Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **553,** 1–13 (2026)
Advance Access publication 2026 October 19

https://doi.org/10.1093/mnras/stag1979

Lattice Boltzmann Methods for compressible magnetohydrodynamics

Fedor Bukreev [1,3★†] Adrian Kummerländer[1,2†] and Mathias J. Krause[1,2,3†]

[1]*Lattice Boltzmann Research Group (LBRG), Straße am Forum 8/Englerstraße 2, D-76131 Karlsruhe, Germany*
[2]*Institute for Applied and Numerical Mathematics (IANM), Karlsruhe Institute of Technology (KIT), Englerstraße 2, D-76131 Karlsruhe, Germany*
[3]*Institute of Mechanical Process Engineering and Mechanics (MVM), Karlsruhe Institute of Technology (KIT), Straße am Forum 8, D-76131 Karlsruhe, Germany*

Accepted 2026 October 15. Received 2026 October 14; in original form 2026 June 17

ABSTRACT

Compressible, magnetized plasma flows pervade astrophysics, from stellar winds to planetary magnetospheres and the interstellar medium. Their high-fidelity simulation at scale is limited by the non-local operations that conventional solvers require for flux reconstruction, gradient and stress evaluation, and magnetic-divergence control, which bottleneck communication on heterogeneous supercomputers. We present a novel class of fully local Lattice Boltzmann Method (LBM) schemes for compressible magnetohydrodynamics (MHD) in which all non-linear operations are point-wise. The approach exploits the kinetic structure to transport each state variable of the Strang-split conservation laws alongside its characteristics, recovering the compressible macroscopic fluxes, gradients, and viscous stresses through local relaxation rather than finite-difference stencils. Divergence control enters as a local source term within the same kinetic update, with no separate global solve. The formulation is partial-differential-equation agnostic and encompasses hydrodynamic limits such as the compressible Euler and incompressible Navier–Stokes equations. We discretize both ideal compressible and resistive incompressible MHD, and validate the scheme across multiple resolutions against the canonical astrophysical MHD test suite: the Brio-Wu shock tube resolves compound waves and slow-fast shocks, the MHD rotor captures strong torsional Alfvén waves, and the Orszag–Tang vortex reproduces current-sheet formation and the transition to MHD turbulence. Implemented in the performance-portable framework OpenLB, the scheme reaches up to 98.9 per cent of the hardware roofline on GPU. Finally, as a capability demonstration, we simulate a rotating, surface-resolved magnetized body modelled after 16 Psyche in a supersonic early solar wind, yielding an emergent bow shock, flux-rope magnetosheath, and reconnecting magnetotail.

Key words: magnetohydrodynamics (MHD) – plasmas – shock waves – methods: numerical.

1 INTRODUCTION

Most of the baryonic Universe is a magnetized, compressible plasma, and its large-scale dynamics are well described in the magnetohydrodynamic (MHD) limit, in which the Navier–Stokes and Maxwell equations couple through an electrically conducting fluid (A. A. Schekochihin & S. C. Cowley 2007; J. M. Stone 2011). MHD flows shape systems across astrophysics, from the solar and stellar winds and the bow shocks of planetary magnetospheres to the turbulent magnetized interstellar medium and accretion flows around compact objects. Capturing them at high fidelity is computationally severe. Beyond the strong non-linear coupling, the wide reconstruction stencils and characteristic decompositions of conventional finite-volume MHD, together with the auxiliary operations needed to control the magnetic divergence, all entail substantial neighbour communication. As demand for resolution grows, it is increasingly this non-local data movement, rather than arithmetic, that bottlenecks simulations on modern heterogeneous high-performance-computing (HPC) architectures (S. Williams, A. Waterman & D. Patterson 2009; D. Unat et al. 2017). The Lattice Boltzmann Method (LBM) offers an alternative whose algorithmic profile is well matched to this constraint, decomposing the update into a perfectly parallel, point-local collision step and a neighbour-local streaming step (T. Krueger et al. 2017). While uniquely suited to massively parallel execution, extending standard scalar LBM to complex, high-Mach-number multiphysics introduces severe numerical challenges. G. Wissocq, Y. Liu & R. Abgrall (2025) notes that recovering the correct macroscopic fluxes often requires complex Hermite expansions or excessive artificial dissipation that compromises physical accuracy. To circumvent these limitations, F. Dubois (2014) and later G. Wissocq et al. (2025) introduced the vectorial LBM (VLBM), which assigns a distinct set of distribution functions to each variable of the macroscopic state vector.

★ E-mail: fedor.bukreev@kit.edu
† These authors contributed equally to this work.

© The Author(s) 2026.
Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

Building on this, and inspired by the unified kinetic schemes of C. Liu & K. Xu (2017), H. Grad (1949), and H. Struchtrup & M. Torrilhon (2003), our approach treats the fully evolved equation system directly at the mesoscopic level. Instead of relying on non-local finite differences to approximate spatial derivatives, we apply Strang splitting in time (G. Strang 1968) and explicitly transport the state variables of the conservation equations alongside their characteristics. This yields a unified kinetic scheme capable of natively solving a wide range of transport equations, including the fully compressible Euler, incompressible Navier–Stokes, and linear Navier-Cauchy systems, as demonstrated by O. Boolakee, M. Geier & L. De Lorenzis (2025), and the compressible Navier-Stokes-Fourier equations (F. Bukreev, A. Kummerländer & M. J. Krause 2026). In this work, we develop this class of schemes into a solver for astrophysical MHD. The formulation is fully local: the compressible fluxes, the magnetic field gradients, and the viscous stresses are all recovered through point-wise kinetic relaxation rather than finite-difference stencils. Magnetic-divergence control is handled in the same spirit, entering as a local source-term correction (K. G. Powell et al. 1999) with the divergence read from the locally evolved field gradients, so that no separate elliptic or global solve is required. To assess accuracy, robustness, and shock-capturing fidelity, we validate the scheme against the standard astrophysical MHD test problems. The Brio-Wu shock tube, applied for instance by H. Huang et al. (2024), probes the resolution of compound waves, slow-fast shocks, and non-equilibrium kinetic effects at sharp interfaces. The MHD rotor, used by P. Londrillo & L. Del Zanna (2000) and E. Gaburov & K. Nitadori (2011), tests the handling of strong torsional Alfvén waves and rapid rotational discontinuities. The Orszag–Tang vortex, analysed by J. M. Stone et al. (2008) and P. J. Dellar (2013), evaluates fidelity in capturing magnetic reconnection, current-sheet formation, and the transition to MHD turbulence. Finally, we apply the scheme to a fully three-dimensional plasma-body interaction: a magnetized, rotating asteroid embedded in a supersonic early solar wind, in which the bow shock, magnetosheath, and reconnecting wake are resolved together with the moving solid boundary and its remnant magnetic field.

The scheme is implemented in the OpenLB software library (M. J. Krause et al. 2021), which provides a hardware-abstracted framework for efficient execution on heterogeneous exascale supercomputers (A. Kummerländer et al. 2023, 2026a). Section 2 introduces the PDE-agnostic framework for general conservation laws and our fully local discretization, then applies it to ideal and resistive MHD and to static and moving boundaries. Section 3 validates predictive accuracy and shock-capturing fidelity against the canonical benchmarks. Section 4 evaluates hardware efficiency on a modern GPU. Section 5 presents the three-dimensional solar-wind interaction.

2 METHODOLOGY

In this section, we present a novel, partial differential equation (PDE) agnostic Lattice Boltzmann framework designed for highly efficient and scalable transport simulations. Subsequently, we detail the application of this methodology to the target MHD equation systems, including the required evolution of spatial gradients and the realization of boundary conditions for the presented simulation cases.

2.1 General conservation laws

The proposed numerical framework is designed to solve generic systems of conservation laws that can be written in the abstract vector form:

$$\partial_t \boldsymbol{Q} + \nabla \cdot \boldsymbol{\Phi} = \boldsymbol{S}, \tag{1}$$

where $\boldsymbol{Q}$ represents an arbitrary macroscopic state vector, $\boldsymbol{\Phi}$ is the corresponding flux tensor governing the transport of the state variables, and $\boldsymbol{S}$ encapsulates any localized source or sink terms. By formulating the numerical method around this generic abstraction, the resulting solver remains largely PDE-agnostic and can be readily adapted to various complex multiphysics problems by redefining the components of $\boldsymbol{Q}$, $\boldsymbol{\Phi}$, and $\boldsymbol{S}$.

2.2 Lattice Boltzmann Methods

To discretize the generic conservation laws defined above, we employ a fully local VLBM. Unlike finite difference or finite volume methods that rely on extended spatial stencils for flux and gradient evaluation, our approach confines these non-linear operations to strict point-wise updates. This algorithmic locality inherently avoids the wide, irregular memory access patterns of traditional solvers, rendering the scheme suited to the saturation of modern heterogeneous HPC architectures.

The fully local VLBM is formulated as

$$f_{k,i}(\boldsymbol{x} + \boldsymbol{c}_i \Delta t, t + \Delta t) = f_{k,i}(\boldsymbol{x}, t) + \Omega^C\Big(\tau_{\mathrm{LB},k},\, f_{k,i}(\boldsymbol{x}, t),\, f^{eq}_{k,i}(\boldsymbol{x}, t)\Big) + \Omega^S\Big(\tau_{\mathrm{LB},k},\, f_{k,i}(\boldsymbol{x}, t),\, f^{eq}_{k,i}(\boldsymbol{x}, t),\, S_{k,i}\Big), \tag{2}$$

where $f_{k,i}$ is the distribution function for a state vector component k in the lattice direction i, Ω^C is the collision operator, $\tau_{\mathrm{LB},k}$ is the lattice relaxation time for a component, $f^{\mathrm{eq}}_{k,i}$ is the equilibrium distribution function, and $\boldsymbol{c}_i = c\hat{\boldsymbol{e}}_i$ is the discrete lattice velocity vector, where $c = \Delta x / \Delta t$ is the lattice speed and $\hat{\boldsymbol{e}}_i$ are dimensionless Cartesian direction vectors (e.g. $\hat{\boldsymbol{e}}_i \in \{(0,0), (\pm 1, 0), (0, \pm 1)\}$ for D2Q5). Ω^S is the source operator with the $S_{k,i}$ source term. For the standard Bhatnagar-Gross-Krook (BGK) collision (P. L. Bhatnagar, E. P. Gross & M. Krook 1954) and direct source injection, these operators read:

$$\Omega^C\Big(\tau_{\mathrm{LB},k},\, f_{k,i},\, f^{eq}_{k,i}\Big) = -\tfrac{1}{\tau_{\mathrm{LB},k}}\Big(f_{k,i}(\boldsymbol{x}, t) - f^{eq}_{k,i}(\boldsymbol{x}, t)\Big), \tag{3}$$

$$\Omega^S\Big(\tau_{\mathrm{LB},k},\, f_{k,i},\, f^{eq}_{k,i},\, S_{k,i}\Big) = \Delta t\, w_i\, S_{k,i}, \tag{4}$$

where w_i are the directional lattice weights. The universality of the applied scheme results in the ability to use any possible collision, such as the regularized BGK (RLB) (J. Latt & B. Chopard 2006), two-relaxation-time collision (TRT) (I. Ginzburg, F. Verhaeghe & D. d'Humières 2008), and multiple relaxation time collision (MRT) (D. d'Humières 2002), among others. Furthermore, distinct source operators can be applied, including direct injection, the scheme proposed by Z. Guo, C. Zheng & B. Shi (2002), and the method developed by Z. Chai & T. S. Zhao (2013). Although the update reads as an explicit forward step, the scheme is second-order in both space and time. The streaming term $f_{k,i}(\boldsymbol{x} + \boldsymbol{c}_i\Delta t, t + \Delta t)$ is an exact shift along a lattice direction, so no spatial derivative is evaluated by the algorithm. The spatial order is instead fixed by the equilibrium of equation (5), whose first two moments match the target exactly. The Chapman-Enskog derivation and its method-of-manufactured-solutions verification for this class of schemes, the resistive MHD system included, are given in the companion work (A. Kummerländer, F. Bukreev &

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

M. J. Krause 2026b), where the measured orders of convergence are near two.

In the shock-dominated benchmarks of Section 3, the lattice relaxation time is also used to add an artificial viscosity. On D2Q5 the relaxation diffusion reduces to $\nu^{\text{eff}} = \nu_{\text{num}}$ when $\tau_{\text{LB},k} = 1/2 + 3\nu_{\text{num}}\Delta t/\Delta x^2$, so setting one sets the other. The ideal-MHD benchmarks below hold the dimensionless numerical diffusivity fixed, giving $\nu_{\text{num}} = \mathcal{O}(\Delta x)$ and hence formally first-order accuracy where that stabilizer is active. This is the cost of shock capture and is in any case the order attainable at a discontinuity. The choice is stated quantitatively for each case in Section 3.

In contrast to standard LBM, the non-linear fluxes are embedded directly in the first moment. For this reason, the equilibrium function requires only a linear, first-order expansion:

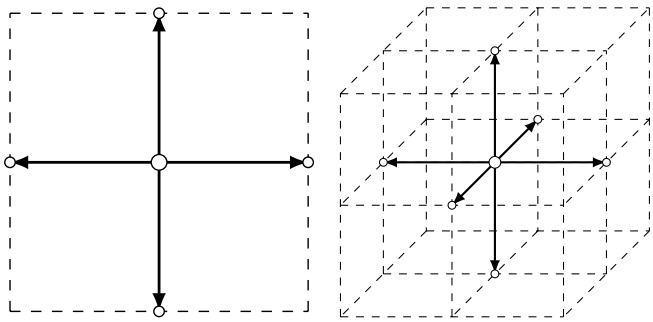
$$f_{k,i}^{eq} = w_i \left(Q_k + \frac{\boldsymbol{c}_i \cdot \boldsymbol{\Phi}_k}{c_s^2} \right), \tag{5}$$

where w_i are the lattice Gauss–Hermite quadrature weights and c_s is the lattice speed of sound. Here, $\boldsymbol{c}_i \cdot \boldsymbol{\Phi}_k = \sum_\alpha c_{i,\alpha}\, \Phi_{k,\alpha}$ contracts over the spatial dimensions α the components of the discrete velocity $\boldsymbol{c}_i$ in lattice direction i with those of the macroscopic flux $\boldsymbol{\Phi}_k$ of state component Q_k, and is therefore a scalar for each i. The constant c_s is a globally fixed numerical reference speed set by the quadrature weights rather than the physical speed of sound, and takes the lattice-specific values $c_s = c/\sqrt{3}$ on D2Q5 and $c_s = c/2$ on D3Q7. The true physical wave propagation speeds (sound speed a, Alfvén speed v_A, and fast/slow magnetosonic speeds v_f, v_s) are embedded directly within the non-linear flux tensor $\boldsymbol{\Phi}$ and equation of state. Stability requires the lattice sound speed c_s to exceed the fastest physical signal speed. This is the subcharacteristic condition of kinetic relaxation schemes (S. Jin & Z. Xin 1995; F. Bouchut 1999; D. Aregba-Driollet & R. Natalini 2000; K. Guillon, R. Hélie & P. Helluy 2024; A. Kummerländer et al. 2026b, c), which for ideal MHD reduces to a bound on the Courant number (CFL) alone, $\text{CFL} = v_{\max}\Delta t/\Delta x \leq c_s/c$, with $v_{\max} = \max(|\boldsymbol{u}| + v_f)$ the maximum signal speed. This ceiling is $1/\sqrt{3}$ on D2Q5 and $1/2$ on D3Q7. The lattice speed of sound is therefore not a proxy for any physical wave speed, but the admissible Courant ceiling. All simulations reported below are run inside this bound.

To recover the state vector variables and their fluxes, moments of the distribution functions must be taken. The zeroth moment directly recovers the conserved variables, while the first moment exactly recovers the macroscopic fluxes $\boldsymbol{\Phi}_k$:

$$Q_k = \sum_i f_{k,i}, \qquad \boldsymbol{\Phi}_k = \sum_i f_{k,i}\boldsymbol{c}_i. \tag{6}$$

Unlike standard scalar LBM, where the macroscopic fluxes must be recovered from the second-order $\mathcal{O}(\epsilon^2)$ moment expansion, requiring high-order lattice isotropy and complex Hermite polynomials, the present approach strictly embeds the physical non-linear fluxes into the first moment. M. Anandan & S. Raghurama Rao (2024) shows that this yields the correct transport equations independent of the underlying discrete velocity set. The two discrete velocity sets used in this work, D2Q5 and D3Q7, are shown schematically in Fig. 1, with their velocities and weights listed in Table 1. Furthermore, this PDE-agnostic formulation can be naturally extended to address Navier–Cauchy and other complex multiphysics systems.



(a) D2Q5 (b) D3Q7

Figure 1. Schematic view of the discrete velocity sets. The discrete velocities are depicted as arrows (zero velocity not shown).

Table 1. Lattice discretization parameters.

Directions i	Normalized lattice velocities $\boldsymbol{c}_i\Delta t/\Delta x$	Lattice weights w_i
	D2Q5	
0	$(0, 0)$	1/3
1,2,3,4	$(\pm1, 0), (0, \pm1)$	1/6
	D3Q7	
0	$(0, 0, 0)$	1/4
$1, 2, \ldots, 6$	$(\pm1, 0, 0), (0, \pm1, 0), (0, 0, \pm1)$	1/8

For more details and mathematical analysis of VLBMs, we recommend the works of G. Wissocq et al. (2025), O. Boolakee et al. (2025), F. Dubois (2014), M. Anandan & S. Raghurama Rao (2024), and K. Guillon et al. (2024). For the aspects of standard LBM and advection-diffusion equation approximation, T. Krueger et al. (2017) provides a comprehensive overview.

2.2.1 *Lattice relaxation time*

In the case of the standard transport equation approximation, the relaxation time is connected to the diffusion constant D as

$$D = \left(\tau_{\text{LB}} - \frac{1}{2}\right) c_s^2 \Delta t. \tag{7}$$

The molecular diffusion (or viscous effects) is embedded directly into the flux term, which is why the lattice relaxation time remains only a stabilization parameter. In the ideal case, $\tau_{\text{LB},k}$ can be set to 0.5, which means zero added artificial diffusion. Due to the challenge of obtaining a stable simulation with this lattice relaxation time, a stabilizing diffusion is imposed in the following simulations. By mesh refinement, this stabilization diffusion can be scaled to zero.

2.2.2 *Unit conversion*

In contrast to standard LBM, where the SI units are scaled to lattice units with respect to the lattice speed of sound, lattice relaxation time, and a common definition that $\Delta x_{\text{LB}} = \Delta t_{\text{LB}} = 1$, the present approach is no longer bound to this convention. Identical to other numerical discretization schemes, the unit scaling is performed here with the chosen reference quantities, such as reference length, velocity, viscosity, and others. This choice is case specific and has only numerical or computational precision limitations.

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

2.2.3 Local moments evolution

A fundamental feature of this generic framework is its ability to recover higher moments (stress tensor components, heat fluxes) natively through a local relaxation mechanism, entirely avoiding finite-difference stencils. When spatial gradients of a variable are required for dissipative fluxes or source terms, they are evolved as independent variables governed by a conservation equation with a relaxation source term.

Taking a generic spatial gradient tensor $\boldsymbol{G}$ tracking the gradients of an arbitrary field $\boldsymbol{V}$ as an example:

$$\partial_t \boldsymbol{G} + \nabla \cdot \boldsymbol{\Phi}_{\boldsymbol{G}} = -\frac{1}{\tau_R}\boldsymbol{G}. \tag{8}$$

If analysed in a quasi-steady state, the time derivative $\partial_t \boldsymbol{G}$ is negligible. The gradients are thus purely determined by the balance between their flux divergences and their relaxation:

$$\boldsymbol{G} \approx -\tau_R(\nabla \cdot \boldsymbol{\Phi}_{\boldsymbol{G}}). \tag{9}$$

This balance holds when the relaxation time is short against the macroscopic dynamical time-scale $\tau_{\text{dyn}} \sim L/v_{\max}$ over a characteristic length L, and its formal asymptotic derivation is given in the companion work (A. Kummerländer et al. 2026b). By formulating the macroscopic flux $\boldsymbol{\Phi}_{\boldsymbol{G}}$ such that the field $\boldsymbol{V}$ is embedded appropriately (e.g. via a $-\boldsymbol{I} \otimes \boldsymbol{V}$ term), the spatial divergence operator natively acts upon the field. Consequently, the relaxation balance recovers the explicitly scaled spatial gradients ($\boldsymbol{G} \approx \tau_R \nabla \boldsymbol{V}$) without ever querying a neighbouring cell.

First, τ_R is a time, not a dimensionless number, and it is parametrized as $\tau_R = g_\tau\, \Delta t$ with a resolution-independent step count g_τ, following the construction of the companion work (A. Kummerländer et al. 2026b). The neglected transient is then $\mathcal{O}(\Delta t)$ and contracts under refinement rather than persisting as a fixed offset. The values of g_τ used here lie between 41 and 116 steps. Secondly, the ratio that the argument above actually requires is τ_R/τ_{dyn}. At the coarsest grid of each benchmark this is 0.038 for the MHD rotor of Section 3.2, 0.010 for the Brio-Wu tube of Section 3.1, 0.023 for the compressible Orszag–Tang vortex of Section 3.5 and 0.024 for the resistive case of Section 3.4, and it falls in proportion to the cell size thereafter.

Varying τ_R over a sixteenfold range at fixed resolution, from $4\,\tau_R$ to $\tau_R/4$, changes the recovered profiles by less than 0.05 per cent of their range for the MHD rotor and by 0.03–0.17 per cent for the Brio-Wu tube. For the compressible Orszag–Tang vortex a fourfold change in τ_R moves the profile by only 0.001 per cent, below what the comparison can resolve. The Strang splitting method (G. Strang 1968) is applied to handle the calculation of these fluxes from the previous time-step.

2.3 Application to magnetohydrodynamics

We now specialize the framework to compressible and dissipative MHD. The governing equations couple the conservation of mass, momentum, and energy with the magnetic induction equation and map onto the generic vector form of equation (1) with:

$$\boldsymbol{Q} = \begin{pmatrix} \rho \\ \rho\boldsymbol{u} \\ e_{\text{tot}} \\ \boldsymbol{B} \end{pmatrix}, \quad \boldsymbol{\Phi} = \begin{pmatrix} \rho\boldsymbol{u} \\ \rho\boldsymbol{u}\otimes\boldsymbol{u} + p_{\text{tot}}\boldsymbol{I} - \boldsymbol{B}\otimes\boldsymbol{B} - \boldsymbol{\tau} \\ (e_{\text{tot}} + p_{\text{tot}})\boldsymbol{u} - \boldsymbol{B}(\boldsymbol{u}\cdot\boldsymbol{B}) - \boldsymbol{u}\cdot\boldsymbol{\tau} + \boldsymbol{q} + \eta\boldsymbol{j}\times\boldsymbol{B} \\ \boldsymbol{u}\otimes\boldsymbol{B} - \boldsymbol{B}\otimes\boldsymbol{u} - \eta\nabla\boldsymbol{B} \end{pmatrix}.$$

Here, ρ is density, $\boldsymbol{u}$ is velocity, e_{tot} is total energy, $\boldsymbol{B}$ is the magnetic field, and $p_{\text{tot}} = p + \frac{1}{2}|\boldsymbol{B}|^2$ is the total pressure. τ represents the viscous shear stress tensor, $\boldsymbol{q}$ is the heat flux, η is magnetic diffusivity, $\boldsymbol{j}$ is electric current density, and $\boldsymbol{I}$ is the unit tensor.

Pressure is obtained from the ideal equation of state for compressible gas within a magnetic field:

$$p = (\gamma - 1)\Big(e_{\text{tot}} - \frac{1}{2}\rho|\boldsymbol{u}|^2 - \frac{1}{2}|\boldsymbol{B}|^2\Big). \tag{10}$$

2.3.1 Evolution of the MHD equation system

In line with the methodology established in the local gradient evolution framework, the macroscopic state vector $\boldsymbol{Q}$ for the dissipative MHD system is expanded to transport the physical viscous stresses (τ) and the magnetic field gradients ($\boldsymbol{G}$) as independent variables:

$$\boldsymbol{Q} = \begin{pmatrix} \rho \\ \rho\boldsymbol{u} \\ e_{\text{tot}} \\ \boldsymbol{B} \\ \boldsymbol{\tau} \\ \boldsymbol{G} \end{pmatrix}. \tag{11}$$

Here, $\boldsymbol{G}$ tracks the spatial gradients $\nabla\boldsymbol{B}$. To recover the physical magnetic field gradients, the macroscopic flux $\boldsymbol{\Phi}_{\boldsymbol{G}}$ is formulated as

$$\boldsymbol{\Phi}_{\boldsymbol{G}} = \boldsymbol{G}\otimes\boldsymbol{u} - \boldsymbol{I}\otimes\boldsymbol{B}. \tag{12}$$

By injecting the magnetic field directly into the flux tensor via the $-\boldsymbol{I}\otimes\boldsymbol{B}$ term, the spatial divergence operator natively acts upon the magnetic field ($\nabla\cdot(\boldsymbol{I}\otimes\boldsymbol{B}) = \nabla\boldsymbol{B}$). The relaxation balance under the assumption of small relaxation time τ_M, parametrized as in Section 2, yields $\boldsymbol{G} \approx \tau_M\nabla\boldsymbol{B}$. A similar expansion logic is applied to the viscous stress tensor τ in the simulations where resistive MHD is considered.

2.3.2 Ideal MHD

For the most part, we assume the ideal MHD limit where all resistive processes are set to zero, such that viscosity, magnetic diffusivity, and thermal conductivity are null. This leads to a simplified equation system based on the Euler equations:

$$\boldsymbol{Q} = \begin{pmatrix} \rho \\ \rho\boldsymbol{u} \\ e_{\text{tot}} \\ \boldsymbol{B} \end{pmatrix}, \quad \boldsymbol{\Phi} = \begin{pmatrix} \rho\boldsymbol{u} \\ \rho\boldsymbol{u}\otimes\boldsymbol{u} + p_{\text{tot}}\boldsymbol{I} - \boldsymbol{B}\otimes\boldsymbol{B} \\ (e_{\text{tot}} + p_{\text{tot}})\boldsymbol{u} - \boldsymbol{B}(\boldsymbol{u}\cdot\boldsymbol{B}) \\ \boldsymbol{u}\otimes\boldsymbol{B} - \boldsymbol{B}\otimes\boldsymbol{u} \end{pmatrix}. \tag{13}$$

Note that the magnetic-field gradient $\boldsymbol{G}$ is retained in this limit, since it still supplies $\nabla\cdot\boldsymbol{B}$ for the Powell source term. Only the viscous stress τ and the resistive flux contributions are dropped relative to the full system.

2.3.3 Incompressible resistive MHD

For the subsonic low compressible Orszag–Tang vortex test presented by P. J. Dellar (2013) (Section 3.4), the full equation system is used with the viscous stress tensor $\tau = \mu(\nabla u + \nabla u^\intercal)$, where μ is shear viscosity. Heat fluxes are set to zero, and the magnetic diffusivity is treated as a locally changing function of the current density.

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

2.3.4 *Powell correction*

To control the spurious magnetic divergence and stabilize the evolution, we apply the eight-wave source-term correction of K. G. Powell et al. (1999) across all simulations:

$$\boldsymbol{S} = -(\nabla \cdot \boldsymbol{B}) \begin{pmatrix} 0 \\ \boldsymbol{B} \\ \boldsymbol{u} \cdot \boldsymbol{B} \\ \boldsymbol{u} \end{pmatrix}. \tag{14}$$

As is standard for eight-wave formulations, this source term renders the discrete system non-conservative and controls $\nabla \cdot \boldsymbol{B}$ by advecting and damping it with the flow rather than driving it to zero to machine precision, in contrast to constrained-transport schemes. The divergence entering $\boldsymbol{S}$ is obtained locally from the evolved gradient field $\boldsymbol{G}$, so the correction adds no non-local operation.

2.4 Boundary conditions

In the current work, we use two types of boundary condition conventions: for static boundaries and for the homogenized ones for moving solid objects.

2.4.1 *Static boundaries*

All boundary conditions used in further simulations are based on the equilibrium distribution functions. For each state vector variable, the boundaries are either Dirichlet or Neumann. In the first case, the variable value is just defined and in the other case, the needed component is taken from the next cell along the boundary discrete normal. Having information from the actual and neighbouring cells, a new equilibrium distribution function is computed and assigned to the lattice direction.

2.4.2 *Homogenization*

For moving solid objects and complex boundaries, we utilize the homogenized LBM (HLBM) developed by M. J. Krause et al. (2017) and efficiently implemented by A. Kummerländer et al. (2026d, f). In standard continuous formulations, this is equivalent in the limit of vanishing permeability to the Brinkman penalization approach (P. Angot, C.-H. Bruneau & P. Fabrie 1999), where a drag term enforces no-slip conditions by driving the permeability to zero inside the solid domain. In our discrete VLBM framework, instead of applying an explicit Brinkman forcing term, we directly impose the zero permeability limit by blending the state variables.

Specifically, we compute a homogenized target state using a convex combination of the locally computed fluid state and the prescribed solid state. For MHD, this applies not only to the macroscopic velocity $\boldsymbol{u}$, but is naturally extended to the magnetic field $\boldsymbol{B}$:

$$\hat{\boldsymbol{u}}(\boldsymbol{x}, t) = d(\boldsymbol{x}, t)\boldsymbol{u}(\boldsymbol{x}, t) + (1 - d(\boldsymbol{x}, t))\boldsymbol{u}^B(\boldsymbol{x}, t), \tag{15}$$
$$\hat{\boldsymbol{B}}(\boldsymbol{x}, t) = d(\boldsymbol{x}, t)\boldsymbol{B}(\boldsymbol{x}, t) + (1 - d(\boldsymbol{x}, t))\boldsymbol{B}^B(\boldsymbol{x}, t), \tag{16}$$

where $\boldsymbol{u}^B$ and $\boldsymbol{B}^B$ are the velocity and magnetic field of the solid object, and the homogenized variables $\hat{\boldsymbol{u}}$ and $\hat{\boldsymbol{B}}$ replace the standard variables in the equilibrium distribution and macroscopic flux evaluations.

This homogenization is governed by the lattice porosity parameter $d(\boldsymbol{x}, t) \in [0, 1]$, which depends on the signed distance $\phi(\boldsymbol{x}, t)$ to the solid surface. Let $\epsilon_h = \epsilon \Delta x$ be the width of the smooth transition region coupled to the grid resolution Δx. The porosity is defined as

$$d(\boldsymbol{x}, t) := \begin{cases} 0 & \text{if } \phi(\boldsymbol{x}, t) \leq -\frac{\epsilon_h}{2} \\ s(\phi(\boldsymbol{x}, t)) & \text{if } \phi(\boldsymbol{x}, t) \in \left(-\frac{\epsilon_h}{2}, \frac{\epsilon_h}{2}\right) \\ 1 & \text{if } \phi(\boldsymbol{x}, t) \geq \frac{\epsilon_h}{2} \end{cases} \tag{17}$$

A common choice for the transition function s is the linear profile:

$$s(\phi) = \frac{\phi}{\epsilon_h} + \frac{1}{2}, \tag{18}$$

such that the exact solid wall is represented by the level set $d(\boldsymbol{x}, t) = \frac{1}{2}$.

By mapping the macroscopic penalization limits directly on to the discrete lattice porosities, with the fully fluid limit corresponding to $d = 1$ and the fully solid limit to $d = 0$, we bypass the use of an empirical physical permeability model. While this sharp, single-cell transition introduces an $\mathcal{O}(\epsilon_h)$ truncation error near the interface, the geometric representation strictly converges to a sharp boundary in the limit $\Delta x \to 0$ due to the direct coupling of the transition width to the spatial resolution.

Finally, for zero-gradient boundary conditions on homogenized boundaries, no special algorithmic measures are required; the state vectors, flux vectors, and LBM distribution functions are computed uniformly across the entire domain.

3 VALIDATION

In this section, the validation of the MHD-VLBM is presented on different established benchmarks, showing the robustness and sufficient accuracy of the built solver. In all tests, the D2Q5 discrete velocity set of Fig. 1 was employed, that is a two-dimensional lattice with five discrete velocities, one stationary and four pointing to the nearest neighbours. As the collision operator, we use BGK and add source terms using the direct injection scheme.

All results reported in this section were computed in single precision on a single NVIDIA RTX A5000. Every discretization parameter of the scheme is specified by a resolution-independent number, the admissible CFL through the sub-characteristic condition of Section 2 and the gradient trackers through the parametrization $\tau_R = g_\tau \Delta t$ (A. Kummerländer et al. 2026b).

Each reference curve is digitized from the published figure identified in Table 2, which records the numerical set-up of the simulation that produced it. For the rotor, J. M. Stone et al. (2008) apply no smoothing at the surface of the disc, whereas the present runs use the linear taper between $r_0 = 0.1$ and $r_1 = 0.115$ described in Section 3.2.

Systematic selection of numerical dissipation and resistivity: In the ideal MHD cases the physical viscosity and resistivity are zero, and a stabilizing diffusivity ν_{num} is added to control the high-wavenumber content generated at discontinuities. It enters through the relaxation time alone: as in Section 2, on D2Q5 the choice of $\tau_{\text{LB},k}$ reproduces ν_{num} exactly, so the two are

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

Table 2. Provenance and numerical set-up of the reference solutions digitized for the benchmarks of Section 3.

Benchmark	Reference	Resolution	Method, $\nabla \cdot \boldsymbol{B}$ control	Reconstruction, flux	CFL, time-step
Brio-Wu shock tube	M. Brio & C. Wu (1988) fig. 2, $t = 80$ ($\equiv 0.2$)	800	2nd-order upwind	Roe, Harten TVD	CFL 0.8, $\Delta t = 0.2$
MHD rotor	J. M. Stone et al. (2008) fig. 26, $t = 0.15$	400^2	unsplit CTU, CT	PPM 3rd-order, Roe	$C_0 \leq 1$ bound, not stated
Resistive Orszag-Tang	P. J. Dellar (2013) fig. 6, $t = 1.0$	2048^2	Fourier pseudospectral, div-free	Hou-Li filter each step	RK4, diffusion-limited
Compressible Orszag-Tang (pressure)	J. M. Stone et al. (2008) fig. 23, $t = 1/2$	512^2	unsplit CTU, CT	PPM 3rd-order, Roe	$C_0 \leq 1$ bound, not stated
Compressible Orszag-Tang (density)	E. Gaburov & K. Nitadori (2011) fig. 11, $t = 1/2$	10^5 particles	meshless, Dedner cleaning	Balsara linear, HLLD	cfl < 1

interchangeable statements of the same choice. The choice itself is made through a single dimensionless numerical diffusivity

$$C = \frac{\nu_{\text{num}}}{\Delta x \, v_{\max}}, \tag{19}$$

with $v_{\max}$ the maximum signal speed that also sets the time-step. The MHD rotor uses $C = 4.1 \times 10^{-2}$ and the compressible Orszag–Tang vortex $C = 1.1 \times 10^{-2}$, held constant under refinement, so that ν_{num} falls in proportion to the cell size: 3.14, 1.57 and 0.785×10^{-4} for the rotor at $N = 500$, 1000 and 2000, and 2.5×10^{-5} down to 6.25×10^{-6} for the Orszag–Tang vortex at $N = 1000$ to 4000. A single dimensionless constant thus fixes the dissipation for every resolution of a given problem, and its value is of order 10^{-2} in both.

Where the target itself carries dissipation, the absolute diffusivity rather than C is held fixed. The Brio-Wu tube keeps $\nu_{\text{num}} = 0.75 \times 10^{-4}$ at all three resolutions, and the resistive Orszag–Tang vortex of Section 3.4 keeps it at five per cent of the physical viscosity $\nu = 10^{-3}$. In those cases the added diffusivity is a fixed fraction of an already resolved physical scale rather than a quantity to be refined away.

The resistivity is not tuned in any case reported here: it is zero in the ideal cases, and in the resistive Orszag–Tang vortex both $\eta_0 = \nu$ and the current-dependent closure $\eta(j) = \eta_0(1 + (j/j_c)^2)$ with $j_c = 25$ are adopted unchanged from P. J. Dellar (2013), whose solution is the reference. For a problem with no such reference, equation (19) is the quantity to specify, and it is the only free constant of the scheme. A more automatic alternative exists within the same framework. A cell-local shock sensor raises the dissipation only where the flow is compressing, and is validated against the Sod and Becker shock tubes and the compressible Taylor-Green vortex by F. Bukreev et al. (2026).

3.1 Brio-Wu shock tube

The Brio-Wu shock tube is a standard 1D problem performed here in a 2D domain; H. Huang et al. (2024) utilizes this to evaluate the scheme's shock-capturing capabilities in compressible ideal MHD. The domain is initialized with two distinct states separated by a discontinuity. Following M. Brio & C. Wu (1988), the left state is given by $(\rho, p, u, v, B_x, B_y) = (1.0, 1.0, 0, 0, 0.75, 1.0)$, and the right state is $(0.125, 0.1, 0, 0, 0.75, -1.0)$. The artificial diffusivity was set constantly to 0.75×10^{-4} m^2 s^{-1}, at a CFL of 0.30. As detailed by H. Huang et al. (2024), the simulation evolves to produce a complex wave structure consisting of a fast rarefaction wave, a slow compound wave, a contact discontinuity, a slow shock, and a fast rarefaction wave. This test validates the robustness of the proposed VLBM in handling strong gradients and assessing non-equilibrium kinetic moments at shock interfaces.

It can be seen in Fig. 2 that the results of different resolutions (2000, 4000, 8000 cells in the domain length), even without separate shock capturing, show good accordance with the reference solution. At finer cell sizes, the oscillations disappear and the variables become smoother. This test shows the validity of the developed solver for accurate reproduction of shock fronts bound to magnetic field in one dimension.

3.2 MHD rotor

The MHD rotor problem (P. Londrillo & L. Del Zanna 2000) tests the propagation of strong torsional Alfvén waves in a compressible medium. The setup consists of a rapidly rotating, dense fluid cylinder embedded in a uniform static background fluid, both permeated by a constant horizontal magnetic field $B_x = 5/\sqrt{4\pi}$ (with $B_y = 0$). Specifically, within a fully periodic 1×1 spatial domain, a rotor of radius $r_0 = 0.1$ with density $\rho = 10.0$ and initial velocity fields $u = -2y/r_0$ and $v = 2x/r_0$ is placed in a resting background medium ($u = v = 0$) with $\rho = 1.0$. The initial thermodynamic pressure is uniformly set to $p = 1.0$ across the domain, and the fluid is modelled as an ideal gas with an adiabatic index of $\gamma = 1.4$. To reduce initial extreme transients, the density and velocity fields are smoothly connected between the core and the background via a linear taper from $r_0 = 0.1$ to $r_1 = 0.115$. E. Gaburov & K. Nitadori (2011) observed that the strong rotational kinetic energy twists the magnetic field lines, launching outward-propagating Alfvén waves that decelerate the rotor while compressing the surrounding fluid.

For this specific configuration, the VLBM simulation is performed at resolutions of 500×500, 1000×1000, and 2000×2000. The simulation runs up to a final time of $t = 0.15$, utilizing a CFL of 0.30 (resulting in a physical time-step of $\Delta t \approx 3.9 \times 10^{-5}$ at the finest resolution). In the simulation, the artificial diffusivity of 3.14×10^{-4}, 1.57×10^{-4}, 0.785×10^{-4} m^2 s^{-1} is used to ensure stability during the robust flow transients.

In Fig. 3, density and local Mach number distributions show smooth shock fronts, matching corresponding figures in the reference work of J. M. Stone et al. (2008). The plots in Fig. 4 show good agreement with the spectral solution, which is accentuated in the subplots. They underline the convergent behaviour of the refinement with acoustic scaling and reducing stabilization diffusivity.

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

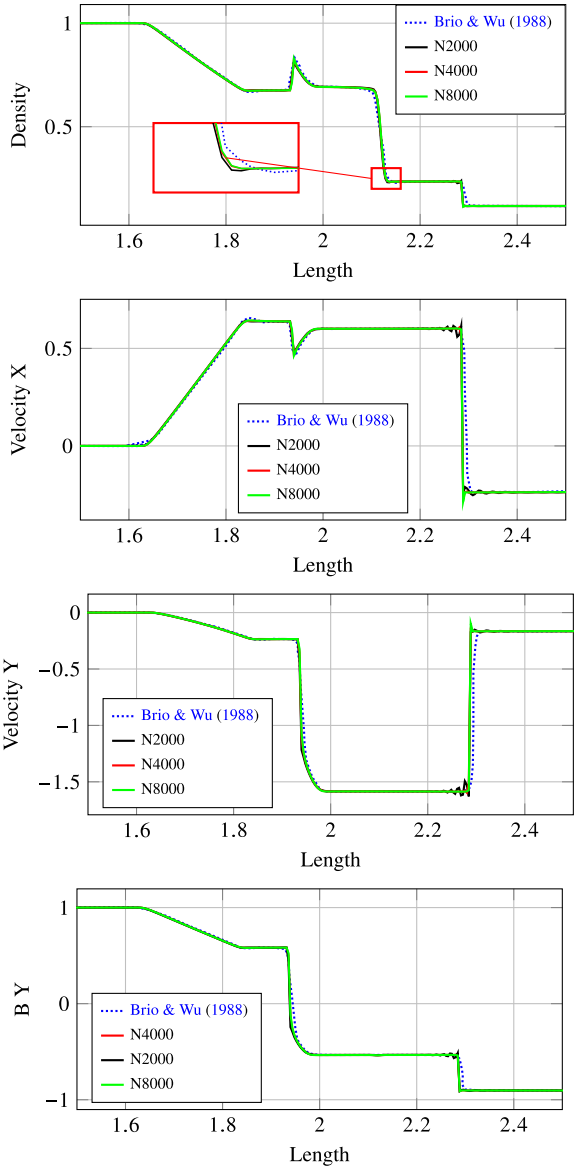


Figure 2. Comparison of reference and simulated data for Brio-Wu benchmark at $t = 0.2$. The reference is digitized from figure 2 of M. Brio & C. Wu (1988), whose abscissa runs over 800 cells of unit width and whose output time $t = 80$ corresponds to $t = 0.2$ in the normalization used here.

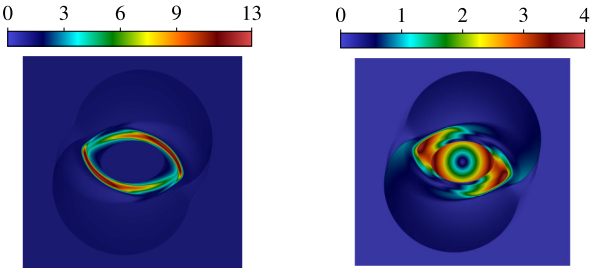


Figure 3. Density and local Mach number distribution at $t = 0.15$ by resolution of 2000. Results can be compared to the reference work of J. M. Stone et al. (2008).

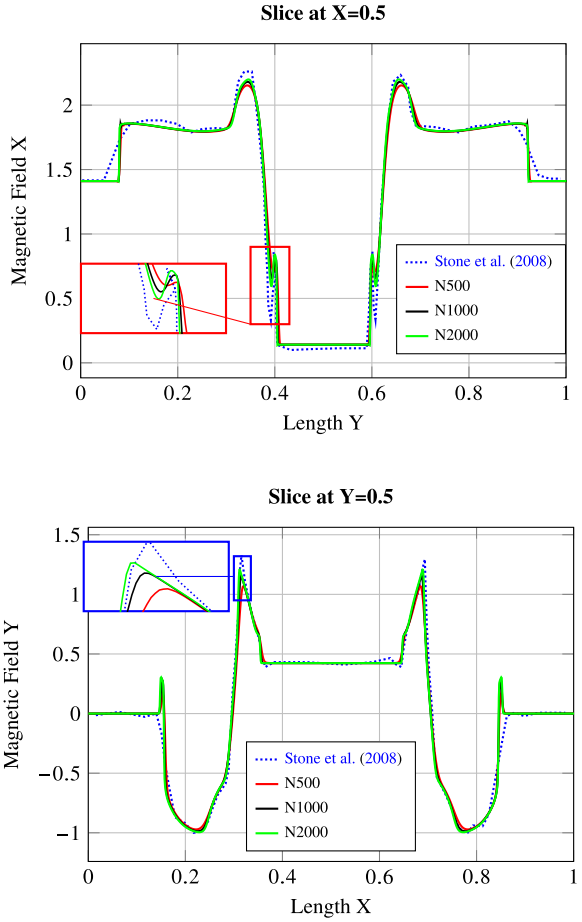


Figure 4. Behaviour of magnetic field X at the $x = 0.5$ slice (top panel) and magnetic field Y in $y = 0.5$ slice (bottom) for the MHD rotor problem at $t = 0.15$.

As we don't apply a separate shock capturer in these validations, at the ends of the domain, the shocks are sharper than the goal reference because of the oscillations occurring there. In future work, we will investigate shock-capturing techniques to handle such situations involving problems with discontinuities.

3.3 Sensitivity to the Courant number

The subcharacteristic condition of Section 2 bounds the CFL by the lattice speed of sound, CFL $\leq c_s/c$, which for the D2Q5 set used here is $1/\sqrt{3}$. First, the bound is imposed on the time-step through an estimate of $v_{\max}$ formed from the initial state, whereas the attained fast magnetosonic speed grows as the flow compresses. The realized CFL therefore exceeds the nominal one, and the usable range is narrower than c_s/c. Secondly, because the relaxation time is set from the imposed diffusivity through the relation of Section 2, the numerical dissipation of the scheme is fixed by that diffusivity alone and does not change with the CFL. The accuracy of a solution should therefore be insensitive to the time-step, and only stability should limit it.

Both expectations are confirmed by a CFL study on the MHD rotor of Section 3.2, repeating the 500×500 configuration at CFL values from CFL $= 0.01$ up to 0.30 and comparing each result against the CFL $= 0.01$ solution. Table 3 lists the outcome. The deviation is flat at approximately 0.12 per cent in relative L_1 over

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

Table 3. Courant-number study for the MHD rotor at 500 × 500, run to $t = 0.15$. Deviations are relative to the CFL = 0.01 solution along the line $x = 0.5$, normalized by the dynamic range of B_x.

CFL	Δt	Steps	L_1 (per cent)	L_∞ (per cent)
0.01	5.21×10^{-6}	28 807	–	–
0.05	2.60×10^{-5}	5761	0.098	0.197
0.10	5.21×10^{-5}	2880	0.110	0.589
0.20	1.04×10^{-4}	1440	0.116	0.907
0.30	1.56×10^{-4}	960	0.118	1.377

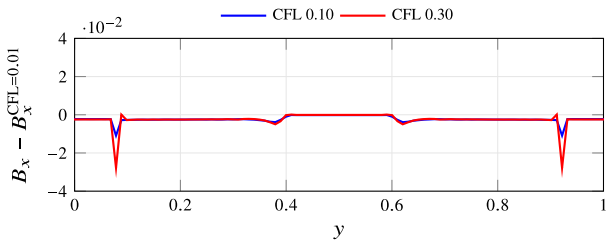


Figure 5. Difference between the MHD rotor solution at raised CFL and the CFL = 0.01 solution, along $x = 0.5$ at $t = 0.15$ and 500 × 500 cells. The differences are confined to the two contact discontinuities near $y = 0.40$ and $y = 0.60$ and remain below 0.035 against a field spanning 2.0. The smooth regions are unaffected.

the whole range, a fifth of the difference between the 500 × 500 and 1000 × 1000 solutions. Fig. 5 shows that what little difference exists is confined to the two contact discontinuities and absent from the smooth regions. Measured against the reference data of Section 3.2 rather than against our own coarse-step solution, the agreement is unchanged to within 0.01 percentage points (4.40 against 4.41 per cent in relative L_1). Raising the time-step from CFL = 0.01 to 0.30 therefore has no measurable effect on the L_1 agreement with the reference. The CFL is limited by stability rather than by accuracy.

The MHD rotor of Section 3.2 is accordingly reported at CFL = 0.30, chosen with a margin below the point at which the simulations lose stability. The remaining cases of this Section are reported at the values stated with each.

3.4 Low compressible resistive Orszag–Tang vortex

The Orszag–Tang vortex is a canonical benchmark for studying the transition to 2D MHD turbulence, characterized by strong velocity-magnetic field coupling and current-sheet formation. To study the incompressible limit, the initial conditions are prescribed via a fluid streamfunction $\phi = 2\cos(x) - 2\sin(y)$ and a magnetic flux function $\psi = 2\cos(x) - \cos(2y)$, which yield the velocity profiles $\boldsymbol{u} = [2\cos(y), -2\sin(x)]$ and magnetic profiles $\boldsymbol{B} = [-2\sin(2y), -2\sin(x)]$. The system is initialized with a uniform density $\rho_0 = 1$ and a sufficiently small Mach number to suppress spurious compressibility effects, matching the setup of P. J. Dellar (2013). The evolution is tracked using a kinematic viscosity of $\nu = 10^{-3}$ (corresponding to an effective Reynolds number of $Re = 4000\pi$), monitoring the peak electric current density and vorticity to evaluate the behaviour of a current-dependent resistivity model $\eta(j) = \eta_0(1 + (j/j_c)^2)$ as the system generates complex, fine-scale turbulent structures. η_0 is set equal to the kinematic viscosity, and $j_c = 25$. To match the results of P. J. Dellar (2013), the stabilization diffusivity was set to 5 per cent of the physical viscosity, at a CFL of 0.30.

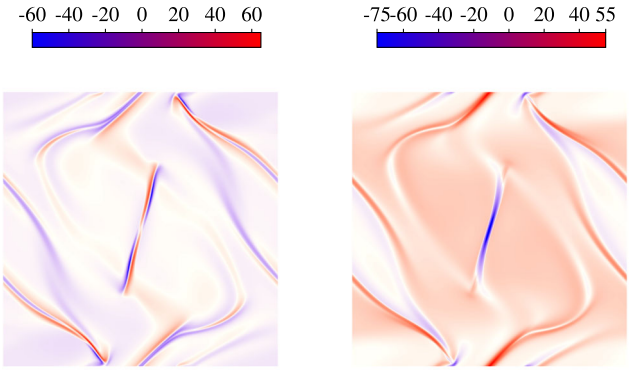


Figure 6. Vorticity and electric current density at a resolution of 2048 cells in the side length at $t = 1$ by incompressible Orszag–Tang vortex. Results can be compared to P. J. Dellar (2013).

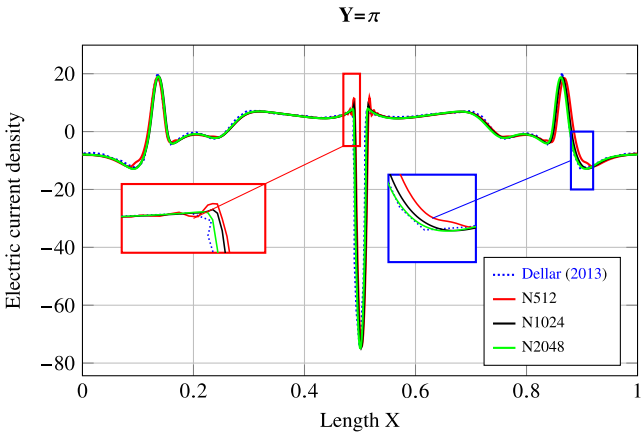


Figure 7. Electric current density j along the horizontal line at the height of $y = \pi$ compared to the spectral simulation taken from P. J. Dellar (2013).

The visual plots (Fig. 6) match those in the referenced publication and show smooth vorticity and electric current density distributions in the domain. Fig. 7 shows how the values are converging towards the spectral simulation values through acoustic spatiotemporal scaling with constant stabilization diffusion. It can be seen that the numerical diffusion is reducing with the refinement of the cell size, and the finest resolution matches the reference spectral values as well as the standard LBM simulation results of P. J. Dellar (2013).

3.5 Compressible non-resistive Orszag–Tang vortex

To evaluate the scheme's ability to capture complex compressible dynamics and multiple shock formations, the adiabatic non-resistive Orszag–Tang vortex for ideal MHD is simulated following the standard validation setup described by J. M. Stone et al. (2008) and E. Gaburov & K. Nitadori (2011). The problem is initialized on a periodic domain of $[0, 1] \times [0, 1]$ with an adiabatic equation of state ($\gamma = 5/3$). The initial density and pressure are set uniformly to $\rho = \frac{25}{36\pi}$ and $p = \frac{5}{12\pi}$, yielding a nominal sound speed of $a = 1.0$.

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

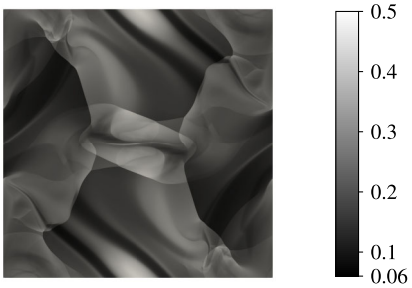


Figure 8. Density distribution in the adiabatic Orszag–Tang vortex benchmark by the resolution of 4000 cells in the side length taken at $t = 0.5$, can be compared to J. M. Stone et al. (2008) and E. Gaburov & K. Nitadori (2011).

The velocity and magnetic field profiles are defined as $\boldsymbol{u} = [-\sin(2\pi y), \sin(2\pi x)]$, $\boldsymbol{B} = [-B_0 \sin(2\pi y), B_0 \sin(4\pi x)]$ with $B_0 = 1/\sqrt{4\pi}$. The stabilization diffusion is set to 2.5×10^{-5} at a resolution of 1000 cells per side length and is decreased linearly together with the cell size refinement. The scaling is performed acoustically by treating a constant CFL of 0.20.

In this compressible regime, the interacting vortices rapidly evolve to produce a system of complex shock waves and strong local current sheets. This configuration serves as a rigorous benchmark for ideal MHD solvers, specifically testing their capacity to sharply resolve fine-scale structures, magnetic reconnection events, and discontinuities. The resulting density distribution is shown in Fig. 8 and the pressure values are provided in Fig. 9.

The results show good agreement with the reference values, whereby small oscillations occur at the local shocks at the coarsest resolution due to the chosen stabilization diffusion. At finer cell sizes, these oscillations disappear, whereby attempts to decrease the stabilization with a higher factor than the cell size refinement factor lead to the divergence of the simulations. In the upper pressure plot in the zoomed snippet of the curve, it can be noted that the resolutions of 2000 and 4000 show lower numerical diffusion than the reference simulation.

4 EFFICIENT IMPLEMENTATION

The unique amenability of classical LBMs to efficient parallel execution directly translates to our present approach. Realizing this potential on diverse target hardware requires an implementation capable of transparently abstracting platform-specific memory and execution hierarchies. The present approach is implemented within the open-source multi-physics framework OpenLB (M. J. Krause et al. 2021), which utilizes a hardware-agnostic C++ abstraction layer to achieve performance portability across SIMD CPUs and diverse GPU architectures (A. Kummerländer et al. 2023, 2026a). By expressing models against abstract cell-level concepts, this design enables compile-time binding to platform-specific data structures.

To bridge the gap between abstract mathematical implementation in C++ and high hardware efficiency, OpenLB employs an automatic optimization pipeline based on *Common Subexpression Elimination* (CSE). As detailed by A. Kummerländer et al. (2026a), the framework repurposes its differentiable architecture to extract the underlying mathematical expression tree which is then automatically processed in the SymPy computer algebra system developed by A. Meurer et al. (2017) in order to minimize the number of arithmetic operations.

Table 4 and Fig. 10 demonstrate the impact of this step on an exemplary NVIDIA A5000 GPU. For the 2D configuration,

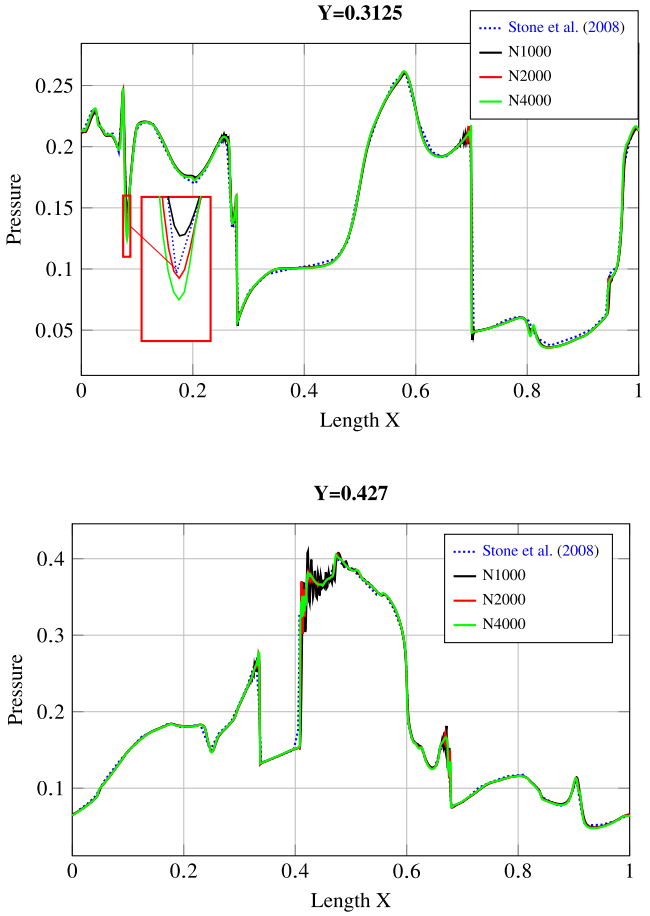


Figure 9. Validation of the gas pressure values along the horizontal lines at the heights of $y = 0.3125$ and $y = 0.427$ in the adiabatic Orszag–Tang vortex benchmark taken at $t = 0.5$.

Table 4. Performance and roofline analysis of the validated 2D and 3D homogenized MHD collision kernels on a NVIDIA RTX A5000 GPU. Compared to measured peak single-precision memory bandwidth 680 GB/s (BabelStream). Throughput is reported in MLUPs, million lattice updates per second, one update being one cell of the computational grid.

	2D		3D Homogenized	
Metric	**Unopt.**	**Opt.**	**Unopt.**	**Opt.**
Arithmetic complexity (FLOPs)	1066	978	1387	1117
Operation reduction (per cent)	–	8.26 per cent	–	19.47 per cent
Memory bandwidth (Bytes/cell)	624	624	608	608
Arithmetic intensity (FLOP/Byte)	1.708	1.567	2.281	1.837
Throughput (MLUPs)	711 ± 8	$1,077 \pm 10$	754 ± 3	847 ± 4
Performance (GFLOPs)	757.5	1053.7	1046.2	946.4
Roofline efficiency (per cent)	65.2 per cent	98.9 per cent	67.4 per cent	75.8 per cent

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

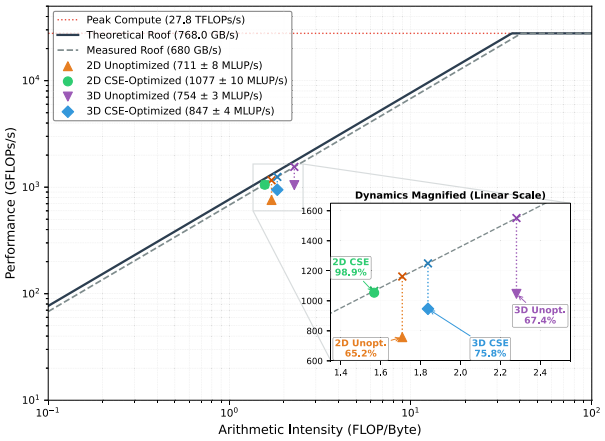


Figure 10. Roofline analysis of both MHD dynamics on NVIDIA A5000 GPU.

the CSE pipeline reduces the arithmetic complexity from 1066 to 978 FLOPs per cell (an 8.26 per cent reduction), shifting the arithmetic intensity from 1.708 to 1.567 FLOP/Byte at a constant memory bandwidth requirement of 624 Bytes/cell. Consequently, while the unoptimized 2D kernel achieves 65.2 per cent of the maximum roofline performance (711 ± 8 million lattice updates per second (MLUPs), one update being one cell of the computational grid), the optimized kernel peaks at 98.9 per cent hardware efficiency relative to the roofline peak, processing at 1077 ± 10 MLUPs. Similarly, the 3D homogenized MHD model sees a 19.47 per cent reduction in operations (1387 down to 1117 FLOPs), reducing the arithmetic intensity from 2.281 to 1.837 FLOP/Byte. Even when resolving the complex, moving porosity fields representing the solid asteroid geometry at 608 Bytes/cell, the CSE optimized 3D model scales throughput from 754 ± 3 to 847 ± 4 MLUPs, achieving a roofline efficiency of 75.8 per cent.

The two-dimensional kernel measured here is the extended MHD collision, which carries the tracked gradients and the resistive fluxes throughout. The ideal benchmarks of Section 3 run it with the magnetic diffusivity and viscosity set to zero rather than through a separate reduced kernel, so the throughput reported above is that of the resistive discretization.

4.1 Computational cost relative to finite-volume solvers

A throughput comparison against established astrophysical codes is meaningful only once the differing hardware, arithmetic precision, physics and accuracy targets are accounted for. We therefore confine the quantitative statements below to measurements made on the present implementation, and treat the comparison with finite-volume methods structurally rather than numerically. At 624 Bytes per cell update the optimized 2D kernel sustains 672 GB s^{-1} of the 680 GB s^{-1} attainable on this device, and the 3D kernel 515 GB s^{-1} at 608 Bytes per cell. Because every operator is point-local and the update is single-stage, this traffic is the state itself, with no multi-cell reconstruction stencil and no intermediate Runge-Kutta or corner-transport stage that a high-order finite-volume MHD code writes and re-reads (U.-H. Zhang, H.-Y. Schive & T. Chiueh 2018; P. Grete, F. W. Glines & B. W. O'Shea 2021). The subdomain overlap is likewise one cell deep, against the four ghost cells of a third-order finite-volume reconstruction (J. M. Stone et al. 2008).

Set against this, the present scheme is bounded at CFL $\leq 1/\sqrt{3}$ in two dimensions, whereas the finite-volume integrator of J. M. Stone et al. (2008) is stable up to a CFL of one there, so a lattice update advances less physical time per step than one of that integrator. This ceiling is a property of that integrator rather than of finite-volume schemes in general, whose stable Courant numbers vary with reconstruction order and the degree of dimensional coupling. A fair ranking would require both codes on the same device, in the same precision and to the same error tolerance, which is challenging and left to future work. Such comparisons have been made for this LBM framework in other demanding applications, an IC-engine flow and an agricultural airflow benchmarked against finite-volume solvers (M. Haussmann et al. 2020; B. Moreau et al. 2026). No speed-up factor is claimed here.

5 SOLAR WIND INTERACTION WITH A MAGNETIZED ASTEROID

We close with a fully three-dimensional plasma-body application: the interaction of a highly magnetized early solar wind with a rotating metallic asteroid modelled after 16 Psyche, using observational data from M. K. Shepard et al. (2021) and NASA Science (2025). Such solar wind interactions with planetary bodies are of significant interest in computational astrophysics, J. Müller et al. (2011) and L. Mejnertsen et al. (2018) have traditionally addressed them via global MHD or adaptive hybrid kinetic schemes.

This specific scenario was selected as a challenging computational stress test. Early stellar environments display extreme plasma conditions characterized by high-Mach coronal mass ejections and enhanced magnetic flux densities. W. Exner et al. (2018) and C. Koenders et al. (2015) demonstrate that these highly energetic plasma flows drive discontinuities and complex spatial features, testing the numerical stability of our approach.

By embedding a rotating, strongly magnetized solid body within a high-Mach compressible flow using our established homogenized fluid-structure interaction (FSI) approach developed by M. J. Krause et al. (2017), A. Kummerländer et al. (2026d), and A. Kummerländer et al. (2026f), the setup encompasses moving boundary conditions, localized dipole fields, and moving shock fronts simultaneously. The basic setup is illustrated by the schematic overview in Fig. 11.

To mimic complex crustal remanent features on the metallic asteroid, an analytical ensemble of four independent, shifted, and smoothed magnetic dipoles is embedded inside the reference lattice onto which the Psyche STL geometry provided by NASA Science (2025) is voxelized. This arrangement scales to a local peak surface field intensity of 225 nT. Furthermore, the body is rotated around a normalized arbitrary tilt vector ($\vec{r} = [1, 0, 1]/\sqrt{2}$) with an artificially accelerated angular frequency. This tumbling behaviour continuously twists the frozen-in *interplanetary magnetic field* (IMF) lines (initially oriented along the $-Y$ direction at 5 nT), forcing local magnetic reconnection events, dynamic bow shock structures, and localized Alfvén transients. The resulting complex plasma topologies and field line structures for a 100 million cell simulation executed on a dual NVIDIA A5000 system are visualized in Fig. 12.

The inflow boundary (X_{min}) is formulated via an equilibrium projection operator to enforce the continuous, supersonic, and moderately compressible early solar wind fluid state defined in Table 5. The outflow and lateral boundaries (X_{max}, $Y_{\text{min/max}}$, $Z_{\text{min/max}}$) are governed by a directed MHD zero-gradient bound-

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

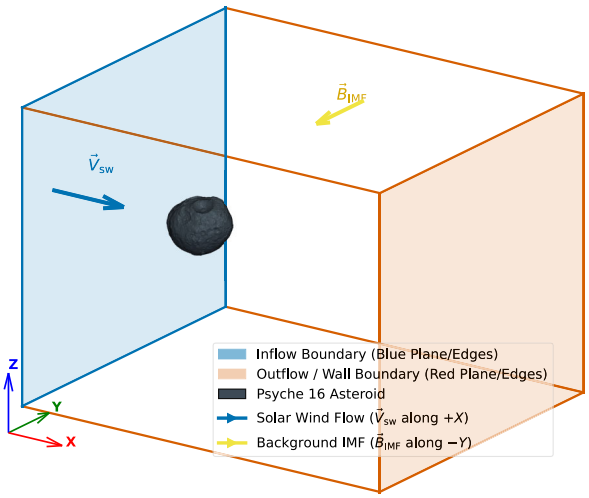
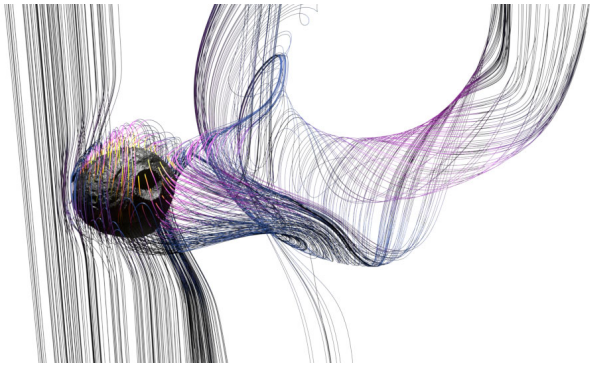
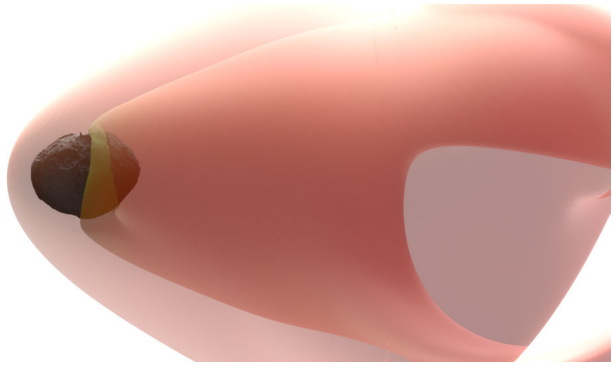


Figure 11. Schematic overview of the 3D computational domain for simulating the interaction of an early solar wind plasma with a magnetized, rotating asteroid. The blue face indicates the supersonic inflow plane ($X_{\min}$), orange faces denote the open boundaries ($X_{\max}$, $Y_{\min/\max}$, $Z_{\min/\max}$), and the solid central structure represents the homogenized discretization of the asteroid body.



(a) Magnetic field line streamlines showing the complex interaction with internal crustal remanent fields.



(b) Normalised plasma density contours capturing the bow shock and wake structures.

Figure 12. Illustrative visualization and rendering of the early solar wind interaction with a magnetized, tumbling asteroid modelled after 16 Psyche (NASA Science 2025). The supersonic (Mach 3.5) plasma inflow interacts with the internal crustal remanent fields to produce: (I) a distinct, high-density bow shock and upstream compression region where the frozen-in interplanetary magnetic field (IMF) lines pile up (II) a magnetosheath of twisted flux ropes deflected around the magnetic obstacle (III) an asymmetric, turbulent downstream magnetotail and low-density wake cavity showing magnetic reconnection events. Rendered in Blender using J. Marín et al. (2025).

Table 5. Simulation parameters for the 16 Psyche early solar wind interaction.

Parameter	Value
Asteroid diameter	220 km
Solar wind velocity	400 km s^{-1}
Solar wind density	8.35×10^{-21} kg m^{-3}
Interplanetary magnetic field (IMF)	5 nT
Maximum crustal magnetic field	225 nT
Simulation Mach number	3.5
CFL	0.2
Lattice relaxation time	0.51
Cell count	100×10^6
Lattice velocities set	D3Q7
Collision operator	BGK
Source scheme	Direct

ary condition to allow the plasma flow and transient features to exit the domain without unphysical reflections. The crustal field maxima drive this configuration into the magnetically dominated regime, well below the plasma beta reached by any of the benchmarks of Section 3. The scheme is kept robust there by bound-preserving safeguards on the recovered state. Like every other operator in the framework, these act on a single cell and therefore leave the strictly local structure of the update intact. The same principle underlies the local shock-capturing treatment of F. Bukreev et al. (2026). Their role and their limitations are discussed in Section 6.

Due to the translation of the porosity-based HLBM approach to MHD, OpenLB's established FSI module (A. Kummerländer et al. 2026d, f) can be directly utilized both to efficiently project the rotating geometry and its remnant magnetic field into the fluid lattice, and to integrate various quantities of interest at every time-step. For the present setup we obtained per-timestep surface integrals of:

(i) Net magnetic flux (Wb) passing through the asteroid.
(ii) Forces (N), comprising both aerodynamic and Maxwell electromagnetic forces.
(iii) Torques (N · m) acting on the rotating asteroid.
(iv) Electromagnetic work (J) transferred to/from the solar wind.

While this configuration is designed as an algorithmic capability and high-performance computing showcase rather than a phenomenologically validated planetary study, it, together with the previous rigorous benchmarks, demonstrates the viability of our strictly local scheme for such applications.

6 DISCUSSION AND CONCLUSION

6.1 Astrophysical extensions and regimes

6.1.1 Extension to relativistic MHD

Because the framework is formulated around generic hyperbolic conservation laws $\partial_t \boldsymbol{Q} + \nabla \cdot \boldsymbol{\Phi} = \boldsymbol{S}$, it extends naturally to Special and General Relativistic Magnetohydrodynamics (SRMHD/GRMHD) (S. S. Komissarov 1999; C. F. Gammie, J. C. McKinney & G. Tóth 2003; L. Del Zanna et al. 2007). In units with $c = 1$, with Lorentz factor Γ, specific enthalpy h and laboratory-frame field $\boldsymbol{B}$, the conserved state $\boldsymbol{Q} = (D, \boldsymbol{S}, E, \boldsymbol{B})^T$ has $D = \Gamma\rho$ and

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

$$\boldsymbol{S} = \left(\rho h \Gamma^2 + B^2\right) \boldsymbol{v} - (\boldsymbol{v} \cdot \boldsymbol{B}) \boldsymbol{B}, \tag{20}$$
$$E = \rho h \Gamma^2 - p + \tfrac{1}{2} B^2 + \tfrac{1}{2} \left[v^2 B^2 - (\boldsymbol{v} \cdot \boldsymbol{B})^2 \right]. \tag{21}$$

The spatial fluxes are algebraic in the primitive variables, so they map directly into the VLBM equilibria and all collision steps remain point-local. The same construction already runs the ideal ultrarelativistic fluid (A. Kummerländer et al. 2026b), so the kinetic formulation is not restricted to Newtonian regimes.

The relativistic case is not pursued here. The primitive-variable recovery is no longer closed-form and needs a non-linear root find in every cell, which stays point-local but changes the cost of the update. The physical signal speeds approach the speed of light, so with c a hard ceiling the subcharacteristic condition of Section 2 tightens the admissible CFL. And positivity of the recovered pressure and of Γ at high Lorentz factor is not guaranteed by the linear equilibrium, requiring the bound-preserving treatment discussed below.

6.1.2 *Applicability in low-β plasma regimes*

Astrophysical environments such as the solar corona and pulsar magnetospheres reach extremely low plasma beta ($\beta = 2p/|\boldsymbol{B}|^2 \ll 1$), where the internal energy is a small difference of large numbers and truncation error can drive the recovered pressure negative. The benchmarks of Section 3 already reach the magnetically dominated range $\beta < 1$, with a lowest attained β of 0.081 (Brio-Wu), 0.125 (rotor) and 0.178 (compressible Orszag-Tang) and positive recovered pressure at every cell without any floor. The asteroid demonstration of Section 5 reaches a local $\beta = \mathcal{O}(10^{-3})$ near the crustal field maxima and remains stable, but it carries explicit cell-local safeguards, a pressure floor, a density floor and a velocity cap, which is one reason it is presented as a demonstration rather than a validated study.

Reaching $\beta \ll 1$ as a validated regime would require positivity preservation beyond these floors. The local structure of the scheme accommodates this, since such treatments act cell by cell, for example the shock sensor of F. Bukreev et al. (2026) or the bound-preserving limiters for vectorial LBM of G. Wissocq et al. (2025).

6.2 Conclusion

In this work, we introduced a novel, fully local, PDE agnostic LBM framework for solving generic systems of conservation laws. We applied it to compressible ideal and incompressible resistive MHD, the regime underlying shock-dominated and reconnecting astrophysical plasma flows. By expanding the macroscopic state vector to transport spatial gradients and viscous stresses as independent variables, the scheme natively recovers the required derivatives while completely avoiding non-local finite-difference operations. In the context of MHD, the same local machinery supplies the magnetic divergence for an eight-wave cleaning correction, so that divergence control remains a point-local source term rather than a global solve. Validation across established astrophysical MHD benchmarks, including the Brio-Wu shock tube, the MHD rotor, and the Orszag–Tang vortex, confirms the framework's high accuracy, stability, and robust discontinuity resolution. Furthermore, implementation within the OpenLB framework, augmented by automated common subexpression elimination, translates this localized algorithmic structure into exceptional computational efficiency, achieving up to 98.9 per cent of the hardware roofline limit on a modern GPU architecture. Finally, a three-dimensional application to the interaction of an early solar wind with a tumbling, magnetized asteroid, resolving the bow shock, magnetosheath, and reconnecting wake together with the moving body, shows that the approach extends naturally to large-scale fluid-structure interaction, establishing this PDE-agnostic scheme as a scalable and robust tool for astrophysical MHD.

A natural next step is to reconstruct the transported gradients from the non-equilibrium part of the distribution, which removes the tracker components from the state vector. This is developed for general conservation laws in A. Kummerländer et al. (2026c) and for the compressible Navier-Stokes-Fourier system in A. Kummerländer et al. (2026e), and its application to the MHD system studied here is left to future work.

ACKNOWLEDGEMENTS

During the preparation of this work, the authors used Google Gemini and Anthropic Claude for code generation and text-formatting assistance. The final manuscript was reviewed, edited, and approved by all authors, who take full responsibility for its content.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

FUNDING

This research received no specific grant from any public, commercial, or not-for-profit funding agency.

DATA AVAILABILITY

The numerical method presented in this article will be made publicly available in a future release of the open-source OpenLB library (M. J. Krause et al. 2021; https://www.openlb.net). Ahead of that release, the implementation can be shared with interested researchers upon reasonable request, within the scope of joint research projects. The simulation data underlying this article are available from the corresponding author upon reasonable request.

REFERENCES

Anandan M., Raghurama Rao S., 2024, Comput. Fluids, 280, 106348
Angot P., Bruneau C.-H., Fabrie P., 1999, Num. Math., 81, 497
Aregba-Driollet D., Natalini R., 2000, SIAM J. Num. Anal., 37, 1973
Bhatnagar P. L., Gross E. P., Krook M., 1954, Phys. Rev., 94, 511
Boolakee O., Geier M., De Lorenzis L., 2025, Comput. Methods Appl. Mech. Eng., 433, 117469
Bouchut F., 1999, J. Stat. Phys., 95, 113
Brio M., Wu C., 1988, J. Comput. Phys., 75, 400
Bukreev F., Kummerländer A., Krause M. J., 2026, preprint (arXiv:e-prints)
Chai Z., Zhao T. S., 2013, Phys. Rev. E, 87, 063309
Del Zanna L., Zanotti O., Bucciantini N., Londrillo P., 2007, A&A, 473, 11
Dellar P. J., 2013, J. Comput. Phys., 237, 115
Dubois F., 2014, Int. J. Mod. Phys. C, 25, 1441014
Exner W., Heyner D., Liuzzo L., Motschmann U., Shiota D., Kusano K., Shibayama T., 2018, Planet. Space Sci., 153, 89
Gaburov E., Nitadori K., 2011, MNRAS, 414, 129

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

Gammie C. F., McKinney J. C., Tóth G., 2003, ApJ, 589, 444
Ginzburg I., Verhaeghe F., d'Humières D., 2008, Commun. Comput. Phys., 3, 427
Grad H., 1949, Commun. Pure Appl. Math., 2, 331
Grete P., Glines F. W., O'Shea B. W., 2021, IEEE Trans. Parallel Distr. Syst., 32, 85
Guillon K., Hélie R., Helluy P., 2024, ESAIM: Proc. Surv., 77, 46
Guo Z., Zheng C., Shi B., 2002, Phys. Rev. E, 65, 046308
Haussmann M. et al., 2020, Computation, 8, 43
Huang H., Jin K., Li K., Zheng X., 2024, Phys. Plasmas, 31, 093902
Jin S., Xin Z., 1995, Commun. Pure Appl. Math., 48, 235
Koenders C., Glassmeier K. H., Richter I., Ranocha H., Motschmann U., 2015, Planet. Space Sci., 105, 101
Komissarov S. S., 1999, MNRAS, 303, 343
Krause M. J., Klemens F., Henn T., Trunk R., Nirschl H., 2017, Particuology, 34, 1
Krause M. J. et al., 2021, Comput. Math. Appl., 81, 258
Krueger T., Kusumaatmaja H., Kuzmin A., Shardt O., Silva G., Viggen E., 2017, The Lattice Boltzmann Method: Principles and Practice. Graduate Texts in Physics. Springer, Cham, Switzerland
Kummerländer A., Dorn M., Frank M., Krause M. J., 2023, Concurr. Comput.: Pract. Exp., 35, e7509
Kummerländer A., Bukreev F., Krause M. J., 2026b, preprint (arXiv:e-prints)
Kummerlander A., Ito S., Schecher M., Dapelo D., Simonis S., Krause M. J., Bukreev F., 2026d, Int. J. Num. Methods Heat Fluid Flow, 36, 2649
Kummerlander A., Bukreev F., Krause M. J., 2026c, preprint (arXiv:e-prints)
Kummerlander A., Bukreev F., Krause M. J., 2026e, preprint (arXiv:e-prints)
Kummerlander A., Bukreev F., Dorneles L., Dorn M., Ito S., Krause M. J., 2026a, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC26). Chicago, USA
Kummerlander A., Tur B., Haase M., Bukreev F., Döllinger M., Krause M. J., Kniesburges S., 2026f, Comput. Methods Appl. Mech. Eng., 457, 119009
Latt J., Chopard B., 2006, Math. Comput. Simul., 72, 165
Liu C., Xu K., 2017, Commun. Comput. Phys., 22, 1175
Londrillo P., Del Zanna L., 2000, ApJ, 530, 508
Marín J., Baptiste T. M., Rodero C., Williams S. E., Niederer S. A., García-Fernández I., 2025, Comput. Graph., 130, 104264
Mejnertsen L., Eastwood J. P., Hietala H., Schwartz S. J., Chittenden J. P., 2018, J. Geophys. Res.: Space Phys., 123, 259
Meurer A. et al., 2017, PeerJ Comput. Sci., 3, e103
Moreau B., Caiazzo A., Amon T., Chen L., Janke D., 2026, Results Eng., 32, 112822
Müller J., Simon S., Motschmann U., Schüle J., Glassmeier K.-H., Pringle G. J., 2011, Comput. Phys. Commun., 182, 946
NASA Science, 2025, Asteroid Psyche 3D Model and Geometry, Washington DC, USA, https://science.nasa.gov/solar-system/asteroids/16-psyche/ (5th of May 2026)
Powell K. G., Roe P. L., Linde T. J., Gombosi T. I., De Zeeuw D. L., 1999, J. Comput. Phys., 154, 284
Schekochihin A. A., Cowley S. C., 2007, in Molokov S., Moreau R., Moffatt H. K., eds, Fluid Mechanics and Its Applications, Vol. 80, Magnetohydrodynamics: Historical Evolution and Trends. Springer, Dordrecht, Netherlands, p. 85
Shepard M. K. et al., 2021, Planet. Sci. J., 2, 125
Stone J. M., 2011, Bull. Astron. Soc. India, 39, 129
Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, ApJS, 178, 137
Strang G., 1968, SIAM J. Num. Anal., 5, 506
Struchtrup H., Torrilhon M., 2003, Phys. Fluids, 15, 2668
Unat D. et al., 2017, IEEE Trans. Parallel Distr. Syst., 28, 3007
Williams S., Waterman A., Patterson D., 2009, Commun. ACM, 52, 65
Wissocq G., Liu Y., Abgrall R., 2025, SIAM J. Sci. Comput., 47, A3276
Zhang U.-H., Schive H.-Y., Chiueh T., 2018, ApJS, 236, 50
d'Humières D., 2002, Phil. Trans. R. Soc. London. Ser. A: Math. Phys. Eng. Sci., 360, 437

This paper has been typeset from a TEX/LATEX file prepared by the author.

Downloaded from academic.oup.com/mnras/article/553/1/stag1979/8943796 by Inst F Entscheidungstheorie and Unternehmensforschung user on 10 November 2026

© The Author(s) 2026.
Published by Oxford University Press on behalf of Royal Astronomical Society. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.