

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/179652/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Talbot, Rosie Y., Pakmor, Rüdiger, Pfrommer, Christoph, Springel, Volker, Werhahn, Maria, Bieri, Rebekka and van de Voort, Freeke 2025. Anisotropic thermal conduction on a moving mesh for cosmological simulations. Monthly Notices of the Royal Astronomical Society , staf1073. 10.1093/mnras/staf1073

Publishers page: <https://doi.org/10.1093/mnras/staf1073>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Anisotropic thermal conduction on a moving mesh for cosmological simulations

Rosie Y. Talbot¹^{*}, Rüdiger Pakmor¹, Christoph Pfrommer², Volker Springel¹, Maria Werhahn¹,
Rebekka Bieri³, Freeke van de Voort⁴

¹Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany

²Leibniz-Institute für Astrophysik Potsdam (AIP), An der Sternwarte 16, 14482 Potsdam, Germany

³Department of Astrophysics, University of Zurich, 8057 Zurich, Switzerland

⁴Cardiff Hub for Astrophysics Research and Technology, School of Physics and Astronomy, Cardiff University, Queen's Buildings, Cardiff CF24 3AA, UK

Accepted 2025 June 23. Received 2025 June 23; in original form 2024 October 9

ABSTRACT

In weakly collisional, strongly magnetised plasmas such as the intracluster medium (ICM), hot accretion flows and the solar corona, the transport of heat and momentum occurs primarily along magnetic field lines. In this paper we present a new scheme for modelling anisotropic thermal conduction which we have implemented in the moving mesh code AREPO. Our implementation uses a semi-implicit time integration scheme which works accurately and efficiently with individual timestepping, making the scheme highly suitable for use in cosmological simulations. We apply the scheme to a number of test-problems including the diffusion of a hot patch of gas in a circular magnetic field, the progression of a point explosion in the presence of thermal conduction, and the evolution and saturation of buoyancy instabilities in anisotropically conducting plasmas. We use these idealised tests to demonstrate the accuracy and stability of the solver and highlight the ways in which anisotropic conduction can fundamentally change the behaviour of the system. Finally, we demonstrate the solver's capability when applied to highly non-linear problems with deep timestep hierarchies by performing high-resolution cosmological zoom-in simulations of a galaxy cluster with conduction. We show that anisotropic thermal conduction can have a significant impact on the temperature distribution of the ICM and that whistler suppression may be relevant on cluster scales. The new scheme is, therefore, well suited for future work which will explore the role of anisotropic thermal conduction in a range of astrophysical contexts including the ICM of clusters and the circumgalactic medium of galaxies.

Key words: conduction – methods: numerical – plasmas – magnetic fields – galaxies: clusters: general – galaxies: clusters: intracluster medium

1 INTRODUCTION

Many astrophysical systems, such as hot accretion flows, the intracluster medium (ICM) of galaxy clusters, the solar wind and some phases of the interstellar medium (ISM), are weakly collisional and strongly magnetised. Such plasmas are characterised by a hierarchy of scales whereby the electron and ion gyroradii are much smaller than the electron mean free path of ion Coulomb collisions, which itself is smaller than characteristic lengthscale associated with the system.

As a result of this scale ordering, charged particles gyrate around magnetic field lines much faster than the rate at which they undergo Coulomb scattering. This effectively ties charged particles to field lines, meaning that the transport of heat and momentum occurs preferentially in the direction of the magnetic field and only gradients oriented along the magnetic field can be relaxed. Such systems are often well described by Braginskii magnetohydrodynamics (MHD, Braginskii 1965) which differs from ideal MHD via the addition of anisotropic heat conduction and viscosity.

Anisotropic heat conduction is thought to be important in a wide range of astrophysical processes in diverse contexts. For example: In determining the structure of the hot plasma in the solar corona (see e.g. Yokoyama & Shibata 1997; Bingert & Peter 2011; Bourdin et al. 2013; Ye et al. 2020; Navarro et al. 2022) and in regulating energy transport in supernova remnants (see e.g. Chevalier 1975; Tilley et al. 2006; Balsara et al. 2008a,b). Thermal conduction also affects the stability properties of the plasma in the formation and dynamics of multiphase structure in the ICM, the circumgalactic medium (CGM) and ISM (see e.g. Sharma et al. 2010a,b; Choi & Stone 2012; Brüggén & Scannapieco 2016; Armillotta et al. 2017; Jennings & Li 2021; Brüggén et al. 2023).

Thermal conduction is thought to play a significant role in shaping the properties and dynamics of the ICM. It may, for example, be important for distributing energy from active galactic nuclei (AGN) and offsetting radiative cooling in cool cores of galaxy clusters (Zakamska & Narayan 2003; Voit 2011; Voit et al. 2015; Yang & Reynolds 2016; Jacob & Pfrommer 2017a,b). Anisotropic transport may also be relevant for explaining features such as cold fronts in the ICM (Markevitch & Vikhlinin 2007; ZuHone et al. 2015) which exhibit temperature changes on scales of order (or smaller than) the electron

* E-mail: rosie@mpa-garching.mpg.de (RYT)

mean free path. Simulations find that magnetic fields tend to drape around cold fronts (Lyutikov 2006; Asai et al. 2007; Dursi & Pfrommer 2008; Pfrommer & Dursi 2010) which act to insulate them from the surroundings and potentially protect their coherence by suppressing mixing (Fabian et al. 2003; Dong & Stone 2009; ZuHone et al. 2015).

Additionally, when heat transport occurs predominantly along the magnetic field, the stability properties of the plasma change markedly. The hydrodynamic stability to convection is no longer determined by the sign of the entropy gradient (Schwarzschild 1958) but rather by that of the temperature gradient, in combination with topology of the magnetic field. The resulting instabilities, driven by fast heat conduction, have the ability to affect the outer regions of galaxy clusters (the magneto-thermal instability, MTI; Balbus 2000, 2001) and also the innermost regions of cool-core clusters (the heat-flux-driven buoyancy instability, HBI; Quataert 2008). The MTI generates a turbulent state driven by the background temperature gradient, consisting of density and velocity fluctuations across a broad range of scales. At saturation, the root mean square values of these fluctuations follow distinct power-law relationships with thermal diffusivity and the gradients of gas entropy and temperature (Perrone & Latter 2022a,b).

Weakly collisional plasmas where the ratio of thermal-to-magnetic pressure is large ($\beta \sim 100$) are, however, susceptible to kinetic instabilities, driven by pressure anisotropies and heat fluxes (Schekochihin et al. 2005; Kunz et al. 2014; Komarov et al. 2016). These micro-instabilities act to enhance the scattering of charged particles, potentially leading to suppression of the conductivity below the collisional (Spitzer 1962) value (Riquelme et al. 2016; Berlok et al. 2021). Alternatively, conductivity can be suppressed below the classical Spitzer value in a weakly collisional plasma as a result of streaming electrons, which resonantly excite whistler waves. These waves frequently scatter electrons, causing a near-isotropisation of the electron distribution in the whistler wave frame. As a result, the mean electron transport speed along the magnetic field is reduced from its thermal value v_{th} to the whistler phase speed $\sim v_{th}/\beta_e$, where β_e is the electron plasma beta (Roberg-Clark et al. 2016, 2018; Komarov et al. 2018; Drake et al. 2021), thus, reducing the efficiency of heat conduction. While heat conduction is indeed suppressed for strong whistler suppression so that the MTI turbulence ceases to exist, externally-driven turbulence can revive the MTI turbulence and reestablish efficient heat conduction at a somewhat lower rate because of the intermittency of strong magnetic flux tubes along which most of the electron heat flux is transported (Perrone et al. 2024a,b).

Structure formation and evolution is a highly complex and non-linear problem with a vast range of relevant length- and time-scales that need to be considered. Numerical simulations therefore play a crucial role in exploring anisotropic thermal conduction in such contexts. Over the past two decades, a number of works have used simulations to investigate the effects of heat conduction in galaxy clusters (see e.g. Dolag et al. 2004; Jubelgas et al. 2004; Parrish et al. 2010; Ruszkowski & Oh 2010, 2011; Yang & Reynolds 2016; Kannan et al. 2017; Barnes et al. 2019; Su et al. 2019; Beckmann et al. 2022; Pellissier et al. 2023).

Some of these simulations were run with conduction solvers that employ explicit time-integration schemes which, for stability, require a timestep limit that scales with the square of the cell radius. Implicit or semi-implicit schemes are, therefore, preferable as they do not require this timestep criterion which can markedly reduce the attainable spatial resolution.

When constructing such a scheme it is also important to ensure that energy is not allowed to flow from lower to higher temperatures, which can occur when not explicitly prevented. Sharma & Hammett

(2007) demonstrated, for a Cartesian mesh, how violating this entropy condition can be avoided with appropriate gradient limiters and Pakmor et al. (2016b) then generalised this procedure for an irregular mesh.

In this paper, we introduce a scheme for anisotropic thermal conduction which we have implemented into the moving-mesh code AREPO (Springel 2010; Pakmor et al. 2016a; Weinberger et al. 2020). Our scheme uses a semi-implicit time integration scheme, is compatible with individual timestepping and ensures the entropy condition is not violated. The underlying method is based on that of Pakmor et al. (2016b) which concerns itself with the problem of cosmic ray diffusion and is, itself, based on those of Sharma & Hammett (2007) and Sharma & Hammett (2011).

The anisotropic thermal conduction solver presented in this work is fundamentally different from the solver described in Kannan et al. (2016), which is also implemented in AREPO. Our solver represents an improvement over this earlier method, primarily due to the fact that it supports local timestepping; a feature which is crucial for efficiency and accuracy in cosmological simulations.

This paper is structured as follows. In Section 2, we introduce the continuous form of the equations of Braginskii MHD and then describe our algorithm and its numerical implementation in Section 3. In Section 4, we assess the accuracy of our solver by performing several test problems including the diffusion of a hot patch of gas in a circular magnetic field, the progression of a point explosion in the presence of thermal conduction and the evolution and saturation of buoyancy instabilities in an anisotropically conducting plasma. In Section 5, we demonstrate the efficiency and stability of the solver when used in computationally demanding cosmological zoom simulations of a galaxy cluster and present some first results. Finally, in Section 6, we summarise our results and provide a brief outlook.

2 BASIC EQUATIONS

The Braginskii MHD equations can be used to describe transport in a fully ionised, weakly collisional¹ and strongly magnetised plasma. In such systems the electrons and ions are tied to magnetic field lines, resulting in the anisotropic transport of heat and momentum, governed by:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (1)$$

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot \left[\rho \mathbf{v} \mathbf{v}^T + P_{\text{tot}} \mathbb{I} - \frac{\mathbf{B} \mathbf{B}^T}{4\pi} \right] = -\nabla \cdot \Pi, \quad (2)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left[(E + P_{\text{tot}}) \mathbf{v} - \frac{\mathbf{B}(\mathbf{v} \cdot \mathbf{B})}{4\pi} \right] = -\nabla \cdot (\Pi \cdot \mathbf{v}) - \nabla \cdot \mathbf{Q}, \quad (3)$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}), \quad (4)$$

where we have used Gaussian units. ρ , \mathbf{v} and \mathbf{B} are the local gas density, velocity, and magnetic field, respectively. \mathbb{I} is the unit rank-two tensor and P_{tot} is the total pressure, accounting for thermal gas and magnetic fields

$$P_{\text{tot}} = P + \frac{B^2}{8\pi}. \quad (5)$$

¹ By ‘weakly collisional’ we are referring to systems where the gradient lengthscale is $10 - 10^3$ times larger than the Coulomb-collisional mean free path.

E is the total energy density

$$E = \rho u + \frac{1}{2} \rho v^2 + \frac{B^2}{8\pi}, \quad (6)$$

where u is the specific internal energy. The anisotropic viscosity tensor Π is

$$\Pi = -\Delta P \left(\mathbf{b} \mathbf{b}^T - \frac{1}{3} \mathbb{I} \right), \quad (7)$$

where $\mathbf{b} = \mathbf{B}/|\mathbf{B}|$ and $\Delta P = P_\perp - P_\parallel$ is the pressure anisotropy, i.e. the difference between the perpendicular and parallel pressures with respect to the magnetic field direction.

The heat flow vector \mathbf{Q} is given by

$$\mathbf{Q} = -\chi [\mathbf{b} (\mathbf{b} \cdot \nabla T)], \quad (8)$$

where χ is the conductivity along the magnetic field and T is the temperature, which is related to the specific internal energy via $u = c_v T$. The specific heat capacity at constant volume is given by $c_v = k_B / [(\gamma - 1) \mu m_p]$, where k_B is the Boltzmann constant, μ is the mean molecular weight, m_p is the proton mass and γ is the adiabatic index.

In a collisional theory of transport processes, the diffusive transfer of heat is dominated by electrons and mediated by particle-particle Coulomb collisions. The resulting conductivity is often referred to as the ‘Spitzer conductivity’ (Spitzer 1962) and is given by

$$\chi_{\text{sp}} = 1.84 \times 10^{-5} \frac{T^{5/2}}{\ln C} \text{ erg s}^{-1} \text{ K}^{-1} \text{ cm}^{-1}, \quad (9)$$

where the temperature, T , is measured in Kelvin and $\ln C \approx 37$ is the Coulomb logarithm.

High β , weakly collisional plasmas, however, are susceptible to a variety of kinetic instabilities which act to alter the mean electron transport velocity which may, ultimately, lead to a suppression of electron transport.

At saturation, the whistler instability has been found to establish a marginal heat flux which is suppressed by a factor of $1/\beta_e$, the inverse of the electron plasma beta (Roberg-Clark et al. 2018; Komarov et al. 2018). To account for this, one can assume a functional form for the conductivity (Komarov et al. 2018) that smoothly interpolates between the two regimes

$$\chi_{\text{sat,whist}} = \frac{\chi_{\text{sp}}}{1 + (1/3) \beta_e \lambda_{\text{mfp,e}} / l_{T,\parallel}}, \quad (10)$$

where $l_{T,\parallel} = |\mathbf{b} \cdot \nabla \ln T|^{-1}$ is the temperature gradient scale parallel to the magnetic field and $\lambda_{\text{mfp,e}}$ is the electron mean free path, given by

$$\lambda_{\text{mfp,e}} = \frac{3^{3/2} k_B^2 T^2}{4\pi^{1/2} n_e e^4 \ln C}, \quad (11)$$

where n_e is the electron number density and e is the electron charge.

One can also construct an analogous form for the conductivity that incorporates the saturation of the heat flux due to the free-streaming of electrons (Cowie & McKee 1977)

$$\chi_{\text{sat,free}} = \frac{\chi_{\text{sp}}}{1 + 4.2 \lambda_{\text{mfp,e}} / l_{T,\parallel}}. \quad (12)$$

In the outskirts of galaxy clusters, however, where the plasma beta is expected to be high ($\beta \gtrsim 100$), whistler suppression is likely to be more relevant.

3 IMPLEMENTATION

We now describe how we implemented the anisotropic thermal conduction solver into the AREPO code (Springel 2010; Pakmor et al. 2016a; Weinberger et al. 2020). Note that there is already a first version of an anisotropic conduction solver implemented in AREPO (Kannan et al. 2016). The solver presented in this work, however, is fundamentally different and is compatible with local timestepping. We discuss these differences in more detail in Section 3.2.4. Note also that a solver for Braginskii viscosity has already been implemented in AREPO by Berlok et al. (2020). While the solver described in this work shares a superficial resemblance to those presented in Pakmor et al. (2016b) and Kannan et al. (2016), it differs substantially in critical details, and these differences dramatically change the practical applicability to the targeted science applications.

AREPO solves the equations of ideal MHD on an unstructured Voronoi mesh using a second order finite volume scheme (Pakmor et al. 2011; Pakmor & Springel 2013). The Voronoi mesh is constructed from a set of mesh-generating points that can move with arbitrary velocities, but which are typically set to the local fluid velocity, resulting in quasi-Lagrangian behaviour. AREPO computes self-gravity using a tree-PM method and couples it to MHD with a second-order Leapfrog scheme (Springel 2010; Springel et al. 2021).

The anisotropic thermal conduction scheme, described in this work, is based on that of the cosmic ray diffusion solver presented in Pakmor et al. (2016b). This approach, itself, generalises and extends the flux limiting scheme of Sharma & Hammett (2007) and the semi-implicit time-integration scheme of Sharma & Hammett (2011) for use with unstructured meshes and local timestepping.

We have implemented this solver in such a way that it can be used for both cosmological and non-cosmological simulations. For simplicity, however, we will continue below to use physical coordinates to describe the scheme and explain how to implement the comoving form of the equations in Section 3.3.

We now focus on the treatment of just the conduction term in equation (3),

$$\frac{\partial u}{\partial t} = \frac{1}{\rho c_v} \nabla \cdot \left[\chi \mathbf{b} (\mathbf{b} \cdot \nabla u) \right], \quad (13)$$

where we have re-cast the equation in terms of the specific internal energy, assuming c_v is spatially constant.

Our numerical treatment of thermal conduction necessarily differs from that of cosmic ray diffusion, described in Pakmor et al. (2016b), in a number of ways. Firstly, the treatment of cosmic ray diffusion assumes that the diffusivity is spatially and temporally constant and that χ , therefore, commutes with the gradient operator. This is, however, not a good assumption for thermal conduction and, therefore, requires an additional treatment. Note, however, that this also may not be a good assumption for cosmic ray diffusion (see e.g. Thomas et al. 2023). Some of the methods described in this paper may, therefore, be relevant for future one-moment treatments of cosmic ray diffusion. In addition to this, there is an extra factor of $1/\rho$ before the flux term in the thermal conduction equation (equation 13), meaning that extra care has to be taken in the integration to ensure the scheme is conservative.

3.1 Spatial discretisation

We now spatially discretise equation (13) and begin by integrating over volume

$$V \frac{\partial u}{\partial t} = \frac{1}{\rho c_v} \int_V \nabla \cdot \left[\chi \mathbf{b} (\mathbf{b} \cdot \nabla u) \right] dV, \quad (14)$$

where we have assumed that ρ and u are constant throughout the volume and that $\int dV$ commutes with the time derivative. Using Gauss' theorem then gives

$$\frac{\partial u}{\partial t} = \frac{1}{V\rho c_v} \int_{\partial V} \left[\chi \mathbf{b} (\mathbf{b} \cdot \nabla u) \right] \cdot d\mathbf{A}, \quad (15)$$

where $d\mathbf{A}$ is the area element on the surface ∂V , directed along the outward normal.

We now take the integral to be over one cell in the simulation, indexed by i ,

$$\frac{\partial u_i}{\partial t} = \frac{1}{m_i c_v} \sum_j \left[\chi_{ij} (\mathbf{b}_{ij} \cdot \nabla u_{ij}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \quad (16)$$

where the sum is taken over all faces of the cell, and quantities indexed by ij are calculated in the interface between cells i and j which has area A_{ij} . Note that we have also used the relation $\rho_i = m_i/V_i$, where m_i and V_i are the mass and volume of the cell, respectively.

We have also implemented a solver for *isotropic* conduction. In this case, the discretised equation that determines the evolution of the specific internal energy is

$$\frac{\partial u_i}{\partial t} = \frac{1}{m_i c_v} \sum_j \left[\chi_{ij} (\nabla u_{ij} \cdot \mathbf{A}_{ij}) \right]. \quad (17)$$

3.1.1 Estimating quantities in the interface

To evaluate the flux on the right-hand side of equation (16), the gradient of the specific thermal energy, ∇u_{ij} , and the magnetic field direction, \mathbf{b}_{ij} , in the interfaces of all cells need to be determined.

When calculating the gradient of the thermal energy in the interface, care has to be taken to ensure the resulting solution is physical. When not explicitly prevented in the gradient estimates, the solution may lead to heat flowing from a cold cell to a hot cell. This can, however, be avoided by using gradient limiters, as described in Sharma & Hammett (2007). To calculate the gradients we largely follow the procedure described in Pakmor et al. (2016b) which generalises the procedure of Sharma & Hammett (2007) for a moving mesh. We now briefly summarise the relevant details here.

To estimate the specific thermal energy gradient in the interface, we first determine the gradients at the corners of the Voronoi face. The interface estimate is then built from these corner estimates. In (2D) 3D, every corner of a Voronoi cell corresponds to the centre of the (circumcircle) circumsphere of a (triangle) tetrahedron in the dual Delaunay tessellation. The corners of a Delaunay (triangle) tetrahedron correspond to mesh-generating points in the Voronoi mesh, so every corner has (three) four adjacent cells.

To estimate the specific thermal energy gradient at the corner, we perform a least-squares fit, using the values of the specific thermal energy at the centres of mass² of these adjacent cells (see section 2.1 of Pakmor et al. 2016b, for an in-depth explanation).

If a corner lies outside of the (triangle) tetrahedron formed by the centres of mass of the (three) four neighbouring cells the gradient estimate will be an extrapolation rather than an interpolation. This is illustrated in Fig. 1 where, for simplicity, we consider the 2D Voronoi mesh. The value of a quantity and its gradient are estimated at corner 1 from the values at the centres of mass, a , b and c , of the three adjacent cells. In the figure, the values at these three points are indicated by the colouring of the crosses and the shading within the




Figure 1. A sketch illustrating the gradient estimate at the corners of the Voronoi mesh in 2D. Mesh generating points are indicated by turquoise circles and the centres of mass of cells by blue crosses. The corners of the mesh are indicated by pink circles and are connected by the faces that form the Voronoi mesh, shown in black. A quantity at corner 1 (highlighted by the red circle) is estimated from the values at the centres of mass a , b and c of the three adjacent cells. The values at these three points are indicated by the colouring of each cross and the shading within the triangle corresponds to the least-squares estimate of the gradient at the corner 1. Since the corner, 1, lies outside of the triangle, the value predicted at this point (illustrated by its colouring) is larger than those at a , b and c and has been extrapolated, producing a new extremum. The gradient at corner 2 (highlighted by the green circle), however, is calculated from the values at the centres of mass d , e and f . This corner lies within the triangle def and the value predicted by the least-squares fit lies within the range of values at d , e and f , corresponding to an interpolation.

triangle corresponds to the fit to the gradient at corner 1. As corner 1 lies outside of the triangle abc , the predicted value (indicated by its colouring) is larger than those at points a , b and c , and has been extrapolated. The gradient at corner 2, on the other hand, is estimated from the values at the centres of mass d , e and f of its three adjacent cells. This corner lies within the triangle def and the value predicted by the least-squares fit is within the range of those at d , e and f , corresponding to an interpolation.

We mark corners where extrapolation has occurred as ‘problematic’ and use a different treatment for the contribution of this corner to gradient estimates in the interface; we will describe this alternative treatment shortly.

In the calculation of the fluxes, we split the gradient of the thermal energy into normal $\nabla u_{ij,N}$ and tangential $\nabla u_{ij,T}$ components relative to the normal to the interface.

$$\nabla u_{ij} = \nabla u_{ij,N} + \nabla u_{ij,T}. \quad (18)$$

The contributions to the flux from these two components are then calculated separately and added together to get the total flux across the interface.

In Fig. 2 we show a sketch illustrating why anisotropic diffusion schemes are prone to violations of the entropy criterion. Both panels

² Note that the centre of mass of a Voronoi cell is not necessarily spatially coincident with the associated mesh-generating point.




Figure 2. A sketch illustrating how anisotropic diffusion schemes are prone to violations of the entropy criterion. Mesh generating points are indicated by turquoise circles and the centres of mass of cells by blue crosses. The corners of the mesh are indicated by pink circles and are connected together by the faces that form the Voronoi mesh, shown in black. Both panels show the same configuration of the Voronoi mesh, with the only difference being the direction of the magnetic field estimate in the interface (as indicated by the arrow labelled \mathbf{B}). On the left-hand side, the configuration is such that the projection of the temperature gradient estimate onto the magnetic field direction and the projection of the magnetic field onto the normal to the face are positive. In this case, heat will flow across this face from the hot cell to the cold (see equation 16). In the right hand panel, however, the projection of the magnetic field onto the face normal is now negative, meaning that heat will flow across this face from the cold cell to the hot. This is a result of the numerical discretisation of the diffusion problem and can lead to extrema being accentuated.

show exactly the same configuration of the Voronoi mesh, with the only difference being the direction of the magnetic field estimate in the interface. In the panel on the left, the quantity $(\mathbf{B} \cdot \nabla T)(\mathbf{B} \cdot \hat{\mathbf{n}})$ is positive (where $\hat{\mathbf{n}}$ is the normal to the face). Inspecting equation (16) shows that, in this configuration, heat will flow across the face, from the hot cell to the cold. In the right hand panel, however, the projection of the magnetic field onto the face normal is now negative, meaning that heat will flow across this face in the opposite direction.

To ensure that temperature extrema are not accentuated we use a procedure very similar to that described in Sharma & Hammett (2007) which requires two steps. First, when estimating the normal gradient of the thermal energy in the interface (described below), we ensure that energy associated with the normal flux flows across a face from the hotter cell to the colder. Second, we slope-limit the tangential gradient, $\nabla u_{ij,T}$, using a generalised version of the van Leer limiter (van Leer 1984)

$$\nabla u_{ij,T} = \frac{N}{\sum_k (\nabla u_{k,T})^{-1}}, \quad (19)$$

where the sum runs over all N corners (indexed by k) of the face ij . If the tangential gradients at the corners $\nabla u_{k,T}$ do not all have the same sign, the tangential gradient in the interface is set to zero.

We calculate the normal component of the gradient in the interface by taking an average of the values at the corner, weighted by the fractional area of the Voronoi face that is closest to this corner. This alone does not, however, guarantee that the result satisfies the entropy condition. We, therefore, additionally calculate a ‘simple’ finite difference estimate using the values of the specific thermal energy at the centres of mass of the two cells on either side of the interface. While the corner averaged normal gradient estimate has a larger stencil, this ‘simple’ estimate is more robust. We then compare the signs of these two gradient estimates and, if they disagree, we

use the ‘simple’ finite difference estimate which satisfies the entropy condition by construction. We, additionally, replace the contribution of corners that were previously flagged as ‘problematic’ (due to extrapolation in the gradient estimate at the corner) with this ‘simple’ estimate.

The magnetic field direction in the interface, \mathbf{b}_{ij} , is calculated in a similar way. We use a least-squares fit to get a first estimate at the corners, and then calculate the interface value by taking a weighted average of these corner estimates³. We then calculate c_V assuming the gas is fully ionised with primordial composition. Next, we calculate the conductivity in the interface, χ_{ij} , by taking the arithmetic average of the conductivities of the two cells on either side of the interface

$$\chi_{ij} = \frac{1}{2}(\chi_i + \chi_j). \quad (20)$$

For the scheme outlined in Sharma & Hammett (2007) they suggest instead using the harmonic average for numerical stability. In our scheme, however, we do not observe any stability problems when using the arithmetic average. Additionally, in Appendix A, we use point explosion simulations (analogous to those presented in Section 4.2.1) to explore the effects of these two averaging procedures and show that, for our scheme, the arithmetic mean best reproduces the analytic predictions of the speed of advance of the conduction fronts.

Note that, when using the functional form of the conductivity that interpolates between the whistler suppressed and collisional regimes (see equation 10), we use the unlimited temperature gradient to calculate the temperature gradient lengthscale. Typically, AREPO stores the limited gradients that are used in the flux calculations, which

³ The weights are the same as those described above for the thermal energy gradient in the interface, i.e. the fractional area of the Voronoi face that is closest to the relevant corner.

can be significantly lower than the true gradients in regions where the gradients are not well resolved. Using these limited gradients in the calculation of $l_{T,\parallel}$ could lead to underestimates of the level of suppression. Additionally, since AREPO uses the single-fluid MHD approximation we do not have access to the electron temperature. When calculating the whistler suppression factor we, therefore, use the bulk plasma temperature to calculate the plasma beta.

For the case of isotropic conduction, only the gradient of the specific thermal energy normal to the surface contributes to the flux, which we estimate in the same way as for $\nabla u_{ij,N}$.

3.2 Time integration

We implement three different time integration schemes for the anisotropic thermal conduction solver: an explicit solver, a semi-implicit, linear solver and a semi-implicit, non-linear solver. Note that in the descriptions of these methods in the next three sub-sections, the equations are strictly only valid when using global timesteps. In Section 3.2.4 we discuss the adaptations that need to be made when using local timestepping.

3.2.1 Explicit time integration

For the explicit integration scheme, we evolve the specific thermal energy using two half timesteps updates which are carried out immediately before and after the first and second gravity kicks, respectively:

$$\begin{aligned} u_i^{n+1/2} &= u_i^n + \frac{\Delta t}{2m_i^n c_v} \sum_j \left[\chi_{ij}^n (\mathbf{b}_{ij} \cdot \nabla u_{ij,T}^n) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right] \\ &\quad + \frac{\Delta t}{2m_i^n c_v} \sum_j \left[\chi_{ij}^n (\mathbf{b}_{ij} \cdot \nabla u_{ij,N}^n) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \\ u_i^{n+1} &= u_i^{n+1/2} + \frac{\Delta t}{2m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^{n+1/2} (\mathbf{b}_{ij} \cdot \nabla u_{ij,T}^{n+1/2}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right] \\ &\quad + \frac{\Delta t}{2m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^{n+1/2} (\mathbf{b}_{ij} \cdot \nabla u_{ij,N}^{n+1/2}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \end{aligned} \quad (21)$$

where the superscripts n , $n + 1/2$, and $n + 1$ correspond to the beginning, mid-point and end of a timestep of length Δt , respectively. We use the same procedure for the explicit treatment of isotropic conduction, with the appropriate flux term (equation 17).

For stability, this explicit integration scheme requires a timestep constraint of the form

$$\Delta t < \eta \frac{(\Delta x)^2}{\kappa}, \quad (22)$$

where Δx is the cell diameter, $\eta < 1$ is a constant, and

$$\kappa = \frac{\chi}{c_v \rho}, \quad (23)$$

is the thermal diffusivity. When determining this timestep for each cell, we use the maximum value of χ in any of its interfaces to calculate κ .

Note that this timestep criterion can become prohibitively restrictive due to its quadratic dependence on the cell size, severely limiting the achievable spatial resolution. It is therefore preferable to use semi-implicit or implicit time-integration schemes, which do not require this timestep criterion for stability. Note also that super-timestepping schemes such as that presented in Berlok et al. (2020) can be used to somewhat alleviate this timestep constraint. So far, however, super-timestepping schemes are not able to deal with the

local timestepping in AREPO, and extending them to be capable of doing so is not straightforward.

3.2.2 Semi-implicit, linear solver

The gradient limiter that ensures the entropy condition is not violated (equation 19) introduces non-linearities into the tangential flux estimate, which significantly increase the numerical complexity when using an implicit integrator. The normal gradient estimate on the other hand has an explicit, linear dependence on internal energy. But it has an additional, implicit dependence on temperature via the conductivity (which is often modelled as being temperature dependent). In this ‘semi-implicit, linear solver’ we assume χ_{ij} for each interface is constant during the determination of the normal flux. In the following section we discuss an integration scheme that relaxes this assumption.

When changes to χ_{ij} across a timestep are ignored, a semi-implicit scheme can be formulated which is almost as stable as a fully implicit scheme and in which only one linear implicit problem is solved per timestep (see Sharma & Hammett 2011; Pakmor et al. 2016b). To this end we split the calculation into two parts. First, the flux associated with the tangential component of the specific thermal energy gradient is evolved using a forward-Euler method,

$$u_i^{\tilde{n}} = u_i^n + \frac{\Delta t}{m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^n (\mathbf{b}_{ij} \cdot \nabla u_{ij,T}^n) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \quad (24)$$

where the superscript \tilde{n} indicates intermediate values. In the second step, the thermal energies are advanced due to the flux associated with the normal component of the specific thermal energy gradient using an implicit backward-Euler scheme,

$$u_i^{n+1} = u_i^{\tilde{n}} + \frac{\Delta t}{m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^n (\mathbf{b}_{ij} \cdot \nabla u_{ij,N}^{n+1}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right]. \quad (25)$$

To solve this linear system of equations, we use the HYPRE library (Falgout & Yang 2002) and carry out a two-step procedure, first using a generalised minimal residual (GMRES) solver (Saad & Schultz 1986) iteratively until the residual drops below 10^{-8} . If this condition is not met within 200 iterations, we add an algebraic multigrid preconditioner (Henson & Yang 2002) to GMRES which is then iterated until the residual drops below 10^{-8} .

We apply a flux limiter to the first (explicit) part of this calculation to ensure that cells are not completely drained of thermal energy during the exchange. Such a limiter is, however, not required in the second (implicit) part of the calculation.

We find that this scheme does not require a timestep criterion similar to that required by explicit schemes (see equation 22) and is stable on much longer timesteps, similar to what was observed in Sharma & Hammett (2011) and Pakmor et al. (2016b). The implicit part of the flux calculation is unconditionally stable. As the timesteps become very long, however, the explicit fluxes may become so large that the flux limiter activates and the solution becomes more isotropic. We, however, do not observe such behaviour when cells are integrated on their MHD timestep.

We also treat the case of isotropic conduction similarly and split the flux into two steps of size $\Delta t/2$. We integrate the first step using the explicit scheme and the second using the implicit scheme. This semi-implicit scheme is fully second-order convergent in the case of isotropic conduction (Pakmor et al. 2016b). We could in principle treat isotropic conduction purely implicitly (as only normal gradient estimates are required which are linear in the specific internal energy of the cells), though this scheme would only be first-order accurate in time.

3.2.3 Semi-implicit, non-linear solver

When the conductivity has a temperature dependence (e.g. Spitzer), the fluxes of thermal energy associated with conduction will necessarily lead to changes in the conductivity. To test whether taking into account changes in χ_{ij} across each timestep affects our results, we also implemented a ‘non-linear solver’.

Put simply, this ‘semi-implicit, non-linear solver’ iteratively applies the procedure described in the previous section for the ‘semi-implicit, linear solver’ and updates the conductivity and gradients at each iteration. Specifically, we iteratively calculate

$$u_i^{\tilde{n},m} = u_i^{n,0} + \frac{\Delta t}{m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^{n,m-1} (\mathbf{b}_{ij} \cdot \nabla u_{ij,T}^{n,m-1}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \quad (26)$$

$$u_i^{n,m*} = u_i^{\tilde{n},m} + \frac{\Delta t}{m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^{\tilde{n},m} (\mathbf{b}_{ij} \cdot \nabla u_{ij,N}^{n,m*}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \quad (27)$$

where n indexes the timestep and m indexes the iteration. \tilde{n} denotes an intermediate quantity, after the tangential flux has been evolved but before that of the normal flux. We use the intermediate thermal energies to calculate the conductivity, $\chi_{ij}^{\tilde{n},m}$, that is used in the normal flux (equation 27).

$u_i^{n,m*}$ is the specific thermal energy that is predicted at the end of iteration m . To get the energies used in the conductivity and gradients at the beginning of the next iteration, we interpolate between this value and the specific thermal energy at the beginning of this iteration, $u_i^{n,m-1}$, via

$$\mathbf{u}^{n,m} = \mu^m \mathbf{u}^{n,m*} + (1 - \mu^m) \mathbf{u}^{n,m-1}, \quad (28)$$

where \mathbf{u} is a vector of thermal energies of all cells being considered in the calculation. A good choice of the interpolation parameter, μ^m , can speed up the rate at which the non-linear solver converges. We implemented two different interpolation methods. First, similar to Kannan et al. (2016), we employ a modified version of the unstable manifold corrector scheme described in Smedt et al. (2010) in which the correction vector (i.e. the vector of changes in thermal energy across the timestep) is adaptively over- or underrelaxed depending on the relative direction of successive correction vectors. In the second interpolation method, we take $\mu^m = 1$ but reduce this when necessary to ensure that the relative change of the thermal energy of all cells from its value at the beginning of the iteration is not more than 10%. We found the second procedure to be more robust, while the first method sometimes took much longer to converge, particularly when using local-timestepping.

We use two criteria to determine if the solution is converged, both of which are checked at every iteration. The first criterion we check is

$$\frac{|\mathbf{u}^{n,m} - \mathbf{u}^{n,m-1}|}{|\mathbf{u}^{n,0}|} < 10^{-6}. \quad (29)$$

For the second convergence criterion we begin by estimating the solution at the end of the iteration using a forward-Euler scheme for both the tangential and normal components. The first part of this is identical to the standard treatment of the tangential fluxes (see equation 26), so only the normal component requires calculating at each iteration.

The estimate of the solution at the end of the iteration, $u_i^{\hat{n},m}$, is then

$$u_i^{\hat{n},m} = u_i^{\tilde{n},m} + \frac{\Delta t}{m_i^{n+1} c_v} \sum_j \left[\chi_{ij}^{\tilde{n},m} (\mathbf{b}_{ij} \cdot \nabla u_{ij,N}^{\tilde{n},m}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij}) \right], \quad (30)$$

where $u_i^{\tilde{n},m}$ is given by equation (26). We then calculate the difference between these predictions, $\mathbf{u}^{\hat{n},m}$, and the internal energies at the beginning of the iteration $\mathbf{u}^{n,m}$,

$$\Delta \mathbf{u}^{n,m} = \mathbf{u}^{n,m} - \mathbf{u}^{\hat{n},m}, \quad (31)$$

and the second convergence criterion is then

$$\frac{|\Delta \mathbf{u}^{n,m}|}{|\Delta \mathbf{u}^{n,0}|} < 10^{-3}. \quad (32)$$

If either of these two conditions are satisfied then we stop the iteration and accept the internal energies at the end of the final iteration.

We find that convergence typically occurs within the first ~ 10 iterations but can sometimes take many more, particularly when using local-timestepping (as described below) at synchronisation points where a significant fraction of gas cells are active. We should note, however, that these procedures are not strictly guaranteed to converge to the solution of the equations, in particular, due to the flux and gradient limiters that are used. Additionally, the first criterion essentially measures whether the iteration has stopped and does not make any statement about how close the solution is to solving the equations.

3.2.4 Local timestepping

The scheme presented thus far solves the conduction problem for all gas cells in the simulation at each timestep. Since, when using local-timestepping, the mesh in AREPO is only guaranteed to be complete for active cells, the scheme, as currently presented, can only be used on global timesteps. This makes it unsuitable for use in complex problems with deep timestep hierarchies (such as those typically found in simulations of galaxy clusters). We, therefore, use a method similar to Pakmor et al. (2016b) to adapt the scheme to work with local timestepping.

At each timestep we compute the specific thermal energy of each cell involved in the calculation and use this to solve the conduction problem for each active interface⁴ as described in the previous sections. We then update the thermal energy of the cells due to these conduction fluxes. Note that the mesh at each timestep is guaranteed to be complete for active cells so all information required to calculate the gradients and quantities at the corners will be available. This means that we only have to solve the conduction problem for all active cells and for a layer of inactive cells that share a face with an active cell.

For the ‘semi-implicit, non-linear solver’ we must additionally define the particle set over which we calculate the convergence criteria (see Section 3.2.3). We choose this set to be all cells that have at least one active interface, (i.e. all cells that can potentially have their internal energies changed by the thermal conduction in this timestep). Note that this is the same particle set that is used in the matrix iteration stopping criterion in the implicit step (used in both the linear and non-linear solvers) as the matrix will include the energies of all cells that are involved in the flux calculation (i.e. those that have at least one shared interface).

The scheme presented in this work uses a fundamentally different method to that of the anisotropic thermal conduction solver described in Kannan et al. (2016) which is also implemented in AREPO. In addition, extending our scheme to work with local timesteps using the procedure described above is relatively straightforward. Building the

⁴ An active interface refers to a Voronoi face where at least one of the two neighbouring cells is active.




Figure 3. Specific internal energy slices at early times ($t = 10$) for the problem of anisotropic conduction of a hot patch of gas around a circular magnetic field in 2D. Each panel shows the entire computational domain for five simulations with different resolutions (indicated in the top right) which use local timestepping (see text for explanation). In the lower-right panel we show the analytic solution for comparison. With increasing resolution the diffusion perpendicular to the magnetic field decreases and the numerical solution is seen to converge on the analytic one.

mesh at each sync-point for the active particle set requires inserting the direct neighbours of all active cells into the mesh which, as discussed above, is sufficient for us to compute the fluxes over all active interfaces. The scheme presented in Kannan et al. (2016), however, requires information about a second layer of cells around active cells, which are not guaranteed to be part of the partial mesh. This means that the scheme can only be used on timesteps where all particles are active, making it unsuitable for use in investigating complex problems (such as those found in studies of structure formation) with deep timestep hierarchies.

3.3 Anisotropic thermal conduction in comoving coordinates

Here we briefly discuss the changes that need to be made to adapt the procedure described in the previous sections for use in cosmological simulations. In cosmological simulations it is convenient to introduce spatial coordinates that are comoving with the cosmological expansion and a set of comoving variables (which we indicate with a subscript ‘c’). Those relevant to the transformation of the conduction equation are defined via

$$\begin{aligned} \mathbf{r} &= a\mathbf{x}, \quad \mathbf{u} = \mathbf{v} - \dot{a}\mathbf{x}, \\ \rho &= \rho_c a^{-3}, \quad P = P_c a^{-3} \quad \chi = \chi_c a^{-1}, \end{aligned} \quad (33)$$

where a is the scale factor, \mathbf{r} are physical coordinates, \mathbf{x} are comoving coordinates, the physical velocity is $\mathbf{v} = \dot{\mathbf{r}}$, and the peculiar velocity is $\mathbf{u} = a\dot{\mathbf{x}}$. Note that, under this coordinate transformation, the specific internal energy is unchanged.

The derivatives in comoving coordinates transform according to

$$\left. \frac{\partial}{\partial t} \right|_{\mathbf{r}} = \left. \frac{\partial}{\partial t} \right|_{\mathbf{x}} - H\mathbf{x} \cdot \nabla_{\mathbf{x}}, \quad \nabla_{\mathbf{r}} = \frac{1}{a}\nabla_{\mathbf{x}}, \quad (34)$$

where $H = H(a) = \dot{a}/a$ is the Hubble rate.

The part of the energy equation relevant for anisotropic thermal conduction in comoving coordinates is then

$$\left. \frac{\partial u}{\partial t} \right|_{\mathbf{x}} = \frac{1}{\rho_c c_v} \nabla_{\mathbf{x}} \cdot \left[\chi_c \mathbf{b} (\mathbf{b} \cdot \nabla_{\mathbf{x}} u) \right]. \quad (35)$$

Note that this equation does not correspond to an exact transform of equation (13) but, rather, comes about by transforming the energy equation (3) and isolating the term relevant for anisotropic thermal conduction. The discretised form of this equation is then

$$\left. \frac{\partial u_i}{\partial t} \right|_{\mathbf{x}} = \frac{1}{m_i c_v} \sum_j \left[\chi_{ij,c} (\mathbf{b}_{ij} \cdot \nabla_{\mathbf{x}} u_{ij}) (\mathbf{b}_{ij} \cdot \mathbf{A}_{ij,c}) \right]. \quad (36)$$

3.4 Comparison with other algorithms and implementations

The solver described in this work presents a significant improvement over previous schemes implemented in AREPO which solve diffusion problems. It combines local timestepping, spatially varying diffusion coefficient with dependence on the quantity being diffused, and an optional non-linear solver to solve the resulting non-linear anisotropic diffusion problem.

The cosmic ray diffusion scheme presented in Pakmor et al. (2016b) is closest to that presented in this paper. Their scheme, however, only works for a spatially uniform diffusion coefficient which




Figure 4. L^1 norm for the problem of anisotropic conduction of a hot patch of gas around a circular magnetic field in 2D (as visualised in Fig. 3). The top panel shows the L^1 norm at early times ($t = 10$), and the bottom panel shows it at late times ($t = 200$). The different coloured lines show the results for different timestepping schemes: ‘local’ and ‘global’ correspond to simulations that use local and global timestepping schemes, respectively. ‘global+diff’ corresponds to simulations that use global timestepping with the addition of the timestep criterion required by explicit schemes for stability (equation 22). In general, all configurations show very similar errors and convergence rates, both at early times ($L^1 \propto N^{-0.52}$) and at late times ($L^1 \propto N^{-0.7}$).

leads to a purely linear problem. In contrast, the schemes presented in the previous section for thermal conduction were designed to work stably and accurately also when the diffusivity varies on a cell-to-cell basis according to the local properties of the gas. The steep temperature gradients that are expected in the CGM and ICM lead to correspondingly steep gradients in the diffusivity, making this a taxing problem. They also make the problem at least mildly non-linear, so it is critical to have a solver that can take this non-linearity into account.

The scheme presented in Berlok et al. (2020) treats the anisotropic diffusion of momentum and uses a similar spatial discretisation as Pakmor et al. (2016b). However, their scheme is not compatible with local timestepping, making it unable to cope with the large spatial and temporal dynamic range that arises in cosmological cluster simulations. Finally, the scheme for anisotropic thermal conduction introduced in Kannan et al. (2016) uses a different spatial discretisation which, similar to the time integration scheme in Berlok et al. (2020), is fundamentally unable to be extended to local timestepping. Additionally, it requires a cap on the diffusivity for stability

(Kannan et al. 2017) which does not allow modelling of full Spitzer conductivity in galaxy clusters and precludes studying models with a physically suppressed heat conduction which are suggested by advanced models of MHD that account for electron-whistler scattering (e.g., Drake et al. 2021).

4 RESULTS: NUMERICAL TESTS

In this section we present results from a number of numerical test problems which we use to assess the accuracy and stability of the anisotropic thermal conduction scheme.

4.1 Anisotropic conduction around a ring in 2D

We test the 2D diffusion of a hot patch of gas in a circular magnetic field, following the setup of Parrish & Stone (2005) and Sharma & Hammett (2007). We use this setup to test and quantify the convergence of the anisotropic solver with respect to the analytical solution when applied to a multidimensional problem at fixed conductivity and a constant background. We, therefore, disable hydrodynamics in this test and keep the mesh fixed.

We use a regular, approximately hexagonal mesh, which we create by taking a uniform Cartesian mesh offsetting the points in every other column by 0.49 times the cell size. We initialise the specific thermal energy within a domain of $[-1, 1]^2$ to

$$u(\mathbf{x}) = \begin{cases} 12 & \text{if } 0.5 < r < 0.7 \text{ and } |\phi| < \pi/12, \\ 10 & \text{otherwise,} \end{cases} \quad (37)$$

where $r = \sqrt{x^2 + y^2}$ is the distance from the centre of the box and $\phi = \tan^{-1}(y/x)$ is the angle to the x -axis. The magnetic field is initialised via

$$\begin{aligned} B_x(\mathbf{x}) &= -\frac{y}{r}, \\ B_y(\mathbf{x}) &= +\frac{x}{r}. \end{aligned} \quad (38)$$

We set the parallel diffusivity, $\kappa = \chi/(c_v \rho)$, to 0.01. There is no perpendicular diffusivity so the energy should stay within the ring.

At early times, the resulting behaviour can be considered as a 1D diffusion problem of a double step function, which has an analytic solution given by

$$u(\mathbf{x}) = 10 + \operatorname{erfc} \left[\left(\phi + \frac{\pi}{12} \right) \frac{r}{D} \right] - \operatorname{erfc} \left[\left(\phi - \frac{\pi}{12} \right) \frac{r}{D} \right], \quad (39)$$

where $D = \sqrt{4\kappa t}$ for $0.5 < r < 0.7$ and $u(\mathbf{x}) = 10$ elsewhere. At late times, the energy should be uniformly distributed in the ring, i.e. $u = 10 + 1/6$ for $0.5 < r < 0.7$ and $u = 10$ elsewhere.

In Section 3.2.3 we introduced the non-linear solver as a way to take the temperature dependence of the conduction coefficient into account. In this test the conduction coefficient is fixed and therefore insensitive to the non-linear solver. We therefore defer discussion of the non-linear solver to the following section and, for this test, consider only the semi-implicit, linear solver described in Section 3.2.2.

In the present section we compare simulations that use local timesteps with those that use global timesteps and additionally consider simulations where we impose the timestep criterion required by explicit schemes for stability (equation 22). We refer to these three types of simulation as ‘local’, ‘global’ and ‘global+diff’, respectively.

In the simulations that use local timesteps we impose a timestep hierarchy, as this does not arise naturally in the problem. We leave the timesteps of cells in the lower-left quadrant unaltered and reduce

that of the cells in the upper-left and lower-right quadrant by a factor of 2, and that of the cells in the upper-right quadrant by a factor of 4.

Figure 3 shows the specific internal energy distribution at early times ($t = 10$) for five simulations with different resolutions, and in the bottom-right we show the analytic prediction. The simulations in Fig. 3 were all run with local timestepping (as described above). Reassuringly, it is evident that there are no unphysical features at the boundaries between regions with different timesteps. All simulations reproduce the general behaviour predicted by the analytic solution. With increasing resolution, however, there is less perpendicular diffusion and the numerical solution becomes a better match to the analytic one.

We now examine these simulations more quantitatively. Figure 4 shows the L^1 norm⁵ as a function of resolution of the simulation at early and late times, for simulations that use local timestepping and for runs that use global timestepping, with and without the diffusion timestep constraint being imposed. In general, all three configurations show very similar errors and convergence rates, at both early times ($L^1 \propto N^{-0.52}$) and late times ($L^1 \propto N^{-0.7}$). The faster convergence rates at late times arise due to the fact that, by this point, the solution is primarily sensitive to errors in the conduction perpendicular to the magnetic field whereas any errors in the parallel conduction speed (which the solution will be sensitive to at early times) have been washed out. We note that these convergence rates are clearly worse than first order but are comparable to other implementations of anisotropic transport solvers (Parrish & Stone 2005; Sharma & Hammett 2007; Kannan et al. 2016; Pakmor et al. 2016b).

The behaviour of the errors in the runs with local and global timesteps are very similar. At late times, however, the run with local timesteps has slightly smaller errors at higher resolution. This may arise because we impose smaller timesteps than required in the local-timestepping simulations. The runs with the additional timestep constraint show slightly smaller errors at early times and higher resolution in comparison to the other schemes. At late times and high resolution, however, the error is slightly worse.

4.2 Point explosion in a uniform medium

To assess the accuracy of the coupling between the ideal MHD and thermal conduction solvers we carry out 3D simulations of a point explosion in a uniform background where the conductivity is self-consistently determined by the temperature. We perform three different types of simulation; one where there is no thermal conduction (i.e. ideal MHD), one with the addition of isotropic conduction and one with anisotropic conduction. We refer to these type of simulations as ‘MHD’, ‘+iso’ and ‘+aniso’, respectively.

In all simulations, the mesh is initially Cartesian but undergoes regularisation during the course of the simulation (see Vogelsberger et al. 2012). We carry out simulations at three different spatial resolutions: 64^3 , 128^3 and 256^3 . For the ideal MHD and isotropic conduction simulations, the computational domain has side length 100 pc. As the conduction front in the setup with anisotropic conduction moves much faster than in the isotropic case we use a computational domain with side length 200 pc for the two lower resolution simulations with anisotropic conduction. This ensures that the conduction front does not leave the box during the simulation. For the highest resolution simulation with anisotropic conduction, increasing the boxsize was too expensive and so we ran it in a box of side length 100 pc and

stopped the simulation before the conduction front left the domain. For consistency we will continue to refer to the lower resolution simulations with a box side length of 200 pc as 64^3 and 128^3 , despite the fact that, when doubling the domain length, we also double the number of cells along each edge. This ensures that, across all simulations of a given resolution (i.e. 64^3 , 128^3 or 256^3), the spatial resolution remains the same.

The background medium has constant density $n_0 = 1 \text{ cm}^{-3}$ and temperature $T_0 = 10^4 \text{ K}$. The background magnetic field has strength $1 \mu\text{G}$ and points in the positive x -direction, which corresponds to an initial plasma beta of ~ 35 . In all simulations we inject $E_0 = 3.33 \times 10^{50} \text{ erg}$ into the central region. In the 64^3 resolution runs we inject into the central 8 cells. We keep the volume of the injection region the same in the higher resolution simulations (128^3 and 256^3), and the number of injection cells increases accordingly to 64 and 512, respectively. This means that the total mass into which the energy is injected is constant across all simulations (and thus the peak initial temperature is also the same).

For simulations with conduction, the conductivity is assumed to take the Spitzer value without any form of suppression (i.e. χ is as in equation 9). Unless specified, the simulations were run with the linear thermal conduction solver. We identify those that were run with the non-linear solver explicitly.

In order to capture the very early evolution of the shock and conduction fronts, we employ maximum timesteps that are log-spaced in time. We, additionally, carried out analogous simulations with local timesteps where we relaxed this maximum timestep criterion and confirmed that the results of these simulations are consistent with the late time behaviour described in the rest of this section.

Figure 5 shows thin temperature projections at three different times for the 128^3 resolution simulations without conduction, with isotropic conduction and with anisotropic conduction. At early times, the conduction fronts in the runs with conduction (middle and right columns) advance faster than the shock front in the pure MHD run (left column). In the simulation with anisotropic conduction there is essentially no diffusion perpendicular to the magnetic field and the parallel expansion occurs noticeably faster than the radial expansion seen in the simulation with isotropic conduction. The perpendicular expansion in the anisotropic conduction case is also slower than the classical adiabatic solution, as was also found in Kannan et al. (2016) and Dubois & Commerçon (2016).

At later times, the pure MHD and isotropic conduction simulations have similar radial extents although the MHD simulation retains a strong radial temperature gradient within the shocked region, while the temperature gradients in the simulations with conduction are much flatter.

We now examine the different behaviour observed in these simulations more quantitatively. Analytic solutions to the problem of point explosion in a uniform medium exist for the purely adiabatic case and the case of pure conduction. For an adiabatic point explosion, the radius of the shock front expands according to

$$r_s(t) = \xi(n) \left(\frac{E_0 t^2}{\rho_0} \right)^{1/(n+2)}, \quad (40)$$

where n is the dimension of the problem and $\xi(n)$ is a constant of order unity which also depends on the adiabatic index (Sedov 1959). In three dimensions, with $\gamma = 5/3$, $\xi(3)$ is approximately 1.15.

For the case of a point explosion with heat conduction where the conductivity has a power law dependence on temperature, $\chi =$

⁵ The sum over all cells of the absolute value of the difference between the simulation result and analytic prediction.

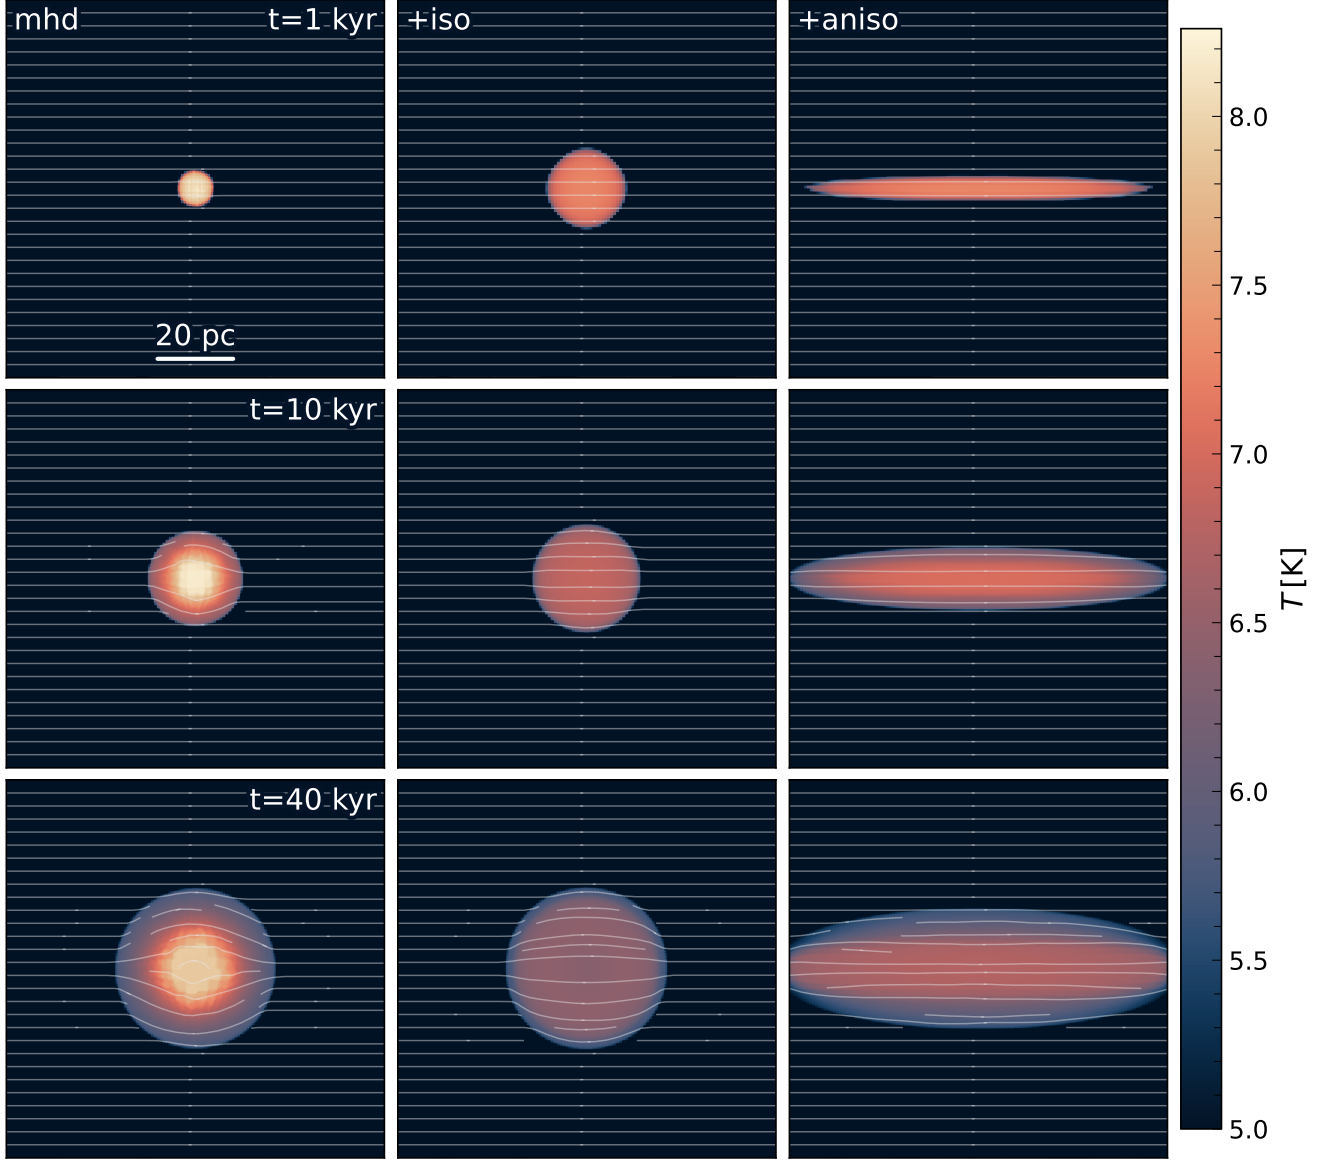


Figure 5. Thin temperature projections of size 100×100 pc for the point explosion test. Each column, from left to right, shows results for the pure MHD case, that with isotropic conduction and that with anisotropic conduction. Each row depicts the state of the system at a different time which is indicated in the top-right corner of each panel in the first column. The grey lines show the magnetic field configuration. Thermal conduction can significantly affect the outcome of a point explosion, particularly when conduction occurs anisotropically along the magnetic field.

$\chi_0 T^\alpha$, the radius of the conduction front evolves according to

$$r_c(t) = \left[\left(\frac{2(n\alpha + 2)}{\alpha} \frac{\chi_0}{c_v \rho_0} \right) \left(\frac{2}{S_n} \frac{1}{\mathcal{B}(\frac{n}{2}, \frac{1}{\alpha} + 1)} \frac{E_0}{\rho_0 c_v} \right)^\alpha t \right]^{1/(n\alpha + 2)}, \quad (41)$$

where $S_n = 1, 2\pi, 4\pi$ for dimensions $n = 1, 2, 3$, and $\mathcal{B}(x, y)$ is the beta function (see e.g. Zel'dovich & Raizer 1967; Barenblatt 1996).

4.2.1 Point explosion with isotropic conduction

We begin by examining the simulations that were run with isotropic conduction. Equation (40) predicts that the 3D shock radius in our

simulations should expand according to

$$r_{s,3D}(t) = 4.75 \left(\frac{t}{\text{kyr}} \right)^{\frac{2}{5}} \text{pc}. \quad (42)$$

From equation (41) the 3D conduction front, assuming Spitzer conductivity ($\alpha = 5/2$), is at radius

$$r_{c,3D}(t) = 10.1 \left(\frac{t}{\text{kyr}} \right)^{\frac{2}{19}} \text{pc}. \quad (43)$$

In Fig. 6 we show the time evolution of the radius of the shock/conduction front for the three different resolution simulations. Additionally, for the lowest resolution, we show the results from a simulation run using the non-linear conduction solver described in




Figure 6. Evolution of the radius of the conduction/shock front in the 3D point explosion test with isotropic conduction (as visualised in the middle column of Fig. 5). Solid, coloured lines show results from simulations run with the linear conduction solver while the dashed, yellow line corresponds to a simulation run with the non-linear solver (see figure legend). The black lines indicate the analytic expectation for the radial evolution of the shock/conduction front in the cases of pure hydrodynamics ($r_{s,3D}$, equation 42) and pure conduction ($r_{c,3D}$, equation 43). At early times conduction is fastest but at ~ 13 kyr, the shock front catches up and overtakes. With increasing resolution the results converge towards the analytic expectation.

Section 3.2.3. We estimate the radius of the shock/conduction front by first calculating a spatial temperature profile of the gas along the x -axis⁶, passing through the centre of the injection region. We then define the conduction/shock front radius to be the distance from the centre to the furthest point where the temperature is above 1%⁷ of the background temperature, T_0 .

At early times, conduction dominates and the radius of the fronts advances $\propto t^{2/9}$, in accordance with analytical expectations. Equating equations (42) and (43) gives a transition time of ~ 13 kyr at which hydrodynamical processes are expected to become dominant. Indeed, this is very close to the time at which we see a distinct change in the evolution of the shock/conduction front in the simulations, which then begins to follow the shock solution. With increasing resolution and at a fixed time, the radius of the measured front is generally smaller, and converges towards the analytic expectation. The behaviour of the shock/conduction front in the simulation run with the non-linear solver (dashed yellow line) almost exactly reproduces that in the analogous simulation run with the linear solver (solid red line). In future work we plan to apply the non-linear solver to other problems (including galaxy cluster zooms) to test whether this result holds generally.

4.2.2 Point explosion with anisotropic conduction

When heat conduction is anisotropic, the problem of point explosion in a uniform medium essentially becomes a 1D problem along the




Figure 7. Evolution of the radius of the conduction/shock front in the 3D point explosion test with anisotropic conduction (as visualised in the right column of Fig. 5). Solid/dot-dash, coloured lines show the radius as measured parallel/perpendicular to the initial magnetic field direction (see figure legend). The black dotted lines show the analytic scalings for the radial evolution of a cylindrical blast wave ($r \propto t^{1/2}$, see equation 40) and a 1D conduction front ($r \propto t^{2/9}$, see equation 41). The expansion of the conduction front along the magnetic field follows the expected scalings for a 1D point explosion with pure conduction. Perpendicular to the initial magnetic field, the shock front expands slower than the conduction front and follows the expected scalings for a cylindrical blast wave.

direction of the magnetic field. Perpendicular to the magnetic field lines, the evolution of the system is determined by MHD processes and the early rapid expansion of the conduction front along the x -axis, as seen in Fig. 5, means that this approximately resembles a 2D cylindrical blast wave.

Equations (40) and (41) show that the radii of the conduction and shock fronts in this scenario should scale as $r_{c,1D} \propto t^{2/9}$ and $r_{s,2D} \propto t^{1/2}$, respectively. In contrast to the isotropic case, the constants E_0 and ρ_0 in equations (40) and (41) are not well defined because the anisotropic case does not decouple into two independent problems along and perpendicular to the magnetic field. We, therefore, just compare the scaling behaviour of the radius of the 2D shock and 1D conduction front, rather than exact values.

In Fig. 7 we show the radii of the conduction/shock front in the point explosion test with anisotropic conduction as a function of time, for the three different resolution simulations. We calculate the radii of the shock/conduction front parallel and perpendicular to the initial magnetic field direction and use the same method as was described in the previous section for the isotropic conduction simulations. The expansion of the conduction front along the magnetic field at early times is much faster than seen for the isotropic case. This is expected and is largely due to the geometry of the problem. The radius of the conduction front follows the analytic scaling well for the higher resolution simulations, but is slightly shallower for the lowest resolution run (red line). At late times, the radial expansion slows with respect to the analytic scaling as hydrodynamical processes that are driving the lateral expansion become important. This behaviour was also seen in Kannan et al. (2016).

The expansion of the shock front perpendicular to the magnetic field is initially slower than that of the conduction front along the magnetic field and follows the expected scaling well at all resolution levels, except perhaps at very early times for the lowest resolution run.

⁶ The results discussed in the remainder of the section are independent of the direction of the axis along which the temperature profile is measured.

⁷ We also tested threshold values of 0.1% and 10% and found that the location of the shock/conduction front radius is sufficiently well resolved in all simulations that the results presented in this section and the one that follows are insensitive to the exact choice of this threshold value.

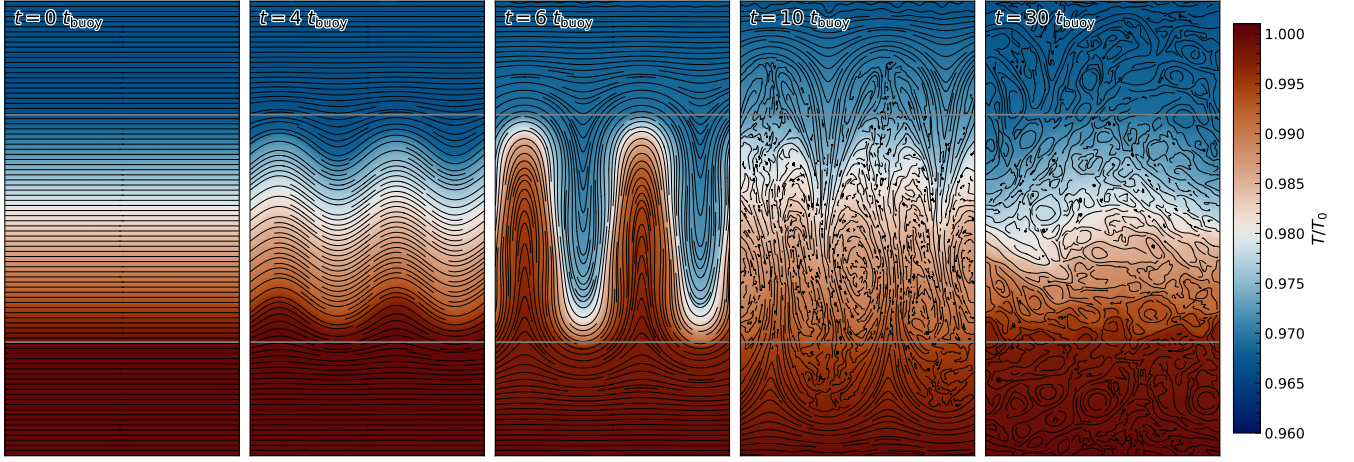


Figure 8. Temperature slices with magnetic field lines overlaid in black that show the time evolution of the MTI for the simulation where the unstable region has resolution 256×256 . The slices show the entire simulation domain ($L \times 2L$), and the grey horizontal lines indicate the boundaries between the unstable and neutral regions. The temperatures are normalised to that at the base of the unstable region. The initial perturbation grows linearly in the first three panels until there is a large vertical component of the magnetic field. This configuration is, however, non-linearly unstable with the end result being a state of sustained turbulence.

Overall, the results in this section have validated the accuracy of the coupling between the hydrodynamics and thermal conduction solvers and have, additionally, demonstrated the accuracy of our treatment of the temperature dependence of the conductivity in the solver.

4.3 The magneto-thermal instability

For both positive and negative temperature gradients, high- β , weakly collisional, stratified plasmas are unstable to magnetically-mediated buoyancy instabilities when anisotropic conduction is rapid compared to the dynamical response of a plasma. Specifically, anisotropic heat conduction causes the slow magnetosonic wave to become buoyantly unstable to the MTI (Balbus 2000, 2001) when the temperature decreases with height and to the HBI (Quataert 2008) when the temperature increases with height.

In regions where $\mathbf{g} \cdot \nabla T > 0$, any misalignment of magnetic-field lines and gravity (i.e. \mathbf{g}) can result in the development of the MTI. In the presence of a sustained temperature gradient, the MTI drives vigorous turbulence that is largely isotropic (see e.g. Parrish & Stone 2005, 2007; Perrone & Latter 2022a,b). In this section, we present simulations which explore the non-linear behaviour of the MTI and compare our results to previous work (such as Parrish & Stone 2005; Parrish & Quataert 2008; McCourt et al. 2011; Kannan et al. 2016).

4.3.1 Initial conditions

We follow the evolution of a 2D region of plasma initially in hydrostatic and thermal equilibrium. The region of plasma we simulate consists of an MTI unstable region, where conduction is anisotropic, sandwiched between two buoyantly neutral regions, where the heat conduction is isotropic. Note, however, that we ignore these buffer regions in the quantitative analysis presented in this section.

The plasma is stratified in density and temperature and is subject to a uniform gravitational field in the vertical direction: $\mathbf{g} = -g_0 \hat{z}$. In the horizontal direction, we apply periodic boundary conditions. At the upper and lower boundaries, we fix the temperature. This enforces a vertical temperature gradient across the box and prevents

the instabilities from saturating by flattening this gradient. We extrapolate the pressure into the ghost cells at these boundaries so as to enforce hydrostatic equilibrium, while all other plasma quantities are treated using reflective boundary conditions.

The size of the simulated domain is $L \times 2L$ and the unstable region lies within $L/2 < z < 3L/2$. In this unstable region we construct an atmosphere where the temperature decreases linearly with height according to

$$T(z) = T_0 \left(1 - \frac{z - L/2}{3H_0} \right), \quad (44)$$

$$\rho(z) = \rho_0 \left(1 - \frac{z - L/2}{3H_0} \right)^2, \quad (45)$$

$$P(z) = P_0 \left(1 - \frac{z - L/2}{3H_0} \right)^3, \quad (46)$$

where H_0 is a scale height. We set $L/H_0 = 0.1$ which means that this setup is ‘local’ in the sense that the size of the simulation domain is much smaller than the pressure scale height.

We initialise the buoyantly neutral layers as isothermal atmospheres where the temperatures are continuous across the boundaries of the unstable region. The pressure and density in these regions vary exponentially with height. We set $g_0 = k_B T_0 / (H_0 \mu m_H)$ to ensure that the system is initially in hydrostatic equilibrium and additionally apply a uniform background magnetic field in the x -direction with magnitude such that the plasma beta at the base of the unstable region is $\beta_0 = 2 \times 10^4$. This means that magnetic tension is negligible on the scales of interest.

The physics of the buoyancy instabilities is independent of the conductivity in the limit that the thermal diffusion time across the spatial scales of interest is short compared to the dynamical time. We, therefore, assume a uniform fixed conductivity and ensure that the simulations are in this fast conduction limit by setting $\chi = 10 \rho_0 c_v \omega_{\text{buoy}} L^2$, where we evaluate the characteristic frequency for the buoyancy instability,

$$\omega_{\text{buoy}} = \left| g_0 \frac{\partial \ln T}{\partial z} \right|^{1/2}, \quad (47)$$

at the base of the unstable region. This means that the thermal diffusion time across the unstable region is $\sim 0.1 \omega_{\text{buoy}}^{-1}$. As the conductivity is fixed, we only present simulations that use the linear solver in this section.

The simulations are initialised with a regular mesh, similar to that described in Section 4.1, which is then allowed to move over the course of the simulation. We, additionally, consider two different spatial resolutions, where the unstable region is resolved by 128^2 or 256^2 cells. The simulations are all run for 50 buoyancy times ($t_{\text{buoy}} = \omega_{\text{buoy}}^{-1}$, where ω_{buoy} is evaluated at the base of the unstable region). We apply single mode transverse perturbations to the velocity field with wave vector $\mathbf{k} = (4\pi/L) \hat{\mathbf{x}}$ and amplitude $10^{-4} c_{s,0}$ where $c_{s,0}$ is the sound speed evaluated at the base of the buoyantly unstable region.

In 2D, when the mesh is allowed to move, we found that not resolving the gradients and the magnetic field topology could lead to the introduction of noise in the simulations. In these simulations we, therefore, applied an extra criterion whereby the tangential fluxes across a face are set to zero if the signs of the normal components of the magnetic field estimates at the two corners of each face are different. This increases the stability of the scheme at the cost of being more diffusive. We find that this condition is not necessary in 3D (where the number of corners that contribute to the interface estimate is significantly higher), or when the resolution is high enough.

4.3.2 Results

Figure 8 shows temperature slices with magnetic field lines overlaid in black for the simulation where the unstable region has resolution 256×256 . The slices show the entire simulation domain ($L \times 2L$), including the upper and lower buoyantly neutral regions and the grey horizontal lines indicate the boundaries between these regions. In the first three panels, the instability is still in the linear growth phase. The perturbations grow exponentially and drive the development of a large vertical magnetic field component.

The MTI growth rate goes to zero when the field lines become vertical. As we see from the final two panels of Fig. 8, however, the MTI does not saturate quiescently when it reaches this linearly stable state but, instead, drives sustained turbulence. This arises due to the existence of zero-frequency modes of the dispersion relation, which correspond to horizontal perturbations that act to drive the plasma away from this equilibrium configuration (Balbus & Reynolds 2010; McCourt et al. 2011). This state of sustained turbulence persists for the rest of the simulation. This qualitative behaviour is consistent with that observed in previous simulation work (see e.g. Parrish & Stone 2005, 2007; McCourt et al. 2011; Kannan et al. 2016; Berlok et al. 2021).

In the final two panels of Fig. 8 we see that the turbulent motions have length scales that are comparable to the size of the computational domain. As shown by McCourt et al. (2011), this prematurely stops the buoyant acceleration and implies that these ‘local’ simulations are not able to accurately capture the saturated state of the MTI and under-predict the turbulent energies.

The left-hand panel of Fig. 9 shows the evolution of the z -component of the magnetic field and highlights the difference between the linear phase of the instability where the vertical component of the magnetic field is exponentially amplified and the saturated state where the magnetic field becomes close to isotropic. We also show the time evolution of the volume averaged kinetic energy density in the right-hand panel of Fig. 9. In both simulations, the kinetic energy is exponentially amplified at early times but then saturates at

$t \approx 6 t_{\text{buoy}}$ as the magnetic field configuration approaches the linearly stable state.

The linear growth rate of a mode of the MTI can be determined from the set of linearised equations that govern the system (see e.g. Quataert 2008; Kunz 2011; Berlok et al. 2021). Assuming locality in the vertical direction, the growth rate for the perturbation seeded in these simulations should be ω_{buoy} . The dashed line in the right-hand panel of Fig. 9 corresponds to this scaling and is in very good agreement with the behaviour observed in our simulations.

Both the lower and higher resolution simulations agree very well on the behaviour of the instability during the linear growth phase ($t \lesssim 6 t_{\text{buoy}}$). At later times, the results differ quantitatively. These differences, however, are marginal and the qualitative behaviour is in good agreement. The evolution of the magnetic field and the kinetic energy densities, shown in Fig. 9, are also very similar to those presented in Kannan et al. (2016), who carried out similar 2D simulations of the MTI, and to those found in the 3D simulations of McCourt et al. (2011).

5 COSMOLOGICAL SIMULATIONS OF GALAXY CLUSTERS WITH ANISOTROPIC THERMAL CONDUCTION

To investigate the effects of anisotropic thermal conduction on structure formation and evolution it is important that the anisotropic thermal conduction solver is able to work efficiently and accurately in a fully cosmological context. In this section, we demonstrate that our code is capable of efficiently running high resolution cosmological zoom simulations of galaxy clusters. An in-depth analysis of the role of anisotropic thermal conduction in the evolution and properties of the cluster will be presented in a follow-up paper. Here we will largely focus on technical details and implementation choices.

In one of our simulations, we set the conductivity to the Spitzer value (equation 9) and in the other we set it to Spitzer with whistler suppression (equation 10).

5.1 Initial conditions and physical model

In this section we present three different cosmological zoom simulations of one halo with mass $M_{200c} = 5.09 \times 10^{14} h^{-1} M_{\odot}$ and radius $R_{200c} = 1.30 h^{-1} \text{ Mpc}$ at $z = 0$. Here, we define M_{200c} and R_{200c} such that the average density within a sphere of radius R_{200c} is 200 times the critical density of the universe and M_{200c} is the mass enclosed by this radius. The simulations use a Planck-2018 cosmology (Planck Collaboration et al. 2020) where $\Omega_m = 0.315$, $\Omega_b = 0.049$, $\Omega_{\Lambda} = 0.684$, and the Hubble constant is $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$ with $h = 0.673$. The simulations presented in this section were carried out as part of the ongoing PICO-Clusters (Plasmas In COsmological Clusters) project, which will be presented in greater detail in forthcoming work.

The particular halo studied here was selected for re-simulation from a parent dark matter only simulation with side of length $1 h^{-1} \text{ cGpc}$. The initial conditions for the zoom simulation were then created using a new code that will be described in Puchwein et al., in prep. We have made sure that the simulations presented here have no low resolution dark matter particles within $2 R_{200c}$ at $z = 0$.

In these simulations, the high resolution dark matter particles have mass $4.0 \times 10^7 h^{-1} M_{\odot}$ while the gas in the high resolution region has target mass resolution $7.4 \times 10^6 h^{-1} M_{\odot}$. The comoving softening length of the high resolution dark matter is set to $3.25 h^{-1} \text{ ckpc}$ with a maximum physical softening length of $1.627 h^{-1} \text{ kpc}$. The




Figure 9. The left-hand panel shows the time evolution of the magnetic field orientation in the two-dimensional MTI simulation visualised in Fig. 8. Specifically, we show the volume average of the square of the vertical component of $\mathbf{b} = \mathbf{B}/|\mathbf{B}|$. A vertical magnetic field will have $b_z^2 = 1$, whereas $b_z^2 = 0.5$ corresponds to an isotropic field in 2D. The right-hand panel shows the time evolution of the volume averaged kinetic energy density. The different coloured lines correspond to simulations with different resolutions (see legend). These quantities were calculated by averaging over all cells within the buoyantly unstable region. The black dashed line shows the theoretical scaling of the growth rate. During the linear phase of the instability ($t \lesssim 6 t_{\text{buoy}}$), the magnetic field is driven towards a vertical orientation and the kinetic energy is exponentially amplified. During the non-linear phase that follows, the instability saturates and the turbulence acts to isotropise the magnetic field.

gas softening is treated adaptively and scales with the radius of the cell, with a minimum comoving value of $0.41 h^{-1} \text{ ckpc}$.

This resolution is about 3 times better than that used in the full-physics MillenniumTNG box (Pakmor et al. 2023), and is slightly better than that used in the TNG300 volume of the IllustrisTNG project (Springel et al. 2018; Marinacci et al. 2018; Nelson et al. 2018; Pillepich et al. 2018b; Naiman et al. 2018) and the recent TNG-Cluster simulations (Nelson et al. 2024), with slight differences arising from the choice of cosmology.

We apply a physics model based on that used in the TNG project, which is described in detail in Weinberger et al. (2017) and Pillepich et al. (2018a). Physical processes in this model include primordial and metal line cooling along with heating from a spatially uniform UV background (Vogelsberger et al. 2013), an effective model for the ISM and star formation (Springel & Hernquist 2003), chemical enrichment of the ISM due to core collapse and thermonuclear supernovae and AGB stars, an effective model for galactic winds (Pillepich et al. 2018a) and the formation, growth of and feedback from supermassive black holes (Weinberger et al. 2017). We additionally initialise a spatially uniform seed magnetic field of comoving strength of 10^{-14} G at the start of the simulations ($z = 127$).

We perform three zoom simulations of this halo, all of which use this physics model as a baseline. The first simulation (which we will refer to as the ‘standard’ simulation) has no additional models for physical processes. In the other two simulations we switch on our model for anisotropic thermal conduction. In both of these simulations we use the ‘semi-implicit, linear solver’. In one of these simulations, we set the conductivity to the Spitzer value (equation 9) and in the other we set it to Spitzer with whistler suppression (equation 10). We refer to these simulations as ‘+spitzer’ and ‘+whistler’, respectively.

For the simulations with thermal conduction, we do not solve the conduction problem across interfaces where one or more of the cells on either side are star forming. The thermal energy of these cells is set by the effective equation of state (Springel & Hernquist 2003) and represents an effective pressure rather than the temperature of the gas,

which should just be considered as a property of the effective subgrid model. In doing this we effectively assume that thermal conduction is unimportant for the transport of energy between the star-forming ISM and the gas surrounding it.

5.2 Results

Figure 10 shows thin projections of various quantities at $z = 0$ for the ‘+whistler’ simulation. The magnetic field in the cluster is clearly highly turbulent and has strength that reaches $\sim 10 \mu\text{G}$ in the central regions. In the outskirts of the cluster, however, the field is significantly weaker. For an in-depth discussion of the amplification of magnetic fields in the PICO clusters, see Tevlin et al. (2024).

Throughout the majority of the cluster, the magnetic field is, however, largely dynamically unimportant, as can be seen in the projection of the plasma beta. The plasma beta typically lies within the range $\sim 10^2 - 10^3$ but can be even higher in the cluster outskirts, particularly in regions where the magnetic field is low. Intermittent magnetic flux tubes can reach plasma beta values of 10.

In the lower-middle panel, we show the thermal diffusivity $\kappa = \chi/(\rho c_v)$, as measured in the simulation (i.e. including the effects of whistler suppression) and in the lower-right panel we show the factor by which the diffusivity is suppressed below the Spitzer value, $(1 + (1/3)\beta \lambda_{\text{mfp},e} / l_{T,\parallel})^{-1}$, see equation (10). The turbulent nature of the magnetic field (in combination with a much smoother temperature gradient) leads to significant variation in the temperature gradient length scale along the direction of the magnetic field, $l_{T,\parallel} = |\mathbf{b} \cdot \nabla \ln T|^{-1}$. This, in combination with the variations in β , can lead to a significant suppression of the conductivity, particularly in the outskirts of the cluster and at the location of shocks (see the temperature projection in the upper-left panel and density in the upper-middle panel).

In Fig. 11 we show radially averaged profiles of various thermodynamical quantities at $z = 0$ for all three simulations. In general, we find good agreement across all three runs. There are, however, some slight differences, particularly in the density, metallicity, en-




Figure 10. Thin ($100 h^{-1}$ kpc) projections of size $(3 h^{-1} \text{ Mpc})^2$ at $z = 0$ for the ‘+whistler’ simulation. From left to right, the top row shows temperature, gas density and magnetic field strength. The bottom row shows the plasma beta, the thermal diffusivity, and the ratio of the (whistler suppressed) thermal diffusivity to the Spitzer value. The temperature projection is weighted by mass, while those of the magnetic field strength, density and thermal diffusivity are volume weighted. The projection of the plasma beta is calculated by dividing the volume weighted projection of the thermal pressure by that of the magnetic pressure; likewise, the projection of the ratio of the thermal diffusivity to the Spitzer value is calculated by dividing the volume weighted projection of the thermal diffusivity by that of the Spitzer value. Dashed circles correspond to R_{200c} . The magnetic field is highly turbulent and this, along with a much smoother temperature gradient leads to significant levels of suppression in the heat flux (see equation 10), particularly in the outskirts of the cluster and at shocks.

tropy and magnetic field strength profiles: the runs with the conduction exhibit higher central densities and magnetic field strengths, lower central entropies and flatter central metallicity gradients. The flattening of the metallicity profile could be explained by enhanced mixing, associated with the effects of thermal conduction. Similar behaviour has also been observed in previous simulation work exploring anisotropic thermal conduction in galaxy clusters (Kannan et al. 2017). We will explore this in more detail in future work with a larger sample of simulated clusters. Additionally, the cluster simulated with full Spitzer conductivity shows larger differences than that simulated with whistler suppressed conduction when compared to the simulation without conduction.

A number of works in the literature which use simulations to explore the effects of anisotropic thermal conduction in galaxy clusters impose a cap on the thermal diffusivity. This is often done when

using conduction schemes which require application of the diffusion timestep criterion (see equation 22) for stability. Such a cap prevents extremely small diffusive timesteps which could otherwise make the simulation computationally intractable (see e.g. Yang & Reynolds 2016; Kannan et al. 2017). Values used for such a cap are typically $\sim 5 \times 10^{31} \text{ cm}^2 \text{ s}^{-1}$. From the thermal diffusivity profiles for our simulations (lower-right panel of Fig. 11) we see that the Spitzer value in the outskirts of the cluster can be significantly higher than this numerical cap. The use of such a cap can, therefore, result in significant numerical suppression of the conductivity in the outskirts of the cluster. The application of this artificial suppression also means that such schemes cannot be reliably used to explore the effects of physically motivated conductivity closures such as whistler suppression.

Whilst the differences in the radially averaged profiles are rather subtle, we do find larger differences between these systems in other




Figure 11. Radial profiles at $z = 0$ for the galaxy cluster zoom simulations. From left to right, the top row shows density, temperature, and magnetic field strength. The bottom row shows entropy (which we define to be $K = k_B T / n_e^{(\gamma-1)}$, where n_e is the electron number density), metallicity, and thermal diffusivity, for the two simulations run with anisotropic thermal conduction. All profiles are in reasonably good agreement with each other, with some deviations seen in the central regions. The cluster simulated with full Spitzer conductivity ('+spitzer') shows larger differences than that simulated with whistler suppressed conduction ('+whistler') when compared to the simulation without conduction ('standard').

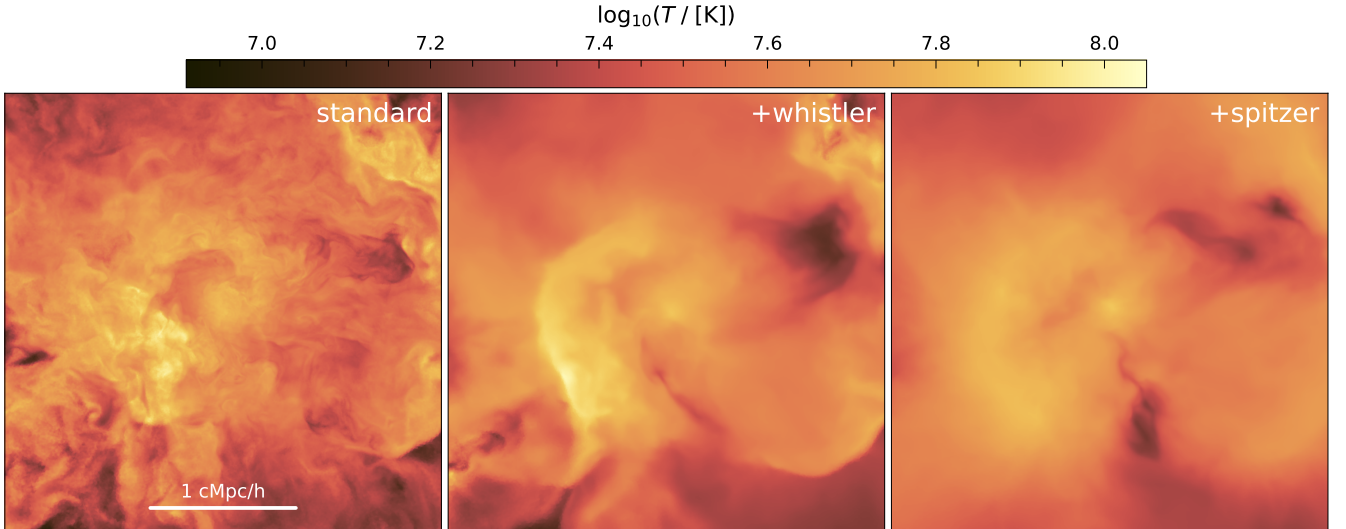


Figure 12. Mass-weighted temperature projections of size $(3 h^{-1} \text{Mpc})^2$ and depth $200 h^{-1} \text{kpc}$ at $z = 0$ for, from left to right, the simulation without conduction ('standard'), that with whistler suppressed anisotropic conduction ('+whistler'), and that with full Spitzer conductivity ('+spitzer'). The 'standard' simulation shows significant small-scale structure in the temperature field. The temperature field in the simulations with conduction, however, is considerably smoother, with the largest effect seen in the '+spitzer' case.

quantities. In Fig. 12 we show temperature projections of the three clusters at $z = 0$. The 'standard' simulation shows significant small-scale structure in the temperature field. In the simulations with conduction, however, this small-scale structure is largely gone and the temperature field is substantially smoother; an effect that is most pronounced in the '+spitzer' run, where the conductivities are generally

higher than in the '+whistler' run (see lower right panel of Fig. 11). These observations are consistent with the expected behaviour of thermal conduction: namely that it acts to flatten temperature gradients.

This can be understood more quantitatively by examining radial profiles of the standard deviation of the temperature within each

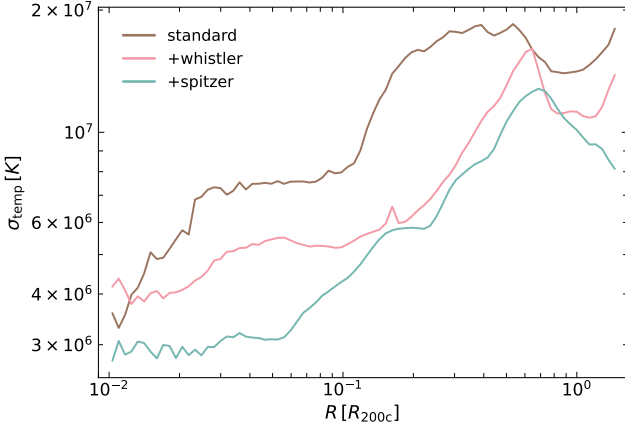


Figure 13. Radial profiles at $z = 0$ of the standard deviation of the temperature distribution in radial bins, weighted by volume, for the simulation without conduction (‘standard’), that with whistler suppressed anisotropic conduction (‘+whistler’), and that with full Spitzer conductivity (‘+spitzer’). The ‘standard’ simulation shows the highest levels of variability in the temperature field, whereas the levels of variability in both simulations with anisotropic conduction are significantly lower.

radial bin, weighted by cell volume, which are shown in Fig. 13. At all radii (except the very centre) the simulation without conduction has the largest temperature variance, while both simulations with anisotropic conduction show significantly lower levels of variability. The ‘+spitzer’ simulation, where the conductivities are generally highest, has the lowest level of variability.

In this section we have shown that the differences between simulations with and without anisotropic conduction can be significant. Additionally, we have shown that the effect of whistler suppression can be considerable. Previous cosmological cluster simulations have also found that conduction can have a substantial effect on the cluster properties (see e.g. Dolag et al. 2004; Jubelgas et al. 2004; Ruszkowski et al. 2011; Kannan et al. 2017; Barnes et al. 2019; Pellissier et al. 2023). In order to carefully assess the extent to which thermal conduction is affecting the cluster properties and in what ways these effects manifest, it is first important to better understand the variability inherent in the standard physics model and which cluster properties it robustly predicts. It is, additionally, important to have more than one data point to draw conclusions from, which requires simulating a larger sample of clusters. This is beyond the scope of this work, where the aim is to introduce our thermal conduction scheme, but will be studied in detail in future work.

It should also be mentioned that the extent to which the thermal conductivity closure used in the ‘+whistler’ simulation (equation 10) accurately captures the effect of whistler waves on electron transport in the ICM is not well understood. The present results, however, demonstrate that it may have a significant impact on the effective conductivity, particularly in the outskirts of the cluster. Future work is needed to better understand the impact of kinetic microinstabilities on electron transport and how these effects manifest on large scales, in a cosmological context.

5.3 Technical considerations

We end this section with a few technical remarks relating to the performance of the anisotropic thermal conduction solver in these cluster zooms. In the two simulations with thermal conduction, about

20% of the run-time is spent in the conduction solver. About a third to half of this time is spent in the matrix solver, and the majority of the remaining time is spent in the computation of the coefficients for the gradient estimates and the preparation of the relevant data structures. The cost associated with the explicit flux calculation itself is negligible in comparison.

Fewer than 10 iterations are typically required to solve the system of equations in the implicit part of the calculation. At later times, and particularly at time steps where a significant fraction of the gas cells are active, the addition of the multigrid preconditioner is required as otherwise the matrix solver is unable to converge within 200 iterations.

As is typical for cluster zooms, our simulations have timestep hierarchies that can be up to ~ 10 levels deep. These simulations would be prohibitively expensive if all cells had to be integrated on the smallest global timestep. The significant difference between the smallest and largest timesteps (a factor of ~ 1024) demonstrates that it is not appropriate to apply the scheme only on global synchronisation points. All this underlines the necessity of our local timestepping scheme for such calculations.

The data structures that store necessary information related to the faces and corners of the Voronoi mesh have the largest memory footprint of all those associated with the conduction solver. In addition, the HYPRE library can have a significant memory footprint in larger simulations which is allocated independently to AREPO’s memory manager. All these structures are, however, not required anywhere else in the code, and are freed upon exit of the thermal conduction routine.

6 SUMMARY AND OUTLOOK

In this paper, we presented a new scheme for modelling anisotropic thermal conduction on a moving mesh. The solver, which we have implemented into the moving-mesh code AREPO (Springel 2010; Pakmor et al. 2016a; Weinberger et al. 2020), is fully conservative, ensures the entropy condition is not violated, and allows for semi-implicit time integration and individual timestepping.

The anisotropic thermal conduction solver introduced in this work represents an improvement over the approach described in Kannan et al. (2016), which is also implemented in AREPO, primarily due to the fact that it supports local timestepping; a feature which is crucial for efficiency and accuracy in cosmological simulations. If conduction is done on global timesteps only, it either becomes too expensive to run, or one loses the coupling to the gas on faster timescales. Accurate treatment of anisotropic thermal conduction is, additionally, more complex than cosmic ray diffusion (Pakmor et al. 2016b), due to the fact that the conductivity is not spatially and temporally constant (as is often assumed in one-moment treatments of cosmic ray diffusion).

We tested our implementation on a variety of numerical problems and demonstrated that our solver is able to reproduce analytic predictions and numerical solutions. The convergence rate of the anisotropic solver is comparable to those of other numerical diffusion solvers, and we showed that the use of local timesteps does not lead to a significant reduction in the accuracy of the solver. We also verified the accuracy of the coupling between the conduction and hydrodynamics and demonstrated that the solver is able to reproduce the speed of the conduction front predicted by analytic arguments. These tests highlighted just a few ways anisotropic thermal conduction can fundamentally change the behaviour of a system.

We also demonstrated that the solver can be applied to highly

non-linear problems with deep timestep hierarchies by performing high-resolution cosmological zoom-in simulations of a galaxy cluster with anisotropic conduction, where we showed that the solver operates efficiently and robustly. In these simulations we found that anisotropic thermal conduction, as well as the presence (or absence) of whistler suppression can have a significant impact on the temperature distribution of ICM and acts to smooth out small-scale structure. We will explore these results in more detail in future work.

The ability to accurately and efficiently capture the effects of anisotropic thermal conduction in a fully cosmological environment is crucial, and the methods presented here will allow us to assess the effects of conduction in a wide range of astrophysical processes in diverse contexts.

In the future, we plan to use high resolution cosmological simulations to better understand the role anisotropic thermal conduction plays in a range of astrophysical contexts, including the ICM of clusters and the CGM of galaxies. This will require a large suite of high resolution cosmological zoom simulations. These studies will be made possible by the fact that our conduction solver is able to operate accurately and efficiently on problems that exhibit significant variations in the relevant spatial and temporal scales.

ACKNOWLEDGEMENTS

We would like to thank Lorenzo Maria Perrone for very helpful discussions and suggestions. CP acknowledges support by the European Research Council under ERC-AdG grant PICO GAL-101019746. RB is supported by the UZH Postdoc Grant, grant no. FK-23116 and the SNSF through the Ambizione Grant PZ00P2_223532. FvdV is supported by a Royal Society University Research Fellowship (URF\R1\191703).

DATA AVAILABILITY

The data underlying this article will be shared upon request to the corresponding author.

REFERENCES

- Armillaotta L., Fraternali F., Werk J. K., Prochaska J. X., Marinacci F., 2017, *MNRAS*, **470**, 114
- Asai N., Fukuda N., Matsumoto R., 2007, *ApJ*, **663**, 816
- Balbus S. A., 2000, *ApJ*, **534**, 420
- Balbus S. A., 2001, *ApJ*, **562**, 909
- Balbus S. A., Reynolds C. S., 2010, *ApJ*, **720**, L97
- Balsara D. S., Tilley D. A., Howk J. C., 2008a, *MNRAS*, **386**, 627
- Balsara D. S., Bendinelli A. J., Tilley D. A., Massari A. R., Howk J. C., 2008b, *MNRAS*, **386**, 642
- Barenblatt G. I., 1996, Scaling, Self-similarity, and Intermediate Asymptotics
- Barnes D. J., et al., 2019, *MNRAS*, **488**, 3003
- Beckmann R. S., Dubois Y., Pellissier A., Polles F. L., Olivares V., 2022, *A&A*, **666**, A71
- Berlok T., Pakmor R., Pfrommer C., 2020, *MNRAS*, **491**, 2919
- Berlok T., Quataert E., Pessah M. E., Pfrommer C., 2021, *MNRAS*, **504**, 3435
- Bingert S., Peter H., 2011, *A&A*, **530**, A112
- Bourdin P. A., Bingert S., Peter H., 2013, *A&A*, **555**, A123
- Braginskii S. I., 1965, Reviews of Plasma Physics, **1**, 205
- Brüggen M., Scannapieco E., 2016, *ApJ*, **822**, 31
- Brüggen M., Scannapieco E., Grete P., 2023, *ApJ*, **951**, 113
- Chevalier R. A., 1975, *ApJ*, **200**, 698
- Choi E., Stone J. M., 2012, *ApJ*, **747**, 86
- Cowie L. L., McKee C. F., 1977, *ApJ*, **211**, 135
- Dolag K., Jubelgas M., Springel V., Borgani S., Rasia E., 2004, *ApJ*, **606**, L97
- Dong R., Stone J. M., 2009, *ApJ*, **704**, 1309
- Drake J. F., et al., 2021, *ApJ*, **923**, 245
- Dubois Y., Commerçon B., 2016, *A&A*, **585**, A138
- Dursi L. J., Pfrommer C., 2008, *ApJ*, **677**, 993
- Fabian A. C., Sanders J. S., Crawford C. S., Conselice C. J., Gallagher J. S., Wyse R. F. G., 2003, *MNRAS*, **344**, L48
- Falgout R. D., Yang U. M., 2002, in Sloot P. M. A., Hoekstra A. G., Tan C. J. K., Dongarra J. J., eds, Computational Science — ICCS 2002. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 632–641
- Henson V., Yang U., 2002, *Applied Numerical Mathematics*, **41**, 155
- Jacob S., Pfrommer C., 2017a, *MNRAS*, **467**, 1449
- Jacob S., Pfrommer C., 2017b, *MNRAS*, **467**, 1478
- Jennings R. M., Li Y., 2021, *MNRAS*, **505**, 5238
- Jubelgas M., Springel V., Dolag K., 2004, *MNRAS*, **351**, 423
- Kannan R., Springel V., Pakmor R., Marinacci F., Vogelsberger M., 2016, *MNRAS*, **458**, 410
- Kannan R., Vogelsberger M., Pfrommer C., Weinberger R., Springel V., Hernquist L., Puchwein E., Pakmor R., 2017, *ApJ*, **837**, L18
- Komarov S. V., Churazov E. M., Kunz M. W., Schekochihin A. A., 2016, *MNRAS*, **460**, 467
- Komarov S., Schekochihin A. A., Churazov E., Spitkovsky A., 2018, *Journal of Plasma Physics*, **84**, 905840305
- Kunz M. W., 2011, *MNRAS*, **417**, 602
- Kunz M. W., Schekochihin A. A., Stone J. M., 2014, *Phys. Rev. Lett.*, **112**, 205003
- Lyutikov M., 2006, *MNRAS*, **373**, 73
- Marinacci F., et al., 2018, *MNRAS*, **480**, 5113
- Markevitch M., Vikhlinin A., 2007, *Phys. Rep.*, **443**, 1
- McCourt M., Parrish I. J., Sharma P., Quataert E., 2011, *MNRAS*, **413**, 1295
- Naiman J. P., et al., 2018, *MNRAS*, **477**, 1206
- Navarro A., Khomenko E., Modestov M., Vitas N., 2022, *A&A*, **663**, A96
- Nelson D., et al., 2018, *MNRAS*, **475**, 624
- Nelson D., Pillepich A., Ayromlou M., Lee W., Lehle K., Rohr E., Truong N., 2024, *A&A*, **686**, A157
- Pakmor R., Springel V., 2013, *MNRAS*, **432**, 176
- Pakmor R., Bauer A., Springel V., 2011, *MNRAS*, **418**, 1392
- Pakmor R., Springel V., Bauer A., Mocz P., Munoz D. J., Ohlmann S. T., Schaal K., Zhu C., 2016a, *MNRAS*, **455**, 1134
- Pakmor R., Pfrommer C., Simpson C. M., Kannan R., Springel V., 2016b, *MNRAS*, **462**, 2603
- Pakmor R., et al., 2023, *MNRAS*, **524**, 2539
- Parrish I. J., Quataert E., 2008, *ApJ*, **677**, L9
- Parrish I. J., Stone J. M., 2005, *ApJ*, **633**, 334
- Parrish I. J., Stone J. M., 2007, *ApJ*, **664**, 135
- Parrish I. J., Quataert E., Sharma P., 2010, *ApJ*, **712**, L194
- Pellissier A., Hahn O., Ferrari C., 2023, *MNRAS*, **522**, 721
- Perrone L. M., Latter H., 2022a, *MNRAS*, **513**, 4605
- Perrone L. M., Latter H., 2022b, *MNRAS*, **513**, 4625
- Perrone L. M., Berlok T., Pfrommer C., 2024a, *arXiv e-prints*, p. arXiv:2402.06718
- Perrone L. M., Berlok T., Pfrommer C., 2024b, *A&A*, **682**, A125
- Pfrommer C., Dursi L. J., 2010, *Nature Physics*, **6**, 520
- Pillepich A., et al., 2018a, *MNRAS*, **473**, 4077
- Pillepich A., et al., 2018b, *MNRAS*, **475**, 648
- Planck Collaboration et al., 2020, *A&A*, **641**, A6
- Quataert E., 2008, *ApJ*, **673**, 758
- Riquelme M. A., Quataert E., Verscharen D., 2016, *ApJ*, **824**, 123
- Roberg-Clark G. T., Drake J. F., Reynolds C. S., Swisdak M., 2016, *ApJ*, **830**, L9
- Roberg-Clark G. T., Drake J. F., Reynolds C. S., Swisdak M., 2018, *Phys. Rev. Lett.*, **120**, 035101
- Ruszkowski M., Oh S. P., 2010, *ApJ*, **713**, 1332
- Ruszkowski M., Oh S. P., 2011, *MNRAS*, **414**, 1493
- Ruszkowski M., Lee D., Brüggen M., Parrish I., Oh S. P., 2011, *ApJ*, **740**, 81
- Saad Y., Schultz M. H., 1986, *SIAM Journal on Scientific and Statistical Computing*, **7**, 856

Schekochihin A. A., Cowley S. C., Kulsrud R. M., Hammett G. W., Sharma P., 2005, *ApJ*, **629**, 139

Schwarzschild M., 1958, Structure and evolution of the stars.

Sedov L. I., 1959, Similarity and Dimensional Methods in Mechanics

Sharma P., Hammett G. W., 2007, *Journal of Computational Physics*, **227**, 123

Sharma P., Hammett G. W., 2011, *Journal of Computational Physics*, **230**, 4899

Sharma P., Parrish I. J., Quataert E., 2010a, *ApJ*, **720**, 652

Sharma P., Parrish I. J., Quataert E., 2010b, *ApJ*, **720**, 652

Smedt B. D., Pattyn F., Groen P. D., 2010, *Journal of Glaciology*, **56**, 257–261

Spitzer L., 1962, Physics of Fully Ionized Gases

Springel V., 2010, *MNRAS*, **401**, 791

Springel V., Hernquist L., 2003, *MNRAS*, **339**, 289

Springel V., et al., 2018, *MNRAS*, **475**, 676

Springel V., Pakmor R., Zier O., Reinecke M., 2021, *MNRAS*, **506**, 2871

Su K.-Y., et al., 2019, *MNRAS*, **487**, 4393

Tevlin L., et al., 2024, *arXiv e-prints*, p. arXiv:2411.00103

Thomas T., Pfrommer C., Pakmor R., 2023, *MNRAS*, **521**, 3023

Tilley D. A., Balsara D. S., Howk J. C., 2006, *MNRAS*, **371**, 1106

Vogelsberger M., Sijacki D., Kereš D., Springel V., Hernquist L., 2012, *MNRAS*, **425**, 3024

Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, **436**, 3031

Voit G. M., 2011, *ApJ*, **740**, 28

Voit G. M., Donahue M., Bryan G. L., McDonald M., 2015, *Nature*, **519**, 203

Weinberger R., et al., 2017, *MNRAS*, **465**, 3291

Weinberger R., Springel V., Pakmor R., 2020, *ApJS*, **248**, 32

Yang H. Y. K., Reynolds C. S., 2016, *ApJ*, **818**, 181

Ye J., Shen C., Lin J., Mei Z., 2020, *Astronomy and Computing*, **30**, 100341

Yokoyama T., Shibata K., 1997, *ApJ*, **474**, L61

Zakamska N. L., Narayan R., 2003, *ApJ*, **582**, 162

Zel'dovich Y. B., Raizer Y. P., 1967, Physics of shock waves and high-temperature hydrodynamic phenomena

ZuHone J. A., Kunz M. W., Markevitch M., Stone J. M., Biffi V., 2015, *ApJ*, **798**, 90

van Leer B., 1984, *SIAM Journal on Scientific and Statistical Computing*, **5**, 1

APPENDIX A: ESTIMATING THE CONDUCTION COEFFICIENT IN CELL INTERFACES

In Section 3.1.1 we explained that we calculate the conductivity in a cell interface by taking the arithmetic average of those of the two cells on each side of the interface (see equation 20). To check that this estimate is reasonable we also carried out isotropic point explosion tests (see Section 4.2) where we estimated the conductivity in the interface by taking the harmonic average (see e.g. Sharma & Hammett 2007) of those of the two cells on either side

$$\chi_{ij,h} = \frac{1}{\frac{1}{2}(\frac{1}{\chi_i} + \frac{1}{\chi_j})}. \quad (\text{A1})$$

From equations (20) and (A1), we can see that, in the case of significant differences in conductivity between the cells, the arithmetic mean will tend towards the larger value of χ , while the harmonic mean tends towards the smaller.

We compare these two averaging procedures by carrying out simulations analogous to those presented in Section 4.2.1, i.e. in the isotropic blast wave setup. We do so because, in this setup, the speed of the conduction front can be predicted using analytical arguments. We show simulations at two different resolutions, 64^3 and 128^3 , for each of these averaging procedures, which we will refer to as ‘arithmetic’ and ‘harmonic’. All details of the simulations are as described in Section 4.2.




Figure A1. Evolution of the radius of the conduction/shock front in the 3D point explosion test with isotropic conduction. Solid/dashed coloured lines show results from simulations where the conductivity in cell interfaces was calculated using the arithmetic/harmonic mean. The black, dot-dash lines show the analytic expectation for the radial evolution of the shock/conduction front in the cases of pure hydrodynamics ($r_{s,3D}$, equation 42) and pure conduction ($r_{c,3D}$, equation 43), respectively. With increasing resolution, the radius of the conduction front converges with the analytic solution when the conductivity is calculated using the arithmetic mean, and diverges from the analytics when the harmonic mean is used.

In Fig. A1 we show the time evolution of the radius of the shock/conduction in these simulations. Solid/dashed coloured lines show results from simulations where the conductivity in cell interfaces was calculated using the arithmetic/harmonic mean. From this figure it is clear that the speed of the conduction fronts in the ‘harmonic’ case are slower than the analytic prediction, $r_{c,3D}$ and those measured in the ‘arithmetic’ simulations.

Additionally, as can be seen in Figures 6 and A1 (and was discussed in Section 4.2.1), with increasing resolution, the size of the conduction front is smaller (at a given time). In the ‘arithmetic’ case, this means that the radius of the conduction front converges towards the analytic prediction, but for the ‘harmonic’ case, the radius diverges further from the analytic solution.

The fact that the results for the ‘harmonic’ case do not converge to the analytic prediction, as the resolution of the initial conditions is increased, likely arises due to the way we set up the initial conditions. Specifically, the fact that the energy is injected into a fixed mass or equivalently volume, regardless of resolution means that, with increasing resolution, the initial conditions do not converge to that of a point explosion (as assumed by the analytic solution) and the initial discontinuity at the edge of the injection region persists. It is this initial discontinuity which is likely responsible for the offset of the results for the ‘harmonic’ case.

In the types of simulations to which we plan to apply this solver (e.g. galaxy cluster zooms), gradients are often unresolved and large discontinuities may arise. We, therefore, choose to calculate the interface conductivities via equation (20) as these results show that it performs better in scenarios typical of what we might expect in our simulations.

From Fig. A1 we see that the late-time behaviour is similar in all simulations. This is due to the fact that hydrodynamical processes dominate at late times (see discussion in Section 4.2.1). If this were

not the case, we would expect that the differences in the sizes of the conduction fronts, seen at early times, would persist.

This paper has been typeset from a \TeX/L\AA\TeX file prepared by the author.