

PHARE : Parallel hybrid particle-in-cell code with patch-based adaptive mesh refinement

Nicolas Aunai^{a,*}, Roch Smets^a, Andrea Ciardi^b, Philip Deegan^a, Alexis Jeandet^a, Thibault Payet^c, Nathan Guyot^a, Loic Darrieumerlou^a

^aLaboratoire de Physique des Plasmas (LPP), CNRS, Observatoire de Paris, Sorbonne Université, Université Paris-Saclay, École polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France

^bSorbonne Université, Observatoire de Paris, Université PSL, CNRS, LERMA, F-75005 Paris, France

^cAix Marseille Univ, CNRS, LIS, Marseille, France

Abstract

Modeling multi-scale collisionless magnetized processes constitutes an important numerical challenge. By treating electrons as a fluid and ions kinetically, the so-called hybrid Particle-In-Cell (PIC) codes represent a promising intermediary between fully kinetic codes, limited to model small scales and short durations, and magnetohydrodynamic codes used large scale. However, simulating processes at scales significantly larger than typical ion particle dynamics while resolving sub-ion dissipative current sheets remain extremely difficult. This paper presents a new hybrid PIC code with patch-based adaptive mesh refinement. Here, hybrid PIC equations are solved on a hierarchy of an arbitrary number of Cartesian meshes of incrementally finer resolution dynamically mapping regions of interest, and with a refined time stepping. This paper presents how the hybrid PIC algorithm is adapted to evolve such mesh hierarchy and the validation of the code on a uniform mesh, fixed refined mesh and dynamically refined mesh.

1. Introduction

The numerical modeling of magnetized space and laboratory plasmas represents an important computational challenge that is generally linked to the need to capture the multi-scale nature of plasma dynamics. While kinetic equations such as the Vlasov-Maxwell system for collisionless plasmas can in principle describe such dynamics, from electron scales to global system scales, their use is in multi-dimensions is in practice restricted to relatively small scales. For example, fully kinetic simulations of magnetic reconnection usually handle scales of a few tens of ion inertial lengths δ_i , quite rarely in three dimensions, and very often with artificially reduced scale separations through modified ion to electron mass ratio and/or reduced speed of light [1, 2]. Modeling larger scales is typically done with more approximate physical formalisms. Among them is the hybrid formalism, in which ions

are treated kinetically but electrons are modeled as a fluid [3, 4]. This allows, in principle, to transfer the computational power required to resolve electron scales into solving larger domains for longer times.

Hybrid codes have been used for modeling large systems such as (small or artificially reduced) planetary magnetospheres. However such large scale models typically poorly resolve the Hall scale [5, 6, 7, 8]. Not resolving such sub-ion scale current sheets may result in errors leading to spurious macroscopic behavior [9]. The main advantage such models have compared to a well resolved yet much lighter magnetohydrodynamics (MHD) counterparts is that they account for collisionless population mixing.

Modeling scales much larger than typical ion scales and in three dimensional systems yet having sub-ion scale space/time resolution remains a very challenging task for hybrid codes. The dispersive nature of the whistler waves supported by the collisionless Ohm's law imposes very strong constraints on the time step, which scales as the square of the

*Corresponding author

Email address: nicolas.aunai@lpp.polytechnique.fr (Nicolas Aunai)

mesh size. Therefore, high resolution large scale models ($\Delta \sim 0.1\delta_i$) remain computationally intensive. In practice, such high resolution hybrid simulations have typically been done in domains that are not vastly larger than those accessible with full-PIC codes using reduced physical normalized constants. The small gain in domain size or evolution time offered by hybrid codes is thus often quickly balanced by the drawback of losing the electron kinetic physics offered by full PIC ones. In practice, hybrid codes have not been so competitive once high resolution is needed.

The hybrid equations can in practice either be solved in a Eulerian way, with so-called Vlasov hybrid codes [10], or with a Lagrangian perspective, with the so-called hybrid Particle-In-Cell (PIC) codes [11]. PIC codes present the advantage over Vlasov ones to be inherently adaptive in their description of the particle distribution. This is more efficient computationally than resolving fine structures developing for different reasons in position space and velocity space and for each of the populations composing the plasma. Some features are also easier and/or lighter to handle in PIC than in Vlasov models, such as the possible interactions between the populations. On the other hand, Vlasov codes present the advantage to be noise-free contrary to PIC results which are inherently based on statistical sampling of the continuous distribution function. Recent advances have been achieved in using the Vlasov hybrid formalism to model large scales in the context of the Earth's magnetosphere [10]. Yet, simulations remain heavier than their PIC counterpart for which noise can moreover be reduced through the use of higher order interpolation kernels and increased number of macroparticles per cell. PIC codes thus constitute a promising approach for developing multi-scale models.

There has been several attempts to make hybrid PIC codes more adapted to model multi-scale systems. Time scale disparities have been for instance addressed by decomposing the domain into different time zones, each evolving with a proper time step [12]. Evolving quantities based on local event triggers rather than at regular and global time step intervals [13, 14, 15] has also been considered.

Spatial scale disparities, on the other hand, are generally handled by adopting a non-uniform meshing. This can be achieved either by refining the grid in the so-called adaptive mesh refinement (AMR) methods [16, 17, 18, 19, 20, 21, 22, 23], or adapting the position of the nodes of a single mesh in the

so-called moving mesh adaptation (MMA) method [24, 25]. In both cases the goal is to focus the meshing in regions developing small scale features where more computational power will be dedicated, and for which one needs to define ad-hoc detection criteria.

Adaptive mesh refinement (AMR) has now become a mainstream feature of fluid codes [26, 27, 28]. It is however much less used in PIC codes. The main reason probably stems from the algorithmic complexity in handling macroparticles across multiple grid levels. Only one AMR fully kinetic electromagnetic code is currently in use in our knowledge [18, 19, 22, 23].

An AMR Hybrid PIC code following a hybrid block AMR method [29] has also been proposed [21]. It is worth noticing however that the code has been mostly used with multiple grids that were fixed in time [30, 31, 32, 33].

In existing AMR implementations for PIC codes, macroparticles interact only with the finest mesh at their location. This method has several drawbacks. The first one, shared with MMA, is that macroparticles see multiple mesh spacings as they evolve. Either because they cross a grid boundary along their trajectory or because the region in which they currently are is being refined or coarsened in time (or stretched/contracted in MMA). Because the compact support of the macroparticle interpolation kernel is usually constant in index-space, these macroparticles will see their physical extent change in time, potentially leading to a spurious evolution of their momentum [11]. A more important drawback of having particles exploring multiple grids is that the number of particles per cell strongly decreases in refined cells if nothing is done. In practice this number is usually kept roughly constant by splitting macroparticles when they enter in refined regions [18, 19, 22, 21]. If the splitting procedure can trivially conserve the velocity distribution of macroparticles, it may not conserve their spatial distribution, and thus their contribution to moments on the mesh. Several splitting procedures have been proposed [34, 23, 35, 21]. They all dispatch refined macroparticles within some ad hoc distance from the original coarse one, usually in the same cell, randomly or not. In any case, these methods do not guarantee the assignment function of the original coarse macroparticle population is equal or as close as possible to that of the refined particles. Thus the contribution of the refined particles to the fine mesh will differ in some uncontrolled way from

what would be that of the coarse population. This typically adds a significant amount of noise on the refined mesh [36]. If macroparticles only interact with the finest mesh at their location, and splitting is used not to decrease their number per cell, the opposite operation, i.e. the merging of macroparticles, needs to be done as well not to stall the simulation. Merging macroparticles is more hazardous than splitting as it generally does not preserve the local structure of the distribution function [34], the very reason why kinetic simulations are done in the first place.

To our knowledge, existing AMR PIC codes use in-house developed code for adaptive meshing and evolve the system with a time step uniform across all mesh levels [19, 21]. In such a case, the time step is thus constrained by the finest grid of the model, which leads to much heavier simulations than necessary.

The Multi-Level-Multi-Domain (MLMD) method has later been proposed as a way to prevent previous issues [35, 37, 38, 39, 40]. Contrary to other AMR codes, the MLMD method solves all equations on all grid levels, as in a patch-based AMR approach. This means that not only electromagnetic fields and moments are defined on all nodes of all levels, but also macroparticles. While this may appear as an overhead to deal with macroparticles in coarse regions where there is also a fine mesh and its associated macroparticles, it comes with several advantages. First, the macroparticles only see one mesh resolution, their shape is thus perfectly constant in time. Then, merging macroparticles is not required anymore since macroparticles can simply be deleted as they exit a refined level because there is a self-consistent kinetic flux in the overlapped region of the next coarser level. As in other AMR PIC codes, refined levels are fed with macroparticles from splitting those from coarser levels at level boundaries. And once the fine solution is obtained, the electromagnetic field is coarsened onto the next coarser level, and in practice it either overwrites the coarser solution in the overlapped region, or is averaged with it. The method, originally proposed with uniform time stepping across grid levels was then updated to consider a proper stepping per level [40]. Coarser levels can evolve much fewer cycles than refined ones, which, considering the dispersive nature of kinetic plasma waves, is much more advantageous than codes based on a uniform and fixed time

stepping. Published results [35, 37, 38, 39, 40] however only demonstrate the method with only one refined level, consisting in a single refined patch with a predefined position which is fixed in time. Such a code is thus useful when the region in which to enhance the resolution is known in advance and does not evolve with time. In complex systems, where critical small scale regions are moving, appear and disappear, the lack of adaptivity imposes the refinement of a substantial part of the domain, which may become prohibitive.

Contrary to MHD codes, AMR kinetic codes cannot have arbitrarily large mesh spacing. Kinetic plasmas include intrinsic particle scales that need to be correctly resolved even in regions where "nothing" happens. Solving explicit fully kinetic equations on a mesh much coarser than the electron Debye length is unstable. Solving hybrid equations with a mesh and time step much coarser than the local ion scales is irrelevant, if not wrong. Such an upper bound to the coarsest mesh resolution make modeling very large domains expensive even with AMR. It thus appears interesting to not only consider refining the mesh and the time step, but also the physical formalism that is resolved. The MLMD method also appears promising in that regard since each refinement level, having its own macroparticles, could in principle be coupled to levels evolving different equations, given that one knows how to transfer information from one to the other.

Coupling a kinetic solver, operating on critical regions, with a fluid solver, evolving less important regions of the domain, has been an important goal over the last decade. The first coupling between MHD and fully kinetic PIC equations was achieved for local simulations of magnetic reconnection [41, 42, 43] along the inflow direction. A 2D coupling was later achieved using anisotropic MHD and Hall MHD and implicit fully kinetic equations using the codes BATS-R-US and iPIC 3D [44]. This method was then extended to 3D [45]. Simulations embedding one or several rectangular full-PIC regions in a global MHD domain were then performed for modeling the Earth's magnetosphere [46], Ganymede's magnetosphere [47, 48], the Mars' magnetosphere [49], or Mercury's magnetosphere [50]. The method has later been implemented between iPIC3D and MPI-AMRVAC [51, 52].

These works provided for the first time a large scale context to otherwise and until then isolated kinetic boxes. However, the choice of using fully kinetic equations, even through an implicit scheme,

forces the resolution of some electron particle scales which prevents a kinetic treatment of ions over large scales. Coupling fluid equations to hybrid ones would instead allow to cover much broader regions with kinetic ions.

In this paper we present the design and implementation of a new hybrid kinetic PIC code with adaptive mesh refinement inspired from the MLMD method. AMR will make high resolution more affordable for large domains and set the code architecture for future multi-formalisms simulations. Section 2 presents the hybrid formalism, its physical assumptions and associated mathematical equations. Section 3 details the AMR method we employ. The proposed code named PHARE, handles an arbitrary number of refinement levels and patches which position and shape is adapted dynamically according to the current state of the solution and refinement criteria. A refined time stepping is used to accommodate the local Courant–Friedrichs–Lewy condition [53] required by our explicit scheme, and use the largest possible time step on each refinement level. Section 4 presents the validation of the hybrid solver with and without AMR. Section 5 explains the main points of the code implementation, in particular related to the AMR mechanism. Finally, section 6 concludes the paper.

2. Hybrid Particle-in-Cell formalism

2.1. Physical assumptions and model equations

In this section we present the physical assumptions concerning the plasma dynamics and the relevant mathematical equations of the hybrid formalism. The magnetic field is evolved from the electric field via Faraday’s equation:

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} \quad (1)$$

The total current density \mathbf{j} is obtained from Ampere’s equation where the displacement current is neglected:

$$\mu_0 \mathbf{j} = \nabla \times \mathbf{B} \quad (2)$$

where μ_0 is the vacuum permittivity. The total ion density, n_i , is the sum of the density n_p of each ion population p , obtained from integration of their distribution function f_p :

$$n_i(\mathbf{r}, t) = \sum_p \int f_p(\mathbf{r}, \mathbf{v}, t) d\mathbf{v} \quad (3)$$

The total ion bulk velocity \mathbf{u} is computed from the particle flux of all populations and the total ion density:

$$\mathbf{u}(\mathbf{r}, t) = \frac{1}{n_i} \sum_p \int \mathbf{v} f_p(\mathbf{r}, \mathbf{v}, t) d\mathbf{v} \quad (4)$$

The hybrid formalism assumes that the plasma dynamics occurs at spatial and temporal scales for which quasi-neutrality holds, i.e. the total ion charge density $\rho_i = \sum_p n_p q_p = \sum_p e Z_p$, where n_p , q_p are the particle density and the electric charge of the population p , respectively, is equal to the electron charge density $e n_e$. With this assumption the electron bulk velocity can be calculated from the known electrical current density \mathbf{j} , and ion moments:

$$\mathbf{v}_e = \mathbf{u} - \frac{1}{e n_e} \mathbf{j} \quad (5)$$

The major physical assumption of the hybrid formalism, that is the less justifiable in collisionless systems, is the closure of the electron moment hierarchy at the pressure level. Currently PHARE uses the isothermal closure:

$$P_e = n_e k_B T_e \quad (6)$$

where k_B is the Boltzmann constant. This closure is the simplest and most commonly used. It assumes that the electron pressure tensor reduces to a simple scalar, P_e , proportional to the electron density. The coefficient T_e defines the temperature which is a constant in both space and time. Since the first three electron moments are known, one can use and rewrite the electron fluid momentum equation to calculate an electric field that is consistent with them. We further assume time and spatial scales at which the electron fluid is accelerated and thus neglect the bulk inertia in the momentum equation. The electric field is then given by:

$$\mathbf{E} = -\mathbf{v}_e \times \mathbf{B} - \frac{1}{e n_e} \nabla P_e \quad (7)$$

In practice, eq. 7 can be problematic because it cannot prevent the possibly infinite thinning of current sheets. In reality, the thickness of any current sheet would be limited by intrinsic electron particle orbit scales, leading to non-negligible non-gyrotropic pressure tensor components and bulk inertia. Since we do not include these scales, we need to add terms in eq. 7 to include a dissipative

scale controlling the minimal current sheet thickness in our systems. Two terms are added, the classical Joule resistive term $\eta \mathbf{j}$, where η is constant, which diffuses large scale structures, and the hyper-resistive term $-\nu \nabla^2 \mathbf{j}$, where ν is constant, which on the contrary is able to dominate over the Hall term and defines a dissipation-dominated small scale [9]. The inclusion of these two terms leads to :

$$\mathbf{E} = -\mathbf{v}_e \times \mathbf{B} - \frac{1}{en_e} \nabla P_e + \eta \mathbf{j} - \nu \nabla^2 \mathbf{j} \quad (8)$$

Finally, in the hybrid formalism, the evolution of the distribution function of each ion population p , used in equations 3 and 4, is described by the Vlasov equation:

$$\frac{\partial f_p}{\partial t} + \mathbf{v} \cdot \frac{\partial f_p}{\partial \mathbf{r}} + \frac{q_p}{m_p} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f_p}{\partial \mathbf{v}} = 0 \quad (9)$$

2.2. Macroparticles and their interaction with the mesh

In practice, equation 9 is not directly solved in PIC codes, but its solution is equivalently obtained by the evolution of a collection of so-called *macroparticles*, which, according to the PIC method, follow the same dynamical equations as real particles:

$$m_p \frac{d\mathbf{v}}{dt} = q_p (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (10)$$

$$\frac{d\mathbf{r}}{dt} = \mathbf{v} \quad (11)$$

The moments of the ion distribution used in 8 are obtained from macroparticles. The distribution function f_p of an ion population p at a given point in phase (\mathbf{r}, \mathbf{v}) space writes:

$$f_p(\mathbf{r}, \mathbf{v}) = \sum_{l=1}^{N_p} S(\mathbf{r} - \mathbf{r}_l) \delta(\mathbf{v} - \mathbf{v}_l) \quad (12)$$

where the sum is performed over the N_p macroparticles of the population p , S is a symmetric kernel normalized to 1 over its compact support, and δ is the Dirac function. Macroparticles thus have a finite spatial extent and a unique velocity. Each of the N_c macroparticles in a cell is assigned a statistical weight $W = V_c n_c / N_c$ representing its contribution to the cell's density n_c of volume V_c .

The function S is chosen to be a b-spline S_ν of order ν . The particle number density at the mesh point \mathbf{r}_{ijk} is obtained from the N_c macroparticles with position \mathbf{r}_l :

$$n_{ijk} = \frac{1}{V_c} \sum_{l=1}^{N_c} W_l S_\nu(\mathbf{r}_{ijk} - \mathbf{r}_l) \quad (13)$$

Similarly, the average particle density flux is obtained from particles velocity \mathbf{v}_l :

$$\Phi_{ijk} = \sum_{l=1}^{N_p} W_l \mathbf{v}_l S_\nu(\mathbf{r}_{ijk} - \mathbf{r}_l) \quad (14)$$

PHARE supports first, second and third order b-splines S_1 , S_2 and S_3 :

$$S_1(x) = \begin{cases} 1 - |x| & \text{if } |x| \leq 1 \\ 0 & \text{elsewhere} \end{cases} \quad (15)$$

$$S_2(x) = \begin{cases} \frac{3}{4} - x^2 & \text{if } |x| \leq 1/2 \\ \frac{1}{2} (\frac{3}{2} + |x|)^2 & \text{if } 1/2 \leq |x| \leq 3/2 \\ 0 & \text{elsewhere} \end{cases} \quad (16)$$

$$S_3(x) = \begin{cases} \frac{|x|^3}{2} - x^2 + 2/3 & \text{if } |x| \leq 1 \\ \frac{4}{3} \left(1 - \frac{|x|}{2}\right)^3 & \text{if } 1 \leq |x| \leq 2 \\ 0 & \text{elsewhere} \end{cases} \quad (17)$$

Increasing the b-spline order uses more mesh nodes. It thus decreases the noise level, at the price of heavier computations and an increased diffusion of gradients that may exist at a scale close to that of the mesh.

2.3. Normalization of physical quantities

The code evolves dimensionless quantities that are obtained with the following normalization. The magnetic field and the particle density are normalized by arbitrary constants B_0 and n_0 , respectively. Masses are normalized to the proton mass m_p and charges to the elementary charge e . It follows that the time is normalized to the inverse proton gyrofrequency $\Omega_0 = (eB_0/m_p)$. Velocities are normalized to the proton Alfvén speed in the reference magnetic field and density $V_{A0} = B_0/\sqrt{m_p n_0 \mu_0}$. The distances are thus normalized to the ion inertial length $\delta_{i0} = V_{A0}/\Omega_{ci0}$




Figure 1: 3D representation of the Yee lattice [54]. The components of the magnetic and electric fields are represented at their position. The magnetic components are located at primal locations in their direction, and dual locations otherwise, whereas the electric components are on the contrary positioned at dual locations in their direction and primal locations otherwise.

2.4. Spatial discretization

The spatial discretization of equations 1 to 8 in PHARE follow the Yee layout [54], which preserves the divergence-free character of the magnetic field. In each dimension, the components of the electric and magnetic field are positioned at integer or half integer multiples of the mesh size, so-called primal and dual positions, respectively. The 3D version of the Yee lattice is represented on Fig. 1. By choice, plasma moments are defined on primal positions.

2.5. Temporal discretization

We adopt the second order iterated Crank-Nicholson method, with predictor-predictor-corrector stages, as used in the code *Pegasus* [55]. The process starts at time step k with known electromagnetic fields \mathbf{B}^k , \mathbf{E}^k , particles position \mathbf{r}^k and velocity \mathbf{v}^k , the electron density is obtained from the ion populations density $n_e^k = \sum_p Z_p n_p^k$, and the total ion bulk velocity from the ion population fluxes $\mathbf{u}^k = \sum_p \Phi_p^k / \sum_p n_p^k$. The first step consists in computing at time $k+1$ the first predicted values (subscript $p1$) of the magnetic field, current density and electric field from Faraday's (eq. 18), Ampere's (eq. 19) and Ohm's (eq. 20) laws respectively:

$$\mathbf{B}_{p1}^{k+1} = \mathbf{B}^k - \Delta t \nabla \times \mathbf{E}^k \quad (18)$$

$$\mathbf{j}_{p1}^{k+1} = \nabla \times \mathbf{B}_{p1}^{k+1} \quad (19)$$

$$\begin{aligned} \mathbf{E}_{p1}^{k+1} &= -\mathbf{u}^k \times \mathbf{B}_{p1}^{k+1} \\ &+ \frac{1}{n_e^k} (\mathbf{j}_{p1}^{k+1} \times \mathbf{B}_{p1}^{k+1} - \nabla P_e^k) \\ &+ \eta \mathbf{j}_{p1}^{k+1} - \nu \nabla^2 \mathbf{j}_{p1}^{k+1} \end{aligned} \quad (20)$$

The time centered value ($k+1/2$) for the predicted electric and magnetic fields is then computed by simple averaging:

$$\mathbf{E}_{p1}^{k+1/2} = (\mathbf{E}^k + \mathbf{E}_{p1}^{k+1}) / 2 \quad (21)$$

$$\mathbf{B}_{p1}^{k+1/2} = (\mathbf{B}^k + \mathbf{B}_{p1}^{k+1}) / 2 \quad (22)$$

These time centered fields are then used to "push" the particles, i.e. to compute the particle velocities $\mathbf{v}^k \rightarrow \mathbf{v}_p^{k+1}$ and positions $\mathbf{r}^k \rightarrow \mathbf{r}_p^{k+1}$ (the particle pusher algorithm is detailed in section 2.6). The new values of velocity and positions are then used to calculate the predicted electron density n_e^{*k+1} and total ion bulk velocity \mathbf{u}^{*k+1} , thus completing the first predictor step.

This is followed by a second prediction, $p2$, of the magnetic field, current density and electric field:

$$\mathbf{B}_{p2}^{k+1} = \mathbf{B}^k - \Delta t \nabla \times \mathbf{E}_{p1}^{k+1/2} \quad (23)$$

$$\mathbf{j}_{p2}^{k+1} = \nabla \times \mathbf{B}_{p2}^{k+1} \quad (24)$$

$$\begin{aligned} \mathbf{E}_{p2}^{k+1} &= -\mathbf{u}^{*k+1} \times \mathbf{B}_{p2}^{k+1} \\ &+ \frac{1}{n_e^{*k+1}} (\mathbf{j}_{p2}^{k+1} \times \mathbf{B}_{p2}^{k+1} - \nabla P_e^{k+1}) \\ &+ \eta \mathbf{j}_{p2}^{k+1} - \nu \nabla^2 \mathbf{j}_{p2}^{k+1} \end{aligned} \quad (25)$$

The new time centered electromagnetic fields

$$\mathbf{E}_{p2}^{k+1/2} = (\mathbf{E}^k + \mathbf{E}_{p2}^{k+1}) / 2 \quad (26)$$

$$\mathbf{B}_{p2}^{k+1/2} = (\mathbf{B}^k + \mathbf{B}_{p2}^{k+1}) / 2 \quad (27)$$

are then used to push particles velocities $\mathbf{v}^k \rightarrow \mathbf{v}^{k+1}$ and positions $\mathbf{r}^k \rightarrow \mathbf{r}^{k+1}$ a second and final time, before computing the moments n^{k+1} and \mathbf{u}^{k+1} .

To finish, a correction step computes the final value of the electromagnetic field and current density:

$$\mathbf{B}^{k+1} = \mathbf{B}^k - \Delta t \nabla \times \mathbf{E}_{p2}^{k+1/2} \quad (28)$$

$$\mathbf{j}^{k+1} = \nabla \times \mathbf{B}^{k+1} \quad (29)$$

$$\begin{aligned} \mathbf{E}^{k+1} &= -\mathbf{u}^{k+1} \times \mathbf{B}^{k+1} \quad (30) \\ &+ \frac{1}{en^{k+1}} (\mathbf{j}^{k+1} \times \mathbf{B}^{k+1} - \nabla P_e^{k+1}) \\ &+ \eta \mathbf{j}^{k+1} - \nu \nabla^2 \mathbf{j}^{k+1} \end{aligned}$$

2.6. Particle pusher

Macroparticles are pushed following the Boris algorithm [56], here split into the 5 following steps. The particle positions and velocities are defined at the same time k but the predictor-predictor-corrector scheme uses the electromagnetic fields defined at the time $k+1/2$, so macroparticles are first push half a time step:

$$\mathbf{x}^{k+1/2} = \mathbf{x}^{(k)} + \frac{\Delta t}{2} \mathbf{v}^k \quad (31)$$

The electric $\mathbf{E}^{k+1/2}$ and magnetic field $\mathbf{B}^{k+1/2}$, known on the mesh, are interpolated at the particles' position using the b-spline assignment functions discussed previously. These interpolated fields $\tilde{\mathbf{E}}^{k+1/2}$ and $\tilde{\mathbf{B}}^{k+1/2}$ are then used to update the particles' velocities following:

$$\mathbf{v}^- = \mathbf{v}^k + \frac{\Delta t}{2m} \tilde{\mathbf{E}}^{k+1/2} \quad (32)$$

$$\mathbf{v}^+ = \mathbf{v}^- + \frac{q\Delta t}{m} \left(\frac{\mathbf{v}^- + \mathbf{v}^+}{2} \right) \times \tilde{\mathbf{B}}^{k+1/2} \quad (33)$$

$$\mathbf{v}^{k+1} = \mathbf{v}^+ + \frac{\Delta t}{2m} \tilde{\mathbf{E}}^{k+1/2} \quad (34)$$

The final particles' position is then given by:

$$\mathbf{x}^{k+1} = \mathbf{x}^{k+1/2} + \frac{\Delta t}{2} \mathbf{v}^{k+1} \quad (35)$$

3. Adaptive Mesh Refinement

3.1. Patch hierarchy

The adaptive mesh refinement in PHARE follows the so-called *patch-based* approach. Equations are solved in rectangular domains of uniform mesh resolution called *patches*. Each patch contains data that is either represented on a grid (e.g field quantities) or that is gridless (e.g macroparticles). Because of spatial derivatives and of the finite extent

of the macroparticle assignment function S_ν , computing the solution in the patch domain requires the knowledge of the solution outside the patch in a so-called *ghost layer*. This ghost layer represents the current state of the solution of the adjacent regions of the simulation domain, and its thickness depends on whether it concerns fields or macroparticles, as explained in more details in sections 3.4 and 3.5.

The union of all patches of the same mesh resolution defines a so-called *patch level*, and the ensemble of patch levels forms the so-called *patch hierarchy*. This is represented on figure 2, in a two-dimensional geometry, for illustrative purposes. Each level increment corresponds to a spatial refinement of the mesh of the enclosed patches by a specific refinement ratio. PHARE adopts an isotropic refinement ratio set to 2. The number of patch levels composing the patch hierarchy can change over time depending on the state of the system, and is bounded by a runtime user supplied parameter.

Cells constituting the ghost layer of a patch can be of two types, represented on Fig. 2. Cells overlapping those of the domain of a neighbor patch in the same patch level are called *patch ghost cells*. The information (field and particles) in these patch ghost cells is just an exact copy of the information existing in the overlapped neighbor domain cells. Cells overlapping the next coarser level mesh are called *level ghost cells*. Their data is obtained from refining next coarser level data. *Patch ghost cells* overlap the cells of the domain of a neighbor patch in the same patch level. The information (field and particles) in these patch ghost cells are just an exact copy of the information existing in the overlapped neighbor domain cells. *Level ghost cells* overlap the next coarser level mesh and their data is obtained from refining next coarser level data.

By definition, the union of the patches of the coarsest patch level cover the whole simulation domain. Therefore ghost layer cells of that level are always patch ghost cells (for periodic domains). The filling of field quantities in ghost nodes and macroparticles in ghost cells for those two types of boundaries need to be handled differently, the details are given in subsections 3.4 and 3.5.

3.2. Refined time stepping evolution

PHARE advances the solution on the patch hierarchy using a refined time stepping procedure, thus avoiding the evolution of the whole domain being constrained by the smallest time step associated




Figure 2: Illustration of a patch hierarchy composed of 3 levels. In this simplified example, the coarsest level L_0 is composed of a single patch represented in black, L_1 has two patches represented in green. Finally L_2 has two blue patches. The two patches on L_1 (resp. L_2) are separated by a patch boundary represented in orange (resp. purple). The other boundaries of L_1 (resp. L_2) are represented in green (resp. blue) and define the so-called level boundaries.

with the finest level. The time stepping is illustrated in Fig. 3 for a patch hierarchy composed of 3 patch levels. The evolution proceeds in a recursive manner as follows. Level L_0 advances the solution on all of its patches by one coarsest time step Δt_0 . Then before it advances another step, L_1 advances its patches by $\Delta t_1 = \Delta t_0/r^2$ where r is the refinement ratio. Then since L_2 is the finest level at that time, it is advanced r^2 steps, of size $\Delta t_1/r^2$, up to the time reached by L_1 . At this point, the solution obtained on L_2 has the best resolution and thus is coarsened to overwrite the one existing on L_1 on overlapped regions, before L_1 is advanced by Δt_1 again. This is repeated r^2 times until L_1 itself reaches the time at which L_0 is and coarsens its solution onto L_0 , which then advances to its second step, etc. This procedure is repeated as many times necessary until L_2 reaches the final desired simulation time.

3.3. Initialization of a patch level

This section explains how data is initialized on a patch level, depending on whether it concerns the initialization of the coarsest level of the hierarchy, a new finest refined level, or the reorganization of an existing one.

3.3.1. Initialization of the coarsest level

The coarsest level uses user-defined functions to initialize fields and particles at the start of a simulation. The macroparticles are loaded on a per-cell basis according to a user supplied density profile and a constant number of macroparticles per cell and per population. The velocity of macroparticles is distributed according to a Maxwellian distribution function with moments locally defined via the




Figure 3: Refined time stepping evolution of the patch hierarchy illustrated for 3 patch levels. On that figure, the evolution of the coarsest (L_0), intermediary (L_i) and finest level (L_3) is represented from left to right. Blue vertical arrows represent the time step taken by each level. Horizontal orange arrows represent the coarsening of L_{i+1} data onto L_i mesh.

user supplied functions. Once the magnetic field and the plasma moments are set, the electric field simply results from Ohm's law (eq. 8).

3.3.2. Initialization of a new refined level

Whether it is at the start time of the simulation or when a new level at the finest resolution is being created during the simulation, a refined level is initialized by refining data existing on the next coarser level. Fields are refined using linear spatial interpolation while macroparticles are refined using the splitting procedure described in Smets et al.[36]. This splitting procedure ensures that the combination of the refined macroparticle b-splines is either exactly or optimally close to the b-spline of the original coarse macroparticle, and thus preserves the contribution of macroparticles to the mesh. Unlike other methods (see [36] and references therein), it ensures that no or minimal spurious noise is added to the refined grid. The method uses pre-computed and tabulated values for the relative position and weights of the refined macroparticles.

3.3.3. Initialization of a level during regridding

The last way a level can be initialized is when regridding occurs. Regridding is the change of the patch hierarchy structure in terms of patch number, position and geometry, and it is often employed to improve the distribution of spatial res-




Figure 4: Illustration of a regridding operation in a patch hierarchy with two levels. The green level is being removed and replaced by the blue one. The new level copies data from the red cells of the old level and refines data from coarser level in the blue cells.

olution around critically important regions as they evolve.

Whenever regridding occurs, a set of levels are removed from the patch hierarchy, and replaced by new ones, as illustrated in Fig. 4. Data is copied (resp. streamed) from local (resp. remote) memory where there is an overlap between old and new patch level regions, and refined from coarser data where there is none. In the latter case, the refinement of fields and particles is performed through the same functions as those used for the initialization of new finest level initialization.

3.4. Macroparticle deposition to the grid and the handling of ghost particles at patch and level boundaries

We now describe how the moments of the ion distribution function are computed on the grid from the macroparticles and introduce the types and use of ghost macroparticles. Without loss of generality an example for a 1D grid and second order interpolation is illustrated in Fig. 5. The first step in the computation of the moments consists in depositing the properties of each macroparticle located in patch domain cells. This is the step **A** of Fig. 5. At this point, nodes 2 to 5 are so-called *complete*, in a sense that all macroparticles that can reach them have contributed to them. For the case of second order interpolation, nodes 0 and 1, as well as nodes 6 and 7, however, are *incomplete*. They received

contributions from domain macroparticles, but lack that of macroparticles located outside the domain. This missing contributions is the main reason why there is a need for so-called *ghost* macroparticles, i.e. macroparticles located beyond the patch domain boundaries that represent the ion populations in these regions. Another reason is that one needs to inject macroparticles into the patch domain consistently with the particle flux existing at the boundaries. This said, we shall now see that two types of ghost macroparticles exist in PHARE: patch and level ghost macroparticles, as described in the two following paragraphs.

3.4.1. Patch ghost macroparticles

Patch ghost macroparticles are those located in patch ghost cells, i.e. those overlapping domain cells of a neighbor patch in the same patch level. By construction, these macroparticles are clones of those positioned in neighbor patches domain cells overlapping the ghost layer. Like those in the domain cells, patch ghost macroparticles are pushed. Those entering the domain contribute to the moments on the mesh like domain macroparticles in step **A** and are copied into the domain particle array. Because domain particles of all neighbor patches have moved, at this point the ensemble of macroparticles in the patch ghost layer is not the exact copy of the neighbor overlapped domain cells anymore. Patch ghost layers are thus refilled from neighbor patches, through either a copy from local memory if the patch and the neighbor are on the same MPI process, or a transfer if they are not. Note that for interpolation of order 2, represented on Fig. 5, ghost macroparticles located as far as $1.5\Delta x$ away from the border can contribute to domain nodes. For interpolation of order 1 and 3, the ghost layer width becomes Δx and $2\Delta x$, respectively. Once received, patch ghost particles are interpolated onto the grid, which in Fig. 5, completes nodes 0 and 1 and is illustrated as step **B** in the figure.

3.4.2. Level ghost macroparticles

At the end of step **B** of Fig. 5, nodes 6 and 7 thus still lack contributions. This other boundary of the patch is a level boundary. There are therefore no neighbor patch cells from which to clone macroparticles in that ghost layer. The kinetic state in this region is only represented in the next coarser level and that is where the information is thus taken from. The level ghost cells are therefore filled by




Figure 5: Representation of three steps **A**, **B** and **C** of the moment deposition phase from macroparticles to the mesh of a patch. For the example, this patch is one that is, on its left, adjacent to another patch of the same level, and at the border of the level on its right. Left and right boundaries are thus a patch boundary and a level boundary, respectively. On each step, primal nodes are represented by colored circles. Nodes colored in blue have, at the given step, received the contribution of all macroparticles that can reach them, and are said "complete". Orange nodes circled in blue have received some but not all contributions and are thus said "incomplete". Macroparticles are represented by their shape factor, here at order 2. On step **A**, blue particles represent domain particles. On step **B**, the green and purple macroparticles represent patch ghost macroparticles that are the closest and farthest from the domain, respectively. The purple particle is just closer to the limit of $1.5dx$ visualized by the purple square, beyond which no particle at this interpolation order can reach a domain node. On step **C**, the color code green and purple as the same meaning. There are however 4 macroparticles, representing the refined macroparticles obtained from splitting next coarser ones at its previous time (solid lines) and current time (dashed lines). α and $1 - \alpha$ are linear time interpolation coefficients weighting the contribution of the future and past sets of macroparticles, respectively.

splitting next-coarser macroparticles, to ensure the continuity of phase space at these boundaries.

It is worth noting that at the time t these operations are performed on level L_i , the next coarser level L_{i-1} already is at a later time, and there is no representation of its distribution at time t . PHARE handles this situation with two different level ghost macroparticle buffers. One represents macroparticles refined from next coarser ones at its previous step, the other represents those refined from the current next coarser state, i.e. future time for level L_i . Nodes 6 and 7 receive the contribution of both buffers, each being weighted via linear time interpolation between the two surrounding next coarser times. This completes nodes 6 and 7 and represents step **C** of Fig. 5.

Like for patch ghost boundaries, the domain must also receive incoming macroparticles at level boundaries. The two level ghost macroparticle buffers described above represent fixed times and cannot change during all steps taken by level L_i to reach the time at which level L_{i-1} is. A third buffer is thus used. At the start of the substepping cycle of level L_i , i.e. when level L_i is at the same time as level L_{i-1} , this buffer is identical to the other buffer defined at this time. However, macroparticles contained in that buffer are pushed at each level L_i substep. Those macroparticles which enter the patch domain are deposited on domain nodes, exactly like is done for entering patch ghost macroparticles.

Since each level has its own set of macroparticles, those which leave a patch domain after being pushed, are simply deleted. In the case they are leaving through a patch boundary, this means that its clone has entered the neighbor patch domain from the patch ghost layer. When a macroparticle leaves through a level boundary, nothing is done except deleting the macroparticle, this kinetic flux being already represented consistently on the next coarser level.

3.4.3. Number of ghost cells for particles

The number of ghost cells filled with macroparticles depends on the interpolation order. At interpolation order 1, macroparticles interact with its two surrounding primal or dual nodes. The border primal node thus cannot be reached by macroparticles outside the domain further than one cell away. We thus only need one ghost cell filled with macroparticles. At order 2, macroparticles as far as 1.5 cell away from the domain can reach the border primal node. The particle ghost layer is thus chosen to be

two cells wide. At order 3, two cells need to be filled with macroparticles and thus the width of the ghost layer is the same as for interpolation order 2.

3.5. Filling field ghost nodes

After the update of the magnetic field, electric field, and current density, the ghost nodes need to be filled with appropriate values. On patch ghost nodes, data is just copied (resp. streamed) from local (resp. remote) memory from overlapped nodes on the same patch level. However for level ghost nodes, values need to be obtained from the next coarser level. However, because levels evolve with different time steps, the values needed do not exist on the next coarser level, which is in fact, at this point of the algorithm, at a later time. Therefore in addition to the spatial linear interpolation, we also perform a time linear interpolation between the previous and current state of the next coarser level domain nodes. In PHARE, only the electromagnetic fields (\mathbf{E} , \mathbf{B}) and the current density \mathbf{j} are refined in this way on ghost nodes. Plasma moments are not refined on ghost nodes but obtained from ghost macroparticles, as described in subsection 3.4. The number of field ghost nodes is primarily constrained by the width of the macroparticle ghost layer and the fact that these ghost particles are pushed. At first order interpolation, we thus need two primal and dual ghost nodes. The same reason implies having 3 and 4 ghost cells for fields of orders 2 and 3, respectively.

3.6. Coarsening of the fine solution

The assumption behind using AMR is that the solution gets better as the mesh gets finer. Thus once a fine patch level has reached the same time as its next coarser level, the coarse region overlapped by the fine level is overwritten by the fine solution via a coarsening operation. Note that in PHARE this coarsening only concerns field quantities (total density, bulk velocity, electric and magnetic fields) and not particles. The distribution function in a patch level is *not* coarsened onto its next coarser level. Macroparticles located in regions overlapped by a finer level always see electromagnetic fields updated by the finer solution before taking a new step. This ensures the dynamics is consistent throughout levels and the whole hierarchy stays synchronized.

The field coarsening operation depends on whether the quantity is positioned at a dual or primal location. Dual coarse nodes receive the equally




Figure 6: Illustration of the coarsening algorithm on a coarse (blue) patch from refined (orange) patch values. Circles and triangles represent primal and dual positions, respectively.

weighted contribution of the two surrounding fine dual values while primal nodes receive the weighted (0.25, 0.5, 0.25) contribution of the three surrounding fine primal values. This operation is represented along a single dimension in figure 6, 2D and 3D coarsening just consist in the outer product of this pattern.

3.7. Refinement strategy

Periodically, the patch hierarchy needs to be updated for the fine levels to map critical regions of the solution. A criterion is therefore defined to "tag" domain cells that should be part of a region to be refined. In general, refinement is needed in regions where the grid is insufficiently fine to resolve the fine structures that would self-consistently appear in the solution. The need for refinement is usually estimated in an empirical manner. For instance the AMR MHD code `Flash` [26] uses by default an adaptation of the Lohner [57] criterion, an estimator based on a modified second derivative, normalized by the average of the gradient over one computational cell. Such criterion has been tested in our code but leads to many regions being sporadically refined as a result of the great sensitivity of this method to the numerical noise in the PIC solution. There is in fact no unique recipe to design a tagging criterion. One could even argue that the best set of criteria depends on the specific application of the code the user has in mind. The code `Flash` provides a way for user to modify or define a custom refinement criterion. The AMR electromagnetic full PIC code in [18] tags cells depending on their size relatively to the local Debye length because it is a severe spatial constraint on the stability of explicit full PIC codes and also on the local amplitude of the electron current, a criterion fine-tuned for the study of magnetic reconnection. In the AMR hybrid PIC code `AIKEF` [21], the local current density is said to "guide" the refinement

but the exact criterion is not specified. In this paper our goal is to provide an example of a criterion that allows to track sharp gradients and that serves to validate our AMR implementation, it is not our aim here to provide a universal tagging recipe. In the future, PHARE will also provide a way to define custom refinement criteria.

We define the local normalized variation of the magnetic field at cell ijk as :

$$r_{ijk} = \max_{\substack{d=\{i,j,k\} \\ \alpha=\{x,y,z\}}} \frac{B_{\alpha}^{d+2} - B_{\alpha}^d}{1 + B_{\alpha}^{d+1} - B_{\alpha}^d} \quad (36)$$

the max is taken along each dimension independently, and then for each magnetic field component.

Finally, the cell tag is set to 1 or 0 depending on whether the ratio r exceeds an empirical threshold σ or not. Results in this paper use $\sigma = 0.1$.

$$c = \begin{cases} 1 & \text{if } r_{ijk} > \sigma \\ 0 & \text{otherwise} \end{cases} \quad (37)$$

Once cells are tagged for refinement, they are grouped into boxes by a so-called tile clustering algorithm [58], used to adapt the number and geometry of patches in a patch level.

Section 4.2.2 will validate this approach is able to capture gradients and show its ability to capture fine structures.

4. Validation of the code

A large number of unit and functional tests have been developed and are executed at each modification of the code. These tests are focused on the behavior of specific components of the code, from an engineering point of view. They decrease a lot the probability components are malfunctioning and above all ensure successive modifications do not alter the working state of the code.

In addition to these tests, we also perform a series of validation tests on physical processes, which are ran during nightly builds when the code has changed. We now discuss these validation tests by starting first with those validating the hybrid solver without refinement, and then including static and dynamic refinement.

4.1. Validation of the hybrid solver

4.1.1. Wave dispersion relations

In this section, we test whether the hybrid solver correctly captures the dispersion of waves propagating parallel to a uniform magnetic field in the




Figure 7: (ω, k) diagram for 1D (larger crosses) and 2D (smaller crosses) computations of the Alfvén-Ion-Cyclotron (red) and Whistler (blue) modes displayed in log-log scales.

dim=1		
setup	"low"	"high"
n_x	128	32
d_x	0.8	0.2
Δt	0.01	0.001
T	200	20
m	(1,2,4,8)	(1,2,4,8)
dim=2		
setup	"low"	"high"
(n_x, n_y)	(200, 10)	(100, 10)
(d_x, d_y)	(0.4, 0.4)	(0.2, 0.2)
Δt	0.00125	0.00025
T	200	20
m	(2,4,8)	(4,8,16)

Table 1: Values of the dimension, setup, number of grid points, grid size, time step, time of duration of the simulation and modes for the runs that are used for the wave dispersion test. These 4 sets of parameters are associated to 8 runs, considering the 2 polarization, left and right in each cases.

hybrid formalism. We initialize both 1D and 2D simulations, each containing an arbitrary set of initial modes (4 and 3 modes for 1D and 2D, respectively) in a Maxwellian electron-proton plasma initially stationary with a density equal to 1, a proton thermal velocity equal 0.01 and a electron thermal velocity equal to 0. For 1D runs, the initial background magnetic field is set to $B_x = 1$ and circularly polarized perturbations are imposed on B_y and B_z . For 2D runs, the initial background magnetic field is oriented along the diagonal of the domain, and the analysis is made on magnetic components perpendicular to this direction. Two setups, hereafter called *low* and *high*, are ran in both 1D and 2D, to capture the wave dispersion at large and small scales, respectively. In each final setup, an arbitrary number of wave modes m are imposed initially. The dimensionality, number of cells, mesh resolution, time step, simulation duration and excited modes numbers are summarized in Tab.1 for all used setups. In all simulations, the hyper resistivity coefficient ν is set to 0.005, the initial amplitude of the excited modes is 0.01 and their phase is random. Each simulation is performed twice, with the excited mode imposed to have either Left- or Right-handed polarization. Then, the Fourier transform of $B_y \mp iB_z$ is computed, the sign depending on the wave polarization. For each mode m associated to the excited wave number k_m , the self-consistent frequency ω_m is found as a local maximum in spectral energy, which is then reported on Fig 7.

The theoretical fluid dispersion relations for the Left-Alfvén-Ion-Cyclotron (eq. 38) and Right-Whistler (eq. 39) modes are also shown on the same figure with the analytical expressions:

$$\omega_L = \frac{k^2}{2} \left(\sqrt{1 + 4/k^2} - 1 \right) \quad (38)$$

$$\omega_R = \frac{k^2}{2} \left(\sqrt{1 + 4/k^2} + 1 \right) \quad (39)$$

The Alfvén-Cyclotron and Whistler modes follow the fluid theory very well in such a cold plasma. After a first visual inspection of the results, the ω_m values are tabulated and the functional tests asserts the code continues finding them during each nightly build.

4.1.2. Right-hand resonant streaming instability

When a beam of charged particles is moving across a magnetized plasma at rest, three kinds of




Figure 8: Time evolution of the first mode of the Fourier transform of $B_y - iB_z$ (black circles) and the associated exponential fit (red curves).

instabilities can develop [59]. The resonant left-hand and right-hand modes, and the non-resonant mode. There is a range of plasma parameters for which only one mode is unstable or where others have a much smaller growth rate[60]. It is thus possible for us to select one of the resonant modes to assess the ability of the code to solve kinetic physics. The right-hand resonant mode is chosen because it results in the most localised structures in phase space, allowing us to test their transport through patch level boundaries (see section 4.2.3). We use a setup already investigated [60] which consists of a main population of protons at rest, with a density equal to 1.0, and a proton beam population, having a density of 0.01 and a bulk velocity of 5.0 embedded in a $B_x = 1$ magnetic field. These two populations are initially Maxwellian with a temperature equal 0.1. The electron temperature is also set to 0.1 so the most unstable mode is numerically predicted with a wavelength ~ 33 and a growth rate ~ 0.09 .

We used a computational domain of length 33 which is the wavelength of the most unstable mode. The simulation is performed with two fixed refinement levels to assess the validity of the multi-grid system as discussed in section 4.2.3. We used 100 macroparticles per cell for the main population, and 1000 for the beam, in order to increase the statistical description of the resonant particles. The unstable mode being right-hand circularly polarized, we compute every $t = 1$ the spatial Fourier transform of $B_y - iB_z$ and then reported with black solid circles in Fig. 8 the modulus of the first mode, that is the one whose wavelength equals the length of

the computational domain, the red curve is an exponential fit of these values. The computational results from several simulations clearly show that, even with a magnetic seed for the magnetic perturbation of amplitude 0.1, the beginning of the linear phase and then the time at which the linear mode is saturating is varying from one run to another, as it depends on the initial microscopic state. We thus have performed 200 runs for which we identify the saturation time, before which the growth of the instability is evaluated by a fit to an exponential. We obtain an averaged growth rate of 0.09 with a standard deviation smaller than 1% of this average, recovering the numerical values previously obtained [60].

4.2. Validation with mesh refinement

Previous tests assessed the ability of the code to obtain solutions expected in the hybrid formalism. In the following parts, we now test the behavior of the code with refined meshes, first refined statically from the initial condition, then dynamically by tagging regions of interest.

4.2.1. Alfvén wave

The goal of this test is to assess the distortion an Alfvén wave may experience as it propagates across patch and level boundaries in 1D. We initialize PHARE with an Alfvén wave of amplitude $\delta B_y = 0.01$ and a wavelength equal to the length of the computational domain. The whole domain at the coarsest level consists of 4000 cells for a total size of $L_x = 1000$ and a time step $dt = 0.001$. The domain is refined between $x = 455$ and $x = 550$ with only one refined level. A single proton population with unit density and thermal velocity $v_{th} = 0.01$ is used. The simulation is ran with 8 MPI processes until $t = 1000$, corresponding to the time at which the wave should have traversed the whole domain. Figure 9 shows the B_y component of the magnetic field at $t=0$, $t=500$ and $t=1000$. One can see the wave is exactly at the expected position at both $t=500$ and $t=1000$. A fit to $B_0 \cos(kx + \phi(t))$ is performed every $\Delta t_{fit} = 0.1$, to obtain the numerical values of B_0 , k and ϕ , the amplitude, wave number and phase of the wave, respectively. This allows us to calculate the phase speed $V_\phi = \dot{\phi}/k$ as a function of time, which is found to be $V_\phi = 1.0031 \pm 0.0292$. After a million time steps, no noticeable artifact appears on the signal although it has crossed many patch boundaries and the two level boundaries. This tests




Figure 9: Alfvén wave moving across a fixed refined region of space. The solid black, blue and orange lines represent the B_y component at times 0, 500 and 1000, respectively. The dashed black line represents $0.01 \cos(2\pi/1000x + 2\pi/1000t)$ for $t=500$. The blue area is the refined portion of the domain.

is automatically executed each time the code is updated and set to fail if the absolute error on the calculated phase velocity is larger than 0.05.

4.2.2. Moving tangential discontinuity

In this test we setup two sharp tangential discontinuities in a globally rightward moving plasma with a velocity $V_x = 2$. In a fluid framework, one expects the two tangential discontinuities to be simply advected at the prescribed velocity and to keep a profile identical to the initial one in the co-moving frame. In this kinetic framework however, the discontinuity is not a kinetic equilibrium and thus can have its structure slightly altered within the few first ion cyclotron times. Finite diffusion can also decrease its sharpness over time. Two simulations are ran, one with adaptive mesh refinement and one without. The non-refined simulation has a mesh size equal to the coarsest mesh size of the refined simulation, set to $\Delta x = 1$. It is enough in regions where the plasma and field are homogeneous, but poor in the discontinuity regions since it equals their half-thickness. The refined run is allowed to create up to 3 levels, i.e. to refine the mesh up to $\Delta x = 0.25$. The goal of this test is twofold : (i) ensure that in the refined case, PHARE refines the solution around the location of the discontinuities at all times and (ii) compare the evolution of the structure of the discontinuity between the refined and non-refined runs. The domain extends over $L_x = 200$. The initial density is $n = 1$, the magnetic field is given by $B_z = 0.5$ and $B_y =$

$-1 + 2(S(x, 0.25L_x, 1) - S(x, 0.75L_x, 1))$ where

$$S(y, y^*, \lambda) = 0.5 \left(1 + \tanh \left(\frac{y - y^*}{\lambda} \right) \right) \quad (40)$$

This produces two discontinuities centered at $x_0 = 50$ and $x_1 = 150$, respectively, of half-width 1. A single proton population is initialized with 100 particles per cell with first order shape function. The ion temperature is chosen to satisfy at each point the total pressure balance: $T(x) = 1 - 0.5B(x)^2$. The hyper-resistivity ν is set to 0.01, the electron temperature is set to 0. The total simulated time is $T = 20$ with 2000 time steps. Figure 10 shows the magnetic field component B_y and the associated current density component J_z along x at $t = 11$ for both runs, in a zoom of the region of the rightmost discontinuity. Given its initial position, and the bulk flow, the discontinuity should in theory be centered at $x = 172$ at this time. One can first notice that PHARE has created 3 patch levels with two refined ones. Each refined level is composed of two patches and is centered in the region of the discontinuity, as expected. Then, one can notice the discontinuity has moved in both runs and roughly arrives at the same location. However if the peak current density is located exactly where expected in the refined solution, it is not the case in the other run. The non-refined run sees its discontinuity profile altered by spurious downwind oscillations and a somewhat decreased slope. This alteration results in a current density which peak value is reduced by half compared to its initial value and its global shape is different. This test is also automatically ran each time the code is updated, it fits the final position of the discontinuity and asserts its width and position are consistent with their initial values.

4.2.3. Continuity of phase space across level boundaries

We previously discussed how the code recovers the growth of the right hand resonant mode associated to the ion streaming instability. This instability will lead beam particles to give their energy to electromagnetic fluctuations, structuring the phase space density in the process. Our goal in this test is to assess how well these self-consistent phase space density structures are transported across (fixed) level boundaries.

In Fig. 11, the particle density in phase space (x, V_x) is displayed at $t = 0$ (upper panel), $t = 64$




Figure 10: Top panel shows B_y for both runs for the whole domain. The orange dashed-line represent the initial profile. Middle panels compares of the B_y component of the magnetic field at $t = 11$ without refinement (black line) and with refinement (blue line). Bottom panel has the same format for the J_z component of the current density. The mesh size of the non-refined simulation is $dx = 1$, which is equal to the lowest resolution of the refined simulation. Intermediary and finest resolutions have $dx = 0.5$ and $dx = 0.25$, respectively. The vertical red line at $x = 172$ marks the theoretical center of the discontinuity at that time. In all panels, light blue rectangles represent the extent of patches, whose refinement level increases vertically.




Figure 11: logarithmic value of the distribution function depending on the X position and V_X velocity for main and beam population at $t = 0$ (upper panel), $t = 64$ (middle panel) and $t = 85.4$ (lower panel). The two vertical black lines are framing the first level of refinement and the two vertical red lines the second one.

(middle panel) and $t = 85.4$ (lower panel) for the same two-stream run that was previously discussed. The upper panel depicts the initial setup. The middle panel shows the beginning of the non-linear phase of the right hand resonant mode where the beam is corrugated with the $m = 1$ mode clearly dominating. We can also notice the decrease of the bulk velocity of the beam, as its associated energy is efficiently converted to magnetic energy. The time in the lower panel has been chosen to emphasize that these complex, non-linear phase space structures cross the level boundaries without any alteration.

4.2.4. Magnetic reconnection

We now show results in two dimensions, with simulations of the magnetic reconnection process. Magnetic reconnection occurs in magnetized plasmas where the magnetic field changes its orientation over a short distance, forming a current sheet. Reconnection is a self-driven process where magnetic tension from the newly connected field lines expels the plasma from the reconnection site, thereby pulling in upstream magnetic flux and the frozen in plasma. Here we setup a double periodic system where current sheets are positioned at $y_1 = 0.3L_y$ and $y_2 = 0.7L_y$, in a domain box that is $L_x = L_y = 40$ wide. The magnetic field profile is

given by

$$B_x(y) = B_1 + (B_2 - B_1) (S(y, y_1, 0.5) - S(y, y_2, 0.5)) \quad (41)$$

and $B_1 = -1$, $B_2 = 1$ a.s S is given by eq. 40. A small magnetic perturbation is superimposed to this magnetic field at the center of the two current sheets. The electron temperature is set to 0. There is only one ion population made of protons, which particle density is given by

$$n(y) = n_b + \cosh^{-2}\left(\frac{y - y_1}{\lambda}\right) + \cosh^{-2}\left(\frac{y - y_2}{\lambda}\right) \quad (42)$$

where the background density is $n_b = 0.4$. The ion temperature is chosen so that the total pressure $K = 0.7$ is initially uniform in the whole domain, and is given by

$$T(y) = \frac{1}{n} (K - 0.5B^2(y)) \quad (43)$$

Three simulations are performed. A first simulation, with AMR enabled, is ran with a coarsest spatial resolution set to $\Delta x = \Delta y = 0.4$ and the associated time step to $\Delta t = 0.005$. The maximum number of levels is set to 3, which means 2 refined levels. A second simulation is performed, so called *equivalent coarsest*, with a uniform mesh at the coarsest resolution of the AMR run and all other parameters identical. The third simulation, called the *equivalent finest*, with the same parameters as the other two except now the spatial resolution is $dx = dy = 0.1$ like the finest one allowed in the AMR run. In all runs the resistivity is set to $\eta = 0.001$ and hyper-resistivity to $\nu = 0.002$.

Figure 12 shows the result of all simulations, zoomed in one of the two reconnection sites. The top panel shows the out of plane magnetic field component B_z , on top of which the borders of L_1 and L_2 patches are shown. The magnetic field reveals the well-known quadrupolar pattern associated with the Hall effect [61]. L_1 patches are located in a narrow band between $y = 24$ and $y = 31$, and L_2 patches are nested in that band, mapping even closer the gradients of the magnetic field structure. The three panels below show the ion density at the same time. All simulations show the expected density pile-up in the downstream X -direction resulting from the expulsion of the initial current sheet with high density plasma downstream towards the periodic boundaries. In all simulations, the density




Figure 12: Top panel shows the out of plane component B_z color coded. Red and black rectangles represent the borders of patches of the L_1 and L_2 refinement levels, respectively. The panels on the second and third row represent, at the same time $t = 3$, the total ion density and the out of plane current density, respectively. On these two rows, the middle panel represent the result of the AMR simulation with 3 mesh levels. The leftmost panel represent the result of a uniform mesh run at the coarsest resolution of the AMR one, so-called *equivalent coarsest* run. The rightmost panel is the so-called *equivalent finest* run, with a uniform mesh at the finest resolution of the AMR run.

is lower on the separatrices, a typical feature of collisionless reconnection [62]. Bottom panels show the out of plane current density. The current density is greatly enhanced around the reconnection site to sustain the large magnetic reversal over the sub-ion scale distance. It is also enhanced on regions connected to the reconnection site, i.e. separatrices. On the left and right side of the reconnection site, the current density increases and forms a broader sheet that results from the pile-up of reconnected magnetic flux against the initial tangential sheet and the periodic boundary. The finest uniform simulation clearly resolves the hyper-resistive current sheet [61]. The coarsest simulation however sees the current sheet collapse down to the grid scale, resulting in many numerical artifacts around the reconnection site. For that simulation to resolve the reconnection current sheet, the hyper-resistivity would need to be set to higher values that would start to dissipate magnetic field at ion scales. The AMR simulation in the middle panels shows the current sheet is well resolved as in the equivalent finest run. The spatial resolution quickly becomes coarse upstream of the reconnection site, however and contrary to the equivalent coarsest run, no numerical artifact is observed. These results confirm that the mesh refinement in PHARE is able to map and resolve fine scale structures at high resolution.

5. Software details

Most of the PHARE code base is written in C++. The high level user input/output interface is written in Python. The code is developed in an open source fashion¹. The management of adaptive mesh refinement intrinsic details is handled by the massively parallel library SAMRAI [58]. SAMRAI performs operations related to the definition of the position and geometry of the patches, through an tile clustering algorithm. It also partitions the patches among the different MPI processes, here the so-called cascade partitioner, also provided by SAMRAI. SAMRAI manipulates abstractions of the data defined on the patches, its transfers and copies, and the determination of geometrical overlaps between patches of the same level or of different ones. Concrete implementation of these data, transfers and geometrical operations are developed in PHARE and follow the concepts explained in above

¹<https://github.com/PHAREHUB/PHARE>

sections. The code is written in such a way that components solving the hybrid equations and those transferring information to the borders of refinement levels, and coarsening data to the next coarser level, are hidden behind abstractions. In this way, despite that here only hybrid PIC equations and data transfers are implemented, the code architecture is already made to support other possible implementations. Future work will in particular focus on implementing a fluid hybrid model, operating on the coarsest levels of the patch hierarchy. The SAMRAI library has already been used in such a multi-formalism simulations [63]. In this way, PHARE could model large fluid domains in which hybrid PIC patches of incrementally fine resolution are created dynamically over critical regions.

6. Summary and perspective

This paper presents the code PHARE, evolving hybrid PIC equations using a patch based AMR approach. Hybrid equations are solved on uniform resolution cartesian patches with a predictor-predictor-corrector temporal scheme and a Yee spatial discretization. Macroparticles are modeled with b-spline interpolation kernels of order 1, 2 or 3 and their trajectory is calculated with the Boris algorithm. Any number of ion populations can be modeled. The AMR methodology adopted here is inspired from the MLMD method where all refinement levels evolve their fields and their macroparticle populations, which allows not to depend on macroparticle merging. The evolution of the patch hierarchy is done through a recursive procedure using a time step divided by 4 between each refinement level, where the spatial resolution is increased by a factor of 2. Validation tests have been presented. They reveal the code successfully captures the expected dispersion of waves in the hybrid kinetic regime. The solver has also been demonstrated to accurately predict the growth rate of the kinetic ion streaming instability. Structures in phase space developing during the nonlinear phase of the instability have been shown to propagate through level borders without alteration. A large scale Alfvén wave was shown to propagate in a system with two grid levels without alteration at level borders and with an accurate phase velocity. The adaptive meshing has first been demonstrated in 1D to be able to capture two tangential discontinuities advected in a global flow. The refinement

was demonstrated to be able to conserve the initial magnetic field profile whereas coarser simulation cannot, resulting in distortions in the current sheet structure. Finally, the code is validated in a 2D system and shown to capture well-established features of collisionless magnetic reconnection. The adaptive meshing adequately maps the structures of the magnetic field and helps resolving the sub-ion scales without needing a uniform high resolution mesh.

This paper validates the method behind the making of a patch-based AMR hybrid PIC code. Future work will now focus on the several following points. First, extend the code to simulate three-dimensional systems. All the code components are already implemented in 3D. Remaining work however remains in its validation and the development of 3D post-processing tooling. 3D runs also call for optimizations of the various components of the code. The code benefits from distributed parallelization through MPI. Multithreading and GPU parallelizations are not yet finished and important for large 3D runs. Work is also being done to design a multi-formalisms patch hierarchy where the coarsest levels would evolve a fluid hybrid system coupled to dynamically created hybrid kinetic refined levels.

Acknowledgements

The authors acknowledge support from the plaspar labex and federation, from the Laboratory of Plasma Physics, from the Centre National d'Etudes Spatiales (CNES) and Programme National Soleil-Terre (PNST). We also acknowledge the SAMRAI developer team for their help.

References

- [1] J. Dargent, N. Aunai, B. Lavraud, S. Toledo-Redondo, F. Califano, Simulation of Plasmaspheric Plume Impact on Dayside Magnetic Reconnection, *Geophysical Research Letters* 47 (4) (2020) e86546. doi:10.1029/2019g1086546. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2020GeoRL...4786546D&link_type=EJOURNAL
- [2] W. Daughton, V. Roytershteyn, H. Karimabadi, L. Yin, B. Albright, B. Bergen, K. Bowers, Role of electron physics in the development of turbulent magnetic reconnection in collisionless plasmas, *Nature Physics* 7 (7) (2011) 539 – 542. doi:10.1038/nphys1965. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2011NatPh...7...539D&link_type=ABSTRACT
- [3] D. S. Filippchev, Hybrid Simulation of Space Plasmas: Models with Massless Fluid Representation of Electrons. V. Reconnection of Magnetic Field Lines, *Computational Mathematics and Modeling* 13 (3) (2002) 215–248. doi:10.1023/a:1016003831330.
- [4] A. S. Lipatov, The Hybrid Multiscale Simulation Technology, An Introduction with Application to Astrophysical and Laboratory Plasmas, *Scientific Computation*, 2002. doi:10.1007/978-3-662-05012-5.
- [5] L. Leclercq, R. Modolo, F. Leblanc, S. Hess, M. Mancini, 3D magnetospheric parallel hybrid multi-grid method applied to planet-plasma interactions, *Journal of Computational Physics* 309 (2016) 295 – 313. doi:10.1016/j.jcp.2016.01.005. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2016JCoPh.309..295L&link_type=EJOURNAL
- [6] S. Hoilijoki, U. Ganse, Y. Pfau-Kempf, P. A. Cassak, B. M. Walsh, H. Hietala, S. v. Alfthan, M. Palmroth, Reconnection rates and X line motion at the magnetopause: Global 2D-3V hybrid-Vlasov simulation results, *Journal of Geophysical Research (Space Physics)* 122 (3) (2017) 2877 – 2888. doi:10.1002/2016ja023709. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2017JGRA...122.2877H&link_type=EJOURNAL
- [7] J. Guo, S. Lu, Q. Lu, Y. Lin, X. Wang, Q. Zhang, Z. Xing, K. Huang, R. Wang, S. Wang, Three-Dimensional Global Hybrid Simulations of High Latitude Magnetopause Reconnection and Flux Ropes During the Northward IMF, *Geophysical Research Letters* 48 (21) (2021). doi:10.1029/2021g1095003.
- [8] Z. Guo, Y. Lin, X. Wang, S. K. Vines, S. H. Lee, Y. Chen, Magnetopause Reconnection as Influenced by the Dipole Tilt Under Southward IMF Conditions: Hybrid Simulation and MMS Observation, *Journal of Geophysical Research: Space Physics* 125 (9) (2020). doi:10.1029/2020ja027795.
- [9] N. Aunai, M. Hesse, C. Black, R. Evans, M. Kuznetsova, Influence of the dissipation mechanism on collisionless magnetic reconnection in symmetric and asymmetric current layers, *Phys Plasmas* 20 (4) (2013) 042901, accepted in *Physics of Plasmas*, 12 pages. doi:10.1063/1.4795727. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2013PhPl...20d2901A&link_type=EJOURNAL
- [10] M. Palmroth, U. Ganse, Y. Pfau-Kempf, M. Battarbee, L. Turc, T. Brito, M. Grandin, S. Hoilijoki, A. Sandroos, S. v. Alfthan, Vlasov methods in space physics and astrophysics, *Living Reviews in Computational Astrophysics* (2018) 1 – 54doi:10.1007/s41115-018-0003-2. URL <https://doi.org/10.1007/s41115-018-0003-2>
- [11] G. Lapenta, Particle simulations of space weather, *Journal of Computational Physics* 231 (3) (2012) 795 – 821. doi:10.1016/j.jcp.2011.03.035. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2012JCoPh.231..795L&link_type=EJOURNAL
- [12] H. Karimabadi, H. Vu, D. Krauss-Varban, Y. Omelchenko, Global Hybrid Simulations of the Earth's Magnetosphere, *Numerical Modeling of Space Plasma Flows: Astronom-2006 ASP Conference Series*

- 359 (2006) 257.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006ASPC..359..257K&link_type=EJOURNAL
- [13] Y. A. Omelchenko, H. Karimabadi, Event-driven, hybrid particle-in-cell simulation: A new paradigm for multi-scale plasma modeling, *Journal of Computational Physics* 216 (1) (2006) 153 – 178. doi:10.1016/j.jcp.2005.11.029.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006JCoPh.216..153O&link_type=EJOURNAL
- [14] H. Karimabadi, A new methodology for multi-scale simulation of plasmas, advanced methods for space simulations (2007) 1 – 9.
- [15] Y. A. Omelchenko, H. Karimabadi, HYPERS: A unidimensional asynchronous framework for multiscale hybrid simulations, *Journal of Computational Physics* 231 (4) (2012) 1766 – 1780. doi:10.1016/j.jcp.2011.11.004.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2012JCoPh.231.1766O&link_type=EJOURNAL
- [16] J. L. Vay, P. Colella, P. McCorquodale, B. v. Straalen, A. Friedman, D. P. Grote, Mesh refinement for particle-in-cell plasma simulations: Applications to and benefits for heavy ion fusion, *Laser and Particle Beams* 20 (4) (2002) 569 – 575. doi:10.1017/s0263034602204139.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2002LPB....20..569V&link_type=EJOURNAL
- [17] J. L. Vay, P. Colella, A. Friedman, D. P. Grote, P. McCorquodale, D. B. Serafini, Implementations of mesh refinement schemes for Particle-In-Cell plasma simulations, *Computer Physics Communications* 164 (1) (2004) 297 – 305. doi:10.1016/j.cpc.2004.06.075.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2004CoPhC.164..297V&link_type=EJOURNAL
- [18] K. Fujimoto, S. Machida, Electromagnetic full particle code with adaptive mesh refinement technique: Application to the current sheet evolution, *Journal of Computational Physics* 214 (2) (2006) 550 – 566. doi:10.1016/j.jcp.2005.10.003.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2006JCoPh.214..550F&link_type=EJOURNAL
- [19] K. Fujimoto, R. D. Sydora, Electromagnetic particle-in-cell simulations on magnetic reconnection with adaptive mesh refinement, *Computer Physics Communications* 178 (12) (2008) 915 – 923. doi:10.1016/j.cpc.2008.02.010.
URL <http://linkinghub.elsevier.com/retrieve/pii/S0010465508000908>
- [20] P. Colella, P. C. Norgaard, Controlling self-force errors at refinement boundaries for AMR-PIC, *Journal of Computational Physics* 229 (4) (2010) 947 – 957. doi:10.1016/j.jcp.2009.07.004.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2010JCoPh.229..947C&link_type=EJOURNAL
- [21] J. Müller, S. Simon, U. Motschmann, J. Schüle, K.-H. Glassmeier, G. J. Pringle, A.I.K.E.F.: Adaptive hybrid model for space plasma simulations, *Computer Physics Communications* 182 (4) (2011) 946 – 966. doi:10.1016/j.cpc.2010.12.033.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2011CoPhC.182..946M&link_type=EJOURNAL
- [22] K. Fujimoto, A new electromagnetic particle-in-cell model with adaptive mesh refinement for high-performance parallel computation, *Journal of Computational Physics* 230 (2) (2011) 8508 – 8526. doi:10.1016/j.jcp.2011.08.002.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2011JCoPh.230.8508F&link_type=EJOURNAL
- [23] K. Fujimoto, Multi-Scale Kinetic Simulation of Magnetic Reconnection With Dynamically Adaptive Meshes, *Frontiers in Physics* 6 (2018) 303 – 8. doi:10.3389/fphy.2018.00119.
URL <https://www.frontiersin.org/article/10.3389/fphy.2018.00119/full>
- [24] G. Lapenta, Automatic adaptive multi-dimensional Particle In Cell, advanced methods for space simulations (2007) 1 – 16.
- [25] G. Lapenta, DEMOCRITUS: An adaptive particle in cell (PIC) code for object-plasma interactions, *Journal of Computational Physics* 230 (1) (2011) 4679 – 4695. doi:10.1016/j.jcp.2011.02.041.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2011JCoPh.230.4679L&link_type=EJOURNAL
- [26] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, H. Tufo, FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes, *The Astrophysical Journal Supplement Series* 131 (1) (2008) 273. doi:10.1086/317361.
- [27] A. Mignone, C. Zanni, P. Tzeferacos, B. v. Straalen, P. Colella, G. Bodo, The PLUTO Code for Adaptive Mesh Computations in Astrophysical Fluid Dynamics, *The Astrophysical Journal Supplement* 198 (1) (2012) 7. doi:10.1088/0067-0049/198/1/7.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2012ApJS..198....7M&link_type=EJOURNAL
- [28] G. Tóth, Y. Ma, T. I. Gombosi, Hall magnetohydrodynamics on block-adaptive grids, *Journal of Computational Physics* 227 (1) (2008) 6967 – 6984. doi:10.1016/j.jcp.2008.04.010.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2008JCoPh.227.6967T&link_type=EJOURNAL
- [29] B. V. d. Holst, R. Keppens, Hybrid block-AMR in cartesian and curvilinear coordinates: MHD applications, *Journal of Computational Physics* 226 (1) (2007) 925 – 946. doi:10.1016/j.jcp.2007.05.007.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2007JCoPh.226..925V&link_type=EJOURNAL
- [30] M. Feyerabend, S. Simon, U. Motschmann, L. Liuzzo, Filamented ion tail structures at Titan: A hybrid simulation study, *Planetary and Space Science* 117 (C) (2015) 362 – 376. doi:10.1016/j.pss.2015.07.008.
URL <http://dx.doi.org/10.1016/j.pss.2015.07.008>
- [31] Y. Vernisse, J. A. Rioussel, U. Motschmann, K. Glassmeier, Stellar winds and planetary bodies simulations: Magnetized obstacles in super-Alfvénic and sub-Alfvénic flows, *Planetary and Space Science* 137 (C)

- (2017) 40 – 51. doi:10.1016/j.pss.2016.08.012.
URL <http://dx.doi.org/10.1016/j.pss.2016.08.012>
- [32] W. Exner, D. Heyner, L. Liuzzo, U. Motschmann, D. Shiota, K. Kusano, T. Shibayama, Coronal mass ejection hits mercury: A.I.K.E.F. hybrid-code results compared to MESSENGER data, *Planetary and Space Science* 153 (2018) 89 – 99. doi:10.1016/j.pss.2017.12.016.
URL <https://doi.org/10.1016/j.pss.2017.12.016>
- [33] H. Arnold, L. Liuzzo, S. Simon, Plasma Interaction Signatures of Plumes at Europa, *Journal of Geophysical Research (Space Physics)* 125 (1) (2020) e27346. doi:10.1029/2019ja027346.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2020JGRA..12527346A&link_type=EJOURNAL
- [34] G. Lapenta, Particle Rezoning for Multidimensional Kinetic Particle-In-Cell Simulations, *Journal of Computational Physics* 181 (1) (2002) 317 – 337. doi:10.1006/jcph.2002.7126.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2002JCoPh.181..317L&link_type=EJOURNAL
- [35] M. E. Innocenti, G. Lapenta, S. Markidis, A. Beck, A. Vapirev, A Multi Level Multi Domain Method for Particle In Cell plasma simulations, *Journal of Computational Physics* 238 (3) (2013) 115 – 140. doi:10.1016/j.jcp.2012.12.028.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2013JCoPh.238..115I&link_type=EJOURNAL
- [36] R. Smets, N. Aunai, A. Ciardi, M. Drouin, M. Campos-Pinto, P. Deegan, A new method to dispatch split particles in Particle-In-Cell codes, *Computer Physics Communications* 261 (10) (2021) 107666. doi:10.1016/j.cpc.2020.107666.
URL <https://linkinghub.elsevier.com/retrieve/pii/S0010465520303246>
- [37] A. Beck, M. E. Innocenti, G. Lapenta, S. Markidis, Multi-level multi-domain algorithm implementation for two-dimensional multiscale particle in cell simulations, *Journal of Computational Physics* 271 (2014) 430 – 443. doi:10.1016/j.jcp.2013.12.016.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2014JCoPh.271..430B&link_type=EJOURNAL
- [38] M. E. Innocenti, C. Tronci, S. Markidis, G. Lapenta, Grid coupling mechanism in the semi-implicit adaptive Multi-Level Multi-Domain method, *Journal of Physics: Conference Series* 719 (1) (2016) 012019. doi:10.1088/1742-6596/719/1/012019.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2016JPhCS.719a2019I&link_type=EJOURNAL
- [39] M. E. Innocenti, A. Beck, S. Markidis, G. Lapenta, Momentum conservation in Multi-Level Multi-Domain (MLMD) simulations, *Journal of Computational Physics* 312 (2016) 14 – 18. doi:10.1016/j.jcp.2016.02.026.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2016JCoPh.312...14I&link_type=EJOURNAL
- [40] M. E. Innocenti, A. Beck, T. Ponweiser, S. Markidis, G. Lapenta, Introduction of temporal sub-stepping in the Multi-Level Multi-Domain semi-implicit Particle-In-Cell code Parsek2D-MLMD, *Computer Physics Communications* 189 (2015) 47 – 59. doi:10.1016/j.cpc.2014.12.004.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2015CoPhC.189...47I&link_type=EJOURNAL
- [41] T. Sugiyama, K. Kusano, Multi-scale plasma simulation by the interlocking of magnetohydrodynamic model and particle-in-cell kinetic model, *Journal of Computational Physics* 227 (2) (2007) 1340 – 1352. doi:10.1016/j.jcp.2007.09.011.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2007JCoPh.227.1340S&link_type=EJOURNAL
- [42] S. Ishiguro, S. Usami, R. Horiuchi, H. Ohtani, A. Maluckov, M. M. Škorić, Multi-scale simulation for plasma science, *Journal of Physics: Conference Series* 257 (1) (2010) 012026 – 9. doi:10.1088/1742-6596/257/1/012026.
URL <https://iopscience.iop.org/article/10.1088/1742-6596/257/1/012026>
- [43] S. Usami, R. Horiuchi, H. Ohtani, M. den, Development of multi-hierarchy simulation model with non-uniform space grids for collisionless driven reconnection, *Physics of Plasmas* 20 (6) (2013) 061208 – 9. doi:10.1063/1.4811121.
URL <http://aip.scitation.org/doi/10.1063/1.4811121>
- [44] L. K. S. Daldorff, G. Tóth, T. I. Gombosi, G. Lapenta, J. Amaya, S. Markidis, J. U. Brackbill, Two-way coupling of a global Hall magnetohydrodynamics model with a local implicit particle-in-cell model, *Journal of Computational Physics* 268 (2014) 236 – 254. doi:10.1016/j.jcp.2014.03.009.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2014JCoPh.268..236D&link_type=EJOURNAL
- [45] G. Tóth, X. Jia, S. Markidis, I. B. Peng, Y. Chen, L. K. S. Daldorff, V. M. Tenishev, D. Borovikov, J. D. Haiducek, T. I. Gombosi, A. Glocer, J. C. Dorelli, Extended magnetohydrodynamics with embedded particle-in-cell simulation of Ganymede's magnetosphere, *Journal of Geophysical Research (Space Physics)* 121 (2) (2016) 1273 – 1293. doi:10.1002/2015ja021997.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2016JGRA..121.1273T&link_type=EJOURNAL
- [46] Y. Chen, G. Tóth, P. Cassak, X. Jia, T. I. Gombosi, J. A. Slavin, S. Markidis, I. B. Peng, V. K. Jordanova, M. G. Henderson, Global Three-Dimensional Simulation of Earth's Dayside Reconnection Using a Two-Way Coupled Magnetohydrodynamics With Embedded Particle-in-Cell Model: Initial Results, *Journal of Geophysical Research (Space Physics)* 122 (1) (2017) 10. doi:10.1002/2017ja024186.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2017JGRA..12210318C&link_type=EJOURNAL
- [47] H. Zhou, G. Tóth, X. Jia, Y. Chen, S. Markidis, Embedded Kinetic Simulation of Ganymede's Magnetosphere: Improvements and Inferences, *Journal of Geophysical Research (Space Physics)* 124 (7) (2019) 5441 – 5460. doi:10.1029/2019ja026643.
URL http://adsabs.harvard.edu/cgi-bin/nph-data_

- query?bibcode=2019JGRA..124.5441Z&link_type=EJOURNAL
- [48] H. Zhou, G. Tóth, X. Jia, Y. Chen, Reconnection-Driven Dynamics at Ganymede's Upstream Magnetosphere: 3-D Global Hall MHD and MHD-EPIC Simulations, *Journal of Geophysical Research (Space Physics)* 125 (8) (2020) 1 – 20. doi:10.1029/2020ja028162. URL <https://onlinelibrary.wiley.com/doi/10.1029/2020JA028162>
- [49] Y. Ma, C. T. Russell, G. Tóth, Y. Chen, A. F. Nagy, Y. Harada, J. McFadden, J. S. Halekas, R. Lillis, J. E. P. Connerney, J. Espley, G. A. DiBraccio, S. Markidis, I. B. Peng, X. Fang, B. M. Jakosky, Reconnection in the Martian Magnetotail: Hall-MHD With Embedded Particle-in-Cell Simulations, *Journal of Geophysical Research (Space Physics)* 123 (5) (2018) 3742 – 3763. doi:10.1029/2017ja024729. URL <http://doi.wiley.com/10.1029/2017JA024729>
- [50] Y. Chen, G. Tóth, X. Jia, J. A. Slavin, W. Sun, S. Markidis, T. I. Gombosi, J. M. Raines, Studying Dawn-Dusk Asymmetries of Mercury's Magnetotail Using MHD-EPIC Simulations, *Journal of Geophysical Research (Space Physics)* 124 (1) (2019) 8954 – 8973. doi:10.1029/2019ja026840. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2019JGRA..124.8954C&link_type=EJOURNAL
- [51] K. D. Makwana, R. Keppens, G. Lapenta, Two-way coupling of magnetohydrodynamic simulations with embedded particle-in-cell simulations, *Computer Physics Communications* 221 (2017) 81 – 94. doi:10.1016/j.cpc.2017.08.003. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=2017CoPhC.221...81M&link_type=EJOURNAL
- [52] K. D. Makwana, R. Keppens, G. Lapenta, Study of magnetic reconnection in large-scale magnetic island coalescence via spatially coupled MHD and PIC simulations, *Physics of Plasmas* 25 (8) (2018) 082904 – 12. doi:10.1063/1.5037774. URL <http://aip.scitation.org/doi/10.1063/1.5037774>
- [53] R. Courant, K. Friedrichs, H. Lewy, On the Partial Difference Equations of Mathematical Physics, *IBM Journal of Research and Development* 11 (1967) 215–234. doi:10.1147/rd.112.0215.
- [54] K. Yee, Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media, *IEEE Transactions on Antennas and Propagation* 14 (3) (1966) 302–307. doi:10.1109/TAP.1966.1138693.
- [55] M. W. Kunz, J. M. Stone, X.-N. Bai, Pegasus: A new hybrid-kinetic particle-in-cell code for astrophysical plasma dynamics, *Journal of Computational Physics* 259 (2014) 154–174. arXiv:1311.4865, doi:10.1016/j.jcp.2013.11.035.
- [56] J. P. Boris, Fourth Conference on Numerical Simulation of Plasmas, 1970. URL [http://scholar.google.com/javascript:void\(0\)](http://scholar.google.com/javascript:void(0))
- [57] R. Löhner, An adaptive finite element scheme for transient problems in CFD, *Computer Methods in Applied Mechanics and Engineering* 61 (3) (1987) 323–338. doi:10.1016/0045-7825(87)90098-3.
- [58] B. T. N. Gunney, R. W. Anderson, Advances in patch-based adaptive mesh refinement scalability, *Journal of Parallel and Distributed Computing* 89 (2016) 65–84. doi:10.1016/j.jpdc.2015.11.005.
- [59] S. P. Gary, *Theory of Space Plasma Microinstabilities*, 1993.
- [60] S. P. Gary, Electromagnetic ion beam instabilities - Hot beams at interplanetary shocks, *Astrophysical Journal* 288 (1985) 342–352. doi:10.1086/162797.
- [61] N. Aunai, M. Hesse, C. Black, R. Evans, M. Kuznetsova, Influence of the dissipation mechanism on collisionless magnetic reconnection in symmetric and asymmetric current layers, *Physics of Plasmas* 20 (4) (2013) 042901. arXiv:1303.0442, doi:10.1063/1.4795727.
- [62] M. Shay, J. Drake, R. Denton, D. Biskamp, Structure of the dissipation region during collisionless magnetic reconnection, *Journal of Geophysical Research* 103 (A) (1998) 9165 – 9176. doi:10.1029/97ja03528. URL http://adsabs.harvard.edu/cgi-bin/nph-data_query?bibcode=1998JGR...103.9165S&link_type=ABSTRACT
- [63] S. Wijesinghe, R. Hornung, A. Garcia, N. Hadji-constantinou, Three Dimensional Hybrid Continuum-Atomistic Simulations for Multiscale Hydrodynamics, in: *APS Division of Fluid Dynamics Meeting Abstracts*, Vol. 55 of *APS Meeting Abstracts*, 2002, p. KA.002.