\text{star}}$

hydrodynamics, magnetic fields, star formation, feedback, radiative cooling) and thereby a plethora of physical, numerical and computational challenges. There is a wide range of spatial and temporal scales in this case. Most of the resolution elements will be in the thin galactic disk and the spatial and temporal scales can shift swiftly; a gas particle in the disk, affected by feedback from a nearby star, may rapidly transition into an outflow. This process transports it from the highly resolved disk to the less resolved circumgalactic medium (CGM). This all makes it a challenge to load balance and scale well. For this scientific case both strong and weak scaling can be conducted.

5.2 Case 2 - Magnetized cloud collapse

The collapse of a rotating magnetized cloud models the early stage of the star formation process. Here, a rotating magnetized cloud is allowed to collapse under its own gravity. During the collapse, the cloud is compressed over several orders of magnitude, testing how the magnetic field develops and interacts with the gas during compression. The large-scale collapse is eventually halted by the formation of a pseudo-disk, which then starts to slowly contract via magnetic braking. The collapse continues within the central region and as the first hydrostatic core starts to form, the magnetic field is twisted until it eventually launches a jet. The formation and subsequent evolution of the first hydrostatic core stalls the collapse and a slow contraction phase begins. The two main jet launching mechanisms are the magneto-centrifugal and the magnetic pressure driven mechanism. With a global poloidal magnetic field, both of these mechanisms play an important role. The resulting magnetic field structure of the jet consists of a poloidal dominated central core with a surrounding toroidal field which produces a strong current along the jet. All these key aspects require the code to have excellent angular momentum conservation, small numerical dissipation and maintain low divergence errors ($\nabla \cdot B$).

This case involves a large span of physical, spatial and temporal scales and proves to be a challenging test for the numerical MHD model. It includes gravity and magnetohydrodynamics but no subgrid physics or special boundary conditions. For this scientific case both strong and weak scaling tests can be conducted. The initial conditions are described in more detail in [66]. The cloud resolutions that we will investigate here is $(50^3, 100^3, 250^3, 500^3)$, employing a mass-to-flux ratio of $\mu = 10$. The mass-to-flux ratio μ is relative to the critical mass-to-flux ratio, $(M_c/\Phi)_{\text{crit}} = \frac{M_c}{\pi R_c^2 B_{0,\text{crit}}}$, in which the cloud is fully supported by magnetic forces against gravity. Here Φ is the magnetic flux which threads the cloud and M_c is the mass of the cloud.

5.3 Case 3 - Sedov-Taylor blast wave

The Sedov-Taylor blast wave is a solution to the problem of a strong explosion occurring at a point in an ambient medium and is one of the most fundamental problems in hydrodynamics and astrophysics. It is particularly

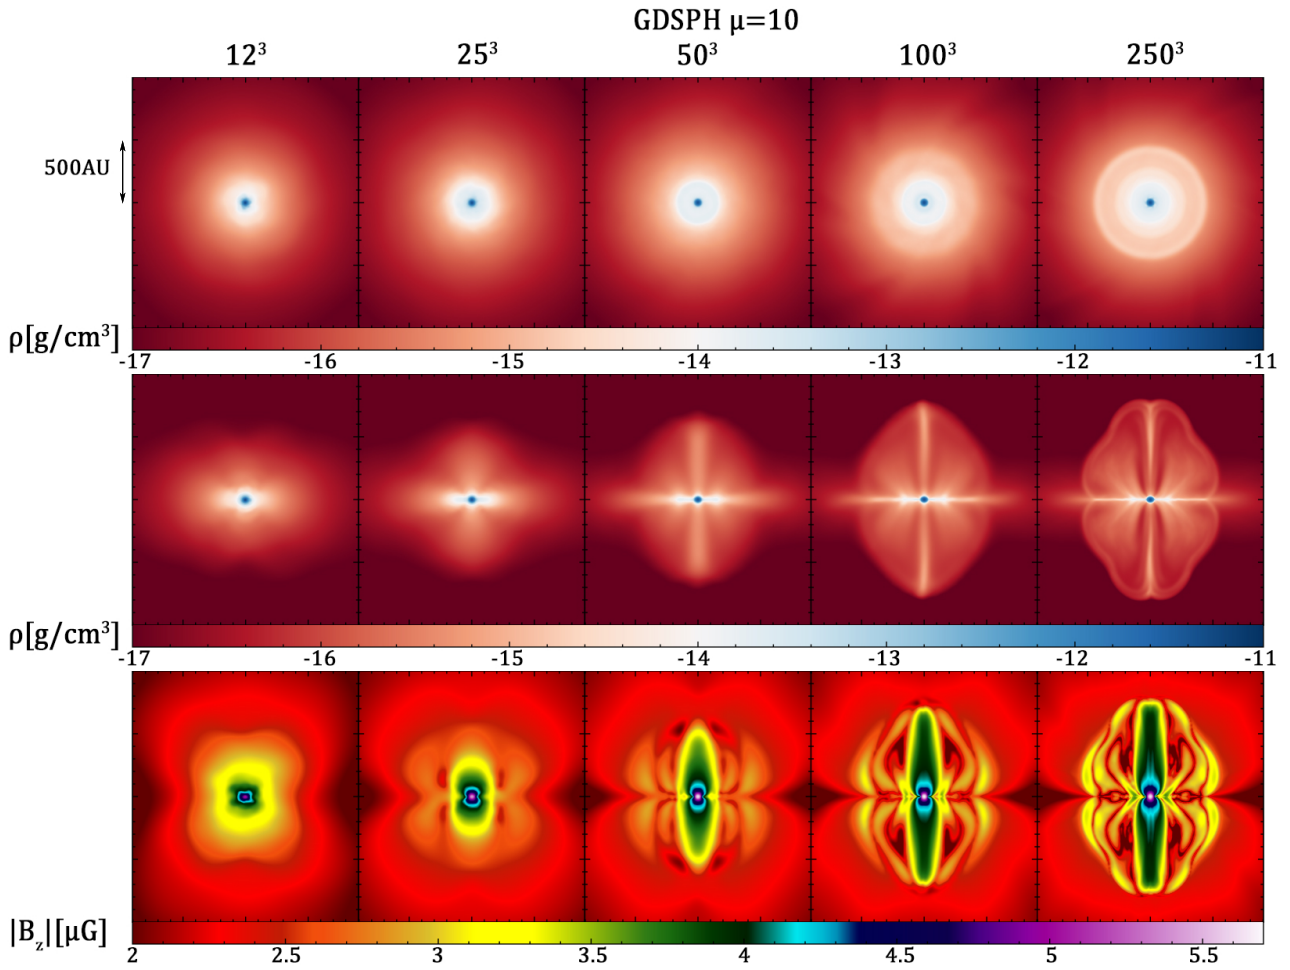


Figure 13: Resolution study of the magnetized cloud collapse with a mass-flux ratio of $\mu = 10$ from [2]. Here we vary the resolution from left to right, in the initial cloud ($12^3, 25^3, 50^3, 100^3, 250^3$). Renderings show the time of jet formation (around the free fall time of the cloud), which occurs due to the winding of the magnetic field during collapse, which produces a magnetic tower structure. The top row shows a rendered face-on slice ($L_{xy} = [2000AU, 2000AU]$) of the density [g/cm^3], the rest of the rows show rendered slices through the rotation axis ($L_{xz} = [2000AU, 2000AU]$), where the second shows density [g/cm^3], and the third shows the absolute poloidal magnetic field [μG]. All quantities are shown in logarithmic scale. For the scientific cases in this project we will perform an even higher resolution version (500^3) and not include the lowest resolution ones ($12^3, 25^3$).

relevant in understanding supernova explosions.

The problem assumes that an immense amount of energy is suddenly released at a point in a homogeneous medium of density ρ . This results in the formation of a blast wave that propagates radially outward from the point of the explosion. There is a self-similar solution to this problem, meaning the shape of the blast wave is preserved at different times, except for a scaling factor. The radius and velocity of the shock front at any time after the explosion can be described by the following scaling law:

$$R(t) = \left(\frac{Et^2}{\rho} \right)^{1/5},$$

$$u_s(t) = \left(\frac{Et^2}{\rho} \right)^{1/5}.$$

Here E is the energy released, t is the time after the energy has been released and ρ is the ambient density. The Sedov-Taylor blast wave is modelled as an "infinite" Mach shock where the upstream pressure/sound speed

is negligible compared to the internal blast pressure. In an ideal gas, the density jump across a shock front is determined by the Rankine-Hugoniot conditions, which describe the relationship between the pre-shock and post-shock states. The compression ratio with an adiabatic index $\gamma = 5/3$ then becomes:

$$\frac{\rho_2}{\rho_1} = \frac{\gamma - 1}{\gamma + 1} = 4$$

The Rankine-Hugoniot conditions also give us the velocity of the flow behind the shock, which can be related to the speed of the shock front and the compression ratio of the gas:

$$u_2 = u_s \left(1 - \frac{\rho_1}{\rho_2}\right) = \frac{3}{4} u_s$$

This is a purely hydrodynamic simulation and is an excellent case to test energy conservation, gradient accuracy and scaling of the hydrodynamics module. For this scientific case both strong and weak scaling can be conducted. For the initial conditions we setup a uniform box with a homogeneous density ($\rho = 1$) and inject $E = 10$ thermal energy to a few particles in the centre of the box, while the rest of the particles have negligible thermal energy. The resolutions we will perform in this case range from $N = 64^3$ to $N = 2048^3$.

5.4 Case 4 - Cosmological boxes

Cosmological boxes are cubic regions of space that are used to simulate the evolution of a patch of the Universe through cosmic time. The main focus in these simulations are usually to capture the large-scale structure and statistical properties of the Universe. Cosmological boundary conditions are required to model effect of the rest of the Universe on the cosmological box. The relevant physics for this kind of system can be immense, but heavily depends on what one seeks to study. For capturing the general large-scale structure of the Universe modeling gravity is enough, but for accurately modeling structures like galaxies and galaxy clusters more physics is required. This allows us to explore patches of the Universe with different physics "enabled", from pure gravity simulations to simulations with a plethora of physics modules. Computationally these simulations can be very challenging and involve a large span of physical, spatial and temporal scales. A more deep and descriptive explanation of the numerical aspects can be read in Section 2.1.

The initial conditions of this case will be the same as outlined in Section 2.1. This allow for strong, weak and physical scaling to be conducted. Additionally, our study incorporates a classic ChaNGa benchmark test case, representing the final state of a dark matter cosmological simulation encompassing a volume of 70 Mpc [67]. This comes in two resolutions ($N = (80 \cdot 10^6, 640 \cdot 10^6)$). This case is quite well load balanced and therefore aims to benchmark the other parts of the gravity solver.

5.5 Case 5 - Zoom-in simulations

While the cosmological boxes mainly focuses on the large-scale, zoom-in simulations focuses on resolving individual objects (galaxy, galaxy clusters, etc.) in a cosmological environment. These simulations use a parent simulation (a cosmological box), to pick out regions of interest. These regions are then re-simulated at much higher resolution together with the surrounding low resolution version of the parent simulation. Thereby the name "zoom-in" simulation. The benefit of these simulations is the modeling of individual objects in a realistic cosmic environment, compared to the isolated environment in Section Case 1. This case poses one of the most formidable challenges for both the computational and the numerical modeling aspects of simulation codes. In addition to the vast amount of physical processes involved, they also include the largest span of physical, spatial and temporal scales. A more deep and descriptive explanation of the numerical aspects can be read in Section 2.2.

The initial setup for this case is as detailed in Section 2.2. In addition, this case also includes a well-established ChaNGa benchmark test case, which is a zoom-in simulation of a dwarf galaxy with a halo mass of approximately $M_{halo} \approx 10^{11} M_{\odot}$. This simulation is a dark matter-only version of the DWF1 model, studied in [68], that demonstrates how disk galaxies can form in a cosmological context. This comes in two resolutions ($N = (50 \cdot 10^6, 400 \cdot 10^6)$). This case allows for both strong and weak scaling to be performed.

6 FIL

The detection of gravitational waves (GW) in 2015 ushered in a new era of astronomy allowing a window into some of the most extreme and violent events in the Universe: black hole and neutron star collisions. The mergers of these compact objects release huge amounts of energy both as electromagnetic (EM) and gravitational radiation; GW15094, the first GW event detected by LIGO released more energy during the micro seconds of the collision than the combined energy output of all the stars in the observable Universe [69] over the same time period. Modeling these systems accurately is a central problem in theoretical astrophysics that has proven incredibly challenging, requiring the development of state of the art numerical relativity codes and large amounts of computing time. The numerical codes need to simultaneously solve Einstein’s equations and the coupled general relativistic magneto-hydrodynamic (GRMHD) equations for magnetized fluids. The GRMHD equations tell the fluids how to move according to the spacetime curvature and the field equations tell the spacetime how to curve based on the mass distribution of the fluids.

Whilst the equations of General Relativity (GR) were written down in 1917 it took until 1999 – 80 years – for the first simulation of binary neutron star collisions and until 2005 for the first simulation of binary black hole mergers [69]. In 2014 the Illinois numerical relativity (ILNR) group released IllinoisGRMHD, an open source user friendly numerical relativity code that uses second order finite difference methods to approximate solutions to the GRMHD equations [70] and simulate compact object collisions. Researchers at Goethe University under Elias Most built upon the work done by the ILNR group and in 2019 developed the Frankfurt Illinois GRMHD code (FIL) [71]. FIL uses fourth order methods to approximate the GRMHD equations and incorporates tabulated equations of state (EOS), allowing users to define their own EOS. This is significant as the EOS of neutron stars is an open question. A visual difference between FILs fourth order method compared to IllinoisGRMHDs second order method can be seen in Figure 14. Constraining the EOS of neutron stars has applications beyond astrophysics and is also of interest to nuclear physicists. The extreme energies, temperatures and densities inside neutron stars are not accessible to us within laboratories and so finding out how matter behaves in these extreme environments will further our collective understanding of nuclear physics[72].

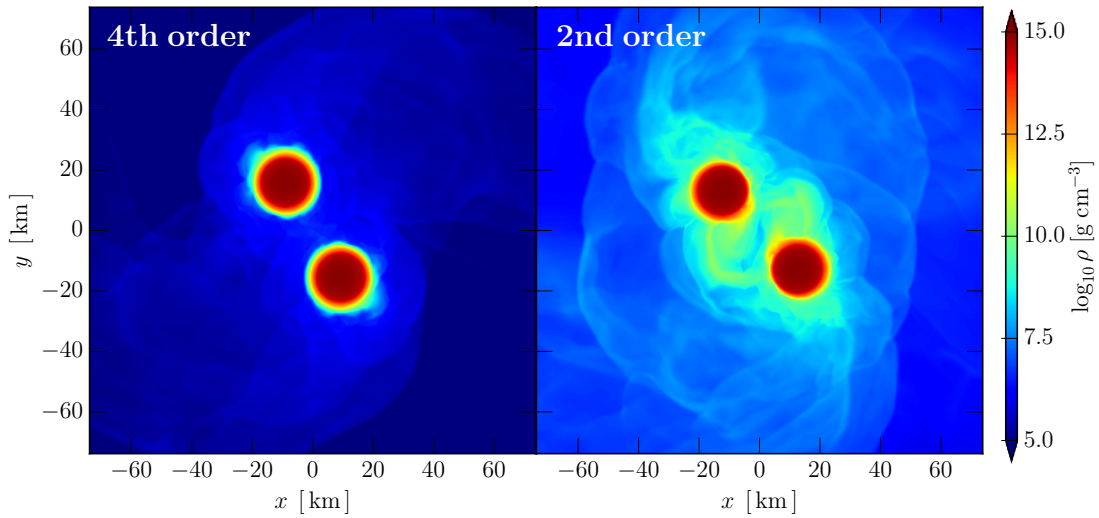


Figure 14: Comparison of equal mass binary neutron star merger simulations computed with the second order IllinoisGRMHD code (right) and with the fourth order FIL code (left). On the y axis is log density ρ , the image is of a snapshot of the in spiral. The Fourth order FIL is able to capture sharp features better.

FIL is coupled with the Einstein toolkit (ET), a publicly available evolution framework that is designed, maintained and extended to enable numerical relativity simulations. FIL is a module of the ET and works in conjunction with other modules such as Carpet, responsible for the AMR grid simulations take place on; Antelope, the spacetime evolution module and the method of lines module which provides the numerical time integrators. FIL and the ET are able to accurately model mergers of binary compact objects and produce the GW waveforms key to GW astronomy. FIL is primarily used to model binary neutron star (BNS) and black hole neutron star (BHNS) collisions. Observations of BNS and BHNS GW and their EM counterparts (multi messenger astronomy) will allow astronomers and theoretical astrophysicists to measure the Hubble constant,

determine the mechanism behind short gamma ray bursts (GRB) and give insights into the origin of heavy elements in the Universe produced through r-process nucleosynthesis as well as develop our understanding of nuclear physics. Observational GW astronomy is only possible with the theoretical understanding that numerical codes like FIL provide.

6.1 Case 1 - Head on Collision

FIL has two important use cases that can be leveraged to ensure the numerical algorithms are properly stressed and to check the accuracy of the numerical results. The first use case is that of head-on collisions where two compact objects either start at rest or have some initial linear velocity directed at the companion. The second and most relevant to multi-messenger astronomy is the in-spiral, merger, and post-merger of two neutron stars. Head on collisions are not expected to be found in nature the head on case contains all the same physics as an in-spiral whilst having the advantage of running in a reasonable amount of computing time. Head on collision's also stresses one of FIL's most time consuming algorithm the conservative to primitive solver, the solver takes longer to converge to a solution in strong magnetic field regimes which are present in the head on case. To test the numerical algorithms the head on collision case is preferable.

To construct and evolve the head on collision system the Einsteins field equation needs to be reformulated. Einstein's field equation is given as

$$G^{\mu\nu} = 8\pi T^{\mu\nu},$$

where G is the Einstein tensor, T is the stress-energy tensor and μ and ν are indices between 0-3 for the four dimension of spacetime, 0 being the time dimension and 1-3 being the spatial dimensions. The Einstein tensor describes the spacetime and the stress-energy tensor describes the matter and radiation within the spacetime [69]. There are multiple reformulations that can be used to rewrite the field equations, the most common being the BSSNOK, CCZ4 and the Z4 [69, 70, 72]. Each framework follows the same basic outline; Einstein's field equation is decomposed using a 3+1 decomposition in which the 4 dimensional spacetime is split into 3 dimensional space hyper-surfaces, described by the spatial metric, γ_{ij} , which is embedded in the full spacetime, specified by the extrinsic curvature K_{ij} and the gauge functions α (lapse) and β^i specify a coordinate frame (see Figure 15). Depending on the framework different gauge conditions, further decomposition and certain




Figure 15: Foliation of the spacetime manifold. The lapse function α measures the proper time along the normal vector n^a to the hyper-surface σ_t which is equipped with an induced metric γ_{ij} . The shift vector β^i measures the displacement, on consecutive hyper-surfaces, between the observer time lines t_a and the normal lines n_a [3].

constraints are used but the end result is the same, the field equations are re-written as a set of constraint and evolution equations. Solving the constraint equations for a given system will provide the spatial metric, γ_{ij} and the extrinsic curvature, K_{ij} , of said system. The evolution equations describe how the spatial hyper-surfaces change between time t and $t + \Delta t$, with the matter and radiation acting as source terms. The evolution equations are evaluated numerically using 4th order Runge Kutta (RK4) methods to integrate forward in time and finite difference methods to evaluate the spatial derivatives.

To evolve the head on collision system first values for γ_{ij} and K_{ij} have to be defined everywhere on the grid, meaning the constraint equations have to be solved. For the BSSNOK formulation used in the head on case, a Hamiltonian and momentum constraint equation have to be solved. The momentum constraint equations are simplified by the head on case not having to solve for orbital angular momentum, only linear momentum. Isolated neutron star solutions can be found by solving the Tolman–Oppenheimer–Volkoff (TOV) equations numerically using the TOV module in the ET. Solving the TOV equations requires the EOS, for the head on case a polytropic EOS was used. The isolated system solutions can be combined linearly if the separation of the two neutron stars is great enough that non linear terms can be treated as negligible. For the head on case a separation of 90,000 km was used.

As well as providing the initial data for the head on collision the constraint equations can also be used to measure the numerical error in the simulation. Analytically if the constraint equations are satisfied at some time they will be satisfied for all time [69]. Constraint violations should be zero everywhere. Because the constraint equations are solved numerically and the resulting set of variables are evolved numerically, numerical errors will naturally be introduced into the solution meaning the constraint violations will not be zero everywhere. The magnitude of the constraint violations is a measure of the accuracy of numerical relativity simulations. The different numerical formulations are designed around constraint violations, minimising them or even constructing the formulae in such a way that the constraint violations dissipate away from the regions of interest.

FIL does not solve the constraint equations FIL focuses on evolving the matter and radiation source terms in parallel. The variables that FIL evolves are; pressure, P ; the fluid 3 velocity, v^i ; the fluid rest mass density, ρ_0 ; and the magnetic field as measured by an eulerian observer, B^i . The variables are evolved by rewriting the conservation equations for the stress energy tensor into the numerically well behaved flux conservation form. Rewriting the conservation equation in this form guarantees the conservation of mass, energy and momentum to truncation error and allows for the use of high order high-resolution shock capturing numerical schemes. This formulation requires two sets of variables; conserved variables, constructed variables which are evolved numerically; and primitive variables, variables that have physical meaning such as density, pressure, magnetic field strength etc. The conservatives are related to the primitive variables by a non linear set of equations. It is easy to get the primitives from the conservative variables but to recover the primitive variables from the conservative variables, something that needs to be done at each substep of the RK4 routine, a Newton–Raphson-based root finder is needed. This step is one of the most time consuming in the simulations. The root finder takes longer to converge to a solution the stronger the magnetic fields. In a neutron star collision the magnetic fields are strongest right before and after merger. Regions of high magnetic field strength are not only of computational interest but also of scientific interest. The behaviour of the magnetic fields is believed to be responsible for many phenomena such as short GRBs. Head on collisions allow greater focus on this regime than in-spirals as a greater percentage of the simulation will be taken up by the short binary separation high magnetic field strength arrangement. The impact of binary separation, resolution and magnetic field strength on the code can all be studied for head on collisions without any loss of physics. For the head on case a 128,128,64 grid was used where the grid extended 2048M in every direction (note geometrized units have been used here ¹) with 4 refinement levels. Two refinement regions exist which are centered on the two neutron stars.

The ultimate goal of the Relativistic Astrophysics group at Goethe University is to study the dynamical properties of BNS post merger remnants over 1 to 10 seconds. This regime is where processes such as the ejection of matter, the development of a globally oriented magnetic field, or the launching of a jet from the merger remnant take place[73]. Currently second long post merger simulations take about a year of wall time to run. The goal is to improve the speed of binary neutron star collisions simulations. Improving the speed of the head on case will also improve the speed of the inspiral simulation, the observationally relevant case, as both use the same evolution framework and contain the same physics. It is hoped that speed up can be found by accelerating the codes using GPUs. Improving the scaling of FIL will also bring down wall time with the new generation of exascale computing.

¹For the Geometrized units the gravitational constant, G ; the speed of light, c ; and the mass of the sun M_{sun} are set equal to 1. Expressed in terms of SI units mass and length are equal to $[M] = M_{sun}$ and $[L] = [M](G/c^2)$ respectively.

7 iPic3D

The particle-in-cell code iPic3D [74] self-consistently solves the kinetic equations for the evolution of the distribution functions $f_s(x, v, t)$ for each plasma species s (electrons, ions) along with Maxwell's equations on a 3D grid:

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \nabla f_s + \frac{q_s}{m_s} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_v f_s = 0$$

The Vlasov equation describes the time evolution of the distribution function f_s for a species s in a plasma. It considers the motion due to velocity \mathbf{v} , and the force exerted by electric (\mathbf{E}) and magnetic (\mathbf{B}) fields on charged particles with charge q_s and mass m_s .

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0}$$

The Gauss law for electricity, relates the electric field \mathbf{E} to the distribution of electric charge ρ , stating that the divergence of \mathbf{E} is proportional to the electric charge density.

$$\nabla \cdot \mathbf{B} = 0$$

The Gauss law for magnetism states that the magnetic field \mathbf{B} has no divergence, implying there are no magnetic monopoles.

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

The above equation shows that a changing magnetic field creates a circulating electric field.

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}$$

The above equation relates the circulation of the magnetic field \mathbf{B} around a closed loop to the electric current \mathbf{J} and the rate of change of the electric field, accounting for the displacement current which arises from changing electric fields.

As the code self consistently simulates the plasma conditions, we plan to perform the following scientific cases.

7.1 Case 1 - Properties of Turbulence in Various Planetary Magnetospheres

Observations have revealed magnetospheric turbulence across terrestrial planets, but many attributes of the turbulence remain poorly constrained or differential between the planets [75, 76, 77]. Properties like intensity, spectral characteristics, spatiotemporal evolution and dissipative behaviors are not well quantified from measurements alone. This hampers the development of first-principles turbulence models needed to generalize predictions for space weather disturbances.

Comparative simulations employing common modeling frameworks could systematically examine differences in the onset of turbulence and dynamics stemming from the driving conditions around each planet. For instance, the lower hybrid drift and other wave-particle interactions expected to mediate Venus' induced magnetotail turbulence are less developed towards fully developed turbulence, than reconnection physics in Earth's magnetosphere. Predictive capabilities would benefit from quantifying the role of plasma instabilities and self-organization behaviors under various solar wind conditions and intrinsic magnetic field strengths.

Multi-scale simulations also present opportunities to probe outstanding questions regarding the coupling between macro-scale and kinetic-scale turbulence - an area lacking detailed constraints from observations. Resolution of the energy cascades and dissipation channels in the collisionless magnetospheric plasmas could have significant implications for internal dynamics and erosion processes. Comparative modeling can therefore guide understanding of the most influential turbulence properties in driving space weather disturbances and atmospheric escape.

In this scientific case study focusing on magnetospheric turbulence, both strong and weak scaling are well-suited approaches due to the diverse requirements of Particle-In-Cell (PIC) simulations. Strong scaling, which involves increasing computational resources while maintaining a constant problem size, is ideal for the high-resolution modeling necessary to capture the fine-scale kinetic processes such as wave-particle interactions and

micro-instabilities. This approach allows the simulation to efficiently process detailed and computationally intensive tasks within the confines of a fixed simulation domain. On the other hand, weak scaling, where the problem size increases in proportion to the computational resources, perfectly aligns with the need to model the extensive range of spatial scales of planetary magnetospheres. As the simulation expands to encompass larger regions of space, weak scaling ensures that the increased data volume and complexity are managed effectively, maintaining computational efficiency.

The precise initialization of this case and the two below is given by two large binary files that range up to tens of GB for the fields and 1TB or even more for the particles. These are generated from observational data or from other simulations at the fluid level. The software for generating these files is provided with the code.

This case is the simplest in terms of computational demand the next two will progressively require a larger computational effort because of the need to resolve smaller scales.

7.2 Case 2 - Regions of Energy Exchange in Ions and Electrons

Turbulence generates numerous reconnection sites at the magnetopause facilitating the explosive transfer of energy from the magnetic field to plasma particles within localized diffusion regions [78], as shown in Figure 16. As revealed in high-resolution Magnetospheric Multiscale data, magnetic islands [79] and electron jets at X-lines are embedded within larger ion-scale flux ropes [80], generating overlap across scales. This drives interactions




Figure 16: In this visualization, we observe a turbulent region characterized by the presence of numerous reconnection sites, depicted as yellow areas. Electron flow paths weave through this tumultuous landscape, intersecting multiple reconnection points. Notably, these electron trajectories are color-coded to reflect the strength of the local electric field encountered along their journey. This electric field is instrumental in the energy exchange process between the magnetic field and the electrons, a key factor in their acceleration. The vibrant color gradient along the flow lines thus serves as a dynamic map of the electric field intensity, highlighting zones of significant energy transfer and acceleration within the turbulent region [4].

between multiple closely-spaced electron- and ion-scale structures. Complex 3D magnetic geometries arise with electric field variations over tiny spatial scales. The resulting nonuniform E field topology provides a site for

direct acceleration of both ions and electrons. For instance, distinct ion and electron diffusion region jets have been observed, indicative of particle decoupling.

The turbulence also stimulates dispersive Alfvén waves and parallel electric fields along separatrices that can energize ions perpendicular and parallel to the magnetic field [81]. Additionally, electrons encountering gradients and curvature drifts display evidence of Fermi acceleration, producing suprathermal power-law spectra.

By regulating acceleration, heating, and mass transport, the hierarchical interactions between micro-scale reconnection electric fields and mesoscale flux structures represent the fundamental pathway for the explosive conversion of magnetic energy to plasma particle energies within turbulent dayside reconnection layers [81]. Further analysis of distribution functions and field-particle correlations is enabling the quantification of this mechanism.

Both strong and weak scaling can be conducted in this scientific case. For strong scaling, the fixed problem would be defined by specifying a three-dimensional simulation box on the order of 100-1000 electron skin depths with sufficient resolution to model electron and ion kinetics. This domain would be distributed across an increasing number of cores during scaling studies, targeting the reduction of overall walltime while maintaining resolution finer than the electron gyroradius.

Weak scaling can complement via proportional expansion of the simulation volume as processor count grows. For example, longer flux tubes capturing the full extent of interconnected X-lines and secondary islands can mitigate boundary effects. Domain decomposing the stretched grid over more cores preserves per-process workload. This investigation of larger systems maintains detail already validated in the baseline calculation used for strong scaling.

7.3 Case 3 - 3D Anti-Parallel Reconnection

This kinetic simulation study concentrates predominantly on anti-parallel magnetic reconnection dynamics at the dayside magnetopause due to the much greater energy conversion efficiency compared to parallel or ortho-parallel orientations [5]. The primary interest lies in quantifying the mechanisms facilitating explosive transfer of electromagnetic energy to particles under dominant southward Interplanetary Magnetic Field (IMF) driving conditions [82].

Crucially, anti-parallel reconnection supports stronger in-plane Hall currents, compared with the cases above, out-of-plane currents, and intense non-ideal electric fields localized over electron scales as evident in the Figure 17. These drive more extensive decoupling of ions and electrons than in the cases above, enabling thermalization by particle crescent distributions and acceleration by parallel/stochastic potentials near X-points. The anti-parallel geometry also stimulates faster reconnection jet front speeds up to the local Alfvén speed, critically controlling bursty flux transport rates [83, 81, 84].

Through massively parallel 3D particle-in-cell simulations, this study aims to examine in detail the particle and field dynamics in the electron diffusion region and separatrices of asymmetric anti-parallel collisions between the geomagnetic field and magnetosheath draped IMF field lines. This scientific case is also well suited for strong and weak scaling. Strong scaling is conducted by fixing a simulation domain spanning 10²-10³ electron skin depths with sufficient kinetic particles per cell. This domain captures the electron diffusion and separatrix regions and is partitioned over increasing processor counts to reduce time-to-solution. Weak scaling initiates from a baseline configuration verified by strong scaling studies. Subsequent simulations proportionally expand the domain dimensions and processor allocation.

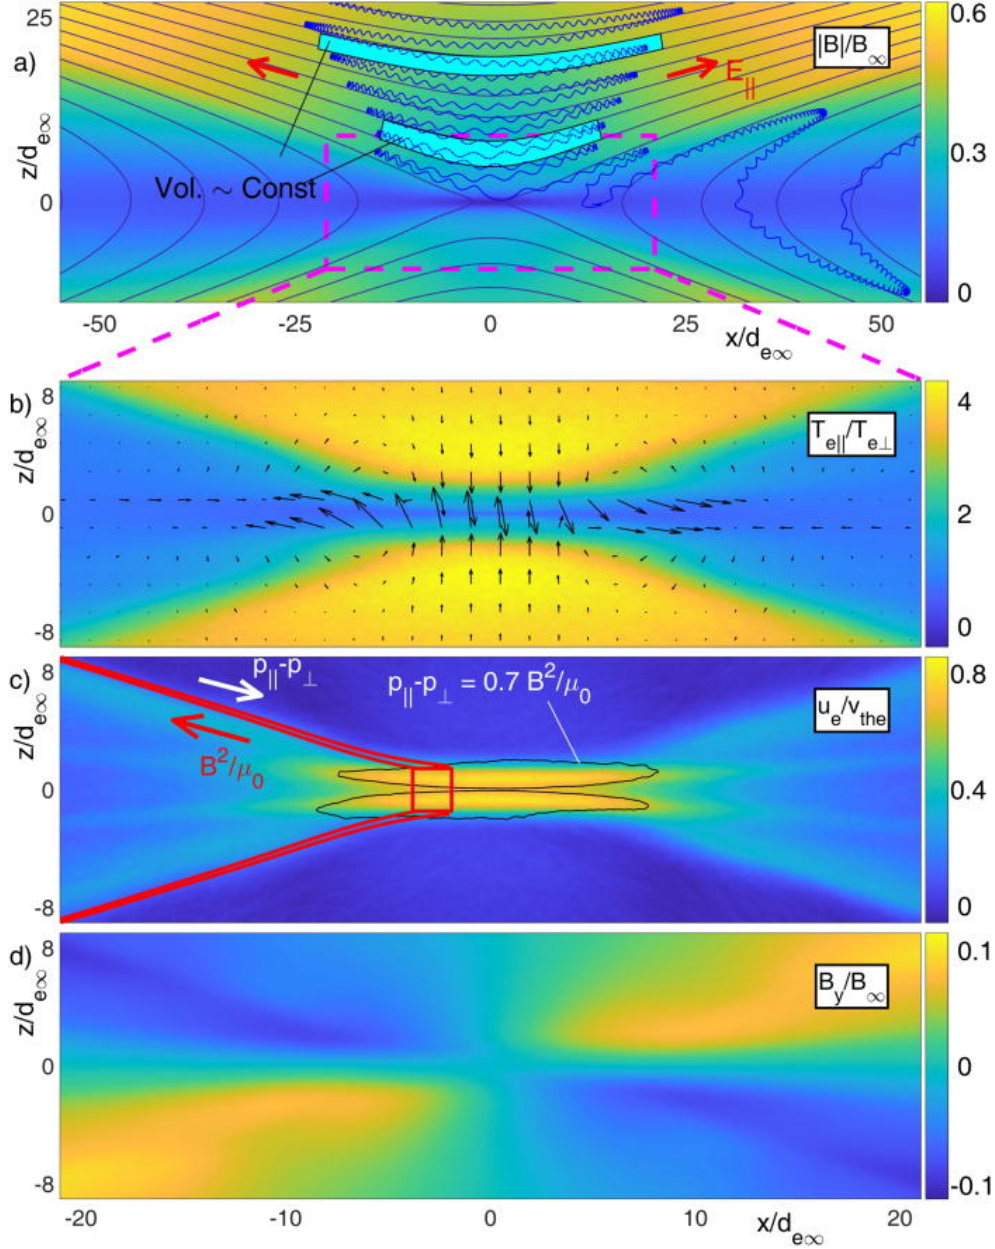


Figure 17: Anti-parallel reconnection depicted in the kinetic simulation [5].

8 RAMSES

RAMSES [85, 86] is an Adaptive-Mesh-Refinement (AMR) code which is used to study astrophysical fluid dynamics and the formation of structures in the Universe. It is based on an oct-tree structure, where parent cells are refined into children cells on a cell-by-cell basis following some user-defined criteria. RAMSES can deal with 1D, 2D and 3D Cartesian grids.

For hydrodynamics, RAMSES integrates the equations of fluid dynamics in their conservative form. This system can be written in the canonical form

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbb{F}(\mathbf{U}) = \mathbb{S}(\mathbf{U}). \quad (16)$$

Here the vector $\mathbf{U} = (\rho, \rho \mathbf{u}, E)$ contains the conservative variables: density ρ , velocity \mathbf{u} , and total energy $E = e + 1/2 \rho u^2$ with e the internal energy. The flux vector $\mathbb{F}(\mathbf{U}) = (\rho \mathbf{u}, \rho \mathbf{u} \otimes \mathbf{u} + P \mathbb{I}, \mathbf{u}(E + P))$ is a linear function of \mathbf{U} , and uses the primitive variables ρ , \mathbf{u} and pressure P . $\mathbb{S}(\mathbf{U})$ represents the source terms, e.g. the

gravitational force contribution. This system is closed using a perfect gas equation of state $e = P/(\gamma - 1)$ with γ the adiabatic index.

RAMSES uses an explicit second-order predictor-corrector finite-volume Godunov scheme to integrate the conservative system of equations. The hydrodynamic solver consists of computing flux at cell interfaces, including coarse-to-fine and fine-to-coarse interfaces. The discretized scheme goes as (here in 1D for simplicity)

$$\frac{\mathbb{U}_i^{n+1} - \mathbb{U}_i^n}{\Delta t} \times V_i = \mathbb{F}_{i+1/2}^{n+1/2} S_{i+1/2} - \mathbb{F}_{i-1/2}^{n+1/2} S_{i-1/2}, \quad (17)$$

where \mathbb{U}_i^n is the state variable \mathbb{U} at time n averaged in cell i of volume V_i , $\mathbb{F}_{i-1/2}^{n+1/2}$ is the flux at the interface of surface $S_{i-1/2}$ between cells i and $i - 1$, computed using a linear Riemann solver. The initial values of the Riemann problems are obtained from the primitive variables of cells $i - 1$ and i , extrapolated in time and space at the interface in the predictive step. Altogether, the scheme follows the same steps as in the PLUTO code.

RAMSES can also handle the evolution of particles, such as stars, dark matter (DM) and sink particles, whose evolution is integrated using a Cloud-In-Cell (CIC) interpolation. The same CIC interpolation is used to deposit the mass of the particles onto the grid to solve for the gravitational potential. For gravity, RAMSES can either use a Conjugate Gradient algorithm or a multigrid solver [87]. For each solver, the gravitational potential is solved level-by-level. RAMSES global time integration can be accelerated using the adaptive-time-step implementation, in which each AMR level ℓ evolves with its own time-step which satisfies a global synchronisation point at the end of the coarser time-step $\Delta t_\ell = 2\Delta t_{\ell+1}$.

For this project, we have identified three use cases described in the following sub-sections, which are representative of the current limitations of RAMSES for what concerns scaling efficiency, load balancing and MPI communications. Besides the technical aspect, these three cases are also representative of a large fraction of the RAMSES community, from cosmology to star formation.

8.1 Case 1 - Sedov-Taylor blast

For the first scientific case, we select the classic Sedov-blast wave test. This test is a standard hydrodynamical test used for verification and benchmarking of gas dynamic codes (see Case 3 for the CHANGA/GASOLINE code, Section 5.3). We run this test in 3D with a uniform fixed grid to focus on the three selected regions of D2.1, without any overhead due to the adaptive mesh refinement. The physical setup consists of a uniform density ($\rho = 1$) and low-pressure ($P = 10^{-5}$) medium at rest (zero velocity) in which an internal energy pulse of amplitude 0.4 is put in the corners of the computational domain in order to check the good behaviour of the periodic boundary conditions. The box size is 0.5. We use resolutions ranging from 64^3 to 2048^3 grid cells and integrate the code over a fixed number of iterations (typically 10 timesteps).

The internal energy is progressively converted into kinetic energy, while the total energy is conserved. A time evolution sequence of the density field is shown in Figure 18.

Since this test is without gravity and therefore purely hydrodynamical, it enables testing the performances of the hydrodynamical kernel and of the communication kernels of RAMSES without the overheads due to the adaptive mesh refinement. In addition, since this tests has an analytic solution, it can also test the accuracy of the code. Here we use a uniform grid at level ℓ and the Lax-Friedrich Riemann solver. A unique load-balancing of the MPI domains is performed at the initialization of the run and no further load balancing is done during the simulation run. In this use case, all time steps (or iterations) are equivalent in terms of computational cost. Therefore, this scientific case is well-suited for memory-bound runs. In summary, this test provides an ideal environment to optimize the communications between MPI domains and we can conduct both strong and weak scaling studies.

8.2 Case 2 - Cosmological box

Cosmological boxes have high scientific relevance in gaining a theoretical understanding and interpretation of large cosmological surveys and the evolution of structure in our Universe. They are the largest simulations run with RAMSES and are typically “flagship” simulations (see e.g. [88]). Since we are now limited by cosmological simulations running on a few tens of thousands of cores, but still taking months to finish, cosmological boxes are a very relevant scientific case for method improvements. To obtain higher resolutions for current-size volumes, larger volumes for current-resolution runs, and shorter wall times for both cases, a good scaling relation is necessary. Cosmological boxes scale relatively well with the number of cores, due to the natural division of

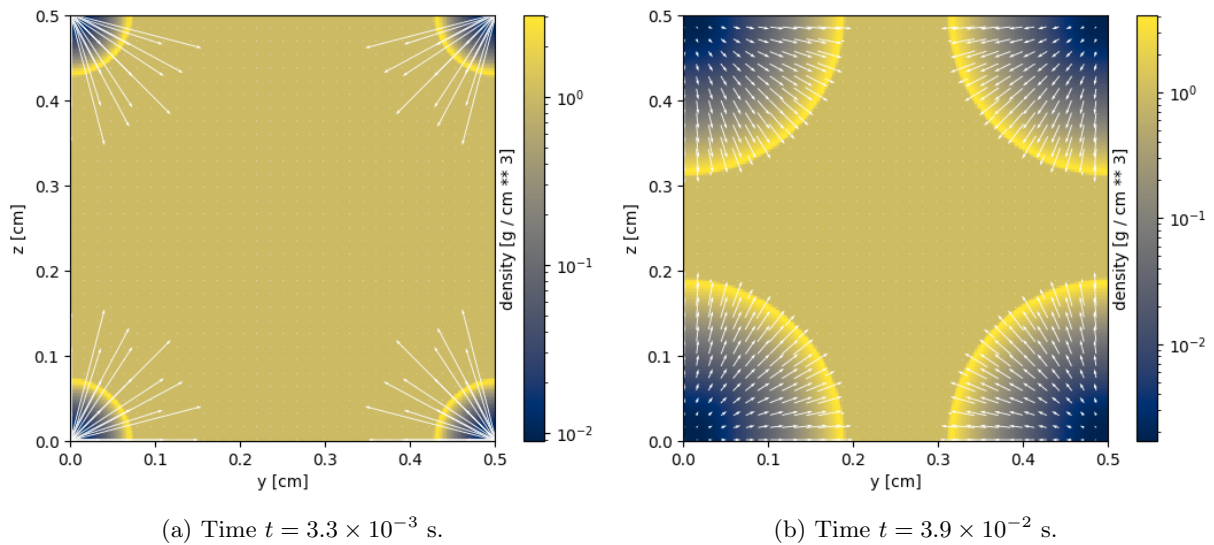


Figure 18: Sedov-Taylor blast wave: projection of the dark matter density at different times and resolutions.

a cosmological volume into different galaxies and their hosting dark matter halos, which makes them a good choice for a scientific case.


A cosmological volume represents a statistically accurate description of a patch of the Universe, in terms of the number density of dark matter halos as well as their spatial and mass distribution. This is achieved by realistic gravitational dynamics of dark matter (to the first order) and initial conditions that are statistically identical to the observed power spectrum of density and temperature fluctuations in the cosmic microwave background. The combination of these two ingredients leads to the collapse into the so-called cosmic web. The fact that these ingredients can reproduce a realistic Universe more than 13 billion years later [89] is among the greatest achievement of cosmological simulations and has made them an essential component in theoretical studies of structure formation from inter-galactic scales to the interstellar medium. For more details about cosmological boxes see the description in Sections 2 and 5 involving OpenGADGET and CHANGA/GASOLINE.

For the initial conditions we use the publicly available MUSIC software package [90], which is commonly used for most cosmological codes. MUSIC takes as parameters the co-moving width of the square volume, L_{box} , the number of DM particles in each dimension, N_{part} , such that the total number of particles, and coarse-grid cells, is N_{part}^3 , and the cosmological parameters representing the Hubble constant, the amplitude of the power spectrum on a scale of 8 Mpc, the matter fraction of the Universe, its dark energy fraction, and baryonic fraction. Additionally, one provides MUSIC a random number seed for the generation of density fluctuations, i.e. by changing only the random seed one can produce multiple “instances” of a Universe with the same cosmological properties.

Starting at these initial conditions a few tens of million years after the Big Bang a cosmological simulation is typically run for more than 13 billion years until the present epoch, where we have the largest wealth of observations to compare to theory. However, as it is too expensive to run the full simulation for each scaling test, we restrict ourselves to the first ten timesteps, starting at redshift 150 when the Universe has almost no structure.

The cosmological runs we use for this science case contain only dark matter particles, i.e. no hydrodynamics, gas cooling, or star formation. Our justification is that these additional physics can make the exact comparisons between different resolutions somewhat tricky – for example, stars may form at different times and at different rates, depending on the resolution, making the amount of computation very uneven between different resolution runs. Additionally, as we only run a few initial timesteps, there is insufficient time for structures to form and therefore non-DM physics remains negligible. For the same reasons, we do not trigger the adaptive mesh refinement. The focus of the resulting uniform-grid benchmarks is therefore to test the particle-in-cell (PIC) algorithm, gravity, load balancing, and domain communication overheads.

While changing the resolution, i.e. N_{part} , by a factor of two, resulting in a change of the total number of particles and cells by a factor of 8 we run all tests with the same L_{box} and initial conditions. Given the same initial conditions, with a factor two increase of N_{part} , we expect to cover the same number of timesteps with



(a) Resolution 128^3 , $N_{\text{step}} = 100$. (b) Resolution 128^3 , $N_{\text{step}} = 200$. (c) Resolution 256^3 , $N_{\text{step}} = 100$.

Figure 19: 2D dark matter density projection in runs with different resolutions and at different timesteps.

an increase in the computational load of a factor 8 Table 4 shows the resolutions adapted for these benchmark tests as well as the cosmological parameters used.

N_{part}	L_{box}	z_{start}	Ω_{m}	Ω_{Λ}	H_0	σ_8	seed
128, 256, 512 1024, 2048	20 cMpc	150	0.3175	0.6825	$67.11 \text{ km s}^{-1} \text{ Mpc}^{-1}$	0.83	4721

Table 4: Resolution and IC parameters for the cosmological benchmark tests. From left to right, the parameters are: number of particles N_{part} per dimension, box width L_{box} in co-moving Mpc, starting redshift z_{start} , matter fraction Ω_{m} , dark energy fraction Ω_{Λ} , Hubble constant H_0 , power spectrum amplitude σ_8 on an 8 Mpc scale, and seed for random number generation with MUSIC.

8.3 Case 3 - Isolated galaxy

This setup of a single pre-defined rotating galactic disk is very often used in the galaxy evolution community to explore the physics that regulate galaxy evolution and generate galactic winds, such as star formation, supernovae, stellar winds, radiation, the growth and feedback of active galactic nuclei, magnetic fields, cosmic rays, and conduction. Isolated galaxy disks are very good for performing controlled experiments of these physics for two main reasons. Firstly, these simulations are much less expensive to run than cosmological simulations. Secondly, they are not affected by the chaotic nature of galaxy evolution in the cosmological context due to mergers and accretion, making it much easier to perform a clean apples to apples comparison between two simulations of the same isolated galaxy.

However, this advantage of not having a cosmological context is also a limitation. The simulated galaxy is isolated and does not experience mergers with other galaxies nor does it accrete gas from the inter-galactic medium. Without these environmental factors, galaxies cannot grow over time since they don't have a fresh supply of baryons. Hence isolated disk galaxy simulations are typically only run for a few rotation times, corresponding to a few hundred million years, a short enough time that one can ignore environmental effects such as the lack of fuel for galaxy growth.

These simulations enable adaptive mesh refinement on a criterion based on mass. Therefore the galaxy placed at the center of the box is highly refined but its volume-filling low-density environment is not. As a consequence, most of the refinement is performed on a small fraction of the computational domain. These simulations are thus challenging in terms of scaling and load balancing, and typically isolated disk simulations described in the literature run only on a few hundred cores before the scaling curve becomes horizontal. There would be a significant gain in improving the scaling in the isolated galaxy setup since it would allow for improved resolution in galaxies, typically limited to cell widths of a few parsecs in Milky-Way mass galaxies at the moment, which is too large for a realistic description of molecular clouds, the birthplace of most of the stars.

We use the isolated disk setup described in paper II of the AGORA code comparison project ([6], <https://sites.google.com/site/santacruzcomparisonproject/>), where it was used to compare several simulation codes including RAMSES, CHANGA, and Gadget-3, which are all three participating in the SPACE CoE. The

initial conditions (ICs) are generated with MAKEDISK, originally written by Volker Springel and adapted to make ICs readable for RAMSES (see [91] for details). The IC generation code is written in C and is part of the public version of the RAMSES code (under `ramses/utls/ic/galaxy_ic/galics`). Here one chooses the mass of the DM halo, its concentration and spin, the number of dark matter particles N_{part} , and hence the particle mass, the baryonic mass, the fraction of the baryonic mass in stars (with the remainder in gas), the number of stellar particles, and hence the stellar particle mass, the height and length scale of the galactic disk, the mass and radius of the stellar bulge, the initially uniform temperature of the gas disk, the temperature of the surrounding circum-galactic medium gas, and its density. One also needs to choose the size of the box (large enough to encompass the dark matter halo particles), the coarse level resolution (via the `levelmin` parameter), and the finest level resolution (via the `levelmax` parameter). To crudely mimic a galaxy with a history of star formation creating heavy elements, one also sets a homogeneous metal mass fraction (i.e. the mass fraction of all elements heavier than helium) for both the gas and the stars.

We list all fixed IC parameters in Table 5. The only aspect that varies between tests is the resolution, via changes to the `levelmin`, `levelmax`, and the numbers of DM and stellar particles. We list the different resolution runs in Table 6. In each run, we allow adaptive refinement of five levels beyond the coarse resolution, using the widely adapted scheme of refining a cell if it contains 8 DM particles or the scaled equivalent mass in baryons, or if the local Jeans length is less than four cell widths. Each test is run for 10 coarse timesteps, and we allow AMR sub-cycling for the two deepest levels, i.e. 10 coarse timesteps correspond to 40 finest-level timesteps. We include the basic thermo-chemistry that comes with RAMSES, but omit star formation or feedback, due to the stochasticity of these processes, which is amplified by the short duration of each run. This scientific case benchmarks mainly the load-balancing and domain communication overheads of the AMR structures on top of the hydro, gravity and Particle-In-Cell (PIC) algorithms.

In Fig. 20 we show face-on and edge-on projections of the intermediate resolution run at 21 Myrs, in gas density, refinement level, DM surface density, and stellar surface density.

	Dark matter halo	Stellar disk	Gas disk	Stellar bulge
Density profile	Navarro et al. (1997)	Exponential	Exponential	Hernquist (1990)
Structural properties	$M_{200} = 1.1 \times 10^{12} M_{\odot}$, $v_{c,200} = 150$ km/s, $R_{200} = 205$ kpc, $c=10$, $\lambda = 0.04$, $L_{\text{box}} = 400$ kpc	$M_{d,*} = 3.4 \times 10^{10} M_{\odot}$, $r_d = 3.43$ kpc, $z_d = 0.1 r_d$	$M_{d,\text{gas}} = 8.6 \times 10^9 M_{\odot}$, $f_{\text{gas}} = 0.2$	$M_{b,*} = 4.3 \times 10^9 M_{\odot}$, $M_{b,*}/M_d = 0.1$

Table 5: Fixed isolated disk setup parameters, copied from AGORA paper II [6]. The parameters are virial mass M_{200} , rotational velocity $v_{c,200}$, virial radius R_{200} , concentration parameter c , spin λ , box width L_{box} , disk stellar mass $M_{d,*}$, disk scale radius r_d , disk scale height z_d , disk gas mass $M_{d,\text{gas}}$, disk gas fraction f_{gas} , and bulge stellar mass $M_{b,*}$.

N_{DM}	N_{SD}	N_{SB}	ℓ_{min}	ℓ_{max}	Δx_{max}	Δx_{min}
1.25×10^5	1.25×10^5	1.25×10^4	7	12	3125 pc	98 pc
10^6	10^6	10^5	8	13	1562 pc	49 pc
8×10^6	8×10^6	8×10^5	9	14	781 pc	24.5 pc

Table 6: Varying isolated disk parameters. The parameters are, from left to right, number of DM particles N_{DM} , number of stellar disk particles N_{SD} , number of stellar bulge particles N_{SB} , coarse refinement level ℓ_{min} , fine refinement level ℓ_{max} , coarse cell width Δx_{max} , and fine cell width Δx_{min} .

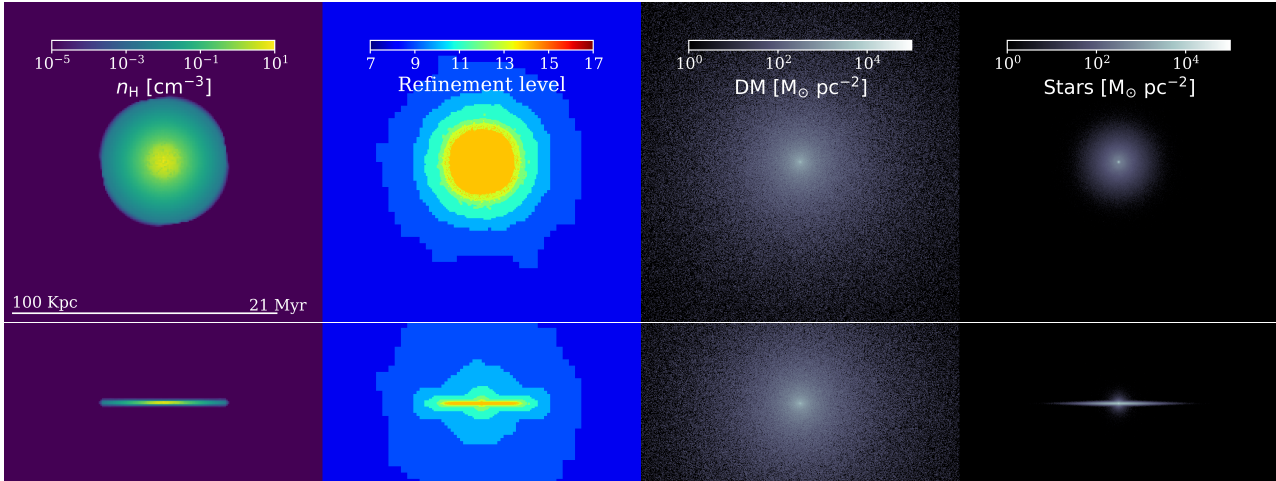


Figure 20: Face-on and edge-on maps of the isolated disk RAMSES galaxy, in the upper and lower rows respectively, after running for a simulated 21 Myrs. The projected maps show, from left to right, maximum hydrogen density along the line of sight, maximum refinement level along the line of sight, DM surface density, and stellar surface density. Note that the maps show only a fraction of the total volume, which is 320 kpc in width.

9 Conclusions

In this document, we have presented and discussed the scientific cases that will be used throughout the project to (i) extract the performance metrics at successive development stages, and (ii) constantly validate the evolved codes. As for the **cosmological codes**, the following cases have been presented:

1. Cosmological boxes (OpenGADGET, ChaNGa, RAMSES)
2. Zoom-in simulations (OpenGADGET, ChaNGa)
3. Sedov-Taylor blast wave (ChaNGa, RAMSES)
4. Isolated Galaxy (ChaNGa, RAMSES)
5. Magnetized cloud collapse (ChaNGa)

We note that many cases are common (in their rationale, not in their realization) among the codes, and most of those not in common could be run with all the codes. Starting from that, in the forthcoming activity, whenever possible and not at the cost of the main activity, we deem it possible to build a common ground where the scientific cases are run in common with the codes for the purpose of testing and comparing. More specifically, the collected scientific use cases are summarized in what follows. The **PLUTO** team has individuated three cases, namely *Resistive Relativistic Connection*, *Plasma Column Instability* and *Synthetic Spectra from 3D Relativistic Jets*. The GR codes **BHAC** and **FIL** have four cases in total: *Magnetised torus around a spinning black hole* in the former, and *Binary Neutron Star Collision* *Black Hole Neutron Star Collision* and *Head on Collision* the latter. Finally, the **iPic3D** code brings three cases: *Properties of Turbulence in Various Planetary Magnetosphere*, *Regions of Energy Exchange in Ions and Electrons* and *Anti-Parallel Reconnection*. **A total of twenty (20)** cases have been selected and realized for this work. Their instances, i.e. the set of files needed to actually run the simulations, are stored in the IDA's folder relative to each code.

These scientific cases are among the major objectives that the SPACE CoE will try to address in the forthcoming three years of activities while porting the codes to exa-scale capabilities.

References

- [1] E. Puzzoni, A. Mignone, and G. Bodo, “On the Impact of the Numerical Method on Magnetic Reconnection and Particle Acceleration – I. The MHD case,” *Monthly Notices of the Royal Astronomical Society*, vol. 508, no. 2, pp. 2771–2783, Oct. 2021.
- [2] R. Wissing and S. Shen, “Smoothed particle magnetohydrodynamics with the geometric density average force expression,” *Astronomy & Astrophysics*, vol. 638, p. A140, Jun. 2020.
- [3] C. Palenzuela, “Introduction to Numerical Relativity,” *Front. Astron. Space Sci.*, vol. 7, p. 58, 2020.
- [4] G. Lapenta, “Power to the particles,” *Nature Physics*, vol. 19, no. 2, pp. 159–160, 2023.
- [5] J. Egedal, H. Gurram, S. Greess, W. Daughton, and A. Lê, “The force balance of electrons during kinetic anti-parallel magnetic reconnection,” *Physics of Plasmas*, vol. 30, no. 6, 2023.
- [6] J.-h. Kim, O. Agertz, R. Teyssier, M. J. Butler, D. Ceverino, J.-H. Choi, R. Feldmann, B. W. Keller, A. Lupi, T. Quinn, Y. Revaz, S. Wallace, N. Y. Gnedin, S. N. Leitner, S. Shen, B. D. Smith, R. Thompson, M. J. Turk, T. Abel, K. S. Arraki, S. M. Benincasa, S. Chakrabarti, C. DeGraf, A. Dekel, N. J. Goldbaum, P. F. Hopkins, C. B. Hummels, A. Klypin, H. Li, P. Madau, N. Mandelker, L. Mayer, K. Nagamine, S. Nickerson, B. W. O’Shea, J. R. Primack, S. Roca-Fàbrega, V. Semenov, I. Shimizu, C. M. Simpson, K. Todoroki, J. W. Wadsley, J. H. Wise, and AGORA Collaboration, “The AGORA High-resolution Galaxy Simulations Comparison Project. II. Isolated Disk Test,” *APJ*, vol. 833, no. 2, p. 202, Dec. 2016.
- [7] G. S. Karademir, R.-S. Remus, A. Burkert, K. Dolag, T. L. Hoffmann, B. P. Moster, U. P. Steinwandel, and J. Zhang, “The outer stellar halos of galaxies: how radial merger mass deposition, shells, and streams depend on infall-orbit configurations,” *MNRAS*, vol. 487, no. 1, pp. 318–332, Jul. 2019.
- [8] N. Katz and S. D. M. White, “Hierarchical Galaxy Formation: Overmerging and the Formation of an X-Ray Cluster,” *ApJ*, vol. 412, p. 455, Aug. 1993.
- [9] G. Tormen, F. R. Bouchet, and S. D. M. White, “The structure and dynamical evolution of dark matter haloes,” *MNRAS*, vol. 286, no. 4, pp. 865–884, Apr. 1997.
- [10] S. Borgani, F. Governato, J. Wadsley, N. Menci, P. Tozzi, T. Quinn, J. Stadel, and G. Lake, “The effect of non-gravitational gas heating in groups and clusters of galaxies,” *MNRAS*, vol. 336, no. 2, pp. 409–424, Oct. 2002.
- [11] K. Dolag, S. Borgani, G. Murante, and V. Springel, “Substructures in hydrodynamical cluster simulations,” *MNRAS*, vol. 399, no. 2, pp. 497–514, Oct. 2009.
- [12] D. Martizzi, Jimmy, R. Teyssier, and B. Moore, “Brightest cluster galaxies in cosmological simulations with adaptive mesh refinement: successes and failures,” *MNRAS*, vol. 443, no. 2, pp. 1500–1508, Sep. 2014.
- [13] D. J. Barnes, S. T. Kay, Y. M. Bahé, C. Dalla Vecchia, I. G. McCarthy, J. Schaye, R. G. Bower, A. Jenkins, P. A. Thomas, M. Schaller, R. A. Crain, T. Theuns, and S. D. M. White, “The Cluster-EAGLE project: global properties of simulated clusters with resolved galaxies,” *MNRAS*, vol. 471, no. 1, pp. 1088–1106, Oct. 2017.
- [14] L. Bassini, E. Rasia, S. Borgani, G. L. Granato, C. Ragone-Figueroa, V. Biffi, A. Ragagnin, K. Dolag, W. Lin, G. Murante, N. R. Napolitano, G. Taffoni, L. Tornatore, and Y. Wang, “The DIANOGA simulations of galaxy clusters: characterising star formation in protoclusters,” *A&A*, vol. 642, p. A37, Oct. 2020.
- [15] A. Mignone, G. Bodo, S. Massaglia, T. Matsakos, O. Tesileanu, C. Zanni, and A. Ferrari, “PLUTO: A Numerical Code for Computational Astrophysics,” *ApJs*, vol. 170, no. 1, pp. 228–242, May 2007.
- [16] A. Mignone, C. Zanni, P. Tzeferacos, B. van Straalen, P. Colella, and G. Bodo, “The PLUTO Code for Adaptive Mesh Computations in Astrophysical Fluid Dynamics,” *ApJs*, vol. 198, no. 1, p. 7, Jan. 2012.
- [17] A. Mignone and G. Bodo, “An HLLC Riemann solver for relativistic flows - I. Hydrodynamics,” *MNRAS*, vol. 364, no. 1, pp. 126–136, Nov. 2005.

-
- [18] —, “An HLLC Riemann solver for relativistic flows - II. Magnetohydrodynamics,” *MNRAS*, vol. 368, no. 35, pp. 1040–1054, May 2006.
- [19] A. Mignone, M. Ugliano, and G. Bodo, “A five-wave Harten-Lax-van Leer Riemann solver for relativistic magnetohydrodynamics,” *MNRAS*, vol. 393, no. 4, pp. 1141–1156, Mar. 2009.
- [20] D. Mukherjee, G. Bodo, P. Rossi, A. Mignone, and B. Vaidya, “Simulating the dynamics and synchrotron emission from relativistic jets - II. Evolution of non-thermal electrons,” *MNRAS*, vol. 505, no. 2, pp. 2267–2284, Aug. 2021.
- [21] L. Sironi and A. Spitkovsky, “Relativistic Reconnection: An Efficient Source of Non-thermal Particles,” *ApJ*, vol. 783, no. 1, p. L21, Mar. 2014.
- [22] V. Berta, A. Mignone, M. Bugli, and G. Mattia, “A 4th-order accurate finite volume method for ideal classical and special relativistic MHD based on pointwise reconstructions,” *arXiv e-prints*, p. arXiv:2310.11831, Oct. 2023.
- [23] L. Del Zanna, E. Papini, S. Landi, M. Bugli, and N. Bucciantini, “Fast reconnection in relativistic plasmas: The magnetohydrodynamics tearing instability revisited,” *MNRAS*, vol. 460, pp. 3753–3765, Aug. 2016.
- [24] C. Palenzuela, L. Lehner, O. Reula, and L. Rezzolla, “Beyond ideal MHD: Towards a more realistic modelling of relativistic astrophysical plasmas,” *MNRAS*, vol. 394, pp. 1727–1740, Apr. 2009.
- [25] S. Selvi, O. Porth, B. Ripperda, F. Bacchini, L. Sironi, and R. Keppens, “Effective Resistivity in Relativistic Collisionless Reconnection,” *The Astrophysical Journal*, vol. 950, no. 2, p. 169, Jun. 2023.
- [26] G. R. Werner, D. A. Uzdensky, B. Cerutti, K. Nalewajko, and M. C. Begelman, “The Extent of Power-law Energy Spectra in Collisionless Relativistic Magnetic Reconnection in Pair Plasmas,” *ApJ*, vol. 816, p. L8, Jan. 2016.
- [27] Y. Mizuno, Y. Lyubarsky, K.-I. Nishikawa, and P. E. Hardee, “Three-dimensional relativistic magnetohydrodynamic simulations of current-driven instability. i. instability of a static column,” *The Astrophysical Journal*, vol. 700, no. 1, p. 684–693, Jul. 2009. [Online]. Available: <http://dx.doi.org/10.1088/0004-637X/700/1/684>
- [28] B. Vaidya, A. Mignone, G. Bodo, P. Rossi, and S. Massaglia, “A particle module for the pluto code. ii. hybrid framework for modeling nonthermal emission from relativistic magnetized flows,” *The Astrophysical Journal*, vol. 865, no. 2, p. 144, oct 2018. [Online]. Available: <https://dx.doi.org/10.3847/1538-4357/aadd17>
- [29] P. Rossi, G. Bodo, A. Capetti, and S. Massaglia, “3D relativistic MHD numerical simulations of X-shaped radio sources,” *A&A*, vol. 606, p. A57, Oct. 2017.
- [30] S. Massaglia, G. Bodo, P. Rossi, S. Capetti, and A. Mignone, “Making Faranoff-Riley I radio sources. I. Numerical hydrodynamic 3D simulations of low-power jets,” *A&A*, vol. 596, p. A12, Nov. 2016.
- [31] O. Porth, H. Olivares, Y. Mizuno, Z. Younsi, L. Rezzolla, M. Moscibrodzka, H. Falcke, and M. Kramer, “The Black Hole Accretion Code,” *Computational Astrophysics and Cosmology*, vol. 4, 2017.
- [32] H. Olivares, O. Porth, J. Davelaar, E. R. Most, C. M. Fromm, Y. Mizuno, Z. Younsi, and L. Rezzolla, “Constrained transport and adaptive mesh refinement in the black hole accretion code,” *Astronomy & Astrophysics*, vol. 629, p. A61, 2019.
- [33] “The Black Hole Accretion Code – Documentation,” <https://bhac.science>.
- [34] “Repository of BHAC on GitLab,” <https://gitlab.itp.uni-frankfurt.de/BHAC-release/bhac>.
- [35] R. Keppens, Z. Meliani, A. van Marle, P. Delmont, A. Vlasis, and B. van der Holst, “Parallel, grid-adaptive approaches for relativistic hydro and magnetohydrodynamics,” *Journal of Computational Physics*, vol. 231, no. 3, pp. 718–744, 2012, special Issue: Computational Plasma Physics. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0021999111000386>

- [36] O. Porth, C. Xia, T. Hendrix, S. P. Moschou, and R. Keppens, “Mpi-amrvac for solar and astrophysics,” *The Astrophysical Journal Supplement Series*, vol. 214, no. 1, p. 4, aug 2014. [Online]. Available: <https://dx.doi.org/10.1088/0067-0049/214/1/4>
- [37] “The MPI - Adaptive Mesh Refinement - Versatile Advection Code – Documentation,” <https://amrvac.org/>, accessed: 2023-11-01.
- [38] “Repository of MPI-AMRVAC on GitHub,” <https://github.com/amrvac/amrvac>, accessed: 2023-11-01.
- [39] Y. Mizuno, Z. Younsi, C. M. Fromm, O. Porth, M. De Laurentis, H. Olivares, H. Falcke, M. Kramer, and L. Rezzolla, “The current ability to test theories of gravity with black hole shadows,” *Nature Astronomy*, vol. 2, pp. 585–590, Apr. 2018.
- [40] The Event Horizon Telescope Collaboration et al., “First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole,” *The Astrophysical Journal Letters*, vol. 875, no. 1, p. L1, apr 2019. [Online]. Available: <https://dx.doi.org/10.3847/2041-8213/ab0ec7>
- [41] H. Olivares, Z. Younsi, C. M. Fromm, M. De Laurentis, O. Porth, Y. Mizuno, H. Falcke, M. Kramer, and L. Rezzolla, “How to tell an accreting boson star from a black hole,” *Monthly Notices of the Royal Astronomical Society*, vol. 497, no. 1, pp. 521–535, 07 2020. [Online]. Available: <https://doi.org/10.1093/mnras/staa1878>
- [42] Z. Younsi, O. Porth, Y. Mizuno, C. M. Fromm, and H. Olivares, “Modelling the polarised emission from black holes on event horizon-scales,” *Proceedings of the International Astronomical Union*, vol. 14, no. S342, p. 9–12, Apr. 2020. [Online]. Available: <http://dx.doi.org/10.1017/S1743921318007263>
- [43] L. G. Fishbone and V. Moncrief, “Relativistic fluid disks in orbit around Kerr black holes.” *ApJ*, vol. 207, pp. 962–976, Aug. 1976.
- [44] C. R. Evans and J. F. Hawley, “Simulation of Magnetohydrodynamic Flows: A Constrained Transport Model,” *ApJ*, vol. 332, p. 659, Sep. 1988.
- [45] SPACE, “Deliverable D2.1: Performance profiling and benchmarking,” https://www.space-coe.eu/files/SPACE_D2.1_Performance_profiling_and_benchmarkingfinal.pdf, December 2023.
- [46] “Repository of the changa code on github,” <https://github.com/N-BodyShop/changa>, accessed: 28.10.2023.
- [47] P. Jetley, F. Gioachin, C. Mendes, L. V. Kale, and T. R. Quinn, “Massively parallel cosmological simulations with ChaNGa,” in *Proceedings of IEEE International Parallel and Distributed Processing Symposium 2008*, 2008.
- [48] H. Menon, L. Wesolowski, G. Zheng, P. Jetley, L. Kale, T. Quinn, and F. Governato, “Adaptive techniques for clustered N-body cosmological simulations,” *Computational Astrophysics and Cosmology*, vol. 2, p. 1, Mar. 2015.
- [49] J. W. Wadsley, B. W. Keller, and T. R. Quinn, “Gasoline2: a modern smoothed particle hydrodynamics code,” *MNRAS*, vol. 471, no. 2, pp. 2357–2369, Oct. 2017.
- [50] J. G. Stadel, “Cosmological N-body simulations and their analysis,” Ph.D. dissertation, University of Washington, Seattle, Jan. 2001.
- [51] A. Boulares and D. P. Cox, “Galactic Hydrostatic Equilibrium with Magnetic Tension and Cosmic-Ray Diffusion,” *A*, vol. 365, p. 544, Dec. 1990.
- [52] R. Beck and P. Hoernes, “Magnetic spiral arms in the galaxy NGC6946,” *Nature*, vol. 379, no. 6560, pp. 47–49, Jan. 1996.
- [53] R. Pakmor and V. Springel, “Simulations of magnetic fields in isolated disc galaxies,” *MNRAS*, vol. 432, no. 1, pp. 176–193, Jun. 2013.
- [54] B.-I. Jun, M. L. Norman, and J. M. Stone, “A Numerical Study of Rayleigh-Taylor Instability in Magnetic Fluids,” *ApJ*, vol. 453, p. 332, Nov. 1995.

-
- [55] M. McCourt, R. M. O’Leary, A.-M. Madigan, and E. Quataert, “Magnetized gas clouds can survive acceleration by a hot wind,” *MNRAS*, vol. 449, no. 1, pp. 2–7, May 2015.
- [56] M. Uhlig, C. Pfrommer, M. Sharma, B. B. Nath, T. A. Enflin, and V. Springel, “Galactic winds driven by cosmic ray streaming,” *MNRAS*, vol. 423, no. 3, pp. 2374–2396, Jul. 2012.
- [57] R. Pakmor, C. Pfrommer, C. M. Simpson, and V. Springel, “Galactic Winds Driven by Isotropic and Anisotropic Cosmic-Ray Diffusion in Disk Galaxies,” *ApJ*, vol. 824, no. 2, p. L30, Jun. 2016.
- [58] I. S. Butsky and T. R. Quinn, “The Role of Cosmic-ray Transport in Shaping the Simulated Circumgalactic Medium,” *ApJ*, vol. 868, no. 2, p. 108, Dec. 2018.
- [59] A. A. Schekochihin, J. L. Maron, S. C. Cowley, and J. C. McWilliams, “The Small-Scale Structure of Magnetohydrodynamic Turbulence with Large Magnetic Prandtl Numbers,” *ApJ*, vol. 576, no. 2, pp. 806–813, Sep. 2002.
- [60] C. Federrath, G. Chabrier, J. Schober, R. Banerjee, R. S. Klessen, and D. R. G. Schleicher, “Mach Number Dependence of Turbulent Magnetic Field Amplification: Solenoidal versus Compressive Flows,” *Physical Review Letter*, vol. 107, no. 11, p. 114504, Sep. 2011.
- [61] J. Schober, D. Schleicher, C. Federrath, R. Klessen, and R. Banerjee, “Magnetic field amplification by small-scale dynamo action: Dependence on turbulence models and Reynolds and Prandtl numbers,” *Physical Review E*, vol. 85, no. 2, p. 026303, Feb. 2012.
- [62] J. Schober, D. R. G. Schleicher, C. Federrath, S. Bovino, and R. S. Klessen, “Saturation of the turbulent dynamo,” *Physical Review E*, vol. 92, no. 2, p. 023010, Aug. 2015.
- [63] G. Ruediger and L. L. Kichatinov, “Alpha-effect and alpha-quenching,” *Astronomy & Astrophysics*, vol. 269, no. 1-2, pp. 581–588, Mar. 1993.
- [64] A. Brandenburg and K. Subramanian, “Astrophysical magnetic fields and nonlinear dynamo theory,” *Physics Reports*, vol. 417, no. 1-4, pp. 1–209, Oct. 2005.
- [65] O. Gressel, D. Elstner, and U. Ziegler, “Towards a hybrid dynamo model for the Milky Way,” *Astronomy & Astrophysics*, vol. 560, p. A93, Dec. 2013.
- [66] R. Wissing and S. Shen, “Numerical dependencies of the galactic dynamo in isolated galaxies with SPH,” *Astronomy & Astrophysics*, vol. 673, p. A47, May 2023.
- [67] D. Reed, J. Gardner, T. Quinn, J. Stadel, M. Fardal, G. Lake, and F. Governato, “Evolution of the mass function of dark matter haloes,” *MNRAS*, vol. 346, no. 2, pp. 565–572, Dec. 2003.
- [68] F. Governato, B. Willman, L. Mayer, A. Brooks, G. Stinson, O. Valenzuela, J. Wadsley, and T. Quinn, “Forming disc galaxies in Λ CDM simulations,” *MNRAS*, vol. 374, no. 4, pp. 1479–1494, Feb. 2007.
- [69] T. W. Baumgarte and S. L. Shapiro, *Numerical Relativity: Starting from Scratch*. Cambridge University Press, 2 2021.
- [70] Z. B. Etienne, V. Paschalidis, R. Haas, P. Mösta, and S. L. Shapiro, “IllinoisGRMHD: An Open-Source, User-Friendly GRMHD Code for Dynamical Spacetimes,” *Class. Quant. Grav.*, vol. 32, p. 175009, 2015.
- [71] E. R. Most, L. J. Papenfort, and L. Rezzolla, “Beyond second-order convergence in simulations of magnetized binary neutron stars with realistic microphysics,” *Mon. Not. Roy. Astron. Soc.*, vol. 490, no. 3, pp. 3588–3600, 2019.
- [72] L. Baiotti and L. Rezzolla, “Binary neutron star mergers: a review of Einstein’s richest laboratory,” *Rept. Prog. Phys.*, vol. 80, no. 9, p. 096901, 2017.
- [73] H. H.-Y. Ng, J.-L. Jiang, C. Musolino, C. Ecker, S. D. Tootle, and L. Rezzolla, “A hybrid approach to long-term binary neutron-star simulations,” 12 2023.
- [74] S. Markidis, G. Lapenta *et al.*, “Multi-scale simulations of plasma with ipic3d,” *Mathematics and Computers in Simulation*, vol. 80, no. 7, pp. 1509–1519, 2010.

- [75] T. A. Bowen, S. D. Bale, R. Bandyopadhyay, J. W. Bonnell, A. Case, A. Chasapis, C. Chen, S. Curry, T. Dudok de Wit, K. Goetz *et al.*, “Kinetic-scale turbulence in the venusian magnetosheath,” *Geophysical research letters*, vol. 48, no. 2, p. e2020GL090783, 2021.
- [76] S. Ruhunusiri, J. Halekas, J. Espley, C. Mazelle, D. Brain, Y. Harada, G. DiBraccio, R. Livi, D. Larson, D. Mitchell *et al.*, “Characterization of turbulence in the mars plasma environment with maven observations,” *Journal of Geophysical Research: Space Physics*, vol. 122, no. 1, pp. 656–674, 2017.
- [77] M. Echim, T. Chang, P. Kovacs, A. Wawrzaszek, E. Yordanova, Y. Narita, Z. Vörös, R. Bruno, W. Macek, K. Mursula *et al.*, “Turbulence and complexity of magnetospheric plasmas,” *Magnetospheres in the solar system*, pp. 67–91, 2021.
- [78] G. Lapenta, F. Pucci, M. Goldman, and D. Newman, “Local regimes of turbulence in 3d magnetic reconnection,” *The Astrophysical Journal*, vol. 888, no. 2, p. 104, 2020.
- [79] S. Huang, F. Sahraoui, A. Retinò, O. Le Contel, Z. Yuan, A. Chasapis, N. Aunai, H. Breuillard, X. Deng, M. Zhou *et al.*, “Mms observations of ion-scale magnetic island in the magnetosheath turbulent plasma,” *Geophysical Research Letters*, vol. 43, no. 15, pp. 7850–7858, 2016.
- [80] K.-J. Hwang, D. Sibeck, J. Burch, E. Choi, R. Fear, B. Lavraud, B. Giles, D. Gershman, C. Pollock, J. Eastwood *et al.*, “Small-scale flux transfer events formed in the reconnection exhaust region between two x lines,” *Journal of Geophysical Research: Space Physics*, vol. 123, no. 10, pp. 8473–8488, 2018.
- [81] R. A. Treumann and W. Baumjohann, “Collisionless magnetic reconnection in space plasmas,” *Frontiers in Physics*, vol. 1, p. 31, 2013.
- [82] K.-I. Nishikawa and S.-i. Ohtani, “Global particle simulation for a space weather model: Present and future,” *IEEE transactions on plasma science*, vol. 28, no. 6, pp. 1991–2006, 2000.
- [83] R. Treumann and W. Baumjohann, “Spontaneous magnetic reconnection: Collisionless reconnection and its potential astrophysical relevance,” *The Astronomy and Astrophysics Review*, vol. 23, pp. 1–91, 2015.
- [84] A. Lazarian, G. Kowal, M. Takamoto, E. de Gouveia Dal Pino, and J. Cho, “Theory and applications of non-relativistic and relativistic turbulent reconnection,” *Magnetic Reconnection: Concepts and Applications*, pp. 409–471, 2016.
- [85] R. Teyssier, “Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES,” *A&A*, vol. 385, pp. 337–364, Apr. 2002.
- [86] “Repository of RAMSES on Bitbucket,” <https://bitbucket.org/rteyssie/ramses/src>, accessed: 2023-11-01.
- [87] T. Guillet and R. Teyssier, “A simple multigrid scheme for solving the Poisson equation with arbitrary domain boundaries,” *Journal of Computational Physics*, vol. 230, no. 12, pp. 4756–4771, Jun. 2011.
- [88] Y. Dubois, R. Beckmann, F. Bournaud, H. Choi, J. Devriendt, R. Jackson, S. Kaviraj, T. Kimm, K. Kraljic, C. Laigle, G. Martin, M.-J. Park, S. Peirani, C. Pichon, M. Volonteri, and S. K. Yi, “Introducing the NEWHORIZON simulation: Galaxy properties with resolved internal dynamics across cosmic time,” *A&A*, vol. 651, p. A109, Jul. 2021.
- [89] V. Springel, C. S. Frenk, and S. D. M. White, “The large-scale structure of the Universe,” *Nature*, vol. 440, no. 7088, pp. 1137–1144, Apr. 2006.
- [90] O. Hahn and T. Abel, “Multi-scale initial conditions for cosmological simulations,” *MNRAS*, vol. 415, no. 3, pp. 2101–2121, Aug. 2011.
- [91] J. Rosdahl, J. Schaye, R. Teyssier, and O. Agertz, “Galaxies that shine: radiation-hydrodynamical simulations of disc galaxies,” *MNRAS*, vol. 451, no. 1, pp. 34–58, Jul. 2015.