



Full length paper

The PLUTO code on GPUs: A first look at Eulerian MHD methods[☆]

M. Rossazza^{a,*}, A. Mignone^a, M. Bugli^{b,a,c}, S. Truzzi^a, L. Riha^d, T. Panoc^d, O. Vydrovsky^d,
N. Shukla^e, A. Romeo^e, V. Berta^{f,g}

^a Dipartimento di Fisica, Università degli Studi di Torino, Via Pietro Giuria 1, I-10125 Torino, Italy

^b Institut d'Astrophysique de Paris, UMR 7095, CNRS & Sorbonne Université, 98 bis bd Arago, 75014 Paris, France

^c INFN - sezione di Torino, Via Pietro Giuria 1, I-10125 Torino, Italy

^d IT4Innovations, VSB - Technical University of Ostrava, Ostrava, Czech Republic

^e CINECA, Department of High Performance Computing, Via Magnanelli 6, 3, 40033 Casalecchio di Reno, BO, Italy

^f Canadian Institute for Theoretical Astrophysics, University of Toronto, ON M5S 3H8, Toronto, Canada

^g Department of Astronomy, Columbia University, 550 W 120th St, NY 10027, New York, USA

ARTICLE INFO

Keywords:

GPU computing

OpenACC

PLUTO code

Parallel computing

High-performance computing (HPC)

Magnetohydrodynamics (MHD)

ABSTRACT

We present preliminary performance results of gPLUTO, the new GPU-optimized implementation of the PLUTO code for computational plasma astrophysics. Like its predecessor, gPLUTO employs a Eulerian finite-volume formulation to numerically solve the equations of magnetohydrodynamics (MHD) in multiple spatial dimensions. Still, this new implementation is a complete rewrite in C++ and leverages the OpenACC programming model to achieve acceleration on NVIDIA GPUs. While a more comprehensive description of the code and its several other modules will be presented in a future paper, here we focus on some preparatory results that demonstrate the code potential and performance on pre exa-scale parallel architectures.

1. Introduction

Magnetohydrodynamics (MHD) simulations play a pivotal role in understanding a vast range of complex physical phenomena. In the context of astrophysical processes, these include the formation and evolution of planets (Lesur, 2021), stars (Teyssier and Commerçon, 2019), and galaxies (Zhang, 2018), the acceleration of high-energy particles (Lazarian et al., 2012; Amato and Blasi, 2018), and the multi-wavelength emission associated to the activity of compact objects such as black holes (Narayan, 2005) and neutron stars (Faber and Rasio, 2012; Uzdensky and Rightley, 2014). There are also numerous terrestrial applications that exploit the predictive power of the MHD framework, such as nuclear fusion (Knaepen and Moreau, 2008; Inghirami et al., 2016), space weather predictions (Shibata and Magara, 2011), and plasma confinement (Doyle et al., 2007).

The computational demands of MHD simulations are inherently tied to the need to describe at once spatial and temporal scales that often span many orders of magnitude. The ability to resolve phenomena occurring at scales much smaller than the typical size of the global system (e.g., the turbulent structures in the plasma constituting an accretion disk or a jet propagating through the interstellar medium) plays a crucial role in enabling a far deeper understanding of their properties and allowing more accurate predictions of their observational/experimental

signatures. Prime examples in the past decade of the qualitative impact of higher computational efficiency in multi-dimensional astrophysical MHD simulations include models of stellar explosions (Mezzacappa, 2023; Janka, 2025), dynamo processes (Charbonneau, 2020; Del Zanna et al., 2022), and accretion/ejection of matter around central compact objects (Mayer, 2019; Cattorini and Giacomazzo, 2024). In recent years, the field of high-performance computing has witnessed a transformative shift thanks to the advent of Graphics Processing Units (GPUs) that offer massive parallelism and computational capabilities. The use of GPUs in scientific computing has indeed opened up new frontiers in numerical simulations, enabling researchers to explore complex physical systems at unprecedented resolutions and speeds. Notable examples are the Eulerian codes H-AMR (Liska et al., 2022), IDEFIX (Lesur et al., 2023), AthenaK (Stone et al., 2024), ECHO (Del Zanna et al., 2024), and AsterX (Kalinani et al., 2025).

In this paper, we present the latest noteworthy advancements in the development of the GPU-ready version of the PLUTO code (Mignone et al., 2007, 2012b) for modeling astrophysical plasmas. PLUTO has been widely used by hundreds of researchers in the astrophysical community over the past two decades. Owing to the variety of physical modules provided (i.e., from classical to relativistic flows, ideal and non-ideal dynamics, different geometries, and equations of state), the

[☆] This article is part of a Special issue entitled: 'HPC in Cosmology and Astrophysics' published in Astronomy and Computing.

* Corresponding author.

E-mail address: marco.rossazza@unito.it (M. Rossazza).

code has been employed in very different astrophysical contexts, ranging from high-energy astrophysics (Mignone et al., 2010a, 2013; Mattia et al., 2023; Pavan et al., 2023) to young stellar objects (Orlando et al., 2011), supernova remnants (Olmi et al., 2018; Olmi and Bucciantini, 2019; Orlando et al., 2019, 2020), plasma physics (Bodo et al., 2022; Bugli et al., 2025), planet formation (Lesur, 2021; Flock and Mignone, 2021), and solar physics (Réville et al., 2015; Perri et al., 2018), to name just a few.

The new GPU-accelerated implementation of the code, gPLUTO (<https://gitlab.com/PLUTO-code/gPLUTO>), is a complete C++ rewrite of the legacy code (<https://plutocode.ph.unito.it>), which employs the OpenACC programming model to harness the vast computing power and energetic efficiency of modern GPU architectures, offering significant advantages in terms of computational speed, power consumption, and cost-effectiveness. gPLUTO has superseded the PLUTO code, which will no longer receive updates. OpenACC was selected for gPLUTO primarily due to the direct support provided by NVIDIA and the model's intrinsic simplicity relative to low-level frameworks. Its directive-based nature enables a minimally intrusive, incremental porting strategy that preserves the original codebase structure while significantly lowering the development effort. Switching to C++ is justified by the need to take advantage of specific language extensions such as function templates, classes and polymorphism. gPLUTO keeps the strengths of the previous versions, such as stability, versatility, and user-friendliness, while also giving access to the most recent high-performance computing resources. This work outlines the code's fundamental design principles and the remarkable results achieved in terms of simulation performance, while a more comprehensive description of the new code and all of its physical modules will appear in a future manuscript.

In this work we present the technical details of gPLUTO by introducing its underlying algorithmic framework, data management and parallelization strategies, and implementation solutions. We also showcase a series of benchmark tests, highlighting the code's effectiveness in reproducing them as well as its notable performance. This paper has a companion, complementary paper that focuses on the Lagrangian Particles module integration and its GPU offloading (Suriano et al., 2025). The paper is organized as follows: in Section 2, we briefly recap the equations of MHD and the finite volume framework. In Section 3, we illustrate technical aspects concerning the OpenACC implementation, while in Section 4, we show the parallel performance on two standard use cases. Finally, we draw our conclusions in Section 5.

2. Numerical methods

2.1. MHD equations

Following its predecessor (Mignone et al., 2007, 2012b), gPLUTO is designed to solve a variety of systems of conservation laws,

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

where \mathcal{U} is an array of conservative quantities, \mathcal{T} is a flux tensor and \mathcal{S} defines the source terms, which accounts, e.g., for geometrical source terms, gravity and rotational effects. Explicit expressions may be found, for instance, in Mignone et al. (2007, 2012a) and Mignone (2014).

Although different sets of equations may be solved, for the present purpose, we only focus on the single fluid classical MHD equations:

$$\mathcal{U} = \begin{pmatrix} \rho \\ \rho \mathbf{v} \\ \mathbf{B} \\ \mathcal{E} \end{pmatrix}, \quad \mathcal{T} = \begin{pmatrix} \rho \mathbf{v} \\ \rho \mathbf{v} \mathbf{v} - \mathbf{B} \mathbf{B} + p_t \mathbf{I} \\ \mathbf{v} \mathbf{B} - \mathbf{B} \mathbf{v} \\ (\mathcal{E} + p_t) \mathbf{v} - (\mathbf{v} \cdot \mathbf{B}) \mathbf{B} \end{pmatrix}^T, \quad (2)$$

where ρ is the fluid density, \mathbf{v} is the fluid velocity, \mathbf{B} is the magnetic field vector, $p_t = p + \mathbf{B}^2/2$ is the total (thermal + magnetic) pressure and

$$\mathcal{E} = \frac{p}{\Gamma - 1} + \frac{\rho \mathbf{v}^2}{2} + \frac{\mathbf{B}^2}{2}, \quad (3)$$

is the total energy density comprising thermal, kinetic and magnetic contributions. In Eq. (3) we have assumed an ideal equation of state with adiabatic index Γ .

The divergence-free constraint, $\nabla \cdot \mathbf{B} = 0$, is not automatically preserved by the underlying discretization method and must be specifically enforced. To this purpose, the code allows different algorithms to be selected, including include Powell's eight wave formulation (Powell et al., 1999), divergence cleaning (Dedner et al., 2002; Mignone et al., 2010b) and the constrained transport method (see Evans and Hawley, 1988; Balsara and Spicer, 1999) in the more recent formulation by Mignone and Del Zanna (2021).

Relativistic and non-ideal extensions of the previous equations will be described in a forthcoming work.

2.2. Discretization

We employ a Cartesian coordinate system with axes defined by the unit vectors $\hat{\mathbf{e}}_x = (1, 0, 0)$, $\hat{\mathbf{e}}_y = (0, 1, 0)$ and $\hat{\mathbf{e}}_z = (0, 0, 1)$. A uniform mesh with coordinate spacing Δx , Δy and Δz is assumed to cover the computational domain.

gPLUTO employs a finite volume (FV) framework to evolve cell-averaged conservative variables in time, given specified initial and boundary conditions. The FV method is based on the integral form of the governing differential equations, obtained by integrating Eq. (1) over a control volume corresponding to a single computational zone. Integrating over space within a single cell yields the semi-discrete form of the conservation equations:

$$\begin{aligned} \frac{d \langle \mathcal{U} \rangle_c}{dt} = & - \frac{\Delta t}{\Delta x} \left(\hat{F}_{i+\frac{1}{2},j,k} - \hat{F}_{i-\frac{1}{2},j,k} \right) \\ & - \frac{\Delta t}{\Delta y} \left(\hat{G}_{i,j+\frac{1}{2},k} - \hat{G}_{i,j-\frac{1}{2},k} \right) \\ & - \frac{\Delta t}{\Delta z} \left(\hat{H}_{i,j,k+\frac{1}{2}} - \hat{H}_{i,j,k-\frac{1}{2}} \right), \end{aligned} \quad (4)$$

where $c = (i, j, k)$ is a short-hand notation for a single computational cell centered at (x_i, y_j, z_k) , $\{i, j, k\}$ is a triplet of zone indices in the three coordinate directions while cell interfaces are located at $c \pm \hat{\mathbf{e}}_x/2$, $c \pm \hat{\mathbf{e}}_y/2$, and $c \pm \hat{\mathbf{e}}_z/2$ in the three directions, respectively. gPLUTO also supports cylindrical and spherical coordinates using the same approach presented in Mignone (2014) and we remind to a future work for a more general discussion.

In Eq. (4), $\langle \mathcal{U} \rangle_c$ represents the volume-averaged conserved quantities within each cell, while \hat{F} , \hat{G} , and \hat{H} are the surface-averaged numerical fluxes across the corresponding cell interfaces, computed via a Riemann solver.

At the numerical level, Eq. (4) is treated as a system of ordinary differential equations (ODEs) in time, and discretized using a strong-stability-preserving (SSP) Runge–Kutta (RK) method (Shu and Osher, 1988). For example, a third-order SSP Runge–Kutta scheme proceeds as:

$$\begin{aligned} \langle \mathcal{U} \rangle^* &= \langle \mathcal{U} \rangle^n + \mathcal{L}^n \\ \langle \mathcal{U} \rangle^{**} &= \frac{3 \langle \mathcal{U} \rangle^n + \langle \mathcal{U} \rangle^* + \mathcal{L}^*}{4} \\ \langle \mathcal{U} \rangle^{n+1} &= \frac{\langle \mathcal{U} \rangle^n + 2 \langle \mathcal{U} \rangle^{**} + \mathcal{L}^{**}}{3}, \end{aligned} \quad (5)$$

where, for brevity, the cell index subscript c has been omitted. In the expression above, $\mathcal{L}^n \equiv \mathcal{L}(\langle \mathcal{U} \rangle^n)$ represents the right-hand side of Eq. (4), evaluated using the solution at time level n , and similarly for \mathcal{L}^* and \mathcal{L}^{**} at intermediate stages.

For each Runge–Kutta stage, the right-hand side \mathcal{L} is typically computed by first performing a reconstruction step, in which zone-centered primitive variables are extrapolated to cell interfaces. For instance, in the x -direction:

$$\mathcal{V}_{i+\frac{1}{2}}^L = \mathcal{R}_x^+(\mathcal{V}_c), \quad \mathcal{V}_{i-\frac{1}{2}}^R = \mathcal{R}_x^-(\mathcal{V}_c) \quad (6)$$

where \mathcal{R}_x^\pm denotes a non-oscillatory, high-order polynomial reconstruction operator (see, e.g., Colella and Woodward, 1984; Suresh and Huynh, 1997; Borges et al., 2008), applied in the positive (+) or negative (−) x -direction.

Here, $\mathcal{V}_c = \{\rho, \mathbf{v}, \mathbf{B}, p\}$ represents the array of primitive variables (density, velocity, magnetic field, and pressure), which are preferred over conservative variables during the reconstruction process. This choice is made because primitive variables typically exhibit smoother variations across the computational domain — except at shocks or contact discontinuities — thereby reducing numerical errors and mitigating spurious oscillations near sharp gradients. Employing primitive variables during reconstruction helps limit the occurrence of unphysical states, such as negative pressures, especially in regions near strong gradients. In MHD, primitive variables are recovered from the conservative ones via straightforward inversion:

$$\mathbf{v} = \frac{\rho \mathbf{v}}{\rho}, \quad p = (\Gamma - 1) \left(\mathcal{E} - \frac{\rho \mathbf{v}^2}{2} - \frac{\mathbf{B}^2}{2} \right) \quad (7)$$

where Γ is the adiabatic index, and \mathcal{E} is the total energy density.

The reconstructed values at cell interfaces serve as input states to the Riemann solver, which computes a stable, upwinded numerical flux. For example:

$$\hat{F}_{i+\frac{1}{2}} = \frac{F(\mathcal{V}_{i+\frac{1}{2}}^L) + F(\mathcal{V}_{i+\frac{1}{2}}^R)}{2} - D_{i+\frac{1}{2}} \cdot \left(\mathcal{U}_{i+\frac{1}{2}}^R - \mathcal{U}_{i+\frac{1}{2}}^L \right) \quad (8)$$

where D is a dissipation matrix whose form depends on the specific Riemann solver employed. The subscripts j and k have been omitted here for clarity, as the context is now unambiguous.

Magnetic fields can be evolved either in a zone-centered formalism, as described so far, or using the constrained transport (CT) method (Evans and Hawley, 1988; Londrillo and del Zanna, 2004). The advantage of CT is its ability to maintain the divergence-free condition for the magnetic field to machine precision throughout the simulation. Further details on the numerical implementation can be found in the original PLUTO code paper (Mignone et al., 2007), and in recent developments related to CT techniques in Mignone and Del Zanna (2021).

As its predecessor, gPLUTO offers a variety of algorithms with different levels of accuracy and computational cost to choose from. In the present work, we employ the third-order SSP Runge–Kutta time-stepping scheme (Eq. (5)), combined with either parabolic reconstruction (Mignone et al., 2005) or the WENOZ reconstruction method (Borges et al., 2008).

3. PLUTO design for GPUs

Modern Graphics Processing Units (GPUs) are specialized, highly parallel processors that have become cornerstones of high-performance scientific computing. Their computational power originates from an architecture containing hundreds or thousands of simple cores designed to execute the same instruction simultaneously across large sets of data. This paradigm makes GPUs exceptionally efficient for algorithms characterized by high data parallelism, such as the mesh-based solvers used in our work. Computations are offloaded from the host CPU to the GPU (called in this context device or accelerator) by launching functions called kernels, which are managed and executed by the device’s hardware scheduler. The execution is organized hierarchically. A thread is a single stream of instructions of that runs in parallel with other threads and is mapped to the GPU cores. Individual threads that execute the same instruction are grouped into warps and multiple warps are contained in blocks, that are scheduled for execution on the GPU’s fundamental hardware units, the Streaming Multiprocessors (SMs). Each SM is an independent multiprocessor containing its own execution cores, scheduler, and a high-speed shared memory, allowing it to execute one or more thread blocks concurrently.

3.1. OpenACC

The programming model chosen for GPU porting of PLUTO is OpenACC, a simpler and higher-level alternative to CUDA. OpenACC is a directive-based programming model for GPU computing developed by Cray, CAPS, NVIDIA, and PGI in 2011. It relies on information provided by the programmer to the compiler through directives, library routines, and environment variables. These annotations instruct the compiler in porting the code to the device. The complex and low-level aspect of the parallelization is left to be handled by the compiler, facilitating the user’s work. Another advantage of choosing a directive-based approach is interchangeability, an important feature allowing the code to be compiled and run, ignoring the OpenACC directives, making it suitable for both CPU and GPU computing. For every GPU-accelerated application, the computing work is divided among CPUs and GPUs. In this context, “host” refers to the primary processor in a system (the CPU) while the “device” refers to the secondary processor (the GPU) responsible for the acceleration.

In broad terms, in gPLUTO the entire computational workload is handled by the GPUs, while the CPUs oversee overall system operations, including running the operating system, managing system memory, and handling I/O operations. Data movement is strictly limited to two scenarios: staging data to the host for disk output at specified intervals, and exchanging boundary conditions between sub-domains. For the latter, the implementation leverages CUDA-aware MPI, which accepts device pointers directly. This bypasses the host memory during halo exchanges and enables direct data transfer between GPUs across nodes. The reason for this strict division is performance. Executing all the computations on the device permits, in fact, the exploitation of its potential without introducing data movement between the memories of the two processors. The alternate use of host and device would introduce memory copies to assure coherence between data and, hence, correct results. OpenACC directives are indicated with the keyword `pragma` followed by the `acc` sentinel, the directive itself, and optional clauses:

```
#pragma acc < directive > < clause , clause >
```

The directive is applied to the instruction, loop, or structured code block that follows. The base approach when accelerating an application with OpenACC is using compute directives (we mainly use `parallel loop`) that indicate to the compiler regions of the code to execute on the device, the kernels. Data directives, on the other hand, such as `enter data copyin` and `exit data delete`, permit explicit control of the data movement. This is crucial since, in many cases, the performance of the kernels is limited by the memory bandwidth of the device, and thus, reducing the amount of data needed to be moved gives a huge benefit.

3.2. Compute directives

A typical compute construct of gPLUTO is reported in Fig. 1, where `NX`, `NY`, `NZ` define domain size, `data->Vc` is the array of the primitive quantities (i.e., density, pressure and velocities), that contains `NVAR` physical quantities for each point of the 3-dimensional domain.

The first directive is at the core of the acceleration of the application: `#pragma acc parallel loop` directive tells the compiler to parallelize the loop, distributing its iterations across the threads of the target accelerator. The loop body is then executed concurrently, significantly improving the performance of the code. The `collapse(3)` clause applies in this case to three tightly nested loops. Hence, the compiler treats the construct as a single loop of `NX x NY x NZ` iterations. This makes it easier for the compiler to map the construct onto the device’s computing resources, potentially improving parallelization efficiency. The present clause implies that data is already present on the GPU and that a copy from the main memory is not necessary. If this is not true a fatal error at the execution would be generated.

```

#pragma acc parallel loop collapse(3) present(data)
for (k = 0; k < NZ; k++){
for (j = 0; j < NY; j++){
for (i = 0; i < NX; i++){
...
#pragma acc loop seq
for (nv = 0; nv < NVAR; nv++){
data->Vc(i,j,k,nv) = ...;
}
...
}}}}

```

Fig. 1. Kernel accelerated with a compute directive.

The innermost loop, over NVAR elements, is marked with `seq`, meaning that the operation on different values of `nv` is executed sequentially by the same thread. Note that while parallelization of this inner loop could be possible, it usually tends to lower the performance, as the number of variables NVAR is typically much smaller than the number of grid points. Moreover, having the threads corresponding to loops with many iterations makes sense for optimal utilization and memory access (see Section 3.3).

3.3. Memory management

To manage data transfer to the device in gPLUTO, we combine explicit OpenACC data directives with CUDA Unified Memory, enabled via the compile-time flag `-gpu=managed`. Unified Memory provides a shared address space accessible by both the CPU and GPU through a single pointer, making it particularly suitable for dynamically allocated data. In most cases, the NVIDIA driver automatically handles data movement. However, we employ explicit data directives to ensure correct deep copying of structured and nested data that are widely used within the code.

GPUs employ various techniques to hide memory latency that can represent an important performance bottleneck. Most of these techniques, such as warp scheduling¹ and prefetching² are handled by the device itself but the programmer can actively act on other important aspects such as coalesced memory access.

Coalesced memory access. Coalesced memory access is a memory access pattern in parallel computing, particularly on GPUs, where multiple threads access consecutive or contiguous memory locations in a single memory transaction. This pattern is essential for achieving efficient memory bandwidth utilization and improving the overall performance of parallel applications. In the pseudo-code of Fig. 1, coalesced memory access is achieved, granting optimal parallelization. Here, consecutive values of the index `i` correspond to consecutive threads but also to adjacent memory access (`i` is the fastest index of `data->Vc`). The advantage of this memory pattern can be traced back to the thread memory access. When a thread needs data for an operation, it receives access to many adjacent memory locations. Neighbor threads can be exploited immediately if they need data received in this memory transaction, hence the advantage of them accessing adjacent memory locations.

¹ When some threads within a warp need to access memory, the scheduler can switch to other warps with active threads, allowing the GPU to continue executing instructions without waiting for the memory access to complete.

² GPUs can predict and preload data into caches before it is actually needed, reducing the time it takes to fetch the data when it is required by the computation.

```

#pragma acc parallel loop collapse(2)
for (k = 0; k < NZ; k++){
for (j = 0; j < NY; j++){
long int offset = NVAR*NX*(j + NY*k);
Ary2D flux(&d->sweep.flux[offset],NX,NVAR);
Ary2D rhs(&d->sweep.rhs[offset],NX,NVAR);

#pragma acc loop vector
for (i = 0; i < NX; i++){
#pragma acc loop seq
NVAR_LOOP(nv) {
rhs(i,nv) = -dtdx*(flux(i+1,nv) - flux(i,nv));
}
}}}}

```

Fig. 2. Computation of the right-hand side contribution along the x -direction.

C++ classes. Optimal memory access pattern is one of the main reasons behind the profound change of many arrays of the code together with the migration to C++. We introduced classes to manage memory for multi-dimensional arrays, replacing the legacy pointer-to-pointer approach with flattened, contiguous memory blocks. This abstraction hides index arithmetic and encapsulates memory layout strategies, allowing element ordering to adapt automatically based on the array type or OpenACC compilation flags. We now discuss a technique we adopted widely, using the class for 2D arrays to improve memory access.

In Fig. 2 `rhs` is a 2D array of size `NX x NVAR` created from the single, contiguous block of memory `d->sweep.rhs` of size `NX x NY x NZ x NVAR`, exploiting the discussed class method and based on an offset defined by `j` and `k`. This assures that each thread writes on a unique memory address (privatization³), which is necessary to grant correct results. Moreover, class definition assures that neighboring threads corresponding to consecutive values of `i` access adjacent memory locations.

Then, when computing the y -direction contribution, the same technique is applied to the same memory pool `d->sweep.rhs`, changing the order of write and access to the elements to obtain again optimal access pattern again. Other arrays such as `Vc` and `Uc` have fixed ordering: `i, j, k, nv` in order of decreasing index speed.

Template function. Another useful feature of C++ that has been adopted is function templates. Many kernels operate on arrays whose indexing depends on the direction of integration, which can vary at runtime. This runtime variability prevents the compiler from knowing in advance which registers will be needed, often resulting in these variables being placed in slower memory spaces such as local or shared memory instead of registers. By using templated functions, we can specialize the code based on template parameters—in our case, the integration direction `dir`. This enables the compiler to inline the specialized functions and propagate constant values, eliminating branches and improving performance. Fig. 3 shows an example of the Riemann solver with `dir` as a template parameter. In this case, the array `v` can be allocated in registers at compile time.

3.4. Multi GPU communication

Multi-process communication at each time step is necessary in gPLUTO to assure correct results when the domain is split between processes. Communication is based on the message passing interface (MPI) protocols. While in PLUTO the data exchange was based on

³ Thread privatization is the process of giving each thread within a kernel its own private copy of a variable, ensuring that parallel computations do not interfere with each other by creating race conditions.

```

template<int dir>
void Riemann_Solver (Data *d, Grid *grid, RBox *box);

void Riemann_Solver_T (Data *d, Grid *grid, RBox *box){
  if(dir==IDIR) Riemann_Solver<IDIR>(d, grid, box);
  else if(dir==JDIR) Riemann_Solver<JDIR>(d, grid, box);
  else if(dir==KDIR) Riemann_Solver<KDIR>(d, grid, box);
}

template<int dir>
void Riemann_Solver (Data *d, Grid *grid, RBox *box)
{
  int VXn = VX1+dir;

  #pragma acc parallel loop collapse(2)
  for (k = kbeg; k <= kend; k++){
    for (j = jbeg; j <= jend; j++){
      #pragma acc loop
      for (i = ibeg; i <= iend; i++){
        int v[NVAR];
        v[VXn] = ...;
      }
    }
  }
}

```

Fig. 3. Templated Riemann Solver. Three version of the routine are compiled with different values of the direction of integration dir: IDIR, JDIR, KDIR.

the library ArrayLib (Mignone et al., 2012), it has here been rewritten using a more basic scheme based on standard buffer arrays and MPI_Isend/MPI_Irecv calls. To exploit the characteristics of GPUs, we employ an asynchronous approach for data exchange. This requires each processor to communicate potentially simultaneously with its 8 neighbors in 2D or 26 neighbors in 3D (this includes also nearby processors in the diagonal directions, necessary for various algorithms such as Constrained Transport with UCT methods Mignone and Del Zanna, 2021). The pseudo-code in Fig. 4 describes the dedicated routine, considering one direction of exchange (IDIR).

Initially, the code initializes non-blocking receives (MPI_Irecv) to get data from neighboring processes. Then, data to be sent to neighboring processes are packed into buffers allocated with acc_malloc,⁴ and each kernel has an async clause with a different number: each kernel is independent, and they could ideally run at the same time on distinct streams (sequences of operations). After the code initializes non-blocking send calls (MPI_Isend) to send data to neighboring processes and then it waits for all receive operations to complete. When this happens, the received data is unpacked into the appropriate variables, again concurrently for each recv_buf. Finally, it waits for all non-blocking send operations to complete and ensures all asynchronous operations on the GPUs are terminated before proceeding.

4. Numerical tests and discussion

Here, we present numerical benchmarks conducted on CPU and GPU based partitions of three different HPC platforms: MareNostrum 5 Accelerated Partition (GPU), Leonardo Booster (GPU), MeluXina GPU nodes (GPU), and Leonardo DCGP (CPU). Every run is performed employing the entire computational power of each node, meaning 4 GPUs per node on GPU partitions and 112 CPU-cores on Leonardo DCGP.

Leonardo is hosted by Cineca and has two main partitions: Booster module and Data-centric General Purpose (DCGP) module. The Booster partition is equipped with Atos BullSequana X2135 “Da Vinci” single-node GPU blade with four NVIDIA Ampere GPUs/node (64 GB HBM2e).

⁴ OpenACC function that allocates memory directly in the accelerator’s global memory.

```

MPI_Irecv (recv_buf_2, ..., nrnks[IDIR][1], ...);
MPI_Irecv (recv_buf_1, ..., nrnks[IDIR][0], ...);

#pragma acc parallel loop collapse(4) async(1)
LOOP(...){
  send_buf_1[ind] = Vc((ibeg+i),j,k,nv);
}

#pragma acc parallel loop collapse(4) async(2)
LOOP(...){
  send_buf_2[ind] = Vc((iend-(nlg-1)+i),j,k,nv);
}

#pragma acc wait

MPI_Isend(send_buf_1, ..., nrnks[IDIR][0], ...);
MPI_Isend(send_buf_2, ..., nrnks[IDIR][1], ...);

MPI_Waitall(..., &recv[2], ...);
MPI_Waitall(..., &recv[1], ...);

#pragma acc parallel loop collapse(4) async(1)
LOOP(...){
  Vc(i,j,k,nv) = recv_buf_1[ind];
}

#pragma acc parallel loop collapse(4) async(2)
LOOP(...){
  Vc((iend+1+i),j,k,nv) = recv_buf_2[ind];
}

#pragma acc wait

MPI_Waitall(..., &send[1], ...);
MPI_Waitall(..., &send[2], ...);

```

Fig. 4. Asynchronous implementation of the parallel data exchange routine.

The DCGP module has instead Atos BullSequana X2140 three-node CPU blade, with dual socket 56-core Intel Sapphire Rapids Intel Xeon Platinum 8480+. MeluXina belongs to LuxProvide and the GPU nodes provide Atos BullSequana XH2000 platform with four NVIDIA Ampere GPUs (40 GB HBM). Finally, MareNostrum 5 is at the Barcelona Supercomputing Center (BSC), and the Accelerated Partition (ACC) we used is a BullSequana XH3000 model with four NVIDIA Hopper GPUs (64 GB HBM2). See Table 1 for a summary of the systems specifics.

The two chosen benchmarks are a 3D Orszag–Tang vortex Section 4.1 (which involves propagation of discontinuities) and a 3D circularly polarized Alfvén wave test Section 4.2 (which presents a smooth wave propagation problem).

4.1. 3D Orszag–Tang vortex

The first benchmark is a 3D version of the well-known Orszag–Tang test problem, with a similar to that of Helzel et al. (2011). We consider a periodic Cartesian domain of size $L_x = L_y = L_z = 1$ with constant initial pressure $p = 5/3$ and density $\rho = 25/9$. Velocity and magnetic field are initialized as follows

$$\mathbf{v} = -\zeta(z) \sin(2\pi y)\hat{e}_x + \zeta(z) \sin(2\pi x)\hat{e}_y + 0.2 \sin(2\pi z)\hat{e}_z, \quad (9)$$

$$\mathbf{B} = -\sin(2\pi y)\hat{e}_x + \sin(4\pi x)\hat{e}_y,$$

where $\zeta(z) = 1 + \sin(2\pi z)/5$. Computations are stopped at $t = 0.5$ and the test employs the HLLD Riemann solver (Miyoshi and Kusano, 2005), WENOZ reconstruction (Borges et al., 2008), and the 3rd-order Runge–Kutta (RK3) time-stepping scheme. All computations are performed in double-precision arithmetic.

Results validation. Numerical correctness of the benchmark on gPLUTO was verified by ensuring that maximum and minimum flux quantities, as well as the maximum Mach number, matched the legacy code to

Table 1
Summary of computational hardware resources utilized.

Detail	Leonardo		MeluXina	MareNostrum 5
	Booster	DCGP	GPU	ACC
Proprietary company	CINECA		LuxProvide	BSC
Active from year	2022		2021	2023
Node model	BullSequana X2135	BullSequana X2140	BullSequana XH2000	BullSequana XH3000
CPUs (per node)	Intel Ice Lake 8358	2x 56-core Intel Sapphire Rapids 8480+	2x AMD EPYC Rome 7452	2x Intel Sapphire Rapids 8460Y+
GPUs (per node)	4x NVIDIA Ampere (64 GB HBM2e)	N/A (CPU-only)	4x NVIDIA Ampere (40 GB HBM)	4x NVIDIA Hopper (64 GB HBM2)

Table 2
Weak scaling resolution (per node) for the 3D Orszag–Tang test.

System	Resolution per node
Leonardo Booster	$704 \times 704 \times 352$
Leonardo DCGP	$704 \times 704 \times 350$
MareNostrum 5 ACC	$704 \times 704 \times 352$
MeluXina GPU nodes	$512 \times 512 \times 512$

sixteen decimal places. Additionally, a visual comparison of the density profile at $t = 0.5$ in ($y = 0.3, z = 0.5$) is displayed in the lower panel of Fig. 5 while the upper panel reports the correspondent density map. While numerical results are strictly consistent, CPU performance remains inconclusive; We conducted comparison tests with PLUTO on various systems and the execution time ratio between the two codes fluctuates significantly around unity depending on the system architecture, necessitating further detailed analysis.

Weak scaling. We first assess the parallel performance using weak scale metric, implying that grid resolution per node — reported in Table 2 — is kept constant. For these runs, we employ the constrained transport method (with UCT HLLD average, Mignone and Del Zanna, 2021) to enforce the divergence-free condition of magnetic field. On the GPU partitions, we chose a resolution that maximized device occupancy to achieve the best performance. A $704 \times 704 \times 352$ resolution was used to nearly fill the device memory on both Leonardo Booster and Marenostrum 5 ACC. A lower resolution was necessary for MeluXina to accommodate its smaller GPU memory. The runs on Leonardo DCGP used a slightly different resolution to allow for domain decomposition across 112 processes while maintaining a nearly identical computational load per process as: $\frac{N_{t,CPU\text{node}}}{N_{t,GPU\text{node}}} \sim 0.99$, with N_t the total number of grid points.

Fig. 6(a) presents the parallel efficiency expressed as the ratio of the wall clock times on 1 node and N nodes, T_1/T_N . We have repeated the test on all 4 HPC systems up to 32, 128, 256 and 512 nodes for MeluXina GPU, Leonardo DCGP, MareNostrum 5 ACC and Leonardo Booster respectively. Among the tested systems, MareNostrum 5 ACC delivered the best results, maintaining a parallel efficiency above 95% even at 256 nodes. On Leonardo Booster, we scaled up to 512 nodes, achieving 88% efficiency at the largest scale. In contrast, MeluXina GPU nodes showed a steeper efficiency drop, reaching 91% at just 32 nodes. The CPU runs on Leonardo DCGP maintain an efficiency above 95% for up to 64 nodes, but this value drops to 84% at 128 nodes. Further investigation of the reasons for this low efficiency at 128 nodes will be needed.

In Fig. 6(b), we show the corresponding performance of the weak scaling runs, measured in terms of the number of cells updated per second per node (computed as the number of grid points multiplied by the number of time steps divided by the execution time and the number of nodes).

From the same runs we retrieved the acceleration factors between runs on the Leonardo Booster and those on Leonardo DCGP. Note in the table in Fig. 6(c) that CPU runs were approximately 9 times slower than GPU runs.

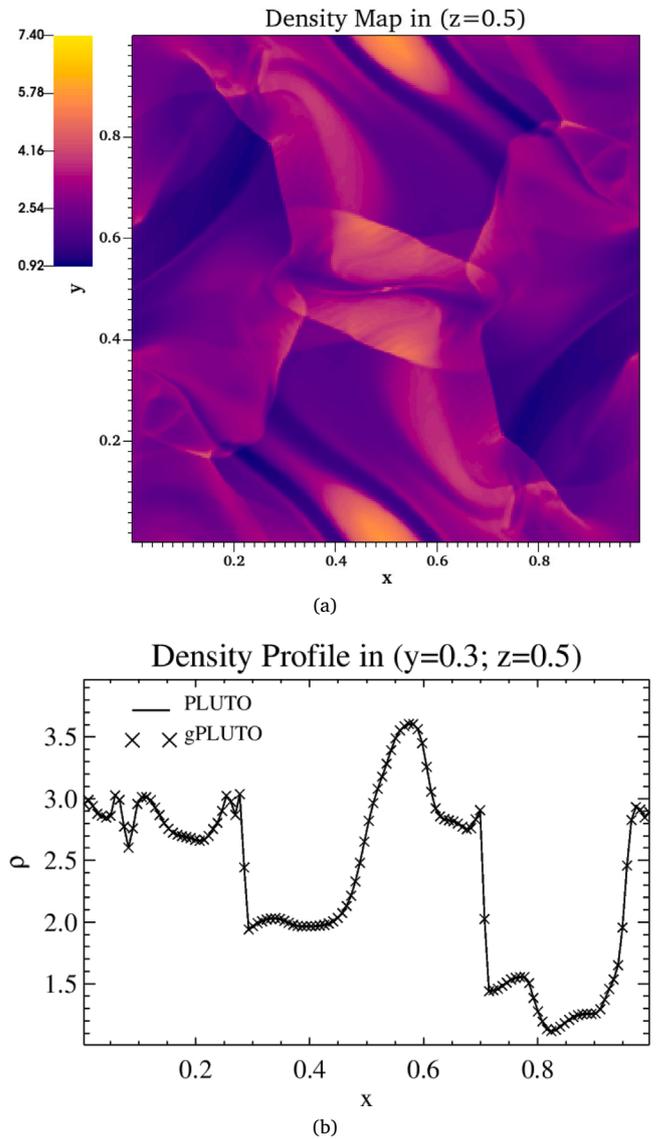
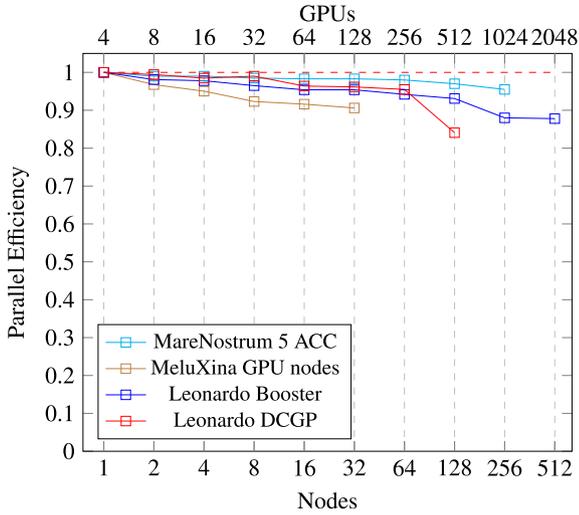
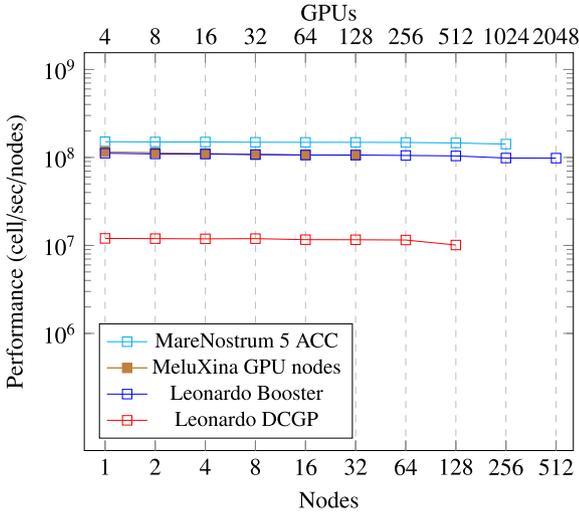


Fig. 5. Upper panel: 2D Density map at $t = 0.5$ and in $z = 0.5$ (gPLUTO). Lower panel: Density profile at $t = 0.5$ in ($y = 0.3, z = 0.5$) comparing PLUTO and gPLUTO.

Strong scaling. We then present, in Fig. 7, strong scaling performance results of gPLUTO, obtained by repeating the test using the divergence cleaning method for the control of the magnetic field divergence. This allowed for choosing a slightly larger grid than that used in the weak scaling test because this algorithm requires less device memory than the constrained transport method (Table 3). In this test, the domain



(a) Weak scaling efficiencies.



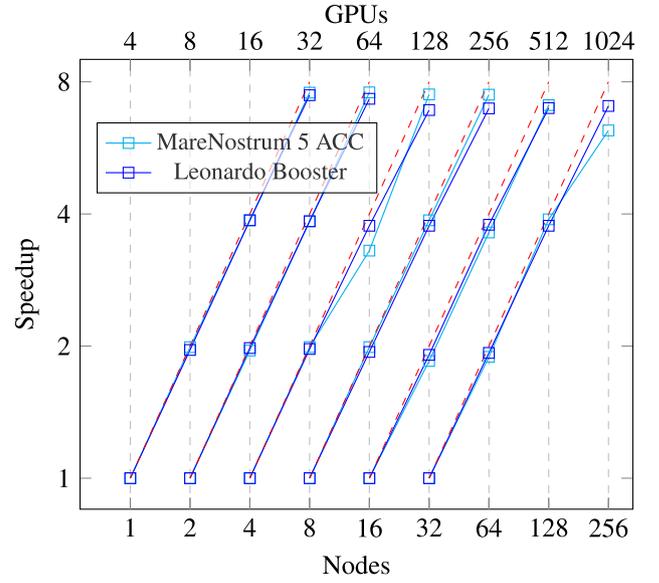
(b) Performance values.

Nodes	T_{GPUs} (sec)	T_{CPUs} (sec)	Acceleration ($T_{\text{CPUs}}/T_{\text{GPUs}}$)
1	312	2885	9.24
2	318	2901.4	9.12
4	319	2918.6	9.14
8	323	2903.2	8.98
16	327	2981.6	9.11
32	327	2989	9.14
64	331	3009.6	9.02
128	335	3430.8	10.24

(c) CPU–GPU comparison on Leonardo DCGP (4 A100 GPUs per node) and Booster partition (112 Intel Xeon 8480+ CPUs cores per node), respectively.

Fig. 6. Parallel Efficiencies (a), Performance (b) and CPUs to GPUs acceleration (c) of weak scaling runs with the 3D Orszag–Tang setup on MareNostrum 5 ACC, MeluXina GPU nodes and Leonardo (Booster and DCGP).

size remains fixed while the number of processes varies. As a result, GPU utilization is optimal in the initial runs but degrades, along with performance, as more nodes are employed. To address this effect, we performed six sets of four runs each, maintaining a fixed resolution

**Fig. 7.** Strong scaling speedup with the 3D Orszag–Tang setup on MareNostrum 5 ACC and Leonardo Booster.

as detailed in Table 3. The resolution was selected to ensure that the grid size nearly filled the GPU VRAM at the lowest node count within each group. We then successively doubled the number of nodes and processes three times per group, obtaining strong scaling efficiencies. The results obtained on MareNostrum 5 ACC and Leonardo Booster largely overlap, with average efficiencies of the runs at higher numbers of nodes of each group at 0.90 and 0.89, respectively. Table 3 provides the exact scaling efficiencies, computed by comparing the measured values of speedup, plotted in Fig. 7, to the ideal ones that correspond to the red lines in figure (speedup = 1, 2, 4, 8).

4.2. 3D circularly polarized Alfvén

The propagation of circularly polarized Alfvén waves is a widely employed benchmark as it provides an exact nonlinear solution of the ideal MHD equations. Following Mignone and Tzeferacos (2010) and Mignone et al. (2010c), we rotate a wave propagating along the x axis to obtain a 3D configuration, with an angle γ around the y axis and an angle α around the z axis. The grid size is the same of Section 4.1 for both strong and weak scaling. The domain is a Cartesian cube with $L_x = L_y = 1$ and $L_z = 1/2$ with periodic boundaries. The initial speed \mathbf{v} and magnetic field \mathbf{B} are chosen as

$$\begin{aligned} \mathbf{v} &= \epsilon \sin(\phi) \hat{\mathbf{e}}_y + \epsilon \cos(\phi) \hat{\mathbf{e}}_z, \\ \mathbf{B} &= \hat{\mathbf{e}}_x + \epsilon \sin(\phi) \hat{\mathbf{e}}_y + \epsilon \cos(\phi) \hat{\mathbf{e}}_z, \end{aligned} \quad (10)$$

where $\epsilon = 0.1$ is the wave amplitude, while $\phi = \mathbf{k} \cdot \mathbf{x}$ with $\mathbf{k} = 2\pi(1/L_x, 1/L_y, 1/L_z)$. Density and gas pressure are initially constant and equal to $\rho_0 = 1.0$ and $p_0 = 0.1$, respectively. We employ an ideal equation of state with adiabatic index $\Gamma = 5/3$.

We solve the MHD equations using the RK3 scheme with WENOZ reconstruction and the HLLD Riemann solver. The divergence-free condition is enforced using divergence cleaning (for strong scaling) and constrained transport (for weak scaling) to differentiate the algorithms employed. All computations are performed in double-precision arithmetic. The tests have been performed on Leonardo Booster and MareNostrum 5 ACC partitions, using the same grid resolution (Table 2) of 4.1,

Results validation. The wave period T is given by:

$$T = \frac{1}{\sqrt{1 + \tan^2 \alpha + \tan^2 \beta}}, \quad (11)$$

Table 3
Strong scaling efficiencies with the 3D Orszag Tang setup on MareNostrum 5 ACC and Leonardo Booster.

Nodes	MareNostrum efficiency	Leonardo efficiency
Group 1, Base: 1 node, Resolution: $832 \times 832 \times 416$		
1	1.00	1.00
2	0.98	0.98
4	0.97	0.97
8	0.94	0.93
Group 2, Base: 2 nodes, Resolution: $832 \times 832 \times 832$		
2	1.00	1.00
4	0.99	0.99
8	0.96	0.96
16	0.94	0.92
Group 3, Base: 4 nodes, Resolution: $1664 \times 832 \times 832$		
4	1.00	1.00
8	0.99	0.99
16	0.83	0.94
32	0.94	0.86
Group 4, Base: 8 nodes, Resolution: $1664 \times 1664 \times 832$		
8	1.00	1.00
16	0.99	0.97
32	0.97	0.94
64	0.93	0.87
Group 5, Base: 16 nodes, Resolution: $1664 \times 1664 \times 1664$		
16	1.00	1.00
32	0.93	0.95
64	0.91	0.94
128	0.89	0.87
Group 6, Base: 32 nodes, Resolution: $3328 \times 1664 \times 1664$		
32	1.00	1.00
64	0.94	0.97
128	0.97	0.94
256	0.77	0.88

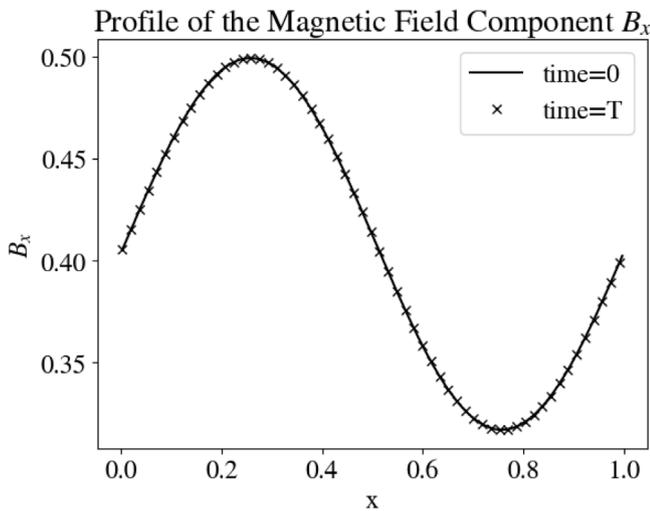


Fig. 8. Plot of the x -component of magnetic field for the 3D Circularly polarized Alfvén wave along the x -axis. The continuous line represents the reference initial solution while x markers show the magnetic field after one period T .

with $\tan \alpha = k_y/k_x = 1$, $\tan \beta = k_z/k_x = 2$ (see Mignone and Tzeferacos, 2010; Mignone et al., 2010c for more details).

We demonstrate the correctness of the test running a simulation for one period T and showing the profile of the B_x component of the magnetic field return to the initial condition (Fig. 8), but the same could have been done for each component of \mathbf{v} and \mathbf{B} .

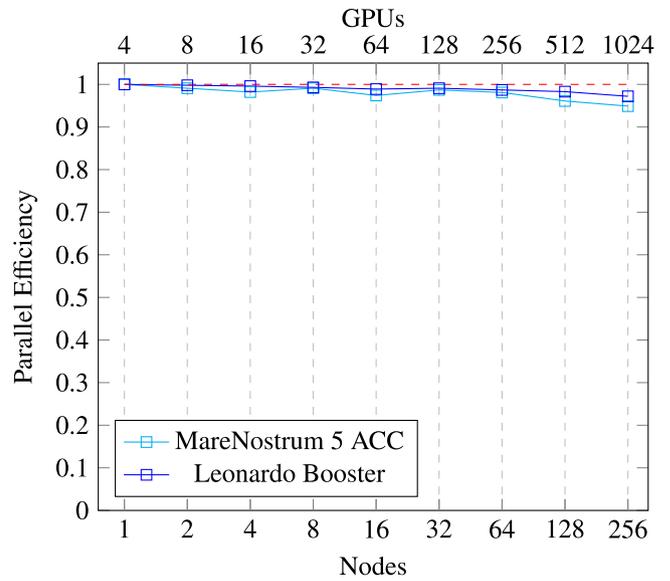


Fig. 9. Weak scaling tests of circularly polarized Alfvén waves on MareNostrum 5 ACC and Leonardo Booster.

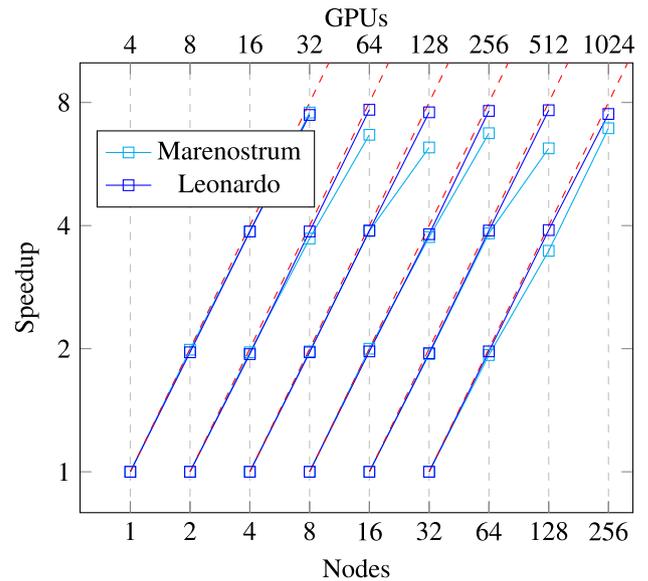


Fig. 10. Strong Scaling speedup with the 3D Circularly Polarized Alfvén setup on MareNostrum 5 ACC and Leonardo Booster.

Weak scaling. Fig. 9 presents the parallel efficiency expressed as the ratio of the wall clock times on 1 node and N nodes, T_1/T_N . The efficiency decreases slightly as the number of nodes grows due to the increase in communications and physical distances between the devices. The best results are obtained, this time, with Leonardo with an efficiency of 0.97 for 256 nodes (1024 GPUs), while we obtain 0.95 on MareNostrum 5.

Strong scaling. Strong scaling test results are presented in Table 4 and in Fig. 10. As for the Orszag–Tang test, six strong scaling tests are presented for four groups of runs with fixed resolution. As expected, the strong scaling efficiency slightly decreases as the number of nodes rises. Leonardo Booster’s results are marginally better than MareNostrum 5 ACC for this test, with an average efficiency of 0.97 and 0.93, respectively.

Table 4

Strong scaling tests with the 3D circularly polarized Alfvén setup on MareNostrum 5 ACC and Leonardo Booster.

Nodes	MareNostrum 5 efficiency	Leonardo efficiency
Group 1, Base: 1 node, Resolution: $832 \times 832 \times 416$		
1	1.00	1.00
2	0.98	0.97
4	0.93	0.97
8	0.83	0.96
Group 2, Base: 2 nodes, Resolution: $832 \times 832 \times 832$		
2	1.00	1.00
4	0.98	0.98
8	0.97	0.98
16	0.78	0.95
Group 3, Base: 4 nodes, Resolution: $1664 \times 832 \times 832$		
4	1.00	1.00
8	1.00	0.99
16	0.94	0.95
32	0.84	0.95
Group 4, Base: 8 nodes, Resolution: $1664 \times 1664 \times 832$		
8	1.00	1.00
16	0.97	0.98
32	0.95	0.97
64	0.81	0.96
Group 5, Base: 16 nodes, Resolution: $1664 \times 1664 \times 1664$		
16	1.00	1.00
32	0.97	0.99
64	0.96	0.97
128	0.77	0.94
Group 6, Base: 32 nodes, Resolution: $3328 \times 1664 \times 1664$		
32	1.00	1.00
64	0.96	0.99
128	0.87	0.97
256	0.87	0.94

4.3. Code accuracy

Although the focus of this work is on assessing code performance on HPC platforms, we emphasize that the numerical accuracy of the code has already been extensively validated in prior studies. In particular, both the 2nd- and 4th-order algorithms have been carefully benchmarked by porting and reproducing all standard test problems presented in earlier works, achieving in most cases almost binary identical results.

These validation efforts include, among others, the cell-centered MHD methods of Mignone and Tzeferacos (2010), the UCT schemes of Mignone and Del Zanna (2021), and the genuinely fourth-order finite-volume algorithms described in Berta et al. (2024) for classical and relativistic ideal MHD, as well as in Mignone et al. (2024) for resistive relativistic MHD. We therefore refer the reader to the above references for a detailed discussion of the code’s accuracy and verification.

4.4. Energy efficiency

Energy efficiency for an application is defined as performance in Floating point Operations Per Second (FLOPS) per unit of power consumption (Watt). The total energy consumed (in Joules) is approximated from power samples (in Watts) obtained at a specific sampling frequency (in Hertz) from a power monitoring system, as depicted in the following equation:

$$\text{Energy}(t) = \int_0^t \text{Power}(t') dt' \approx \frac{\sum_{i=0}^n P_i}{f_s},$$

where P_{s_i} is i th power sample, f_s is sampling frequency.

4.4.1. Methodology

For energy measurements, we used GPU-accelerated nodes of the EuroHPC Karolina cluster (IT4Innovations, 2025) equipped with two AMD EPYC 7763 CPUs and eight NVIDIA A100-SXM4 GPUs. The energy consumption of the GPU is measured using performance counters of the GPU via NVIDIA Management Library (NVML) (NVIDIA Corp., 2025). For CPU, we use AMD RAPL (AMD, 2025; Schone et al., 2021). The power consumption of the remaining components of the compute node is estimated using a simple mathematical model (see Table 5).

To investigate the energy efficiency of the application, we have tuned the frequency of the GPUs’ Streaming Multiprocessors (SMs), which is analogous to the CPU core frequency. For this purpose, we use NVML function `nvmlDeviceSetApplicationsClocks()` that sets a specific clock speed of a GPU SMs. Since A100-SXM4 uses HBM2 memory, its frequency cannot be tuned; on GPUs with GDDR memory, on the other hand, the frequency can be scaled using the same function.

We use static frequency tuning, meaning that a single GPU SMs frequency is applied at the start of the application execution and remains constant until the end. The key advantage of static tuning is that it can be performed by a standard HPC job scheduler like SLURM (Yoo et al., 2003) or PBS (Nitzberg et al., 2004) at the job launch time. The disadvantage of this method is that the applied configuration may not be optimal for all executed kernels. Thus, dynamic tuning can bring more energy savings for the same or lower performance penalty (Catthoor et al., 2020).

4.4.2. Platform analysis

The key properties of the Karolina GPU-accelerated node from the energy consumption point of view are shown in Table 5. We break the hardware into 3 key groups: CPUs, GPUs, and the rest of the server (including motherboard, memory DIMMs, network cards, GPU board with NVlink switches, etc.). For each group, we show the “Range of power consumption”, where the lower number is the power consumption of the component when the server is idle (no computations are running) and the higher value presents maximum power consumption under full load. When all values for idle power consumption for all components are summed, we get the “Static power consumption” of the entire server. In this case, just to keep the server turned on, it consumes 1.12 kW. The “Dynamic power consumption” then denotes the additional power that can be consumed when the server is under full load. This is a maximum/theoretical value if a workload is capable of full utilization of the CPUs or GPUs. By summing the static and dynamic consumption, we obtain the server’s peak power consumption. Our methodology can only control and reduce dynamic power consumption. From the server user perspective, the static power cannot be controlled, and it is the main reason why energy consumption increases when runtime is extended, as energy is the integral of power over time. Finally, the last two rows in the Tab. present the platform’s behavior when gPLUTO runs the 3D Orszag–Tang vortex benchmark. We can see that CPUs and remaining server components consumption is mostly static (because we do not tune CPU frequency), while the key component with predominantly dynamic power consumption are the 8 GPUs. In summary, each compute node can consume between 1.6 and 2.8 kW depending on how many nodes are used for the experiment (we use strong scaling) and the selected frequency of the GPU SMs, resulting in 1.2 kW dynamically controllable power consumption.

The compute node behavior is presented in Table 6 in more detail. We can see that: (a) by changing the SM frequency from the default 1.41 GHz to 0.99 GHz the power consumption of all 8 GPUs per node is reduced from 2.0 to 1.2 kW (when the benchmark runs on one node); (b) by performing the strong scaling from 1 to 32 compute nodes, we also reduce GPUs power consumption per node from 2.0 to 1.2 kW. This is caused by the strong scaling because the problem size per GPU is halved with every step of strong scaling. This results in lower utilization of the GPUs and, therefore, lower average power consumption. For 32 nodes, if we scale the SM frequency down to 0.99 GHz, we can

Table 5

Power consumption of the key components of the GPU accelerated compute node of the Karolina supercomputer. Static power consumption is always present, even when a server is idle. Dynamic power consumption is additional power consumption when a server is under load (CPUs and GPUs perform some computations) and it is workload-dependent. The large gPLUTO power consumption range for GPUs is caused by both lowering the SM frequency to 990 GHz as well as lowering GPU utilization for the high number of compute nodes in the strong scaling test.

Karolina GPU node	CPUs	GPUs	Other parts	Total
Units per server	2× AMD EPYC 7763	8× NVIDIA A100	–	–
Range of power consumption	90–280 W per CPU	55–400 W per GPU	520–625 W per server	–
Static power consumption	180 W	420 W	520 W	1120 W
Dynamic power consumption	max. 380 W	max. 2780 W	max. 105 W	max. 3265 W
PLUTO power consumption range	115–120 W per CPU	99–248 W per GPU	600–610 W	1622–2834 W
PLUTO dynamic pow. consumption	5 W per CPU; 10 W total	149 W per GPU; 1192 W total	10 W	1212 W

Table 6

Power consumption per node for the 3D Orszag–Tang Vortex benchmark. The top panel shows absolute values in [kW]. The bottom one shows the relative values with respect to the entire node power consumption in [%].

SM freq. [MHz]	1 node 8 GPUs			2 nodes 16 GPUs			4 nodes 32 GPUs			8 nodes 64 GPUs			16 nodes 128 GPUs			32 nodes 256 GPUs					
	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver			
1410	2,0	0,2	0,6	1,9	0,2	0,6	1,7	0,2	0,6	1,6	0,2	0,6	1,4	0,2	0,6	1,2	0,2	0,6	1,1	0,2	0,6
1350	1,8	0,2	0,6	1,7	0,2	0,6	1,5	0,2	0,6	1,5	0,2	0,6	1,2	0,2	0,6	1,1	0,2	0,6	1,0	0,2	0,6
1290	1,6	0,2	0,6	1,5	0,2	0,6	1,4	0,2	0,6	1,3	0,2	0,6	1,2	0,2	0,6	1,0	0,2	0,6	1,0	0,2	0,6
1230	1,5	0,2	0,6	1,4	0,2	0,6	1,3	0,2	0,6	1,2	0,2	0,6	1,1	0,2	0,6	1,0	0,2	0,6	0,9	0,2	0,6
1170	1,4	0,2	0,6	1,3	0,2	0,6	1,2	0,2	0,6	1,2	0,2	0,6	1,0	0,2	0,6	0,9	0,2	0,6	0,9	0,2	0,6
1110	1,3	0,2	0,6	1,2	0,2	0,6	1,1	0,2	0,6	1,1	0,2	0,6	1,0	0,2	0,6	0,9	0,2	0,6	0,9	0,2	0,6
1050	1,2	0,2	0,6	1,2	0,2	0,6	1,1	0,2	0,6	1,0	0,2	0,6	0,9	0,2	0,6	0,8	0,2	0,6	0,8	0,2	0,6
990	1,2	0,2	0,6	1,1	0,2	0,6	1,0	0,2	0,6	1,0	0,2	0,6	0,9	0,2	0,6	0,8	0,2	0,6	0,8	0,2	0,6

SM freq. [MHz]	1 node 8 GPUs			2 nodes 16 GPUs			4 nodes 32 GPUs			8 nodes 64 GPUs			16 nodes 128 GPUs			32 nodes 256 GPUs		
	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver	GPU	CPU	Ser- ver
1410	70	8	21	69	9	22	67	9	24	66	10	24	62	11	27	58	12	30
1350	68	9	23	67	9	24	65	10	25	64	10	26	60	11	29	57	12	31
1290	66	10	24	65	10	25	63	10	27	62	11	28	58	12	30	55	12	33
1230	64	10	26	63	10	27	61	11	28	60	11	29	57	12	31	54	13	33
1170	62	11	27	61	11	28	59	11	29	58	12	30	55	13	33	52	13	35
1110	61	11	28	60	11	29	58	12	30	57	12	31	53	13	34	51	14	36
1050	59	12	29	58	12	30	56	12	31	55	13	32	52	13	35	49	14	37
990	58	12	30	57	12	31	56	13	32	54	13	33	51	14	35	49	14	37

further reduce GPU power consumption to 0.8kW. It is important to realize that, for this configuration, the static power consumption of the non-GPU components amounts to half the total. When this effect is combined with the extended runtime due to SM frequency reduction (for 32 nodes, reducing the SM frequency from 1.4 to 0.99 GHz yields a runtime extension to 117.6% - see Table 7 top) it becomes dominant in terms of energy consumption of the computation. This explains why the highest energy savings (see bottom panels of Tables 7 and 8) are not achieved for the lowest SM frequency we used.

4.4.3. Energy consumption optimization results

We have performed strong scaling from 1 to 32 GPU-accelerated compute nodes of Karolina for both 3D circularly polarized Alfvén wave and 3D Orszag–Tang vortex benchmarks. For each number of nodes, we scaled down the SM frequency from 1.4 to 0.99 GHz and measured the runtime and energy consumption of all compute nodes employed for the benchmark. The measurement results are in Tables 7 and 8. We do not present results for frequencies lower than 0.99 GHz because, for all cases, the energy savings for 0.99 GHz are smaller than the savings for 1.11 and 1.05 GHz.

If we consider acceptable an increase in runtime up to 6% (values close to this threshold are highlighted in red in both figures), we see that we can save up to 8% of energy for the 3D Orszag–Tang vortex and up to 11% for the 3D circularly polarized Alfvén wave benchmark. We also see that by increasing the number of compute nodes we gain lower savings. As mentioned previously, this is due to a high ratio of static power consumption and extended runtime.

We want to highlight that these servers are air-cooled, and therefore, the energy consumption of the remaining components is relatively

Table 7

Impact of the static tuning of the GPU SM frequency on the runtime (top) and energy consumption (bottom) of the 3D Orszag–Tang vortex benchmark. The top panel shows runtime variations with respect to the default execution time shown in the first line. The bottom one shows relative energy consumption with respect to the energy consumption of the default execution, as shown in the first line of the table.

SM freq. [MHz]	1 node 8 GPUs	2 nodes 16 GPUs	4 nodes 32 GPUs	8 nodes 64 GPUs	16 nodes 128 GPUs	32 nodes 256 GPUs
1410/	129 s	68,2 s	38,1 s	21,6 s	15,5 s	11,2 s
default	100%	100%	100%	100%	100%	100,0%
1350	102,8%	102,3%	102,8%	102,6%	100,5%	100,8%
1290	105,8%	105,4%	105,3%	105,2%	102,1%	101,1%
1230	108,9%	108,4%	108,4%	108,0%	104,8%	103,7%
1170	112,6%	112,1%	111,4%	113,1%	108,8%	107,5%
1110	116,2%	115,3%	115,5%	115,9%	112,9%	110,0%
1050	120,7%	119,9%	119,7%	121,2%	116,3%	115,2%
990	125,7%	125,0%	124,1%	123,9%	120,5%	117,6%

SM freq. [MHz]	1 node 8 GPUs	2 nodes 16 GPUs	4 nodes 32 GPUs	8 nodes 64 GPUs	16 nodes 128 GPUs	32 nodes 256 GPUs
1410/	365 kJ	368 kJ	387 kJ	423 kJ	543 kJ	719 kJ
default	100%	100%	100%	100%	100%	100%
1350	95,7%	94,9%	96,2%	96,4%	95,7%	96,4%
1290	92,0%	92,2%	93,5%	93,7%	93,7%	93,5%
1230	90,1%	90,4%	91,7%	91,8%	91,9%	93,7%
1170	88,8%	89,4%	90,5%	91,8%	91,8%	93,0%
1110	87,6%	88,4%	90,2%	90,9%	92,1%	92,8%
1050	87,4%	88,5%	90,4%	91,7%	92,4%	94,5%
990	88,9%	90,5%	92,1%	92,5%	94,3%	95,5%

high. By using modern hot water-cooled servers and racks, the static power consumption can be reduced, with a corresponding increase in energy savings.

5. Summary

In this paper we have presented a first look into gPLUTO, a GPU-optimized implementation of the widely used PLUTO code for computational plasma astrophysics. The new code is written in C++ and adopts OpenACC as a programming model to efficiently exploit modern GPU architectures. The new implementation maintains the flexibility and robustness of the legacy code while significantly improving computational performance, power efficiency, and scalability on HPC clusters.

This paper outlines the fundamental design principles behind gPLUTO, including the numerical methods employed, the OpenACC-based parallelization strategy, memory management optimizations, and multi-GPU communication. The performance of gPLUTO has been evaluated through two numerical benchmarks: the 3D Orszag–Tang vortex and the 3D circularly polarized Alfvén wave, conducted on three European HPC platforms: Marenstrum 5, Leonardo, and Meluxina. Both weak and strong scaling tests have been conducted.

Table 8

Same as Table 7 but for the 3D circularly polarized Alfvén test.

SM freq. [MHz]	1 node 8 GPUs	2 nodes 16 GPUs	4 nodes 32 GPUs	8 nodes 64 GPUs	16 nodes 128 GPUs	32 nodes 256 GPUs
1410/	129 s	68,2 s	38,1 s	21,6 s	15,5 s	11,2 s
default	100%	100%	100%	100%	100%	100,0%
1350	102,8%	102,3%	102,8%	102,6%	100,5%	100,8%
1290	105,8%	105,4%	105,3%	105,2%	102,1%	101,1%
1230	108,9%	108,4%	108,4%	108,0%	104,8%	103,7%
1170	112,6%	112,1%	111,4%	113,1%	108,8%	107,5%
1110	116,2%	115,3%	115,5%	115,9%	112,9%	110,0%
1050	120,7%	119,9%	119,7%	121,2%	116,3%	115,2%
990	125,7%	125,0%	124,1%	123,9%	120,5%	117,6%

SM freq. [MHz]	1 node 8 GPUs	2 nodes 16 GPUs	4 nodes 32 GPUs	8 nodes 64 GPUs	16 nodes 128 GPUs	32 nodes 256 GPUs
1410/	365 kJ	368 kJ	387 kJ	423 kJ	543 kJ	719 kJ
default	100%	100%	100%	100%	100%	100%
1350	95,7%	94,9%	96,2%	96,4%	95,7%	96,4%
1290	92,0%	92,2%	93,5%	93,7%	93,7%	93,5%
1230	90,1%	90,4%	91,7%	91,8%	91,9%	93,7%
1170	88,8%	89,4%	90,5%	91,8%	91,8%	93,0%
1110	87,6%	88,4%	90,2%	90,9%	92,1%	92,8%
1050	87,4%	88,5%	90,4%	91,7%	92,4%	94,5%
990	88,9%	90,5%	92,1%	92,5%	94,3%	95,5%

Results demonstrate excellent strong and weak scaling efficiencies, maintaining around 90% performance across multiple nodes, in a situation of high GPU occupancy. Direct comparison between the GPU- and CPU-specialized partitions of Leonardo reveals a factor ≈ 9 acceleration (per node usage). The results of this study confirm that gPLUTO successfully harnesses the power of GPUs to enhance the performance of MHD simulations.

The strong scaling behavior suggests that gPLUTO can efficiently utilize a large number of GPUs, making it well-suited for exascale computing applications.

Energy efficiency analyses highlight the impact of GPU frequency tuning in optimizing computational power consumption. Our experiments indicate that adjusting GPU clock frequencies can lead to significant energy savings with minimal runtime penalties, highlighting potential avenues for optimizing HPC resource utilization. In particular, lowering the SMs frequency can save around 8% and 11% of energy consumption for 3D Orszag–Tang tests and circularly polarized Alfvén waves, respectively, at a cost of 6% longer execution time.

The development of gPLUTO represents a significant step toward enabling more efficient and scalable astrophysical simulations on next-generation, Exascale supercomputers. A more comprehensive description of the code along with all its physical modules (including relativistic and non-ideal MHD extensions, particles and different systems of coordinates) and the implemented numerical methods — including high-order methods (Berta et al., 2024; Mignone et al., 2024) — will be presented in a forthcoming work. We also invite to refer to the companion paper (Suriano et al., 2025) for the Lagrangian Particles module integration in gPLUTO. Future works on HPC porting will also explore additional acceleration frameworks such as OpenMP for execution on AMD-based HPC systems.

CRediT authorship contribution statement

M. Rossazza: Writing – review & editing, Visualization, Validation, Supervision, Data curation. **A. Mignone:** Supervision, Writing – review & editing. **M. Bugli:** Writing – review & editing. **S. Truzzi:** Writing – review & editing. **L. Riha:** Writing – review & editing. **T. Panoc:** Writing – review & editing. **O. Vysocky:** Writing – review & editing. **N. Shukla:** Writing – review & editing. **A. Romeo:** Writing – review & editing. **V. Berta:** Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Andrea Mignone reports financial support was provided by European Union. Marco Rossazza reports financial support was provided by European Union. Stefano Truzzi reports financial support was provided by European Union. Vittoria Berta reports financial support was provided by European Union. Lubomir Riha reports financial support was provided by European Union. Tomas Panoc reports financial support was provided by European Union. Ondrej Vysocky reports financial support was provided by European Union.

Nitin Shukla reports financial support was provided by European Union. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

MR and AM wish to thank F. Spiga from NVIDIA for his active support during the very early stages of this development. MR and AM wish also to thank M. Bettencourt for his technical guidance and precious advices while porting the code to C++.

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 (SPACE).

This project has received funding from the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101064953 (GR-PLUTO).

M.B. acknowledges the support of the French Agence Nationale de la Recherche (ANR), under grant ANR-24-ERCS-0006 (project BlackJET).

This work was supported by the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID: 90254).

We acknowledge ISCRA for awarding this project access to the Leonardo supercomputer and the EuroHPC Joint Undertaking for granting us access to the Leonardo, MareNostrum and MeluXina supercomputer through an EuroHPC [Extreme/Regular/Benchmark/Development/...] Access call.

This paper is supported by the Fondazione ICSC, Spoke 3 Astrophysics and Cosmos Observations and National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR), Project ID CN_00000013 “Italian Research Center on High-Performance Computing, Big Data and Quantum Computing” funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di “campioni nazionali di R&S (M4C2-19)” - Next Generation EU (NGEU).

Data availability

Data will be made available on request.

References

- Amato, E., Blasi, P., 2018. Cosmic ray transport in the galaxy: A review. *Adv. Space Res.* 62 (10), 2731–2749. doi:10.1016/j.asr.2017.04.019, arXiv:1704.05696.
- AMD, 2025. Kernel driver amd_energy. https://github.com/amd/amd_energy. [Online]; (Accessed 28 February 2025).
- Balsara, D.S., Spicer, D.S., 1999. A staggered mesh algorithm using high order Godunov fluxes to ensure solenoidal magnetic fields in magnetohydrodynamic simulations. *J. Comput. Phys.* 149, 270–292. doi:10.1006/jcph.1998.6153.
- Berta, V., Mignone, A., Bugli, M., Mattia, G., 2024. A 4th-order accurate finite volume method for ideal classical and special relativistic MHD based on pointwise reconstructions. *J. Comput. Phys.* 499, 112701. doi:10.1016/j.jcp.2023.112701, arXiv:2310.11831.

- Bodo, G., Mamatsashvili, G., Rossi, P., Mignone, A., 2022. Current-driven kink instabilities in relativistic jets: dissipation properties. *Mon. Not. R. Astron. Soc.* 510 (2), 2391–2406. doi:10.1093/mnras/stab3492, arXiv:2111.14575.
- Borges, R., Carmona, M., Costa, B., Don, W.S., 2008. An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws. *J. Comput. Phys.* 227 (6), 3191–3211. doi:10.1016/j.jcp.2007.11.038.
- Bugli, M., Lopresti, E.F., Figueiredo, E., Mignone, A., Cerutti, B., Mattia, G., Del Zanna, L., Bodo, G., Berta, V., 2025. Relativistic reconnection with effective resistivity: I. Dynamics and reconnection rate. *J. Comput. Phys.* 693, A233. doi:10.1051/0004-6361/202452277, arXiv:2410.20924.
- Cattoor, F., Basten, T., Zompakis, N., Geilen, M., Kjeldsberg, P.G., 2020. System-Scenario-based Design Principles and Applications, vol. 1, Springer International Publishing, doi:10.1007/978-3-030-20343-6.
- Cattorini, F., Giacomazzo, B., 2024. GRMHD study of accreting massive black hole binaries in astrophysical environment: A review. *Astrophys. J.* 154, 102892. doi:10.1016/j.astrophys.2023.102892, arXiv:2401.02521.
- Charbonneau, P., 2020. Dynamo models of the solar cycle. *Living Rev. Sol. Phys.* 17 (1), 4. doi:10.1007/s41116-020-00025-6.
- Colella, P., Woodward, P.R., 1984. The piecewise parabolic method (PPM) for gas-dynamical simulations. *J. Comput. Phys.* 54, 174–201. doi:10.1016/0021-9991(84)90143-8.
- Dedner, A., Kemm, F., Kröner, D., Munz, C.-D., Schnitzer, T., Wesenberg, M., 2002. Hyperbolic divergence cleaning for the MHD equations. *J. Comput. Phys.* 175, 645–673. doi:10.1006/jcph.2001.6961.
- Del Zanna, L., Landi, S., Serafini, L., Bugli, M., Papini, E., 2024. A GPU-accelerated modern fortran version of the ECHO code for relativistic magnetohydrodynamics. *Fluids* 9 (1), 16. doi:10.3390/fluids9010016.
- Del Zanna, L., Tomei, N., Franceschetti, K., Bugli, M., Bucciantini, N., 2022. General relativistic magnetohydrodynamics mean-field dynamos. *Fluidika* 7 (2), 87. doi:10.3390/fluids7020087.
- Doyle, E.J., Houlberg, W.A., Kamada, Y., Mukhovatov, V., Osborne, T.H., Polevoi, A., Bateman, G., Connor, J.W., Cordey, J.G., Fujita, T., Garbet, X., Hahn, T.S., Horton, L.D., Hubbard, A.E., Imbeaux, F., Jenko, F., Kinsey, J.E., Kishimoto, Y., Li, J., Luce, T.C., Martin, Y., Ossipenko, M., Parail, V., Peeters, A., Rhodes, T.L., Rice, J.E., Roach, C.M., Rozhansky, V., Rytter, F., Saibene, G., Sartori, R., Sips, A.C.C., Snipes, J.A., Sugihara, M., Synakowski, E.J., Takenaga, H., Takizuka, T., Thomsen, K., Wade, M.R., Wilson, H.R., ITPA Transport Physics Topical Group, Confinement Database, I., Modelling Topical Group, Pedestal, I., Edge Topical Group, 2007. Chapter 2: Plasma confinement and transport. *Nucl. Fusion* 47 (6), S18–S127. doi:10.1088/0029-5515/47/6/S02.
- Evans, C.R., Hawley, J.F., 1988. Simulation of magnetohydrodynamic flows - a constrained transport method. *Astrophys. J.* 332, 659–677. doi:10.1086/166684.
- Faber, J.A., Rasio, F.A., 2012. Binary neutron star mergers. *Living Rev. Relativ.* 15 (1), 8. doi:10.12942/lrr-2012-8, arXiv:1204.3858.
- Flock, M., Mignone, A., 2021. Streaming instability in a global patch simulation of protoplanetary disks. *J. Comput. Phys.* 650, A119. doi:10.1051/0004-6361/202040104, arXiv:2103.15146.
- Helzel, C., Rossmannith, J.A., Taetz, B., 2011. An unstaggered constrained transport method for the 3D ideal magnetohydrodynamic equations. *J. Comput. Phys.* 230 (10), 3803–3829. doi:10.1016/j.jcp.2011.02.009, arXiv:1007.2606.
- Inghirami, G., Del Zanna, L., Beraudo, A., Moghaddam, M.H., Becattini, F., Bleicher, M., 2016. Numerical magneto-hydrodynamics for relativistic nuclear collisions. *Eur. Phys. J. C* 76 (12), 659. doi:10.1140/epjc/s10052-016-4516-8, arXiv:1609.03042.
- IT4Innovations, 2025. Karolina supercomputer user guide. <https://docs.it4i.cz/karolina/introduction/>. [Online]; (Accessed 28 February 2025).
- Janka, H.T., 2025. Long-term multidimensional models of core-collapse supernovae: Progress and challenges. doi:10.48550/arXiv.2502.14836, ArXiv E-Prints arXiv:2502.14836.
- Kalinani, J.V., Ji, L., Ennoggi, L., Lopez Armengol, F.G., Timotheo Sanches, L., Tsao, B.-J., Brandt, S.R., Campanelli, M., Ciolfi, R., Giacomazzo, B., Haas, R., Schnetter, E., Zlochower, Y., 2025. AsterX: a new open-source GPU-accelerated GRMHD code for dynamical spacetimes. *Classical Quantum Gravity* 42 (2), 025016. doi:10.1088/1361-6382/ad9c11, arXiv:2406.11669.
- Knaepen, B., Moreau, R., 2008. Magnetohydrodynamic turbulence at low magnetic Reynolds number. *Annu. Rev. Fluid Mech.* 40 (1), 25–45. doi:10.1146/annurev.fluid.39.050905.110231.
- Lazarian, A., Vlahos, L., Kowal, G., Yan, H., Beresnyak, A., de Gouveia Dal Pino, E.M., 2012. Turbulence, magnetic reconnection in turbulent fluids and energetic particle acceleration. *Space Sci. Rev.* 173 (1–4), 557–622. doi:10.1007/s11214-012-9936-7, arXiv:1211.0008.
- Lesur, G., 2021. Magnetohydrodynamics of protoplanetary discs. *J. Plasma Phys.* 87, 205870101. doi:10.1017/S0022377820001002.
- Lesur, G.R.J., 2021. Systematic description of wind-driven protoplanetary discs. *J. Comput. Phys.* 650, A35. doi:10.1051/0004-6361/202040109, arXiv:2101.10349.
- Lesur, G.R.J., Baghdadi, S., Wafflard-Fernandez, G., Mouxion, J., Robert, C.M.T., Van den Bossche, M., 2023. IDEFIX: A versatile performance-portable Godunov code for astrophysical flows. *J. Comput. Phys.* 677, A9. doi:10.1051/0004-6361/202346005, arXiv:2304.13746.
- Liska, M.T.P., Chatterjee, K., Issa, D., Yoon, D., Kaaz, N., Tchekhovskoy, A., van Eijndt, D., Musoke, G., Hesp, C., Rohoza, V., Markoff, S., Ingram, A., van der Klis, M., 2022. H-AMR: A new GPU-accelerated GRMHD code for exascale computing with 3D adaptive mesh refinement and local adaptive time stepping. *Astrophys. J. Suppl. Ser.* 263 (2), 26. doi:10.3847/1538-4365/ac9966, arXiv:1912.10192.
- Londrillo, P., del Zanna, L., 2004. On the divergence-free condition in godunov-type schemes for ideal magnetohydrodynamics: the upwind constrained transport method. *J. Comput. Phys.* 195, 17–48. doi:10.1016/j.jcp.2003.09.016, arXiv:astro-ph/0310183.
- Mattia, G., Del Zanna, L., Bugli, M., Pavan, A., Ciolfi, R., Bodo, G., Mignone, A., 2023. Resistive relativistic MHD simulations of astrophysical jets. *Astron. Astrophys.* 679, A49. doi:10.1051/0004-6361/202347126, arXiv:2308.09477.
- Mayer, L., 2019. Super-eddington accretion; flow regimes and conditions in high-z galaxies. In: Latif, M., Schleicher, D. (Eds.), *Formation of the First Black Holes*. pp. 195–222. doi:10.1142/9789813227958_0011.
- Mezzacappa, A., 2023. Toward realistic models of core collapse supernovae: A brief review. In: Bisikalo, D., Wiebe, D., Boily, C. (Eds.), *The Predictive Power of Computational Astrophysics As a Discover Tool*. In: *IAU Symposium*, 362, pp. 215–227. doi:10.1017/S1743921322001831, arXiv:2205.13438.
- Mignone, A., 2014. High-order conservative reconstruction schemes for finite volume methods in cylindrical and spherical coordinates. *J. Comput. Phys.* 270, 784–814. doi:10.1016/j.jcp.2014.04.001, arXiv:1404.0537.
- Mignone, A., Berta, V., Rossazza, M., Bugli, M., Mattia, G., Del Zanna, L., Pareschi, L., 2024. A fourth-order accurate finite volume scheme for resistive relativistic MHD. *Mon. Not. R. Astron. Soc.* 533 (2), 1670–1686. doi:10.1093/mnras/stae1729, arXiv:2407.08519.
- Mignone, A., Bodo, G., Massaglia, S., Matsakos, T., Tesileanu, O., Zanni, C., Ferrari, A., 2007. PLUTO: A numerical code for computational astrophysics. *Astrophys. J. Suppl. Ser.* 170, 228–242. doi:10.1086/513316, arXiv:astro-ph/0701854.
- Mignone, A., Del Zanna, L., 2021. Systematic construction of upwind constrained transport schemes for MHD. *J. Comput. Phys.* 424, 109748. doi:10.1016/j.jcp.2020.109748, arXiv:2004.10542.
- Mignone, A., Flock, M., Stute, M., Kolb, S.M., Muscianisi, G., 2012a. A conservative orbital advection scheme for simulations of magnetized shear flows with the PLUTO code. *J. Comput. Phys.* 545, A152. doi:10.1051/0004-6361/201219557, arXiv:1207.2955.
- Mignone, A., Muscianisi, G., Rivib, M., Bodoc, G., I/O Optimization Strategies in the PLUTO Code.
- Mignone, A., Plewa, T., Bodo, G., 2005. The piecewise parabolic method for multidimensional relativistic fluid dynamics. *Astrophys. J. Suppl. Ser.* 160 (1), 199–219. doi:10.1086/430905, arXiv:astro-ph/0505200.
- Mignone, A., Rossi, P., Bodo, G., Ferrari, A., Massaglia, S., 2010a. High-resolution 3D relativistic MHD simulations of jets. *Mon. Not. R. Astron. Soc.* 402 (1), 7–12. doi:10.1111/j.1365-2966.2009.15642.x, arXiv:0908.4523.
- Mignone, A., Rossi, P., Bodo, G., Ferrari, A., Massaglia, S., 2010b. High-resolution 3D relativistic MHD simulations of jets. *Mon. Not. R. Astron. Soc.* 402, 7–12. doi:10.1111/j.1365-2966.2009.15642.x, arXiv:0908.4523.
- Mignone, A., Striani, E., Tavani, M., Ferrari, A., 2013. Modelling the kinked jet of the crab nebula. *Mon. Not. R. Astron. Soc.* 436 (2), 1102–1115. doi:10.1093/mnras/stt1632, arXiv:1309.0375.
- Mignone, A., Zteferacos, P., 2010. A second-order unsplit godunov scheme for cell-centered MHD: The CTU-GLM scheme. *J. Comput. Phys.* 229, 2117–2138. doi:10.1016/j.jcp.2009.11.026, arXiv:0911.3410.
- Mignone, A., Zteferacos, P., Bodo, G., 2010c. High-order conservative finite difference GLM-MHD schemes for cell-centered MHD. *J. Comput. Phys.* 229 (17), 5896–5920. doi:10.1016/j.jcp.2010.04.013, arXiv:1001.2832.
- Mignone, A., Zanni, C., Zteferacos, P., van Straalen, B., Colella, P., Bodo, G., 2012b. The PLUTO code for adaptive mesh computations in astrophysical fluid dynamics. *Astrophys. J. Suppl. Ser.* 198, 7. doi:10.1088/0067-0049/198/1/7, arXiv:1110.0740.
- Miyoshi, T., Kusano, K., 2005. A multi-state HLL approximate Riemann solver for ideal magnetohydrodynamics. *J. Comput. Phys.* 208 (1), 315–344. doi:10.1016/j.jcp.2005.02.017.
- Narayan, R., 2005. Black holes in astrophysics. *New J. Phys.* 7 (1), 199. doi:10.1088/1367-2630/7/1/199, arXiv:gr-qc/0506078.
- Nitzberg, B., Schopf, J.M., Jones, J.P., 2004. PBS Pro: Grid computing and scheduling attributes. In: *Grid Resource Management: State of the Art and Future Trends*. Kluwer Academic Publishers, USA, pp. 183–190.
- NVIDIA Corp., 2025. NVML API reference. <https://docs.nvidia.com/deploy/nvml-api/nvml-api-reference.html>. [Online]; (Accessed 28 February 2025).

- Olmi, B., Bucciantini, N., 2019. On the origin of jet-like features in bow shock pulsar wind nebulae. *Mon. Not. R. Astron. Soc.* 490 (3), 3608–3615. doi:10.1093/mnras/stz2819, arXiv:1910.01926.
- Olmi, B., Bucciantini, N., Morlino, G., 2018. Numerical simulations of mass loading in the tails of bow-shock pulsar-wind nebulae. *Mon. Not. R. Astron. Soc.* 481 (3), 3394–3400. doi:10.1093/mnras/sty2525, arXiv:1809.03807.
- Orlando, S., Miceli, M., Petruk, O., Ono, M., Nagataki, S., Aloy, M.A., Mimica, P., Lee, S.H., Bocchino, F., Peres, G., Guarrasi, M., 2019. 3D MHD modeling of the expanding remnant of SN 1987A. Role of magnetic field and non-thermal radio emission. *J. Comput. Phys.* 622, A73. doi:10.1051/0004-6361/201834487, arXiv:1812.00021.
- Orlando, S., Ono, M., Nagataki, S., Miceli, M., Umeda, H., Ferrand, G., Bocchino, F., Petruk, O., Peres, G., Takahashi, K., Yoshida, T., 2020. Hydrodynamic simulations unravel the progenitor-supernova-remnant connection in SN 1987A. *J. Comput. Phys.* 636, A22. doi:10.1051/0004-6361/201936718, arXiv:1912.03070.
- Orlando, S., Reale, F., Peres, G., Mignone, A., 2011. Mass accretion to young stars triggered by flaring activity in circumstellar discs. *Mon. Not. R. Astron. Soc.* 415 (4), 3380–3392. doi:10.1111/j.1365-2966.2011.18954.x, arXiv:1104.5107.
- Pavan, A., Ciolfi, R., Kalinani, J.V., Mignone, A., 2023. Jet-environment interplay in magnetized binary neutron star mergers. *Mon. Not. R. Astron. Soc.* 524 (1), 260–275. doi:10.1093/mnras/stad1809, arXiv:2211.10135.
- Perri, B., Brun, A.S., Réville, V., Strugarek, A., 2018. Simulations of solar wind variations during an 11-year cycle and the influence of North-South asymmetry. *J. Plasma Phys.* 84, 765840501. doi:10.1017/S0022377818000880.
- Powell, K.G., Roe, P.L., Linde, T.J., Gombosi, T.I., De Zeeuw, D.L., 1999. A solution-adaptive upwind scheme for ideal magnetohydrodynamics. *J. Comput. Phys.* 154 (2), 284–309. doi:10.1006/jcph.1999.6299.
- Réville, V., Brun, A.S., Matt, S.P., Strugarek, A., Pinto, R.F., 2015. The effect of magnetic topology on thermally driven wind: toward a general formulation of the braking law. *Astrophys. J.* 798, 116. doi:10.1088/0004-637X/798/2/116.
- Schone, R., Ilsche, T., Bielert, M., Veltin, M., Schmidl, M., Hackenberg, D., 2021. Energy efficiency aspects of the AMD Zen 2 architecture. In: 2021 IEEE International Conference on Cluster Computing. CLUSTER, pp. 562–571, abs/2108.00808, URL <https://arxiv.org/abs/2108.00808>.
- Shibata, K., Magara, T., 2011. Solar flares: Magnetohydrodynamic processes. *Living Rev. Sol. Phys.* 8 (1), 6. doi:10.12942/lrsp-2011-6.
- Shu, C.-W., Osher, S., 1988. Efficient implementation of essentially non-oscillatory shock-capturing schemes. *J. Comput. Phys.* 77 (2), 439–471. doi:10.1016/0021-9991(88)90177-5.
- Stone, J.M., Mullen, P.D., Fielding, D., Grete, P., Guo, M., Kempfski, P., Most, E.R., White, C.J., Wong, G.N., 2024. AthenaK: A performance-portable version of the Athena++ AMR framework. doi:10.48550/arXiv.2409.16053, ArXiv E-Prints arXiv:2409.16053.
- Suresh, A., Huynh, H.T., 1997. Accurate monotonicity-preserving schemes with Runge Kutta time stepping. *J. Comput. Phys.* 136, 83–99. doi:10.1006/jcph.1997.5745.
- Suriano, A., Truzzi, S., Costa, A., Rossazza, M., Shukla, N., Berta, V., Mignone, A., Zanni, C., 2025. The PLUTO code on GPUs: Offloading Lagrangian particle methods. submitted for publication and currently under revision (Companion Paper).
- Teyssier, R., Commerçon, B., 2019. Numerical methods for simulating star formation. 6, 51. doi:10.3389/fspas.2019.00051.
- Uzdensky, D.A., Rightley, S., 2014. Plasma physics of extreme astrophysical environments. *Rep. Progr. Phys.* 77 (3), 036902. doi:10.1088/0034-4885/77/3/036902, arXiv:1401.5110.
- Yoo, A.B., Jette, M.A., Grondona, M., 2003. SLURM: Simple linux utility for resource management. In: Feitelson, D., Rudolph, L., Schwiegelshohn, U. (Eds.), *Job Scheduling Strategies for Parallel Processing*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 44–60.
- Zhang, D., 2018. A review of the theory of galactic winds driven by stellar feedback. *Galaxies* 6, 114. doi:10.3390/galaxies6040114.