



A robust and efficient solver based on kinetic schemes for Magnetohydrodynamics (MHD) equations.

Hubert Baty^a, Florence Drui^b, Philippe Helluy^{c,*}, Emmanuel Franck^c,
Christian Klingenberg^d, Lukas Thanhäuser^d

^a Observatoire de Strasbourg, France

^b CEA Saclay, France

^c IRMA, Université de Strasbourg, Inria Tonus, France

^d University of Würzburg, Germany

ARTICLE INFO

Article history:

Received 16 May 2022

Received in revised form 13 October 2022

Accepted 20 October 2022

ABSTRACT

This paper is devoted to the simulation of compressible magnetohydrodynamic (MHD) flows with the Lattice Boltzmann Method (LBM). The usual LBM is limited to low-Mach flows. We propose a robust and accurate numerical method based on the vectorial kinetic construction of [5,25], which allows us to extend the LBM to arbitrary Mach flows. We also explain how to adjust the numerical viscosity in order to obtain stable and accurate results in smooth or discontinuous parts of the flow and reduced divergence errors. The method can handle shock waves and can be made second order in smooth regions. It is also very well adapted to computing with Graphics Processing Unit (GPU). Our GPU implementation in 2D achieves state-of-the-art accuracy, with near-optimal performance. We finally present numerical computations of a tilt instability that demonstrate the capability of the method to handle physically relevant simulations.

© 20XX

1. Introduction

The MagnetoHydroDynamics (MHD) system is a fundamental model used in many fields of physics: astrophysics, plasma physics, geophysics... Indeed, the MHD model is commonly adopted as an excellent framework for collisional plasma environments. The numerical approximation of this system is a difficult task. Compared to other models in fluid mechanics, it contains more conservative unknowns and more physical scales. It is also subject to complex phenomena such as occurrence of shock waves, current sheet formation, magnetic reconnection, instabilities and turbulent behaviors.

Many numerical methods have been developed for solving MHD. The MHD equations enter the category of nonlinear hyperbolic systems. The finite volume method with various upwind numerical fluxes based on approximate Riemann solvers are natural choices, because those methods are known to handle in robust way shock waves arising in nonlinear hyperbolic systems. One of the first works based on upwind schemes for solving the MHD equations is [9]. Other notable works extending the initial idea of Godunov to MHD are, among many others, [7,10,19,40,45]. An additional specificity is that the magnetic field has to satisfy a divergence-free condition, which is generally difficult to be verified by numerical solutions. When this condition is not respected numerically, this generally leads to nonphysical solutions and instabilities that may lead to the crash of the computations. See, for instance, [3,20,44]. In order to

* Corresponding author.

E-mail address: helluy@unistra.fr (P. Helluy).

deal with the divergence-free condition, we adopt here a modified version of the MHD equations by a divergence cleaning term proposed in [20].

We propose a simple scheme, based on an abstract kinetic interpretation, for computing two-dimensional MHD solutions. The kinetic interpretation is vectorial and has been first proposed by Bouchut in [5] and Aregba-Natalini in [1]. In this approach, the original system of conservation laws is represented by a system of transport equations coupled through a so-called collision source term. The idea originates from the Boltzmann kinetic theory of gases, but it is purely abstract and has no physical meaning. It leads to natural numerical methods, where the transport step and the collision step are made separately. Many interesting features arise from this representation. We can mention the possibility to design explicit schemes with very large time steps, simplified parallel implementations [2,14]. It can also be used for establishing simplified stability proofs [5]. Finally, we would like to mention that the vectorial kinetic representation permits to construct schemes for compressible flows, while the usual scalar Lattice Boltzmann approach only handles low-Mach flows.

In this paper, we solve the kinetic representation with a Lattice-Boltzmann Method (LBM), where the transport step is solved exactly on a regular grid. The same approach has already been used in [25] for solving compressible Euler equations. An important aspect of the LBM is the choice of the relaxation parameter in the collision step. The choice of the parameter allows adjusting the numerical viscosity of the LBM scheme. We provide an analysis in a simplified one-dimensional framework which shows that it is possible to adjust more precisely the numerical viscosity with a generalized matrix relaxation parameter. We also show that the divergence cleaning effect is improved if the relaxation parameter is chosen differently for the physical variables and the divergence cleaning potential.

In order to capture fine structures, it is necessary to consider very fine meshes. We have programmed the algorithm in a very efficient way in order to address recent GPUs (Graphic Processing Units) or multicore CPUs (Central Processing Units). GPU hardware is notably known to be especially efficient for LBM computations. This is due to the simplicity of the computations and the regularity of memory access [27,37,54].

We describe a very simple but very efficient implementation, which relies on the OpenCL [30] and PyOpenCL [34] libraries, and the memory optimizations used for reaching high performance. The program allows performing full MHD simulations on grids as fine as 7000×7000 within a few hours.

Finally, we apply the whole approach to several MHD simulations: the classical Orszag-Tang test and a more physical study of ideal MHD instabilities (tilt modes) and associated formation of quasi-singular current sheets.

2. Mathematical model

2.1. MHD equations with divergence cleaning

The MagnetoHydroDynamic (MHD) system is a model used in many fields of physics. It consists of an extension of the compressible Euler equations for taking into account magnetic effects. It reads:

$$\partial_t \begin{pmatrix} \rho \\ \rho \mathbf{u} \\ Q \\ \mathbf{B} \end{pmatrix} + \nabla \cdot \begin{pmatrix} \rho \mathbf{u} \otimes \mathbf{u} + \left(p + \frac{\mathbf{B} \cdot \mathbf{B}}{2} \right) \mathbf{I} - \mathbf{B} \otimes \mathbf{B} \\ \left(Q + p + \frac{\mathbf{B} \cdot \mathbf{B}}{2} \right) \mathbf{u} - (\mathbf{B} \cdot \mathbf{u}) \mathbf{B} \\ \mathbf{u} \otimes \mathbf{B} - \mathbf{B} \otimes \mathbf{u} \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{0} \\ 0 \\ 0 \end{pmatrix},$$

where the notations are defined as follows. The velocity and magnetic field are denoted

$$\mathbf{u} = (u_1, u_2, u_3)^T, \quad \mathbf{B} = (B_1, B_2, B_3)^T,$$

the pressure is given by a perfect-gas law with a constant polytropic exponent $\gamma > 1$

$$p = (\gamma - 1) \left(Q - \rho \frac{\mathbf{u} \cdot \mathbf{u}}{2} - \frac{\mathbf{B} \cdot \mathbf{B}}{2} \right).$$

The other variables are the density ρ , the total energy Q .

A difficulty is that the magnetic field has to satisfy an additional divergence-free condition

$$\nabla \cdot \mathbf{B} = 0. \tag{1}$$

Standard finite volume methods do not guarantee that the numerical magnetic field is divergence-free. More annoying: the divergence errors generally grow with the simulation time, which leads to physically wrong results. For limiting the divergence errors, we adopt the divergence cleaning method described in [20].

The MHD equations with Divergence Cleaning [20] (called the MHD-DC equations in the following) read:

$$\partial_t \begin{pmatrix} \rho \\ \rho \mathbf{u} \\ Q \\ \mathbf{B} \\ \psi \end{pmatrix} + \nabla \cdot \begin{pmatrix} \rho \mathbf{u} \otimes \mathbf{u} + \left(p + \frac{\mathbf{B} \cdot \mathbf{B}}{2}\right) \mathbf{I} - \mathbf{B} \otimes \mathbf{B} \\ \left(Q + p + \frac{\mathbf{B} \cdot \mathbf{B}}{2}\right) \mathbf{u} - (\mathbf{B} \cdot \mathbf{u}) \mathbf{B} \\ \mathbf{u} \otimes \mathbf{B} - \mathbf{B} \otimes \mathbf{u} + \psi \mathbf{I} \\ c_h^2 \mathbf{B} \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{0} \\ 0 \\ \mathbf{0} \\ 0 \end{pmatrix}.$$

The MHD-DC system contains an additional unknown, the divergence cleaning potential ψ . The divergence cleaning velocity is a positive parameter $c_h > 0$.

When the magnetic field satisfies the divergence-free condition (1) and the potential ψ is a constant, then the MHD-DC equations simply reduce to the usual MHD equations. In other words, the MHD-DC equations are a generalization of the MHD equations, where the magnetic field can have a non-vanishing divergence.

It is possible to show that the MHD-DC system, as the MHD system, is hyperbolic [20]. However this system is not strictly hyperbolic, which leads to some difficulties, such as non-uniqueness of the Riemann problems in some situations [51].

The interest of the MHD-DC formulation is for numerical approximations. Indeed, standard approximations of the usual MHD equations suffer from drifting errors along time of the divergence constraint. This has been observed by several authors [3,20,44]. A review on this topic is developed for instance in [52]. Approximations of the MHD-DC generally have a much better behavior. In this generalized model, the divergence errors propagate at the wave speed c_h . The errors are then damped at the boundaries of the computational domain or by the numerical diffusion in the domain. In [20] several divergence corrections are proposed. Another term can for instance be considered: a damping relaxation term of the divergence cleaning potential ψ . With this choice, the MHD system is transformed into a parabolic system of conservation laws, which does not fit into the kinetic framework. Therefore we choose to consider only the correction that leads to a conservative and hyperbolic model, in order to be able to apply the general theory of vectorial kinetic representations given in [1,5]. Theoretically, the parameter c_h can take any value. But in practice it is generally chosen larger than all the wave speeds of the MHD system (see [20]). Let us introduce the conservative variables

$$\mathbf{w} = \mathbf{w}(\mathbf{x}, t) = \begin{pmatrix} \rho \\ \rho \mathbf{u} \\ Q \\ \mathbf{B} \\ \psi \end{pmatrix} \in \mathbb{R}^m, \quad m = 9.$$

and the flux (\mathbf{n} is a vector of \mathbb{R}^3)

$$\mathbf{F}(\mathbf{w}, \mathbf{n}) = \begin{pmatrix} \rho \mathbf{u} \cdot \mathbf{n} + \left(p + \frac{\mathbf{B} \cdot \mathbf{B}}{2}\right) \mathbf{n} - \mathbf{B} \cdot \mathbf{n} \mathbf{B} \\ \left(Q + p + \frac{\mathbf{B} \cdot \mathbf{B}}{2}\right) \mathbf{u} \cdot \mathbf{n} - (\mathbf{B} \cdot \mathbf{u}) \mathbf{B} \cdot \mathbf{n} \\ \mathbf{u} \cdot \mathbf{n} \mathbf{B} - \mathbf{B} \cdot \mathbf{n} \mathbf{u} + \psi \mathbf{n} \\ c_h^2 \mathbf{B} \cdot \mathbf{n} \end{pmatrix} \in \mathbb{R}^m.$$

With this definition, we can make more precise the choice of the divergence cleaning parameter c_h . We denote by $\lambda_i, i = 1 \dots m$ the real eigenvalues of the jacobian of the flux

$$D_{\mathbf{w}} \mathbf{F}(\mathbf{w}, \mathbf{n}).$$

It can be shown that c_h is one of the eigenvalues [20]. Let us decide that $\lambda_m = c_h$. Then we choose c_h to be the largest eigenvalue:

$$c_h > \max_{1 \leq i \leq m-1} |\lambda_i|. \tag{2}$$

In this work, we assume that all the fields do not depend on the x_3 space variable. We are thus computing two-dimensional solutions. If we set

$$\mathbf{n}_1 = (1, 0, 0)^T, \quad \mathbf{n}_2 = (0, 1, 0)^T, \quad \mathbf{F}_1 = \mathbf{F}(\mathbf{w}, \mathbf{n}_1), \quad \mathbf{F}_2 = \mathbf{F}(\mathbf{w}, \mathbf{n}_2),$$

the MHD equations can also be written as a two-dimensional system of nine conservation laws

$$\partial_t \mathbf{w} + \partial_1 \mathbf{F}_1 + \partial_2 \mathbf{F}_2 = 0, \tag{3}$$

where we use the notation ∂_i for the partial derivative $\partial/\partial x_i$. If the fields depend only on the x_1 space variable, the system reduces to

$$\partial_t \mathbf{w} + \partial_1 \mathbf{F}_1 = 0. \tag{4}$$

In this case, the mathematical analysis is simplified. We shall perform an analysis of the numerical viscosity of the kinetic method in this simplified framework.

2.2. Kinetic representation

The Lattice-Boltzmann Method (LBM) originated from the physical kinetic interpretation of the Navier-Stokes equations [12]. In the physical world, the fluid particles can have arbitrary velocities. The main idea of the LBM is that it is possible to construct abstract kinetic interpretations of the Navier-Stokes equations in which the particles velocities can take a few number of given values. This makes it possible to solve the kinetic model directly in an efficient way. We refer to [50] for a history of the LBM.

Initially devised for solving incompressible Navier-Stokes equations, the LBM is based on a single scalar particle distribution function. More recently, it has been extended to other systems of conservation laws. Extensions to MHD can be found in [21,39]. Dellar, in [21] showed that it is not possible to rely on a single scalar kinetic function for approximating the MHD equations. He proposes to represent the magnetic part of the equations with a vectorial kinetic formulation. Another vectorial kinetic interpretation of the MHD equations was also previously proposed in [17] for building a numerical flux for a Finite Volume method.

The resulting hybrid kinetic model (scalar for the fluid part, vectorial for the magnetic part) is, however, limited to low-Mach number flows.

Several attempts have been done to extend the LBM to arbitrary Mach flows with a scalar kinetic function. For instance, in [28,29,36] the Maxwell distribution is approximated with more velocities in the lattice. The additional degrees of freedom allow to computing a discrete Maxwellian distribution that matches exactly the moments of the compressible Euler equations. These approaches require quite large lattices. For instance, in [36] the authors propose a D3Q39 LBM with 39 kinetic velocities for approximating the five-equation, three-dimensional, compressible Euler equations. In addition, the methods proposed in [28,29,36] become unstable when the temperature exits from a given range. Another method is presented in [46] for improving the stability of the scalar LBM for compressible flow. It relies on the classical D2Q9 or D3Q27 lattices. However, the LBM is supplemented with a finite-difference approximation of the temperature equation. This approach is thus not fully an LBM method.

For addressing transonic and supersonic flows, we think that a fully vectorial kinetic model is preferable. The general theory of lattice vectorial kinetic model is discussed in [1,5,25,31]. It is valid for any hyperbolic system of conservation laws and is no more limited to low-Mach number flows. It relies on a relatively small set of kinetic data. For instance, the LBM proposed by Dubois in [25] requires only 30 transported kinetic quantities compared to the 39 kinetic velocities of [36]. In this work Dubois observes that while very robust and promising, his method is quite numerically diffusive. For obtaining accurate results, it is therefore necessary to consider refined meshes.

Let us finally mention that the vectorial kinetic approach is also of interest for low Mach numbers flows and the incompressible limit [6,55].

The vectorial kinetic approach can be used on arbitrary unstructured meshes at any order of approximation [2,13]. When the lattice velocities are aligned with the mesh, it is possible to adopt the simple exact transport step of the LBM. This is the method that we present here.

We consider a real number $\lambda > 0$ and four velocities $\mathbf{v}_k = (v_k^1, v_k^2)^T$, $k = 1 \dots 4$, defined by

$$\mathbf{v}_1 = \lambda \begin{pmatrix} -1 \\ 0 \end{pmatrix}, \quad \mathbf{v}_2 = \lambda \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \mathbf{v}_3 = \lambda \begin{pmatrix} 0 \\ -1 \end{pmatrix}, \quad \mathbf{v}_4 = \lambda \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

We consider vectorial distribution functions $\mathbf{f}_k(\mathbf{x}, t) \in \mathbb{R}^m$, $k = 1 \dots 4$. The conservative variable \mathbf{w} is related to the kinetic data by

$$\mathbf{w} = \sum_{k=1}^4 \mathbf{f}_k.$$

Usually, in the Lattice-Boltzmann philosophy, the kinetic system is a set a transport equations coupled through a relaxation source term and reads (with $\tau > 0$)

$$\partial_t \mathbf{f}_k + \mathbf{v}_k \cdot \nabla \mathbf{f}_k = \mu_k = \frac{1}{\tau} (\mathbf{f}_k^{eq} - \mathbf{f}_k), \quad k = 1 \dots 4. \quad (5)$$

The equilibrium function \mathbf{f}_k^{eq} also called the Maxwellian state is chosen as

$$\begin{aligned} \mathbf{f}_1^{eq}(\mathbf{w}) &= \frac{1}{4} \mathbf{w} - \frac{1}{2\lambda} \mathbf{F}_1(\mathbf{w}), & \mathbf{f}_2^{eq}(\mathbf{w}) &= \frac{1}{4} \mathbf{w} + \frac{1}{2\lambda} \mathbf{F}_1(\mathbf{w}), \\ \mathbf{f}_3^{eq}(\mathbf{w}) &= \frac{1}{4} \mathbf{w} - \frac{1}{2\lambda} \mathbf{F}_2(\mathbf{w}), & \mathbf{f}_4^{eq}(\mathbf{w}) &= \frac{1}{4} \mathbf{w} + \frac{1}{2\lambda} \mathbf{F}_2(\mathbf{w}). \end{aligned}$$

One can check that

$$\mathbf{w} = \sum_{k=1}^4 \mathbf{f}_k^{eq}, \quad \mathbf{F}_i = \sum_{k=1}^4 v_k^i \mathbf{f}_k^{eq}. \quad (6)$$

When the relaxation time $\tau \rightarrow 0$ then from (5) we see that

$$\mathbf{f}_k \simeq \mathbf{f}_k^{eq},$$

and thus from (6) we recover the MHD equations (3).

In order to approximate, (5) we first introduce a time step $\Delta t > 0$. Assume that we know the solution at time $t = (i-1)\Delta t$, where i is an integer. For computing the solution at time $t = i\Delta t$, one first solves the free transport equation

$$\partial_t \mathbf{f}_k + \mathbf{v}_k \cdot \nabla \mathbf{f}_k = 0, \quad (7)$$

for $t \in [(i-1)\Delta t, i\Delta t[$. This gives the solution at time $t = i\Delta t^-$. At this time $t = i\Delta t^-$, it is clear that the distribution functions \mathbf{f}_k have deviated from the equilibrium $\mathbf{f}_k \simeq \mathbf{f}_k^{eq}$, because of the free transport evolution. In order to recover the consistency of the approximation, a classical choice is to force a return (also called a *relaxation* or *projection*) to equilibrium at each time step:

$$\mathbf{f}_k^i = \mathbf{f}_k^{eq}.$$

But it appears that this choice is not the most accurate one. It is better to apply an *over-relaxation*. Let ω be a relaxation parameter $\in [1, 2]$. The approximate solution has a jump in time at times $t = i\Delta t$, and \mathbf{f} is given by

$$\mathbf{f}_k(\mathbf{x}, t^+) = \mathbf{f}_k(\mathbf{x}, t^-) + \omega (\mathbf{f}_k^{eq}(\mathbf{w}(\mathbf{x}, t^-)) - \mathbf{f}_k(\mathbf{x}, t^-)).$$

The computation of the jumps at times $t = i\Delta t$ are called the collision steps.

Let us remark that when $\omega = 1$ the collision step reads

$$\mathbf{f}_k(\mathbf{x}, t^+) = \mathbf{f}_k^{eq}(\mathbf{w}(\mathbf{x}, t^-)).$$

We recover the simple projection on the Maxwellian state associated to the conservative data. This scheme is a first order in time approximation of the original equations. An interesting case is the case of over-relaxation $\omega = 2$. Then the scheme is a second order in time approximation of the original equations. See for instance [22]). See also below for a mathematical analysis in the one-dimensional case. We shall see that the choice of the relaxation parameter ω amounts to introducing some numerical viscosity in the numerical scheme. With this interpretation, $\omega = 1$ corresponds to the maximal viscosity. The case $\omega = 2$ leads to a second order accurate scheme, with third order dispersive terms replacing the diffusive terms.

3. Numerical method

For solving (5) numerically, we first construct a structured grid of the square $\mathcal{D} =]0, L[\times]0, L[$. The space step is given by $\Delta x = \frac{L}{N}$ and the grid-points are then as follows

$$\mathbf{x}_{i,j} = \left(\begin{array}{c} \left(i + \frac{1}{2} \right) \Delta x \\ \left(j + \frac{1}{2} \right) \Delta x \end{array} \right), \quad i, j \in \{0, \dots, N-1\}.$$

For a simpler presentation we can assume periodic boundary condition, which amounts to the following equivalences

$$i + N \equiv i, \quad j + N \equiv j.$$

We denote by $\mathbf{w}_{i,j}^n$ and $\mathbf{f}_{k,i,j}^n$ the approximation of \mathbf{w} and \mathbf{f}_k at the grid points $\mathbf{x}_{i,j}$ and time $t_n = n\Delta t$ just after the collision step. The values of the kinetic vectors just before the collision step are denoted $\mathbf{f}_{k,i,j}^{n-}$. For solving the kinetic system (5) we treat the transport and the relaxation terms separately.

3.1. Transport solver

The transport Eq. (7) admits an exact solution

$$\mathbf{f}_k(\mathbf{x}, t + \Delta t) = \mathbf{f}_k(\mathbf{x} - \Delta t \mathbf{v}_k, t).$$

If we assume that the time step satisfies $\Delta t = \frac{\Delta x}{\lambda}$, the transport operator then reduces to a simple shift. Before the collision step, we have thus

$$\begin{cases} \mathbf{f}_{1,i,j}^{n+1,-} = \mathbf{f}_{1,i+1,j}^n, & \mathbf{f}_{2,i,j}^{n+1,-} = \mathbf{f}_{2,i-1,j}^n, \\ \mathbf{f}_{3,i,j}^{n+1,-} = \mathbf{f}_{3,i,j+1}^n, & \mathbf{f}_{4,i,j}^{n+1,-} = \mathbf{f}_{4,i,j-1}^n. \end{cases} \quad (8)$$

3.2. Relaxation

At the end of the transport step, we can compute the conservative data

$$\mathbf{w}_{i,j}^{n+1} = \sum_{k=1}^4 \mathbf{f}_{k,i,j}^{n+1,-}. \quad (9)$$

The usual projection method would be to write

$$\mathbf{f}_{k,i,j}^{n+1} = \mathbf{f}_k^{eq} \left(\mathbf{w}_{i,j}^{n+1} \right). \quad (10)$$

As stated above, it is more accurate to consider a second order relaxed scheme

$$\mathbf{f}_{k,i,j}^{n+1} = 2\mathbf{f}_k^{eq} \left(\mathbf{w}_{i,j}^{n+1} \right) - \mathbf{f}_{k,i,j}^{n+1,-},$$

but this second order scheme may suffer from numerical instabilities, particularly in shock waves. It is possible to add some dissipation. This dissipation provides a better stability in numerical cases when sharp fronts are generated.

For introducing the dissipative mechanism, we consider a relaxation step of the form

$$\mathbf{f}_{k,i,j}^{n+1} = \omega \mathbf{f}_k^{eq} \left(\mathbf{w}_{i,j}^{n+1} \right) - (\omega - 1) \mathbf{f}_{k,i,j}^{n+1,-}, \quad (11)$$

where $\omega \in [1, 2]$. We see that the whole algorithm is extremely simple. It is a succession of “shifts” (8) and “collisions” (11).

For $\omega = 2$ the scheme is second order but unstable in shocks. For $\omega = 1$ the scheme is very robust, entropy dissipative, but quite diffusive.

With a simple one-dimensional analysis recalled in Section 3.4, we can give an estimate of the numerical viscosity introduced by the relaxation parameter ω . Roughly speaking (see a more precise formulation in Section 3.4), the numerical viscosity ν scales like

$$\nu = \lambda^2 \Delta t \left(\frac{1}{\omega} - \frac{1}{2} \right) \left(1 - \frac{1}{\lambda^2} c_h^2 \right),$$

where c_h is the maximal wave speed of the MHD-DC system, according to (2). We observe that ν indeed vanishes when $\omega = 2$ and that it has the right sign if λ satisfies the *subcharacteristic condition*

$$\lambda > c_h. \quad (12)$$

For the two-dimensional case, it is more difficult to give a rigorous stability condition. In practice, we have observed that the stability is ensured with the following subcharacteristic condition

$$\lambda > 2c_h. \quad (13)$$

It would be interesting to construct a rigorous strategy for choosing locally the optimal value of ω . Let us mention that many finite volume schemes have been designed for solving MHD equations, with specific Riemann solvers. We can mention for instance the solver of [7], based on a relaxation approach, with proven stability and accuracy features. However, the solver presented in [7] is more complicated to program and less computationally efficient. We can thus compensate the slightly lower accuracy of the kinetic scheme by finer meshes, without increasing too much the computational load. For a quantitative comparison of the relative efficiencies of the finite volume/difference methods and the LBM, we refer for instance to [38,53]

3.3. Boundary conditions

For the moment, we have assumed periodic boundary conditions. But Dirichlet or Neumann conditions are also possible. Neumann conditions are a good candidate for handling variables associated with outgoing waves. Indeed, because of the hyperbolic nature of the equations, it is not possible to impose all the physical data at the boundaries. The number of boundary conditions depends on the wave pattern. We refer to [23], for a specific analysis of the boundary conditions in the over-relaxation scheme.

3.4. Analysis of the numerical viscosity in the one-dimensional case

In this section, we state some results about the numerical viscosity of the kinetic relaxation scheme in the one-dimensional case. This one-dimensional analysis will give us simple intuitions for adjusting the relaxation parameter in the two-dimensional case. We consider the one-dimensional MHD-DC system (4). For more simplicity, we will denote $\mathbf{F} = \mathbf{F}_1$ and $x = x_1$. The equations then read

$$\partial_t \mathbf{w} + \partial_x \mathbf{F}(\mathbf{w}) = 0.$$

For the analysis, it is possible to replace the scalar relaxation parameter ω by a matrix Ω , for more generality. In the following, we establish the comparison between matrices in the usual way, by the comparison of the associated quadratic form (the resulting order is thus not total).

It is then possible to prove the following result:

Theorem 1. *If the relaxation matrix satisfies $\mathbf{I} < \Omega < 2\mathbf{I}$ and if $\mathbf{f} = \mathbf{f}^{eq}$ at the initial time, then, up to second-order terms in $O(\Delta t^2)$, \mathbf{w} is a solution of the following system of conservation laws*

$$\partial_t \mathbf{w} + \partial_x \mathbf{F}(\mathbf{w}) = \lambda^2 \Delta t \partial_x \left(\left(\Omega^{-1} - \frac{1}{2} \mathbf{I} \right) \left(\mathbf{I} - \frac{1}{\lambda^2} \mathbf{F}'(\mathbf{w})^2 \right) \partial_x \mathbf{w} \right) + O(\Delta t^2). \quad (14)$$

Remark 1. The proof is based on standard Taylor expansions. For the scalar case $\Omega = \omega \mathbf{I}$, it can be found (for instance) in [15]. The approach is classical in the analysis of the Lattice Boltzmann Method (LBM). See for instance [24,42].

Remark 2. For $\Omega = \mathbf{I}$ the proof, based on a Chapman-Enskog expansion, is also given in [11].

The above analysis allows recovering formally the so-called sub-characteristic condition. Assuming that $\mathbf{I} < \Omega < 2\mathbf{I}$, the second-order (“viscous”) terms have the good sign, which ensures stability of the model, if the following matrix is positive:

$$\mathbf{V}(\lambda, \mathbf{w}) = \mathbf{I} - \frac{1}{\lambda^2} \mathbf{F}'(\mathbf{w})^2 > 0. \quad (15)$$

A fully rigorous mathematical proof of stability of the kinetic model is given by Bouchut in [5], Section 3.2, pp. 140–142. Bouchut’s proof is not based on asymptotic expansions but on fully non-linear entropy estimates. Let us emphasize that the vectorial kinetic construction ensures stability even when shock waves occur and is not limited by a low-Mach assumption.

From (14), we formally observe that the scheme is second-order accurate in time in the over-relaxation case, when

$$\Omega = 2\mathbf{I}.$$

Finally, this analysis provides a way to numerically approximate, up to second order in time, the second-order system of conservation laws

$$\partial_t \mathbf{w} + \partial_x \mathbf{F}(\mathbf{w}) - \partial_x (\mathbf{A}(\mathbf{w}) \partial_x \mathbf{w}) = 0. \quad (16)$$

If the diffusion matrix $\mathbf{A}(\mathbf{w})$ is small, it is natural to take

$$\Omega = 2 \left(\mathbf{I} + \frac{2}{\lambda^2 \Delta t} \mathbf{A}(\mathbf{w}) \mathbf{V}(\lambda, \mathbf{w})^{-1} \right)^{-1}. \quad (17)$$

Let us remark that if λ is large enough then the above matrix is well defined.

In order to check practically the accuracy of the approximation, we apply the above analysis for a simplified system of two conservation laws.

We consider the one-dimensional isothermal Euler equations with a diagonal diffusion of $\epsilon \partial_{xx}(\rho, \rho u)^T$. Hence the full fluid system we are interested in is given by

$$\partial_t \begin{pmatrix} \rho \\ \rho u \end{pmatrix} + \partial_x \begin{pmatrix} \rho u \\ (\rho u^2 + c^2 \rho) \end{pmatrix} = \epsilon \partial_{xx} \begin{pmatrix} \rho \\ \rho u \end{pmatrix}. \quad (18)$$

This system is a very simplified version of the MHD equation, where the magnetic field is assumed to vanish and the gas is supposed to be isothermal. We have chosen to use this specific non-physical diffusion for our first test since it is exactly the type of diffusion that is apparent in a standard finite volume code using a Lax-Friedrichs flux.

The diffusion matrix for (18) is simply given by

$$\mathbf{A}(\mathbf{w}) = \begin{pmatrix} \epsilon & 0 \\ 0 & \epsilon \end{pmatrix}. \quad (19)$$

From (17), the inverse of the relaxation matrix is given by

$$\mathbf{\Omega}^{-1} = \begin{pmatrix} a_{1,1} & a_{1,2} \\ a_{2,1} & a_{2,2} \end{pmatrix} \quad (20)$$

with

$$a_{1,1} = \frac{\epsilon ((\lambda^2 - c^2) - 3u^2)}{\Delta t (u^4 + (c^2 - \lambda^2)^2 - 2u^2 (c^2 + \lambda^2))} + \frac{1}{2},$$

$$a_{1,2} = \frac{2u\epsilon}{\Delta t (u^4 + (c^2 - \lambda^2)^2 - 2u^2 (c^2 + \lambda^2))},$$

$$a_{2,1} = -\frac{2u(u-c)(c+u)\epsilon}{\Delta t (u^4 + (c^2 - \lambda^2)^2 - 2u^2 (c^2 + \lambda^2))},$$

$$a_{2,2} = \frac{\epsilon (u^2 + (\lambda^2 - c^2))}{\Delta t (u^4 + (c^2 - \lambda^2)^2 - 2u^2 (c^2 + \lambda^2))} + \frac{1}{2}.$$

For our first test of the matrix relaxation, we solve (18) comparing the Lattice Boltzmann scheme using (17) as relaxation matrix and a standard explicit centered finite volume scheme for approximating (18). In this centered scheme, the time step is taken very small in such way that the stability condition is satisfied and that the time integration error can be neglected. In other words, the equivalent PDE of both schemes is (16). Hence, given the parameters in both schemes are set to represent the same diffusion ϵ , one should get the same type of diffusion for both schemes.

To test this we take the simple case of a stationary viscous shock.

The initial data for this shock tube problem are

$$w_l = (\rho_l, \rho_l u_l)^T \quad w_r = \left(\frac{\rho_l u_l^2}{c}, \frac{c}{u_l} \right)^T \quad (21)$$

where the sound speed c is set as $c = 1$ while for the left state we have chosen

$$(\rho_l, u_l) = \left(2, \frac{3}{2} \right)^T. \quad (22)$$

For the diffusion we use $\epsilon = 0.1$. For the Lattice Boltzmann scheme, we set $\lambda = 40$. The results on a grid of 128 cells at time $T = 0.1$ are shown in Fig. 1. The test shows that the matrix relaxation lattice Boltzmann scheme provides the correct diffusion and therefore results in the right viscous profile. The conclusion of this section is that the relaxation parameter is related to numerical viscosity. In 1D it can indeed be adjusted to fit the exact viscosity up to second order. For two-dimensional computations the analysis is more complicated, but heuristically we expect a similar behavior.

In the following, for simplicity reasons, we only consider scalar relaxation.

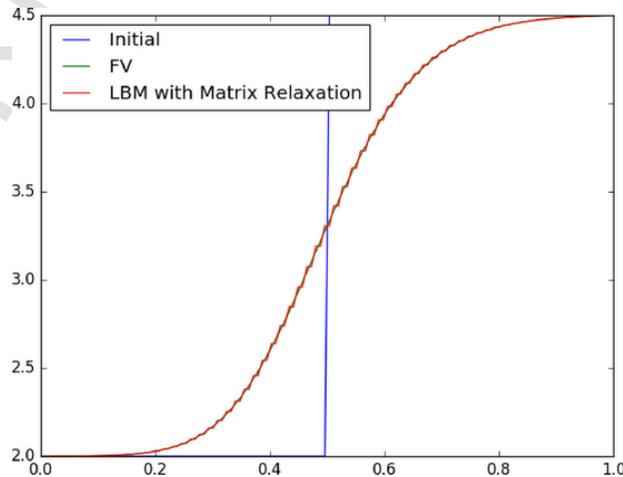


Fig. 1. Viscous Shock Test with $\epsilon = 0.1$: Comparison of the viscous profiles for ρ of the finite volume and the lattice Boltzmann scheme. The small staircase effect is typical of the LBM scheme and is due to the over-relaxation. It can be removed by only plotting every other point.

4. GPU implementation

We have implemented the above LBM algorithm. The LBM is particularly well adapted to parallelism. It is possible to provide very efficient implementations on GPU (Graphic Purpose Units) hardware.

4.1. OpenCL

Today Graphic Processing Units (GPU) have more and more computation power. The Open Computing Language (OpenCL) is a software environment for simplifying the use of the GPUs for general computing. It is also possible to use OpenCL for driving a heterogeneous set of general multicore processors.

In the OpenCL terminology, an *accelerator* is a parallel computing device, such as a GPU or a multicore CPU. A *kernel* is a (generally small) program that is executed on several of the computing cores of the GPU. For instance, the Nvidia GPU GTX 1660 has 448 computing cores. Thanks to the OpenCL command queue management, it is possible to launch several million kernel instances, which are dispatched on the hundreds of cores of the GPU.

The *OpenCL runtime* is a library of C functions, called from the host, in order to drive the GPU. The OpenCL runtime, because it is written in C is quite heavy to use in practice: the verbosity is high, the API is not very user-friendly and memory management is cumbersome. For this reason it is advised to use OpenCL wrappers written in a higher-level language such as C++ or Python. We have used the Python OpenCL wrapper written by Andreas Klöckner, PyOpenCL [34], which makes OpenCL initializations and calls much easier and shorter to program.

The *OpenCL language* is a C-like language for writing the kernels that will be executed on the computing cores.

In order to perform a complex task, a kernel has to be executed many times. Each execution of a kernel is called a *work-item*. Each work-item is identified by a unique global ID p .

For more details on OpenCL, we refer for instance to [16,30,34].

4.2. OpenCL implementation of the LBM algorithm

We have implemented the above algorithm using PyOpenCL.

The different devices used in this paper are listed in Table 1. In this table, we also describe the hardware: size of the accelerator memory, size of the fast-access cache memory, number of compute units and number of processing elements. The AMD processor was used inside a virtual environment, which implies a non-negligible loss of performance. PyOpenCL allows us to select either the CPU or the GPU for the computations. With the OpenCL AMD drivers, when the CPU is selected, it is also possible to choose the number of activated CPU cores through a Linux environment variable. This is useful for estimating (in a crude way) the efficiency of the OpenCL parallelism.

The LBM is very simple. Our implementation is made of two OpenCL kernels, a few C functions and a small Python driver for initializing the memory buffers, launching the OpenCL kernel and plotting the results. The role of the first OpenCL kernel is to compute the initial condition directly into the memory buffer created on the OpenCL accelerator. The role of the second kernel is to perform a time step of the LBM. The time-stepping is driven from the PyOpenCL program.

The most important point is to take care of the organizations of the kinetic data into memory. Indeed ensuring aligned memory access is essential for performance. In practice all the values $f_{k,i,j}$ for kinetic velocities v_k in cells (i,j) are arranged in a single memory buffer $f_n[]$, with the following storage

$$f_{k,i,j} = f_n[\text{imem}], \quad \text{imem} = i + j * N_x + k * N_x * N_y,$$

where N_x is the number of grid points in the x_1 -direction and N_y is the number of grid points in the x_2 -direction. Then, in the OpenCL kernel, each cell (i,j) is associated to the work-item

Table 1

Characteristics of the OpenCL devices tested in this paper. CU stands for ‘‘Compute Units’’ and ‘‘PE’’ for ‘‘Processing Elements’’.

short name	name					
V100	Tesla V100-PCIE-16GB					
Quadro	Nvidia Quadro P6000					
GTX	Nvidia GTX 1660					
AMD	AMD EPYC 7551 32-core					
CPU	Intel Xeon E5-2609 v4					
Iris 640	Intel Iris Graphics 640					
short name	type	frequency	memory	cache	CU	PE
V100	GPU	1.4 GHz	16 GB	48 kB	80	5120
Quadro	GPU	1.5 GHz	24 GB	32 kB	30	3840
GTX	GPU	1.6 GHz	6 GB	48 kB	22	1408
AMD	CPU	2 GHz	48 GB	32 kB	1–24	1–24
Intel	CPU	1.7 GHz	64 GB	32 kB	16	16
Iris 640	GPU	1.0 GHz	4 GB	64 kB	48	192

$$p = i + j * Nx.$$

This numbering ensures that neighboring work-items will access neighboring memory locations during the shift algorithm (8). For instance, if work-item p access data in global memory, for reading or writing, at location i_{mem} , then work-item $p' = p + 1$ access location $i_{mem}+1$.

This property is still true in the relaxation algorithm. First, the kinetic data are copied in processor registers in an aligned way. The computations of the conservative data (9), equilibrium (10) and collisions (11) are done in registers, which ensure very fast memory access. Finally, the kinetic data are copied back to the global buffer $f_n[]$ in a fully aligned way.

We remark that thanks to the chosen organization into memory, we do not have to use the local cache memory for accelerating the algorithm.

In order to measure the efficiency of the implementation, we perform a memory bandwidth test for a 512×512 grid. One time-step of the method implies the read access in the global memory of the set of fields of the previous time-step. The local computations are done in registers. Then there is another write access to global memory for storing the data of the next time-step. The memory size in Gigabyte of one set of fields is

$$n_{GB} = \frac{Nx \times Ny \times prec \times 4 \times m}{1024^3},$$

where $prec$ is the number of bytes for storing one floating-point number ($prec = 4$ for single precision and $prec = 8$ for double precision). We then perform a given number of time iterations n_{iter} and measure the elapsed time $t_{elapsed}$ (in s) in the OpenCL kernels. We perform two kinds of experiments. In the first experiment, we deactivate the numerical computations and only perform the shift operations. The memory bandwidth (in GB/s) of the shift algorithm is then given by

$$b = \frac{2 \times n_{GB} \times n_{iter}}{t_{elapsed}}.$$

In the second experiment, we reactivate the computations and measure how the bandwidth is reduced. This allows to evaluating how the elapsed time is shared between memory transfers and computations. The results are given in Table 2. We observe a good efficiency of the shift algorithm in the shift-only case: the transfer rates are not very far from the maximal bandwidth of the device, at least for the GPU accelerators. From these results we also observe that the LBM algorithm is clearly memory bound. When the single precision computations are activated on the GPU devices (GTX, Quadro, V100), the elapsed time of the shift-and-relaxation test is not very different from the shift-only test. For the double precision computations, we observe that the V100 device outperforms all the other GPUs.

5. Numerical applications to MHD

5.1. Smooth vortex (performance test)

The smooth vortex test is a classical test for MHD codes. It is described for instance in [26]. Because this is an exact solution, it allows us to assess the accuracy of the solver. Here we also used this test to evaluate the efficiency of the parallel implementation. The test case is built upon a single vortex, which is a stationary solution of the MHD system, to which a constant drift velocity is added. In the moving frame centered on $\mathbf{r}_O(t) = t\mathbf{u}_{drift}$, with $\mathbf{u}_{drift} \in \mathbb{R}^2$, the analytical solution reads in polar coordinates

Table 2

Bandwidth efficiency of the LBM algorithm. Comparison of the data transfer rates of the shift-only algorithm and of the shift-and-relaxation algorithm for different computing devices and floating-point accuracies. The resulting bandwidth is compared with the maximal memory bandwidth advertised by the vendor of the hardware device. Let us observed that the maximal measured bandwidth can be smaller than the vendor's specifications. For instance, for the V100 GPU, the practical bandwidth is measured to 825 GB/s in [48]. This has to be compared with the theoretical bandwidth of 900 GB/s.

	prec.	b (GB/s, shift-only)	b (GB/s, shift-relax)	max. b (GB/s)
Intel	float32	17.58	13.38	60
Intel	float64	19.12	17.48	60
Iris 640	float32	26.20	24.98	34
Iris 640	float64	20.08	3.78	34
GTX	float32	147.54	146.94	192
GTX	float64	148.76	49.72	192
Quadro	float32	336.45	329.06	432
Quadro	float64	344.50	127.21	432
V100	float32	692.31	676.44	900
V100	float64	705.88	610.17	900

$$\begin{aligned} \rho(r, \theta) &= \rho_0, \\ \mathbf{u}(r, \theta) &= u_0 [\mathbf{u}_{\text{drift}} + h(r) \mathbf{e}_\theta], \\ \mathbf{B}(r, \theta) &= b_0 h(r) \mathbf{e}_\theta, \\ p(r, \theta) &= p_0 + \frac{b_0^2}{2} (1 - h(r)), \end{aligned}$$

with $b_0 = \rho_0 u_0^2$. The results shown below are obtained with $\gamma = 5/3$ and the parameter set

$$\rho_0 = p_0 = 1, \quad u_0 = b_0 = 0.2, \quad \mathbf{u}_{\text{drift}} = (1, 1)^T, \quad h(r) = \exp\left[\frac{(1 - r^2)}{2}\right].$$

The computational domain is the square $\Omega =]-L/2, L/2[\times]L/2, L/2[$, with $L = 20$. We compute the solution at time $t = 10$. The grid contains N points in x_1 and x_2 directions. For this smooth test case, we can take a relaxation parameter $\omega = 2$ without triggering instabilities. The cleaning parameter is set to $c_h = 6$ and the kinetic speed to $\lambda = 20$ in order to enforce the subcharacteristic condition (13) with a security margin. We compute the error e_N in the L^1 norm at the final time between the exact solution and the numerical solution on the first component of the momentum (the other components of the solution would give similar results):

$$e_N = \int_{\Omega} \left| (\rho u_1)_{\text{num}} - (\rho u_1)_{\text{exact}} \right|.$$

Asymptotically, the order of the scheme is evaluated by

$$\delta \simeq \frac{\ln(e_N/e_{2N})}{\ln 2}.$$

The obtained numerical results are summed up in Table 3, where we give the convergence study and a performance evaluation of the implementation. OpenCL permits to run the same code on a multicore CPU or a GPU. We have tested several CPU or GPU hardware in single or double precision. Table 3 confirms the second order of accuracy of the scheme in the case $\omega = 2$. In addition we observe a good efficiency of the implementation on several types of GPU. On CPU there is also a speedup achieved by the OpenCL parallelism, but it is very sensitive to the OpenCL drivers. For instance, with the same hardware (an Intel Xeon two-CPU system) the program runs almost three times faster with the OpenCL Intel drivers than with the open source POCL drivers. Let us finally mention that when it is run on only one core, the code is very slow. This is due to the fact that our implementation is not really optimized for correctly harnessing the CPU cache. Here, with a more clever tiling strategy, the one-core run could probably be accelerated by an order of magnitude (see for instance [32]).

5.2. Orszag-Tang vortex

The Orszag-Tang test case [18,41,43] is often used to test a numerical method for MHD. It consists in a vortex system where turbulent structures and shocks develop. The domain is $\mathcal{D} =]0, 2\pi[\times]0, 2\pi[$ and the boundary conditions are periodic in $x_1 = x$ and $x_2 = y$. The initial conditions are given in Table 4.

As in the previous test case, we took $c_h = 6$ and $\lambda = 20$. But we have to set the relaxation parameter to $\omega = 1.8$ in order to avoid instabilities in the shock waves. In Figs. 2 and 3 we present several snapshots of the evolution of the vortex. In Fig. 4 we compare the results obtained with several grid refinements. The grid refinement clearly improves the sharpness of the shock profiles.

Table 3

Convergence and performance study. Some tests are done in single precision (float32) and others in double precision (float64). The speedup is a comparison for $N=1024$ with the slowest device. ‘‘CU’’ means ‘‘Compute Units’’: it is the number of activated cores in a CPU computation or of OpenCL compute units for a GPU computation.

	CU	accuracy	N=128	N=256	N=512	N=1024	speedup
AMD	1	float32	11.9 s	159 s	621 s	6396 s	1
AMD	24	float32	1.01 s	9.4 s	153 s	1380 s	5
POCL	16	float32	2.30 s	17.3 s	96.6 s	644 s	10
Intel	16	float32	0.75 s	3.93 s	32 s	226 s	30
Intel	16	float64	0.82 s	5.62 s	53 s	315 s	20
GTX	22	float32	0.04 s	0.31s	2.46 s	19.48 s	330
Quadro	30	float32	0.017 s	0.15 s	1.06 s	8.25 s	780
Quadro	30	float64	0.15 s	0.81 s	5.67 s	45.53 s	140
V100	80	float32	0.015 s	0.084 s	0.54 s	3.93 s	1600
V100	80	float64	0.031 s	0.21 s	1.17 s	8.35 s	770
e_N		float64	0.0507	0.0130	0.0033	0.0008	
δ		float64	-	1.96	1.99	2.00	

Table 4
Initial states for the Orszag-Tang test case.

Variables	States
γ	5/3
ρ	γ^2
p	γ
u_x	$-\sin(y)$
u_y	$\sin(x)$
u_z	0
B_x	$-\sin(y)$
B_y	$\sin(2x)$
B_z	0

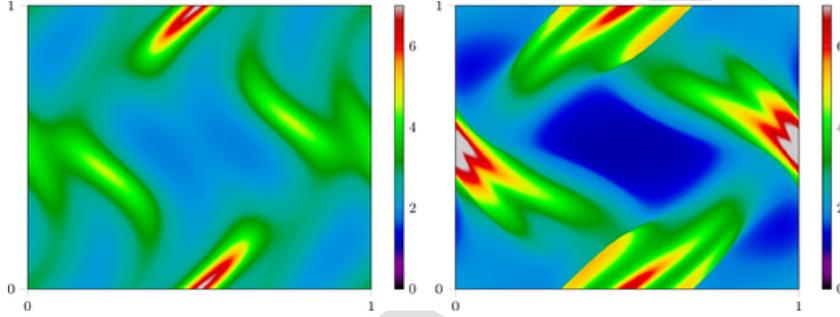


Fig. 2. Snapshots of ρ for the Orszag-Tang configuration recorded at times $t = 0.1$ and $t = 0.2$. Grid size is $N_x \times N_y = 1024 \times 1024$.

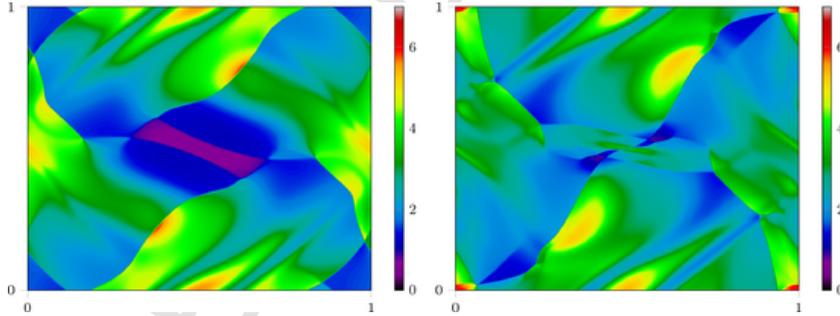


Fig. 3. Snapshots of ρ for the Orszag-Tang configuration recorded at times $t = 0.3$ (left) and $t = 0.4$ (right). Grid size is $N_x \times N_y = 1024 \times 1024$.

5.3. MHD tilt instability

In this test-case, we consider a squared spatial domain $\mathcal{D} =]-3, 3[\times]-3, 3[$. We define $r^2 = x^2 + y^2$.

The initial condition is an MHD equilibrium given by:

$$\rho = 1.0, \quad \mathbf{u} = \mathbf{0}, \quad p_0 = 1.0, \quad \psi = 0, \quad (23)$$

$$k = 3.8317059702, \quad K = \frac{-2}{kJ_0(k)}, \quad \varphi = KJ_1(kr)\frac{y}{r}, \quad (24)$$

$$p = \begin{cases} p_0 & \text{if } r \geq 1, \\ p_0 + \frac{k^2}{2}\varphi^2 & \text{otherwise,} \end{cases} \quad (25)$$

$$B_1 = \begin{cases} \frac{x^2 - y^2}{r^4} - 1 & \text{if } r \geq 1, \\ K \left[\frac{ky^2}{r^2} J_0(kr) + \frac{x^2 - y^2}{r^3} J_1(kr) \right] & \text{otherwise,} \end{cases} \quad (26)$$

$$B_2 = \begin{cases} \frac{2xy}{r^4} & \text{if } r \geq 1, \\ -K \left[\frac{kxy}{r^2} J_0(kr) - \frac{2xy}{r^3} J_1(kr) \right] & \text{otherwise,} \end{cases} \quad (27)$$

where J_0 and J_1 are the Bessel functions of the first kind of orders zero and one respectively.

We add an initial perturbation of this equilibrium of the form:

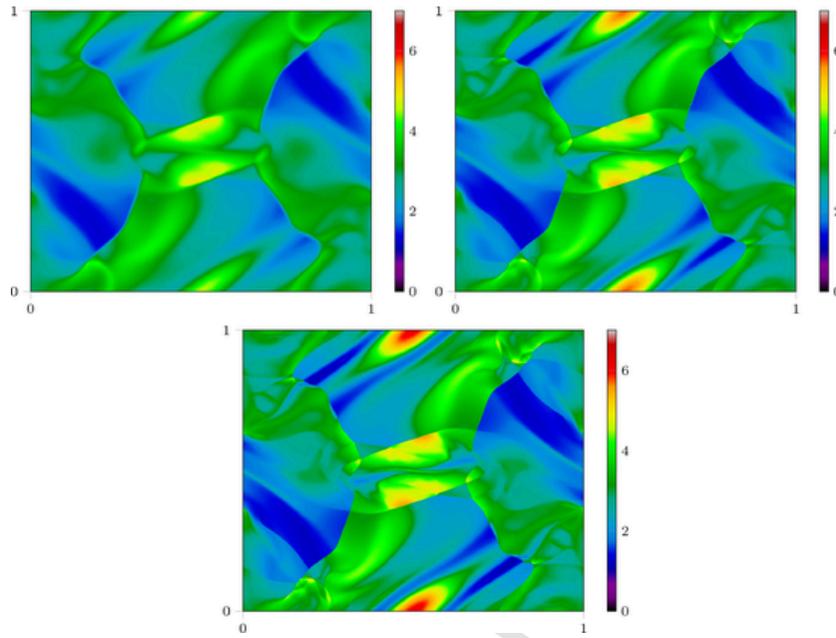


Fig. 4. Snapshots of ρ for the Orszag-Tang configuration recorded at time $t = 0.5$. Grid sizes are $N_x \times N_y = 256 \times 256$ (top left), 1024×1024 (top right) and 4096×4096 (bottom).

$$\mathbf{u} = 2\epsilon \exp(-r^2) \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix},$$

with $\epsilon = 10^{-3}$.

This kind of test-case has been studied in [35,47,49]. The given initial condition is a formal stationary solution of the MHD system. However, when perturbed, an instability develops. A detailed description of the different stages of this instability can be found in [4, 47]. In [33,47], the study of the tilt instability was also made for a compressible MHD system. In [35,49] the equations are incompressible, which leads to slightly different results.

The

For the analysis of the simulation results, we consider the current density given by:

$$\mathbf{C} = \nabla \times \mathbf{B},$$

and we are interested in the third component of \mathbf{C} , *i.e.* C_3 . We will also look at the kinetic energy of the system, which is computed as the integral over the spatial domain of the kinetic energy density:

$$E_{\text{kin}} = \int_{\mathcal{D}} \frac{1}{2} \rho |\mathbf{u}|^2.$$

We apply Dirichlet boundary conditions in this test case: the data in the boundary cells are simply kept constant and equal to the initial condition. Finally, the numerical scheme involves the following parameters

$$\Delta x = L_x / N_x, \quad c_h = 6, \quad \lambda = 20, \quad \Delta t = \Delta x / \lambda,$$

and $\omega = 1.9$ for all variables, except for the kinetic variables associated with ψ where $\omega = 1$ (see Section 5.3). In our numerical tests, we have noticed that choosing $\omega = 2$ leads to an unstable numerical solution, due to the lack of dissipation.

The instability is illustrated in Figs. 5, 6, 7 for different times and the numerical approach presented above. An important numerical feature of the tilt instability is the formation of a current sheet: as time advances, the current density concentrates in a thin layer on the boundary of the magnetic islands.

Current sheets and current peak

Current sheets develop at the edges of the magnetic islands. The width of these current sheets, as well as the maximum intensity of the current density, highly depend on the numerical implementation. We compare our results with those of [35], based on a sophisticated adaptive algorithm (mixing grid refinement and adaptive order of interpolation).

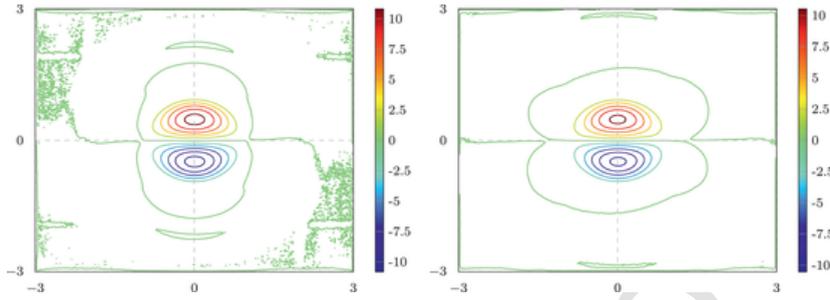


Fig. 5. Snapshots of the magnetic current density C_3 recorded at (non-dimensional) times $t = 1$ and $t = 2$. Grid size is $N_x \times N_y = 1024 \times 1024$.

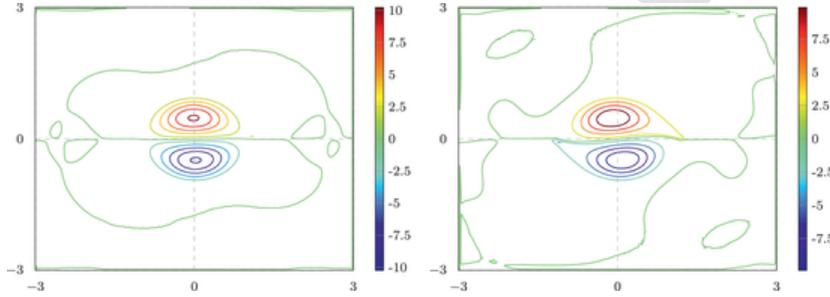


Fig. 6. Snapshots of the magnetic current density C_3 recorded at times $t = 3$ and $t = 4$. Grid size is $N_x \times N_y = 1024 \times 1024$.

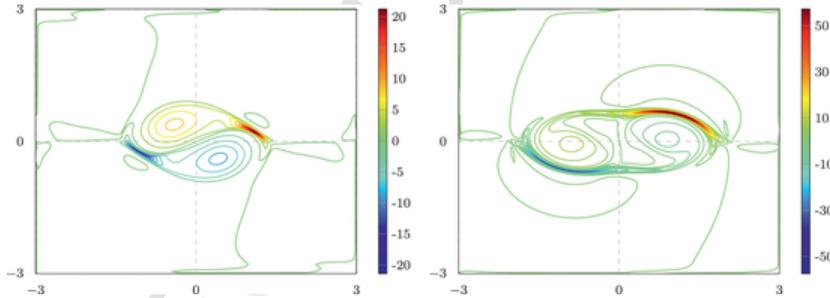


Fig. 7. Snapshots of the magnetic current density C_3 recorded at times $t = 5$ and $t = 6$. Grid size is $N_x \times N_y = 1024 \times 1024$.

In Fig. 8, we zoom on the top right current sheet different grid sizes and at time $t = 6$. The finest current structures are obtained with the finest grid (here the 4096×4096 grid). In Fig. 9, we present the ratio of the maximum current density $C_{3,max}$ over the initial current density $C_{3,0}$. This ratio increases with the grid refinement and should be infinite in the case of a perfectly ideal MHD system [35], *i.e.* without numerical viscosity.

The maximum ratio we can obtain is 14.2 with the 7000×7000 grid. In comparison, [35] obtained a ratio of 5 for the simulation with a first-order method and 32,768 triangles. They improved the ratio up to 41 with their adaptive grid algorithm near the current sheets. In conclusion, our method seems to have a higher numerical viscosity, which limits the maximum current peak, but the computations are certainly much faster. Indeed, adaptive methods are known to suffer from overhead due to the algorithm complexity.

Kinetic energy growth rate

During the development of the instability, the kinetic energy of the system presents an exponential behavior near $t = 4$:

$$E_{\text{kin}}(t) \simeq E_{\text{kin}}(4) \exp(\eta(t - 4)).$$

The growth rate η of the kinetic energy has been studied for instance in [47], where the growth rate was estimated to $\eta = 1.44$. In [35], the study is limited to the incompressible case. Then, the grid-converged value for the growth rate is estimated to $\eta = 1.3$. In [33] for a similar configuration, the growth rate is evaluated to $\eta = 1.498$. This work probably gives the best estimation because the simulation is conducted with a very accurate adaptive scheme.

The simulations of the tilt instability that we perform are post-processed at regular time intervals. In Figs. 10 and 11, we show the time evolution of the total kinetic energy $E_{\text{kin}}(t)$. One can note that the growth of the kinetic energy happens in two stages: in a first stage, the growth is faster than exponentially, then slows down and after two seconds the second stage begins with exponential growth. When the grid resolution is higher, the first stage tends to vanish and the second stage begins sooner.

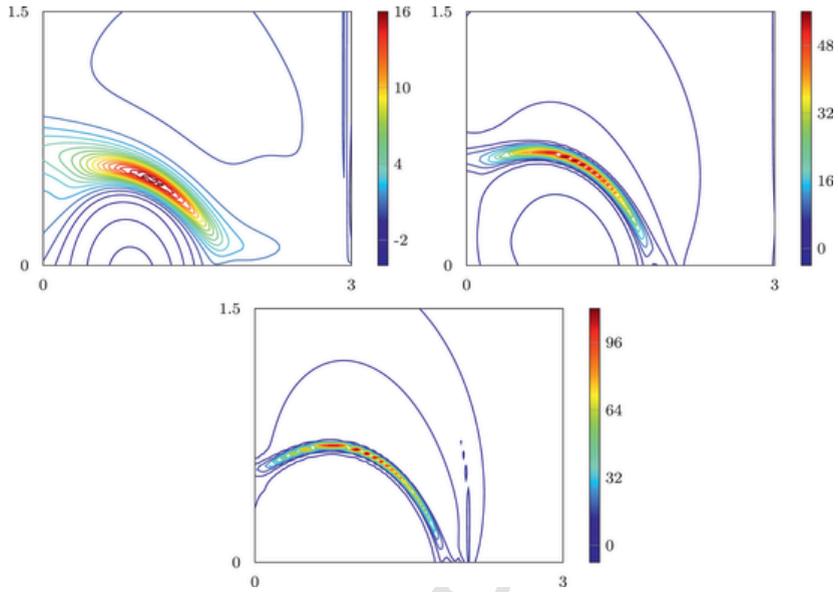


Fig. 8. Snapshots of the magnetic current density near the top right current sheet. The snapshots time is $t = 6$. Grid sizes are $N_x \times N_y = 256 \times 256$ (top left), $N_x \times N_y = 1024 \times 1024$ (top right), $N_x \times N_y = 4096 \times 4096$ (bottom). We observe a sharpening of the current sheet with the grid refinement.

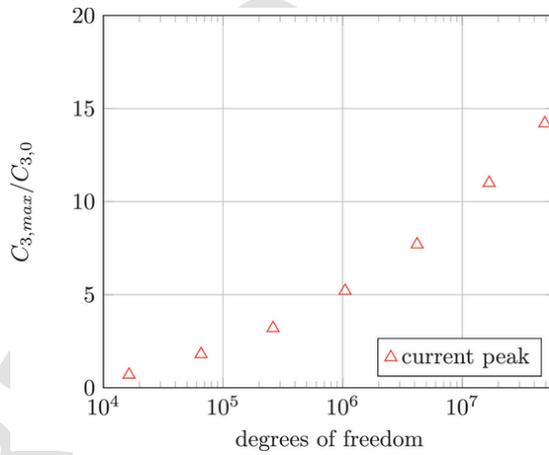


Fig. 9. Ratio of the peak current and the initial current density for simulation grids ranging from 128×128 to 7000×7000 .

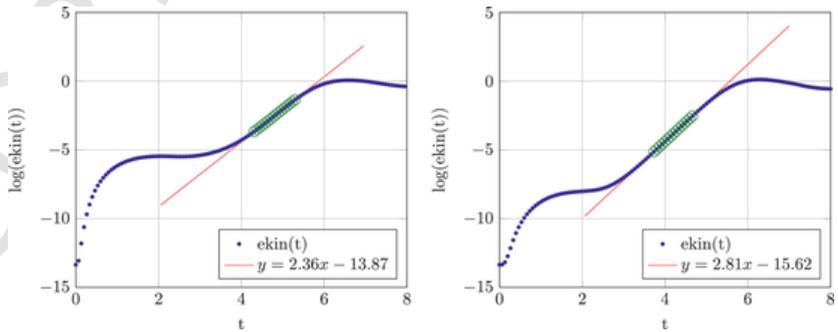


Fig. 10. Growth of the kinetic energy during the simulation. Measures that have been accounted for the linear regression are identified with green marks and the regression line is in red. Results for grid sizes $N_x \times N_y = 512 \times 512$ and $N_x \times N_y = 2048 \times 2048$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

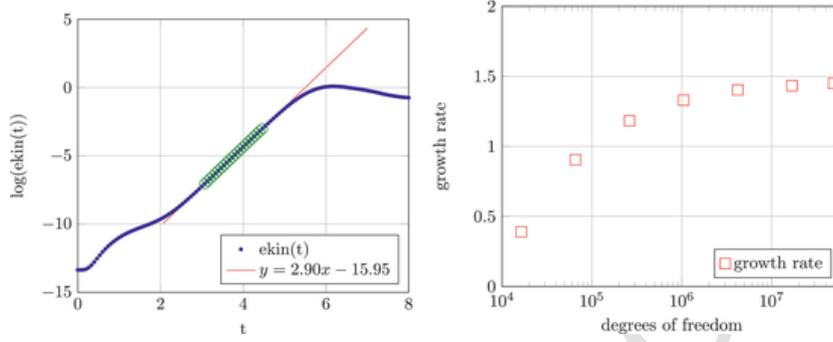


Fig. 11. Growth of the kinetic energy during the simulation. Measures that have been accounted for the linear regression are identified with green marks and the regression line is in red. Results for grid size $N_x \times N_y = 7000 \times 7000$. Growth rates from $N_x = N_y = 128$ (16384 degrees of freedom) to $N_x = N_y = 7000$ (49×10^6 degrees of freedom). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In a logarithmic scale, the linear regression of the kinetic energy growth is performed with the `Scipy` function `linregress` and the regression line is illustrated in Figs. 10 and 11. The results of the linear regression are reported in Table 5 and Fig. 11. The converged growth rate is close to $\eta = 1.45$ for our simulations. It is higher than the results of [47] and [35] but slightly lower than the most accurate results of [33]. Our results seem thus to be quite correct on the finest mesh.

In Table 5, we summarize the characteristics of the linear regression for each mesh size.

Divergence cleaning effect

In order to ensure a magnetic field close to a divergence-free field, we have presented the divergence cleaning procedure in Section 2.1. Here we compare two strategies for the numerical implementation of the divergence cleaning equation:

1. the kinetic variables associated with the conservative variable Ψ are solved numerically like any other kinetic variable. This means that in the relaxation step, the relaxation coefficient ω is set to $\omega = 1.9$ and the order of resolution is close to 2,
2. during the relaxation step, the kinetic variables associated with the macro-variable Ψ are relaxed with a coefficient $\omega = 1$, which comes down to setting

$$\begin{aligned} \mathbf{f}_{k,i,j}^{n+1} &= \omega \mathbf{f}_k^{eq}(\mathbf{w}_{i,j}^{n+1}) - (\omega - 1) \mathbf{f}_{k,i,j}^{n+1,-} \\ &= \mathbf{f}_k^{eq}(\mathbf{w}_{i,j}^{n+1}). \end{aligned}$$

Now, the order of resolution is 1 and the waves associated with the perturbations of the divergence-free constraint are better damped than with $\omega = 1.9$. In [20], several strategies are proposed for damping the divergence errors. One of the methods is to introduce a viscous damping term. In its spirit our proposal is similar, but the viscous term is introduced here in a numerical way.

These two options are compared in Figs. 12, 13, 14 for different times. For obtaining those plots, the divergence of the magnetic field is computed with a simple centered finite difference approximation

$$\nabla \cdot \mathbf{B} \simeq \frac{B_1(x + \Delta x, y) - B_1(x - \Delta x, y)}{2\Delta x} + \frac{B_2(x, y + \Delta y) - B_2(x, y - \Delta y)}{2\Delta y}.$$

At time $t = 0.5$, one can notice the effects of the boundary conditions that have propagated towards the center of the domain. The perturbations of the divergence-free constraint have been more attenuated with the second strategy than with the first one. At time $t = 1$, the divergence-free constraint is mainly violated at the edges of the magnetic vortices and close to the domain boundaries. Once

Table 5

Characteristics of the linear regression for the kinetic energy growth rate for each mesh size: range of values that have been accounted for and growth rate.

Mesh Size	Regression range	Growth rate η
128^2	[5.6,6.7]	0.390
256^2	[4.9,5.9]	0.904
512^2	[4.3,5.3]	1.082
1024^2	[3.9,4.9]	1.330
2048^2	[3.7,4.7]	1.403
4096^2	[3.0,4.5]	1.433
7000^2	[3.0,4.5]	1.450

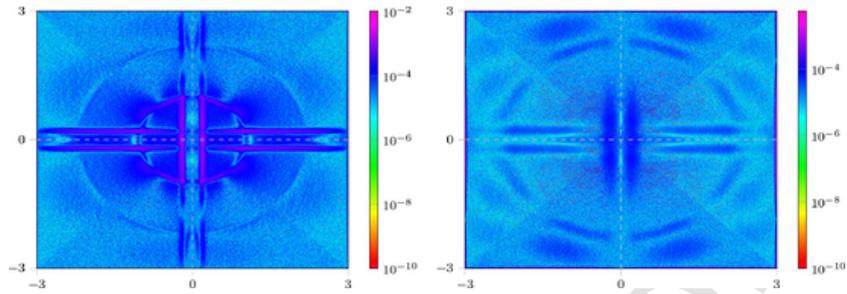


Fig. 12. Snapshots of the divergence of the magnetic field recorded at times $t = 0.5$ for both strategies of divergence cleaning. Grid size is $N_x \times N_y = 1024 \times 1024$.

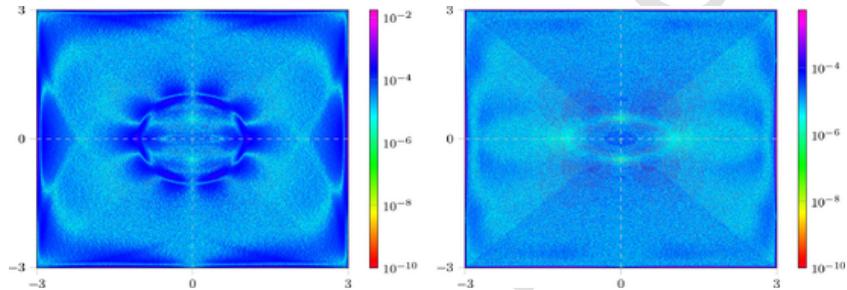


Fig. 13. Snapshots of the divergence of the magnetic field recorded at times $t = 1$ for both strategies of divergence cleaning. Grid size is $N_x \times N_y = 1024 \times 1024$.

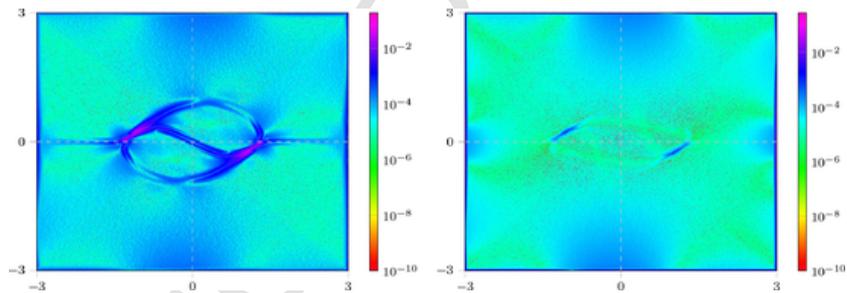


Fig. 14. Snapshots of the divergence of the magnetic field recorded at times $t = 5$ for both strategies of divergence cleaning. Grid size is $N_x \times N_y = 1024 \times 1024$.

more, the results are better with the second divergence cleaning strategy. Finally, at time $t = 5$, when the vortices have begun to align, the divergence-free constraint is strongly perturbed, but results are still better with the second strategy.

Because the second strategy is more precise than the first, all other test cases previously exposed in this paper are computed with the second strategy.

6. Conclusion

In this work we have proposed a fast and robust Lattice-Boltzmann solver for the two-dimensional MHD equations with divergence cleaning. The method is general and can be extended to other systems of conservation laws.

We have provided a preliminary one-dimensional analysis in order to evaluate the numerical viscosity of the numerical scheme. We have shown that in principle, the numerical viscosity can be adjusted to a specific viscosity. An interesting challenge would be to extend the analysis to higher dimensions, but this leads to complex calculations because of cross second-order derivatives. If this analysis is possible, this would lead to a scheme where the physical resistive terms could be approximated properly.

The kinetic representation of the equations is very well adapted to massive parallel computing. The simplicity of the algorithm allows achieving almost optimal efficiency on GPU hardware. This leads to the possibility to conduct computations on very fine uniformly refined meshes.

When conducting the simulations, we observed that the limiting factor was memory rather than computation time. In order to extend the method to even finer meshes or to three-dimensional computations, the memory management has to be improved. An obvious possibility is to distribute the computations on several GPUs, using a task-based runtime system (such as in [8]). This approach could also be mixed with a compression strategy in order to reduce memory occupation and data transfers.

Thanks: this work has been finalized during a one-week stay at the Oberwolfach Research Institute for Mathematics (MFO) in September 2020. The authors are very grateful to the staff of the MFO for their warm welcome and perfect organization despite the constraints caused by the pandemic.

The authors also thank Matthieu Boileau for his helpful contribution in the design of the PyOpenCL GPU code.

References

- [1] D. Aregba-Driollet, R. Natalini, Discrete kinetic schemes for multidimensional systems of conservation laws, *SIAM J. Numer. Anal.* 37 (6) (2000) 1973–2004.
- [2] J. Badwaik, M. Boileau, D. Coulette, E. Franck, P. Helluy, C. Klingenberg, L. Mendoza, H. Oberlin, Task-based parallelization of an implicit kinetic scheme, *ESAIM Proc. Surv.* 63 (2018) 60–77.
- [3] T. Barth, On the role of involutions in the Discontinuous Galerkin discretization of Maxwell and magnetohydrodynamic systems. *Compatible Spatial Discretizations*, Springer, 2006, pp. 69–88.
- [4] H. Baty, FINMHD: an adaptive finite-element code for magnetic reconnection and formation of plasmoid chains in magnetohydrodynamics, *Astrophys. J. Suppl. Ser.* 243 (2) (2019) 23.
- [5] F. Bouchut, Construction of BGK models with a family of kinetic entropies for a given system of conservation laws, *J. Stat. Phys.* 95 (1–2) (1999) 113–170.
- [6] F. Bouchut, Y. Jobic, R. Natalini, R. Ocellini, V. Pavan, Second-order entropy satisfying BGK-FVS schemes for incompressible Navier-Stokes equations, *SMAI J. Comput. Math.* 4 (2018) 1–56.
- [7] F. Bouchut, C. Klingenberg, K. Waagan, A multiwave approximate Riemann solver for ideal MHD based on relaxation II: numerical implementation with 3 and 5 waves, *Numer. Math.* 115 (4) (2010) 647–679.
- [8] B. Bramas, P. Helluy, L. Mendoza, B. Weber, Optimization of a discontinuous Galerkin solver with OpenCL and StarPU, *Int. J. Finite Vol.* 15 (1) (2020) 1–19.
- [9] M. Brio, C.C. Wu, An upwind differencing scheme for the equations of ideal magnetohydrodynamics, *J. Comput. Phys.* 75 (2) (1988) 400–422.
- [10] P. Cargo, G. Gallice, Roe matrices for ideal MHD and systematic construction of Roe matrices for systems of conservation laws, *J. Comput. Phys.* 136 (2) (1997) 446–466.
- [11] G.Q. Chen, C.D. Levermore, T.P. Liu, Hyperbolic conservation laws with stiff relaxation terms and entropy, *Commun. Pure Appl. Math.* 47 (6) (1994) 787–830.
- [12] S. Chen, G.D. Doolen, Lattice Boltzmann method for fluid flows, *Annu. Rev. Fluid Mech.* 30 (1) (1998) 329–364.
- [13] D. Coulette, E. Franck, P. Helluy, M. Mehrenberger, L. Navoret, *Palindromic Discontinuous Galerkin Method*, Springer International Publishing, Cham, 2017, pp. 171–178.
- [14] D. Coulette, E. Franck, P. Helluy, M. Mehrenberger, L. Navoret, High-order implicit palindromic Discontinuous Galerkin method for kinetic-relaxation approximation, *Comput. Fluids* 190 (2019) 485–502.
- [15] C. Courtès, D. Coulette, E. Franck, L. Navoret, Vectorial kinetic relaxation model with central velocity. Application to implicit relaxations schemes, *Commun. Comput. Phys.* 27 (4) (2020).
- [16] A. Crestetto, P. Helluy, J. Jung, Numerical resolution of conservation laws with OpenCL. *ESAIM: Proceedings*, vol. 40, EDP Sciences, 2013, pp. 51–62.
- [17] J.-P. Croisille, R. Khanfir, G. Chanteur, Numerical simulation of the MHD equations by a kinetic-type method, *J. Sci. Comput.* 10 (1) (1995) 81–92.
- [18] R.B. Dahlburg, J.M. Picone, Evolution of the Orszag-Tang vortex system in a compressible medium. I. Initial average subsonic flow, *Phys. Fluids B Plasma Phys.* 1 (11) (1989).
- [19] W. Dai, P.R. Woodward, An approximate Riemann solver for ideal magnetohydrodynamics, *J. Comput. Phys.* 111 (2) (1994) 354–372.
- [20] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, M. Wesenberg, Hyperbolic divergence cleaning for the MHD equations, *J. Comput. Phys.* 175 (2) (2002) 645–673.
- [21] P.J. Dellar, Lattice kinetic schemes for magnetohydrodynamics, *J. Comput. Phys.* 179 (1) (2002) 95–126.
- [22] P.J. Dellar, An interpretation and derivation of the lattice Boltzmann method using Strang splitting, *Comput. Math. Appl.* 65 (2) (2013) 129–141.
- [23] F. Drui, E. Franck, P. Helluy, L. Navoret, An analysis of over-relaxation in a kinetic approximation of systems of conservation laws, *Comptes Rendus Mécanique* 347 (3) (2019) 259–269.
- [24] F. Dubois, Equivalent partial differential equations of a lattice Boltzmann scheme, *Comput. Math. Appl.* 55 (7) (2008) 1441–1449.
- [25] F. Dubois, Simulation of strong nonlinear waves with vectorial lattice Boltzmann schemes, *Int. J. Modern Phys. C* 25 (12) (2014) 1441014.
- [26] M. Dumbser, R. Loubère, A simple robust and accurate a posteriori sub-cell finite volume limiter for the Discontinuous Galerkin method on unstructured meshes, *J. Comput. Phys.* 319 (2016) 163–199, <https://doi.org/10.1016/j.jcp.2016.05.002>.
- [27] Z. Fan, F. Qiu, A. Kaufman, S. Yoakum-Stover, GPU cluster for high performance computing. SC'04: Proceedings of the 2004 ACM/IEEE Conference on Supercomputing, IEEE, 2004. 47–47
- [28] N. Frapolli, S.S. Chikatamarla, I.V. Karlin, Entropic lattice Boltzmann model for compressible flows, *Phys. Rev. E* 92 (6) (2015) 061301.
- [29] N. Frapolli, S.S. Chikatamarla, I.V. Karlin, Theory, analysis, and applications of the entropic lattice Boltzmann model for compressible flows, *Entropy* 22 (3) (2020) 370.
- [30] B. Gaster, L. Howes, D.R. Kaeli, P. Mistry, D. Schaa, *Heterogeneous Computing with OpenCL: revised OpenCL 1*, Newnes, 2012.
- [31] B. Graille, Approximation of mono-dimensional hyperbolic systems: a lattice Boltzmann scheme as a relaxation method, *J. Comput. Phys.* 266 (2014) 74–88.
- [32] P. Helluy, T. Strub, M. Massaro, M. Roberts, Asynchronous OpenCL/MPI numerical simulations of conservation laws. *Software for Exascale Computing-SPPEXA 2013–2015*, Springer, 2016, pp. 547–565.
- [33] R. Keppens, O. Porth, C. Xia, Interacting tilt and kink instabilities in repelling current channels, *Astrophys. J.* 795 (1) (2014) 77.
- [34] A. Klöckner, N. Pinto, Y. Lee, B. Catanzaro, P. Ivanov, A. Fasih, PyCUDA and PyOpenCL: a scripting-based approach to GPU run-time code generation, *Parallel Comput.* 38 (3) (2012) 157–174.
- [35] S. Lankalapalli, J.E. Flaherty, M.S. Shephard, H. Strauss, An adaptive finite element method for magnetohydrodynamics, *J. Comput. Phys.* 225 (1) (2007) 363–381.
- [36] J. Latt, C. Coreixas, J. Beny, A. Parmigiani, Efficient supersonic flow simulations using lattice Boltzmann methods based on numerical equilibria, *Philos. Trans. R. Soc. A* 378 (2175) (2020) 20190559.
- [37] M. Lehmann, Esoteric pull and esoteric push: two simple in-place streaming schemes for the lattice Boltzmann method on GPUs, *Computation* 10 (6) (2022) 92.
- [38] S. Marié, D. Ricot, P. Sagaut, Comparison between lattice Boltzmann method and Navier–Stokes high order schemes for computational aeroacoustics, *J. Comput. Phys.* 228 (4) (2009) 1056–1070.
- [39] D.O. Martínez, S. Chen, W.H. Matthaeus, Lattice Boltzmann magnetohydrodynamics, *Phys. Plasmas* 1 (6) (1994) 1850–1867.
- [40] T. Miyoshi, K. Kusano, A multi-state HLL approximate Riemann solver for ideal magnetohydrodynamics, *J. Comput. Phys.* 208 (1) (2005) 315–344.
- [41] S.A. Orszag, C.-M. Tang, Small-scale structure of two-dimensional magnetohydrodynamic turbulence, *J. Fluid Mech.* 90 (1) (1979) 129–143.
- [42] H. Otomo, B.M. Boghosian, F. Dubois, Two complementary lattice-Boltzmann-based analyses for nonlinear systems, *Physica A* 486 (2017) 1000–1011.
- [43] J.M. Picone, R.B. Dahlburg, Evolution of the Orszag-Tang vortex system in a compressible medium. II. Supersonic flow, *Phys. Fluids B Plasma Phys.* 3 (29) (1991).
- [44] K.G. Powell, An Approximate Riemann Solver for Magnetohydrodynamics (That Works in More Than one Space Dimension). Technical Report NASA/CR-194902 ICASE Report No. 94-24, ICASE-NASA Langley, ICASE, NASA Langley Research Center, 1994.
- [45] K.G. Powell, An approximate Riemann solver for magnetohydrodynamics. Upwind and High-Resolution Schemes, Springer, 1997, pp. 570–583.
- [46] F. Renard, Y. Feng, J.-F. Boussuge, P. Sagaut, Improved compressible hybrid lattice Boltzmann method on standard lattice for subsonic and supersonic flows, *Comput. Fluids* (2021) 104867.
- [47] R.L. Richard, R.D. Szyora, M. Ashour-Abdalla, Magnetic reconnection driven by current repulsion, *Phys. Fluids B Plasma Phys.* 2 (3) (1990) 488–494, <https://doi.org/10.1063/1.85911>.

doi.org/10.1063/1.859338.

- [48] M. Špet'ko, O. Vysocký, B. Janský, L. Říha, DGX-A100 Face to Face DGX-2 Performance, power and thermal behavior evaluation, *Energies* 14 (2) (2021) 376.
- [49] H.R. Strauss, D.W. Longcope, An adaptive finite element method for magnetohydrodynamics, *J. Comput. Phys.* 147 (2) (1998) 318–336, <https://doi.org/10.1006/jcph.1998.6091>.
- [50] S. Succi, *The Lattice Boltzmann Equation: for Fluid Dynamics and Beyond*, Oxford University Press, 2001.
- [51] M. Torrilhon, Non-uniform convergence of finite volume schemes for Riemann problems of ideal magnetohydrodynamics, *J. Comput. Phys.* 192 (1) (2003) 73–94.
- [52] G. Tóth, The $\nabla \cdot B = 0$ constraint in shock-capturing magnetohydrodynamics codes, *J. Comput. Phys.* 161 (2) (2000) 605–652.
- [53] K.-R. Wichmann, M. Kronbichler, R. Löhner, W.A. Wall, A runtime based comparison of highly tuned lattice Boltzmann and finite difference solvers, *Int. J. High Perform. Comput. Appl.* 35 (4) (2021) 370–390.
- [54] W. Xian, A. Takayuki, Multi-GPU performance of incompressible flow computation by lattice Boltzmann method on GPU cluster, *Parallel Comput.* 37 (9) (2011) 521–535.
- [55] J. Zhao, Discrete-velocity vector-BGK models based numerical methods for the incompressible Navier-Stokes equations, *Commun. Comput. Phys.* 29 (2) (2021) 420–444.