

A dissipative extension to ideal hydrodynamics

Marcus John Hatton   and Ian Hawke

Mathematical Sciences and STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK

Accepted 2024 October 30. Received 2024 September 26; in original form 2024 July 29

ABSTRACT

We present a formulation of special relativistic dissipative hydrodynamics (SRDHD) derived from the well-established Müller–Israel–Stewart (MIS) formalism using an expansion in deviations from ideal behaviour. By re-summing the non-ideal terms, our approach extends the Euler equations of motion for an ideal fluid through a series of additional source terms that capture the effects of bulk viscosity, shear viscosity, and heat flux. For efficiency these additional terms are built from purely spatial derivatives of the primitive fluid variables. The series expansion is parametrized by the dissipation strength and time-scale coefficients, and is therefore rapidly convergent near the ideal limit. We show, using numerical simulations, that our model reproduces the dissipative fluid behaviour of other formulations. As our formulation is designed to avoid the numerical stiffness issues that arise in the traditional MIS formalism for fast relaxation time-scales, it is roughly an order of magnitude faster than standard methods near the ideal limit.

Key words: hydrodynamics – relativistic processes – methods: numerical – software: simulations – stars: neutron – neutron star mergers.

1 MOTIVATION

Binary neutron star mergers represent complex astrophysical laboratories that probe supernuclear-density matter, strong-gravity space-time and the origin of the heavy elements in our universe. Merging neutron stars produce multimessenger signals comprised of extraordinary electromagnetic and gravitational wave components, as confirmed by their detection in the GW170817 merger event by the LIGO–VIRGO–Kagra collaboration (Abbott et al. 2017). Since then, observations by these ground-based detectors have been used to put constraints on the mass, radius, and tidal deformability of neutron stars, informing us in turn about their equation of state (Abbott et al. 2018, 2019).

To simulate these events, we require a highly non-linear, general relativistic magnetohydrodynamic (GRMHD) model. This is then coupled to a space–time evolution procedure using numerical relativity, with neutrino transport and cooling schemes often added as well. For *simplicity*, the fluid of the neutron star is often treated as ‘ideal’, in that fluid stresses are purely isotropic.

Recently, however, more attention has been paid to effects resulting from non-ideal fluid behaviour, for example by Shibata, Kiuchi & Sekiguchi (2017), Rezzolla et al. (2018), Bemfica, Disconzi & Noronha (2019), Chabanov, Rezzolla & Rischke (2021), Pandya, Most & Pretorius (2022), and Yang et al. (2024). These dissipative effects arise due to out-of-equilibrium processes that are particularly important shortly after the neutron stars merge. When next-generation, ground-based gravitational wave detectors such as Cosmic Explorer (Reitze et al. 2019), NEMO (Ackley et al. 2020), LIGO-voyager (Berti et al. 2022), and the Einstein Telescope

(Punturo et al. 2010) come online, accurately modelling this next-to-leading-order behaviour will be essential to make precise physical inferences from observations.

Theoretical work (Chugunov & Yakovlev 2005; Manuel & Tolos 2011; Schmitt & Shternin 2018) and numerical investigations (Hammond, Hawke & Andersson 2021) have been undertaken into the out-of-equilibrium state of matter and its transport properties in neutron stars. For example, Urca and reverse-Urca nuclear reactions operate at an atomic scale and may give rise to an effective bulk viscosity at the fluid scale that quantitatively affects the gravitational wave signal we obtain from the merger and its remnant’s ringdown (Alford et al. 2018; Most et al. 2021, 2022; Hammond, Hawke & Andersson 2023). Similarly, work has been done to investigate the possible effects of both shear viscosity (Duez et al. 2004) and heat transport (Alford et al. 2018) in binary neutron star mergers, particularly for modulating the turbulence that ensues post-merger, both in the remnant itself and its associated accretion disc. Viscous braking redistributes momentum in a differentially rotating remnant, removing centrifugal support that can aid in the collapse of the core into a black hole. This produces delayed gravitational wave emission. It is also able to provide thermal support from viscous heating, negating this effect.

One well-established model of non-ideal hydrodynamics is that of Müller–Israel–Stewart (MIS; Israel 1976; Israel & Stewart 1979). Its theoretical properties have received thorough investigations (Molnár, Niemi & Rischke 2010; Biswas et al. 2020; Bemfica et al. 2021; Bemfica, Disconzi & Noronha 2022; Wagner & Gavassino 2024) and it has been used extensively in the context of high-energy, quark-gluon-plasma (QGP) physics to model post-collision fluid evolution (Del Zanna et al. 2013; Du & Heinz 2020), as well as in the astrophysical community for modelling viscous black hole accretion (Chabanov et al. 2021), for example.

* E-mail: mjh1n20@soton.ac.uk

The MIS model includes viscous and heat-conductive effects in the evolved conserved and flux vectors, as well as relaxation-type sources that drive the non-ideal terms to relativistic analogues of their Navier–Stokes forms. A numerical issue arises when the dissipative relaxation time-scales become small and the sources become ‘stiff’. The relaxation time-scales tend to zero in the ideal limit, which is relevant for the majority of the lifecycle of a binary neutron star merger. One must either reduce the time-step of the simulation drastically or adopt implicit time-integrator methods to ensure accurate and stable numerical evolution. See Palenzuela et al. (2009), Dionysopoulou et al. (2013), Miranda-Aranguren, Aloy & Rembiasz (2018), Ripperda et al. (2019), Wright (2020), Wright & Hawke (2020), and Dash et al. (2023) for examples of approaches taken to evolve stiff numerical systems. Both options increase the computational cost of simulations greatly. Sometimes, ‘best-of-both’ implicit–explicit methods (Pareschi & Russo 2005) may be used but in any case, the computational cost increases, potentially by orders of magnitude, when source terms become stiff near the ideal limit.

This in turn limits the spatial resolution of simulations that are performed, leading to coarse numerical grids that represent fluid elements with sizes well above those that ‘should’ be used to satisfy the fluid approximation. That is to say, there is significant variation in fluid properties occurring over length-scales well below that of the grid cells’ size. Estimates of the dissipation length-scale above which structure can form through turbulence suggest that simulations may need to resolve scales below the cm level (Thompson & Duncan 1993; Radice & Hawke 2024). However, the current highest-resolution simulations have fluid elements with sizes ≈ 10 m (Kiuchi et al. 2018). To bridge this gap computationally is impractical for the foreseeable future.

Instead, to address this ‘sub-grid’ behaviour, extensions to existing hydrodynamic models have recently begun being employed. These additions to the model aim to capture, at least in a statistical sense, either genuine sub-grid microphysics or mathematical artefacts resulting from the implicit filtering process introduced by coarse simulations.

Sub-grid models are beginning to see a number of applications in modelling astrophysical systems. The general principle behind these extensions is to include additional terms into the equations of motion, aimed at capturing the effects of unresolved fluid behaviour at scales below that which can be directly resolved in a numerical simulation. The benefit of these models lies in their ability to, without greatly increased computational cost, capture the influence of unresolvable microphysics or fluctuations, at least in a statistical sense.

A common application of sub-grid sources is in the modelling of turbulence. In large-eddy simulations, the equations of motion are explicitly redefined in terms of resolved and unresolved quantities. A closure relation is then applied that allows the sub-grid fields to be formulated in terms of the resolved ones. Using this technique, it is possible to replicate the behaviour that would result, on average, from using more fine-scale numerical grids.

For instance, Radice (2017) first applied an analogue of the classical Smagorinsky closure (Smagorinsky 1963) to the equations of general relativistic hydrodynamics for a merger simulation, showing that by modelling the sub-grid scale turbulence, the collapse of the hypermassive neutron star remnant is altered. Other work by Carrasco, Viganò & Palenzuela (2020) and Viganò et al. (2020) uses a gradient expansion approach to prescribe the unresolved fields in the MHD equations. See Radice & Hawke (2024) for a modern review of the field.

One might ask why these sub-grid models are relevant to the non-ideal hydrodynamic formulation presented here. In Celora et al.

(2021), it is shown that when a linear, covariant filtering operation is applied to an ideal fluid formulation, the fine-scale variation that is spatially averaged over may be described on the coarse scale by algebraic terms that mimic those present in a non-ideal fluid formulation. The corollary of this is that we may use our models of non-ideal hydrodynamics to describe a fluid that does not genuinely exhibit dissipative effects, at least not on the coarse scale at which we simulate it, but instead to capture unresolved effects due to resolution limitations. In effect, our sub-grid closure relation is given by the model’s prescription for the non-ideal dissipation terms within it. Of course, the meaning of the ‘dissipative’ terms changes when we do this. Instead, they now capture the effects of filtering, and follow-up work will be published investigating this.

In this paper, we develop an extension to the special relativistic, ideal hydrodynamic equations that captures the dissipative effects present in full non-ideal fluid descriptions. This extension, dubbed a dissipative extension to ideal fluid dynamics (DEIFY), is derived from first principles arguments and, as such, requires no fine tuning of parameters for different astrophysical scenarios. The rest of the paper is laid out as follows. In Section 2, we introduce the hydrodynamic models we are concerned with: first, zeroth-order ideal hydrodynamics; then, second-order, non-ideal hydrodynamics in the MIS formulation. Section 3 introduces the Chapman–Enskog (CE) expansion we use here and derives a number of simple models to demonstrate the pertinent points. Section 4 presents the full ‘MISCE’ model with its source derived from applying the CE expansion to the MIS model. In Section 5, we show results of simulations that use the MISCE formulation of dissipative hydrodynamics. In particular, we quantitatively compare results and performance with the MIS model. Our appendices cover considerations about initial data and stability for the MISCE model, as well as how one may calculate time derivatives of primitive fluid variables without using lagged updates. Finally, in Section 7, we summarize the findings of the previous sections, discuss how they fit into current astrophysical simulations, and propose the future direction of the project.

2 HYDRODYNAMIC MODELS

In this section, we outline two models of hydrodynamics that are used in relativistic astrophysics. In order to simplify the numerics in later sections and to test the validity of the method, we will limit ourselves to special relativity. In moving to a general relativistic description, only the form of the equations should change, and so the analysis we perform here should still apply. We will also adopt the Einstein summation convention over repeated indices, where Greek letters run over 4 indices (1 temporal and 3 spatial) and Roman letters run over 3 (spatial) indices. δ_j^i is the Kronecker delta (3,3)-tensor. We use units where $c = 1$ throughout.

2.1 Ideal hydrodynamics

The first model we present is that of ideal, non-dissipative hydrodynamics. This is the simplest relativistic model of fluids that one can write down. The stress-energy tensor for such a fluid is

$$T^{\mu\nu} = (\rho + p)u^\mu u^\nu + pg^{\mu\nu}, \quad (1)$$

where ρ is the energy density of the fluid, p is its pressure, u^μ is its 4-velocity, and $g^{\mu\nu}$ is the metric tensor defining the space–time geometry. In this paper, we only work in flat space–time such that the metric is Minkowski and $g^{\mu\nu} \rightarrow \eta^{\mu\nu} = \text{diag}(-1, +1, +1, +1)$.

We enforce conservation of energy, momentum, and conserved current(s) by demanding that the covariant derivative of the stress-

energy tensor and of a conserved number-current $N^\mu = nu^\mu$ both vanish. Mathematically, the equations

$$\nabla_\mu T^{\mu\nu} = 0, \quad (2a)$$

$$\nabla_\mu N^\mu = 0 \quad (2b)$$

lead to the ideal fluid equations of motion. In conservative form, these are

$$\partial_t \begin{pmatrix} D \\ S_j \\ \tau \\ X \\ Y_j \\ Z_{jk} \end{pmatrix} + \partial_i \begin{pmatrix} Dv^i \\ S_j v^i + p\delta_j^i \\ \tau v^i + p v^i \\ \end{pmatrix} = 0, \quad (3)$$

where the three conserved quantities, $\{D, S_j, \tau\}$, correspond to the fluid density, specific momentum in the j^{th} -direction, and kinetic energy density, respectively, and are related to the primitive quantities, $\{n, v_j, p\}$, namely the baryon number density, fluid 3-velocity and hydrodynamic pressure via

$$D = nW, \quad (4a)$$

$$S_j = (\rho + p)W^2 v_j, \quad (4b)$$

$$\tau = (\rho + p)W^2 - D. \quad (4c)$$

Additionally, we have that $\rho = mn(1 + \epsilon)$, where m is the mass per baryon and ϵ the specific internal energy. The hydrodynamic pressure p is given by an equation of state to close the system: numerically, we often write $p \equiv p(n, \epsilon)$ but here it is more convenient to use $p \equiv p(n, \rho)$. Specifically, we will use a Gamma-law equation of state throughout this paper of the form $p = (\Gamma - 1)(\rho - mn)$, which typically describes an ideal relativistic gas.

Two more important thermodynamic quantities of interest are specific enthalpy $h = 1 + \epsilon + p/(mn) = (\rho + p)/(mn)$ and temperature $T = p/n$ (as well as its inverse $\beta = m/T$). Note that we assume a uniform baryon mass so scale it out of the equations such that $m = 1$ throughout this paper. Finally, the spatial three-velocity v_j and the Lorentz factor $W = (1 - v_j v^j)^{-1/2}$ make up the four-velocity $u_\mu = W(1, v_j)$.

2.2 Dissipative hydrodynamics within Müller–Israel–Stewart formalism

The second model we present is that describing a non-ideal fluid with a stress-energy tensor given by

$$T^{\mu\nu} = (\rho + p + \Pi)u^\mu u^\nu + (p + \Pi)g^{\mu\nu} + q^\mu u^\nu + q^\nu u^\mu + \pi^{\mu\nu}, \quad (5)$$

where the new dissipative terms are the bulk viscosity pressure Π , the heat flux vector q^μ , and the shear viscosity tensor $\pi^{\mu\nu}$. The first of these encapsulates isotropic stresses (compression and expansion). The second, momentum transport orthogonal to fluid's velocity. And the third, anisotropic stresses within the fluid. To set their form, and derive the equations of motion of the fluid, we follow the MIS (Israel 1976; Israel & Stewart 1979) formalism. Their approach, in words, involves first performing a gradient expansion of entropy-generating terms to second-order. Then, applying the second law of thermodynamics such that the entropy-generation rate is always non-negative and using this condition to set the form of the non-ideal, dissipative terms. Finally, the additional degrees of freedom this introduces are wrapped-up into six non-ideal coefficients (three dissipation 'strengths' and three time-scales). This gives us the model that is our starting point and the one to which we will apply a CE type expansion, eventually giving us our new 'MISCE' model.

The entire MIS equations, in balance-law form, are given as

$$\partial_t \mathbf{U} + \partial_i \mathbf{F}^i(\mathbf{U}) = \mathbf{S}, \quad (6)$$

where

$$\mathbf{U} = \begin{pmatrix} D \\ S_j \\ \tau \\ X \\ Y_j \\ Z_{jk} \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi)W^2 v_j + W(q_0 v_j + q_j) + \pi_{0j} \\ (\rho + p + \Pi)W^2 + 2q_0 W - (p + \Pi - \pi_{00}) - D \\ nW\Pi \\ nWq_j \\ nW\pi_{jk} \end{pmatrix}, \quad (7a)$$

$$\mathbf{F}^{(i)} = \begin{pmatrix} Dv^i \\ S_j^i - Dv^i \\ Xv^i \\ Y_j v^i \\ Z_{jk} v^i \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad (7b)$$

$$\mathbf{S} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{n}{\tau\Pi}(\Pi_{NS} - \Pi) \\ \frac{n}{\tau_q}(q_{j,NS} - q_j) \\ \frac{n}{\tau_\pi}(\pi_{jk,NS} - \pi_{jk}) \end{pmatrix} \quad (7c)$$

and

$$S_j^i = (\rho + p + \Pi)W^2 v^i v_j + W(q^i v_j + q_j v^i) + (p + \Pi)\delta_j^i + \pi_j^i. \quad (8)$$

The heat flux and shear viscosity are orthogonal to the four velocity on all indices, and the shear viscosity is trace free, implying

$$q_0 = v^k q_k, \quad (9a)$$

$$\pi_{0j} = v^k \pi_{kj}, \quad (9b)$$

$$\pi_{j0} = v^k \pi_{jk}, \quad (9c)$$

$$\pi_{00} = -\pi_k^k. \quad (9d)$$

The first-order, relativistic 'Navier–Stokes' terms to which the dissipative system relaxes are

$$\Pi_{NS} = -\zeta \Theta, \quad (10a)$$

$$q_{j,NS} = -\kappa T(\delta_j \log T + a_j), \quad (10b)$$

$$\pi_{jk,NS} = -2\eta \sigma_{jk}, \quad (10c)$$

where the non-ideal coefficients of bulk viscosity, heat conductivity, and shear viscosity are ζ, κ , and η , respectively, which we may collectively represent as ξ . The following quantities:

$$\Theta = \partial_\mu u^\mu, \quad (11a)$$

$$a_\mu = u^\nu \partial_\nu u_\mu, \quad (11b)$$

$$\sigma_{\mu\nu} = (\partial_\mu u_\nu + \partial_\nu u_\mu - \frac{2}{3}\eta_{\mu\nu}\Theta) \quad (11c)$$

are the expansion, acceleration, and shear of the 4-velocity, respectively.

The first three conserved quantities ($\{D, S_j, \tau\}$), and their associated equations of motion, form the non-stiff subsystem, labelled \mathbf{q} , which reduces to the Euler equations (in the form of equation 3)

in the ideal limit of zero dissipation ($\xi \rightarrow 0$). The remaining three conserved quantities, $\{X, Y_j, Z_{jk}\}$ are labelled \bar{q} . These quantities are evolved with a source that is proportional to the reciprocal of a possibly small time-scale and may therefore represent a stiff subsystem. Note that the terms q, \bar{q} should not be confused with the heat flux, which will intentionally be written in component form, q_j , for clarity.

2.2.1 Asymptotic behaviour

One major difference between the two models of hydrodynamics that are presented here lies in the form of their source terms. Whilst the ideal model's source vector is entirely zero, the Israel-Stewart model's source vector has non-zero sources for the dissipative evolution components. These terms are proportional to the reciprocal of the relaxation time-scales, τ . In the limit as $\tau \rightarrow 0$, any deviation of the dissipative variables from their equilibrium Navier-Stokes form will be instantaneously quenched. This represents a reduction to a first-order theory, essentially a relativistic version of the classical, dissipative Navier-Stokes equations. However, we may link the dissipation time-scales and strengths, for example through thermodynamic relations for a Boltzmann gas as in Israel (1976) where

$$\tau_\Pi = \zeta \beta_0, \quad (12a)$$

$$\tau_q = \kappa \beta_1 T, \quad (12b)$$

$$\tau_\pi = 2\eta \beta_2 \quad (12c)$$

and the β terms are non-negative thermodynamic functions of the enthalpy, temperature, and pressure given therein. Alternatively, there are analytical bounds on the ratio of dissipation strengths: time-scales from enforcing causality due to recent work by Heller et al. (2023). One can see this practically from the fact that dissipation modifies the characteristic propagation speeds of waves travelling in the fluid (sound speeds). For instance, by a factor $\propto \sqrt{\zeta/\tau_\Pi}$ for bulk viscosity as seen in Chabakov et al. (2021). This, of course leads to a divergence when we take the $\tau \rightarrow 0$ limit without taking the $\xi \rightarrow 0$ limit along with it. In reality of course, there are no instantaneous processes. However, there are physically and mathematically motivated reasons why taking the instantaneous-relaxation limit also implies taking the zero-dissipation limit, in which case one recovers ‘zeroth-order’ (ideal) hydrodynamics as in Section 2.1.

2.2.2 Numerical difficulties

When the time-scales that the source acts on are shorter than the time-step of the simulation, $\tau \lesssim \Delta t$, the system is said to be *stiff*. In order to maintain a stable evolution, one may either reduce the size of the time-step used in the simulation, or employ a set of implicit or semi-implicit time integrators such as those seen in Pareschi & Russo (2005). In the first case where the time-step, the execution time will increase by a factor $\approx \Delta t/\tau$, making it impractical for dissipation acting well below the hydrodynamic time-scale. In the latter case where (semi-) implicit methods are used, the time-integrator algorithm is considerably, maybe orders of magnitude, more costly.

Hence, our motivation is to find a source term that captures dissipative behaviour but avoids the numerical difficulties of the stiff MIS system of equations. The following section will derive such a source term using a CE-type analysis.

First we introduce the notation that will be used. We can re-write the conservative form of the MIS equations, equation (7) in the following, more compact way

$$\partial_t q + \partial_i f^i(q, \bar{q}) = s(q, \bar{q}) = 0, \quad (13a)$$

$$\partial_t \bar{q} + \partial_i \bar{f}^i(q, \bar{q}) = \frac{\bar{s}(q, \bar{q})}{\tau}, \quad (13b)$$

where we indicate equations that become stiff as $\tau \rightarrow 0$ with an over-bar. This means that $\bar{q} = \{U, Y_j, Z_{jk}\}$ with the corresponding fluxes, $\bar{f}^i(q, \bar{q})$, and sources, $\bar{s}(q, \bar{q})$, taken from equation (7). The remaining conserved variables are non-stiff in the ideal limit, and denoted $q = \{D, S_j, \tau\}$. We will also denote the vector of primitive variables present in ideal hydrodynamics as $w = \{p, \rho, n, v_x, v_y, v_z\}$ and the dissipative primitive variables as $\bar{w} = \{q_j, \Pi, \pi_{jk}\}$.

3 CHAPMAN-ENSKOG EXPANSION

In this section, we will use a type of CE method of expansion to derive the form of the DEIFY source term. Let us first give a brief introduction to the historical use of the CE expansion for solving the Boltzmann equation.

The Boltzmann equation provides a statistical description of atomic-scale particle kinetics for a thermodynamic system in a state of non-equilibrium. It describes the evolution of the one-particle distribution function, f , and may be written as

$$p^\mu \partial_\mu f(x, p) = \mathcal{C}[f], \quad (14)$$

where p^μ is the 4-momentum of the particle and $f(x, p)$ is the 8D distribution function which depends on the particle's 4-position and 4-momentum. Finally, $\mathcal{C}[f]$ is the collision term, an integral operator (in momentum space) whose precise form depends on the interactions and statistics of the particle ensemble being considered. The distribution function counts the particles and is normalized such that the number density $n(x)$ is given by

$$n(x) = \int \frac{d^3 p}{(2\pi)^3} f(x, p) \quad (15)$$

whilst familiar quantities such as the number current and stress-energy tensor are given by

$$N^\mu = \int_p p^\mu f(x, p), \quad (16a)$$

$$T^{\mu\nu} = \int_p p^\mu p^\nu f(x, p), \quad (16b)$$

where $\int_p \dots \equiv \int d^3 p / [(2\pi)^3 p^0]$ is the Lorentz-invariant integration measure.

In the early 20th century, Sydney Chapman and David Enskog introduced a method for solving the Boltzmann equation (for f), as seen in Chapman & Cowling (1990) and Kumar (1967). Their approach allows one to derive macroscopic descriptions of fluids, including the familiar Euler and Navier-Stokes equations (Cercignani & Kremer 2002), and obtain expressions for various transport coefficients such as thermal conductivity and bulk viscosity in the process.

One of the central assumptions of the CE analysis is that the particle collision duration time is not only far less than the time between collisions, but also far smaller than *extrinsic* time-scales: the time-scales associated with macroscopic fluid behaviour. The solution is in fact obtained as an asymptotic expansion in the

closely related Knudsen number¹, where at zero'th order we have the equilibrium solution.

Work has been done in the field of resistive MHD (Wright & Hawke 2019, 2020) by applying this method to describe a system where a source acts rapidly to relax the solution towards equilibrium (in that case, ideal MHD). A simple but general application of this type of expansion is presented in LeVeque (2002), in which LeVeque demonstrates how a coupled system of balance law equations may be reduced to a single, modified system. One of the coupled equations contains a potentially stiff source term, whilst the other is an advection equation. The reduced system is a balance law with a derivative source term that is non-stiff. Hence, this example represents a simplified version of the system we have here, seen in equation (13). Next, we will show how this method works for a heat-flux model obtained from the MIS equations.

3.1 A simple heat model

To demonstrate the approach, we first apply it to a simple model governing two variables: the temperature, T and the spatial heat flux, q_i . To obtain this model, which will be familiar once derived, we apply a number of simplifying assumptions to the MIS model.

We work with a static fluid such that the 3-velocity and, hence, the bulk and shear viscosity all vanish. All Lorentz factors become unity and the orthogonality relation $v_\mu q^\mu = 0$ means that $q_0 = 0$. We treat our particle number current, n , and hydrostatic pressure, p , as constants and hence the density, ρ , is now purely a function of the temperature i.e. $\rho \equiv \rho(T)$ and may be scaled out of the equations. After setting any remaining constant terms to one, we arrive at

$$\partial_t T + \partial_i q^i = 0, \quad (17a)$$

$$\partial_t q_j = \frac{1}{\tau_q} (q_{j,NS} - q_j). \quad (17b)$$

We note that the acceleration term usually present in the heat-flux's source will vanish so that $q_{j,NS} \rightarrow -\kappa \partial_j T$ and we obtain a rather simple pair of equations where the first has no source and the second, the 'Maxwell–Cattaneo' equation (Cattaneo 1948), has no flux:

$$\partial_t T + \partial_i q^i = 0, \quad (18a)$$

$$\partial_t q_j = -\frac{1}{\tau_q} (\kappa \partial_j T + q_j). \quad (18b)$$

When τ_q is small, the heat flux, q_j , will relax to its equilibrium value, $q_{j,NS}$, rapidly, with small deviations being modulated by the size of τ_q . Hence, we first write the non-ideal variable q_j that we wish to eliminate from the system in terms of its equilibrium value $q_{j,NS} = -\kappa \partial_j T$ and a (small) correction term of order τ_q , so

$$q_j = q_{j,NS} + \tau_q q_j^{(1)}, \quad (19)$$

where $q_j^{(1)}$ is to be determined. We can then write the pair of equations, to first order in τ_q , as

$$\partial_t T + \partial_i q_{NS}^i = -\partial^i (\tau_q q_i^{(1)}), \quad (20a)$$

$$\partial_t q_{j,NS} = -q_j^{(1)}. \quad (20b)$$

¹The Knudsen number (Kn) is a dimensionless number defined as the ratio of the molecular mean free path length to a representative physical length-scale. This length-scale could be, for example, the radius of a rigid body emersed in a fluid.

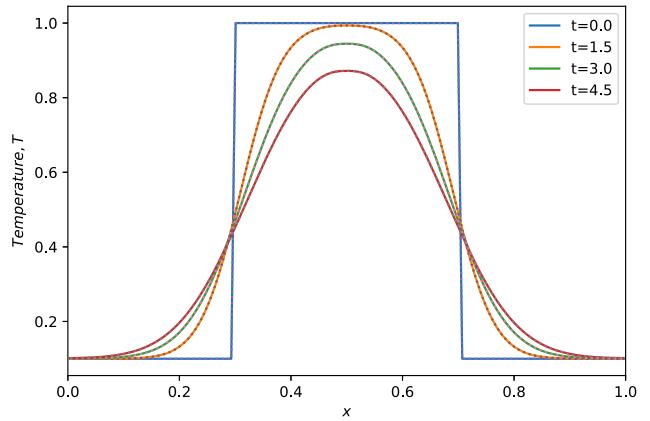


Figure 1. The widening of a top-hat temperature profile through dissipation by the heat flux within our toy model, both with (dotted lines) and without (continuous lines) the CE expansion. There is an excellent agreement between the two, with the small numerical errors being $\mathcal{O}(\tau_q)$.

By using the explicit form for the equilibrium value $q_{j,NS} = -\kappa \partial_j T$, we can write this as

$$\partial_t T = \partial^i (\kappa \partial_i T) - \partial^i (\tau_q q_i^{(1)}), \quad (21a)$$

$$\partial_t (-\kappa \partial_j T) = -q_j^{(1)} \quad (21b)$$

to obtain an expression for $q_j^{(1)}$, but one which includes a temporal derivative.

For simplicity, we will assume that κ and τ_q are constants in time and space. By commuting the temporal and spatial derivatives in equation (21b), we can now substitute the leading order form (zero'th order in τ_q) of the equation of motion for T , equation (21a), into the relaxation equation (21b), to determine the correction $q_j^{(1)} = \kappa^2 \partial_j^{(3)} T$ as purely spatial derivatives. Inserting this result back into equation (21a) and writing the result in one spatial dimension, we finally have the CE form

$$\partial_t T = \kappa \partial_x^{(2)} T - \kappa^2 \tau_q \partial_x^{(4)} T + \mathcal{O}(\tau_q^2). \quad (22)$$

Note that this result is an evolution equation written purely in terms of the temperature, T , and is half the size of the original system given by equation (18). This reduced model is non-stiff (as the source term is multiplied by the small time-scale τ_q , not by its reciprocal), but may lead to other numerical problems due to the higher derivatives. These features will hold true when we apply the CE expansion method to the full MIS equations, whereby dissipation will be modelled using only the primitive variables and their derivatives, and the system size will be reduced significantly.

Equation (22) is essentially a power series expansion in $\{\kappa, \tau\}$ where the leading term alone ($\mathcal{O}(\kappa)$) gives us the 1D heat equation. With the next-to-leading order term, ($\mathcal{O}(\tau_q \kappa^2)$) we have a linear, diffusion-retention equation. That is to say, the second-order derivatives represent diffusive effects that spread heat isotropically and the fourth-order derivatives retain heat locally. The numerical significance of these higher-order terms are described well in (Bevilacqua, Galeão & Costa 2011).

In Fig. 1, we use a 'top-hat' initial temperature profile to compare the CE model equation (22) to its originating equations (equation 18). We see excellent agreement between the two. The performance improvement will be quantified below.

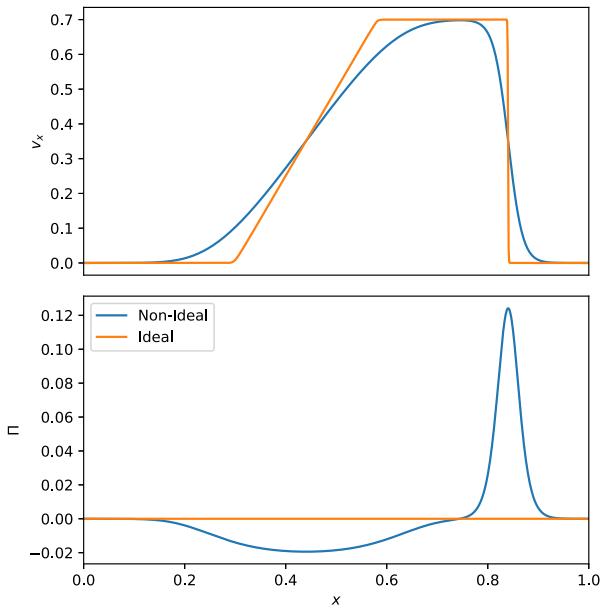


Figure 2. The velocity and bulk viscosity for a 1D test using the model given by equation (24) (with and without the ‘right-hand side’ terms) and an initially discontinuous velocity profile. One can see that bulk viscosity can be both positive and negative and in general acts to smooth gradients in the velocity, as well as accelerating the leading edge of the shock. $t_{\text{final}} = 0.2$ and the non-ideal simulation ran with parameters $\zeta = 1 \times 10^{-2}$ and $\tau_{\Pi} = 1 \times 10^{-3}$.

3.2 Simple viscosity models

Similar ‘toy’ models that govern the evolution of bulk and shear viscosity may be derived from the MIS equations. For the bulk viscosity, in one dimension the relevant equations are

$$\partial_t v_j + \partial_i (v^i v_j + \Pi \delta_j^i) = 0, \quad (23a)$$

$$\partial_t \Pi + \partial_i (\Pi v^i) = -\frac{1}{\tau_{\Pi}} (\zeta \partial_k v^k + \Pi), \quad (23b)$$

which reduces to the single 1D equation

$$\partial_t v + \partial_x (v^2) = \zeta \partial_{xx} v + \tau_{\Pi} \zeta (\partial_x v \partial_{xx} v - v \partial_{xxx} v) \quad (24)$$

when the CE expansion is performed. At leading order ($\tau_{\Pi} = 0$) this is essentially the viscous Burgers equation, only missing a factor of one half in the flux compared to its canonical form. In Fig. 2, we test this model (at next-to-leading order). The initial velocity profile is given by a top-hat function with $v_x = 0.7$ for $0.3 \leq x \leq 0.7$ and $v_x = 0$ otherwise. The bulk viscosity is calculated by its CE-expanded form as $\Pi = \Pi_0 + \tau_{\Pi} \Pi_1$, where $\Pi_0 = \Pi_{NS} = -\zeta \partial_x v_x$ and $\Pi_1 = \zeta [\partial_x (v \partial_x v) - \partial_{xx} v^2]$.

For the shear viscous case, if we work in two spatial dimensions but consider a purely y -directed flow ($v_x = 0$) we have

$$\partial_t v_y + \partial_x \pi_{xy} = 0, \quad (25a)$$

$$\partial_t \pi_{xy} = -\frac{1}{\tau_{\Pi}} (2\eta \sigma_{xy} + \pi_{xy}). \quad (25b)$$

Given that σ_{xy} reduces to $\partial_x v_y$, we have

$$\partial_t v_y = 2\eta \partial_x^{(2)} v_y - 4\eta^2 \tau_{\Pi} \partial_x^{(4)} v_y \quad (26)$$

as the CE form that governs the shear damping of the fluid velocity. This has an analytical solution in the limit $\tau_{\Pi} \rightarrow 0$ given by $v_y(t, x) \sim v_y(0, x) \text{erf}(\frac{x}{\eta})$, which allows us to test convergence of the CE expression as the dissipation time-scale vanishes. In Fig. 3, this

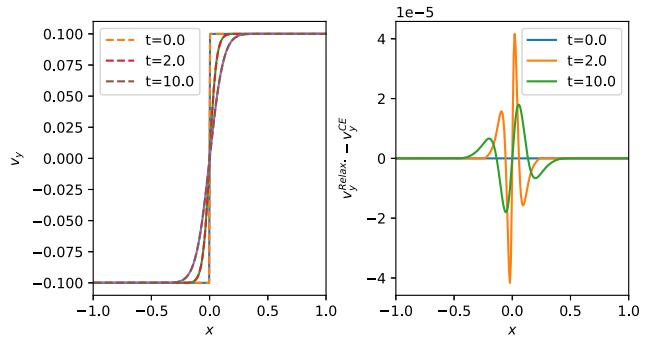


Figure 3. The evolution of the y -directed component of the velocity plotted across the x -domain at times $t = 0.0, 2.0, 10.0$. The initial data for the velocity form a step function and the viscous parameter values are $\eta = 2 \times 10^{-4}$ and $\tau_{\Pi} = 2 \times 10^{-4}$. In the left panel, two models’ results are plotted: first, the MIS-derived simple shear model given by equations (25; solid lines); secondly, the model obtained applying the CE-expansion to it, given by equation (26) (dotted lines). In the right panel, the difference between the two results is plotted. The shear viscosity damps the initial step function, causing the velocity to develop approximately according to the analytical error-function. There is no visual difference seen between the two models. The numerical difference is an order of magnitude smaller than even the value of $\mathcal{O}(\tau_{\Pi})$ and is decreasing in time.

Table 1. Code run-times for our simple bulk viscosity model in one dimension, simulating the initial data seen in Fig. 3. The expected scaling of the run-time, $t_{\text{run}} \propto N_x^2$, can be seen for both the full model and the CE expansion (the latter in parentheses). Crucially, a nearly three-fold speed-up is achieved by using the CE expansion, even with identical numerical methods.

Grid N_x	t_{final}	Model Runtimes [s]			Scaling	Speed-up
		Full	CE			
2048	1.0	5.40	1.60		– (–)	3.4
4096	1.0	21.92	7.77		4.1 (4.9)	2.8
8192	1.0	92.07	33.17		4.2 (4.3)	2.9

test case is plotted, and again excellent agreement is seen between the original relaxation model, the CE-expanded one.

Finally, Table 1 gives a comparison of run-times that indicates a significant speed-up is achievable using the CE form without compromising on accuracy. This speed-up of a factor between 2 and 3 comes purely from the reduced system size – both models are evolved with the same, explicit time-integrators. This is possible for simple systems like this one, but is often not in the case of the MIS model, depending on the chosen parameter values. For the full MIS model where more expensive time integrators would be needed, the speed-up will be larger, as investigated below.

4 GENERAL BALANCE-LAW DERIVATION

We begin with the full MIS model in balance law form, equation (13), recalling that the non-stiff and stiff conserved variables are labelled \mathbf{q} and $\bar{\mathbf{q}}$, respectively. In order to maintain finite solutions in the ideal limit, we require that $\lim_{\tau \rightarrow 0} \bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}) = \mathbf{0}$. This motivates an expansion of the stiff variables in powers of τ , with each increasing order providing a further deviation from the ideal limit.

In fact, because of the mathematical and physical links between the dissipation time-scales (τ) and strengths (ξ) discussed earlier, we choose to perform the expansion in powers of ϵ where both τ and ξ are $\mathcal{O}(\epsilon)$. This reflects that in practice these parameters often take on similar (small) values in numerical simulations. This choice also

means that $\mathcal{O}(\epsilon^0)$ corresponds to ideal behaviour with no dissipation. We could equally well perform the expansion in powers of τ and arrive at the same result, simply with a shifted series definition. Moving on, we now have

$$\bar{\mathbf{q}} = \bar{\mathbf{q}}_0 + \bar{\mathbf{q}}_1 + \bar{\mathbf{q}}_2 + \mathcal{O}(\epsilon^3), \quad (27)$$

where $\bar{\mathbf{q}}_0$ is $\mathcal{O}(\epsilon^0)$, $\bar{\mathbf{q}}_1$ is $\mathcal{O}(\epsilon)$, $\bar{\mathbf{q}}_2$ is $\mathcal{O}(\epsilon^2)$, and so on. To identify the terms in this expansion, we take the form of the stiff source:

$$\bar{s}(\mathbf{q}, \bar{\mathbf{q}}) = \frac{1}{\epsilon}(\bar{\mathbf{q}} - \bar{\mathbf{q}}_{NS}) \quad (28)$$

and simply rewrite it as

$$\epsilon \bar{s}(\mathbf{q}, \bar{\mathbf{q}}) = \bar{\mathbf{q}} - \bar{\mathbf{q}}_{NS}. \quad (29)$$

Noting that \mathbf{q}_{NS} is $\mathcal{O}(\epsilon)$, we have that at zeroth order (the ideal case) $\bar{\mathbf{q}} = \bar{\mathbf{q}}_0 = \mathbf{0}$. At first order we have $\bar{\mathbf{q}} = \bar{\mathbf{q}}_1 = \bar{\mathbf{q}}_{NS}$ and at second order we have $\bar{\mathbf{q}} = \bar{\mathbf{q}}_2$ where $\bar{\mathbf{q}}_2$ is yet to be determined.

At zeroth order, the non-stiff subsystem of equations is given by

$$\partial_t \mathbf{q}_0(\mathbf{w}) + \partial_i \mathbf{f}_0^i(\mathbf{w}) = \mathbf{0}, \quad (30)$$

where

$$\mathbf{q}_0(\mathbf{w}) = \begin{pmatrix} D \\ S_j \\ \tau \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p)W^2 v_j \\ (\rho + p)W^2 - p - nW \end{pmatrix} \quad (31)$$

and

$$\mathbf{f}_0^i(\mathbf{w}) = \begin{pmatrix} nW v^i \\ (\rho + p)W^2 v^i v_j + p \delta_j^i \\ (\rho + p)W^2 v^i - nW v^i \end{pmatrix}. \quad (32)$$

These are simply the relativistic Euler equations. At first order, it can be written as

$$\begin{aligned} \partial_t [\mathbf{q}_0(\mathbf{w}) + \mathbf{H}_{(1)}(\mathbf{w}, \partial_t \mathbf{w}, \partial_i \mathbf{w})] \\ + \partial_i [\mathbf{f}_0^i(\mathbf{w}) + \mathbf{F}_{(1)}^i(\mathbf{w}, \partial_t \mathbf{w}, \partial_i \mathbf{w})] = 0, \end{aligned} \quad (33)$$

where

$$\mathbf{H}_{(1)}(\mathbf{w}) = \begin{pmatrix} 0 \\ W(q_{0,NS} v_j + q_{j,NS}) + \pi_{0j,NS} \\ 2q_{0,NS} W - \Pi_{NS} + \pi_{00,NS} \end{pmatrix} \quad (34)$$

and

$$\mathbf{F}_{(1)}^i(\mathbf{w}) = \begin{pmatrix} 0 \\ \Pi_{NS}(v^i v_j + \delta_j^i)W^2 + W(q_{NS}^i v_j + q_{j,NS} v^i) + \pi_{j,NS}^i \\ W(q_{0,NS} v^i + q_{NS}^i) + \pi_{0,NS}^i \end{pmatrix}. \quad (35)$$

Here, we have separated the dissipative parts of the state and flux vectors and can view $\mathbf{H}_{(1)}$ and $\mathbf{F}_{(1)}$ as $\mathcal{O}(\epsilon)$ perturbations on-top of the ideal $\mathcal{O}(\epsilon^0)$ state and flux vectors, \mathbf{q}_0 and \mathbf{f}_0 . In general, we can rewrite the expanded system as

$$\partial_t \mathbf{q}_0 + \partial_i \mathbf{f}_0^i = \sum_{p=0} \tilde{\mathbf{R}}_{(p)} \equiv \sum_{p=0} (-\partial_t \mathbf{H}_{(p)} - \partial_i \mathbf{F}_{(p)}^i), \quad (36)$$

where each additional term in the series on the RHS of equation (36) represents a source correction of order ϵ^p . Hence, $\tilde{\mathbf{R}}_{(0)} = \mathbf{0}$, $\tilde{\mathbf{R}}_{(1)} = -\partial_t \mathbf{H}_{(1)} - \partial_i \mathbf{F}_{(1)}^i$, and $\tilde{\mathbf{R}}_{(2)} = -\partial_t \mathbf{H}_{(2)} - \partial_i \mathbf{F}_{(2)}^i$ and so on.

Using symbolic Python, we have fully derived the first-order ($\mathcal{O}(\epsilon)$) source terms in $\tilde{\mathbf{R}}_{(1)}$ such that they contain only spatial gradients. At second order, we have derived the flux contribution to $\tilde{\mathbf{R}}_{(2)}$ that is $-\partial_i \mathbf{F}_{(2)}^i$. The presence of high order time derivatives in $-\partial_t \mathbf{H}_{(2)}$, which in turn introduce even higher order spatial derivatives, leads to algebraic terms that rapidly scale in number

and complexity, making it impractical to derive and implement, even using computer algebra packages.

Note that one cannot always directly align powers of ϵ ($\{\zeta, \kappa, \eta\}$ or $\{\tau_{\Pi}, \tau_q, \tau_{\pi}\}$) with the order of spatial derivatives appearing in these source terms. To see this, consider the simple (CE) bulk viscosity and heat flux models from earlier given by equations (22) and (24). In the former, the next-to-leading order correction is $\mathcal{O}(\kappa^2 \tau_q)$ and contains a fourth-order derivative, whilst in the latter it is $\mathcal{O}(\zeta \tau_{\Pi})$ and contains a mixture of first-, second-, and third-order derivatives. However, the leading order correction in each case, $\tilde{\mathbf{R}}_{(1)}$, contains mostly second-order gradients in the primitive variables ($\partial_i \partial_j \mathbf{w}$) with some products of two first-order derivatives ($\partial_i \mathbf{w} \partial_j \mathbf{w}$). To see this, consider that the dissipative variables we move from the state and flux vectors to the new sources contain first-order gradients. If moved from the flux vector, becoming $-\partial_i \mathbf{F}_{(1)}^i$, they pick up another spatial derivative from the flux-gradient. If moved from the state vector, becoming $-\partial_t \mathbf{H}_{(1)}$, they pick up a first-order temporal derivative, which we will show can be swapped for a first-order spatial derivative. Hence, they are always diffusive, second-order gradients as one would expect for dissipation. This can also be seen in equations (22) and (24) at leading order.

We choose to perform the series expansion and truncation such that terms $\mathcal{O}(\epsilon)$ contain no time-scales and are first-order in the dissipation strengths $\{\zeta, \kappa, \eta\}$. Terms considered to be $\mathcal{O}(\epsilon^2)$ are first-order in the time-scales $\{\tau_{\Pi}, \tau_q, \tau_{\pi}\}$ and the strengths. We often choose to work with the first-order ($\mathcal{O}(\epsilon)$) source terms only as we find that including higher orders generally only makes small quantitative differences. However, using the $\mathcal{O}(\epsilon^2)$ source, the effect of varying time-scales for both the MIS and MISCE models will be shown. Finally, despite the inherent instability of first-order theories of relativistic dissipation in fluids (Hiscock & Lindblom 1983), we do not find any instabilities arising with our first-order MISCE model, at least for the test problems and parameter space explored so far.

In order to make it practical to implement the system numerically, we need to replace the time derivatives present in $\tilde{\mathbf{R}}_{(1)}$ and $\tilde{\mathbf{R}}_{(2)}$ with spatial ones. We have two potentially problematic sources of time-derivatives. First, the Navier–Stokes forms of the dissipative variables themselves contain time derivatives. Secondly, the entire dissipative state vector $\mathbf{H}(\mathbf{w})$ is time-differentiated in the equations of motion. Because both \mathbf{H} and \mathbf{F} can be expressed entirely as functions of primitive, non-stiff variables, we need expressions for the time derivatives of the primitive variables. Making use of the chain rule and equation (30), which contains the time derivative of the ideal state vector and hence the primitive variables that constitute it, we have

$$\frac{\partial \mathbf{w}}{\partial t} = \frac{\partial \mathbf{w}}{\partial \mathbf{q}_0} \frac{\partial \mathbf{q}_0}{\partial t} + \frac{\partial \mathbf{w}}{\partial \mathbf{q}_1} \frac{\partial \mathbf{q}_1}{\partial t} + \dots = - \left(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}} \right)^{-1} \partial_i \mathbf{f}_0^i + \mathcal{O}(\epsilon), \quad (37)$$

where we again note that \mathbf{w} is the vector of primitives. This means the term $\frac{\partial \mathbf{w}}{\partial \mathbf{q}_0}$ has a matrix form that is far more easily obtained through an inversion of the matrix $\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}}$. We can use this result to substitute wherever a time-derivative appears in our source such that we then have

$$\tilde{\mathbf{R}}_{(1)} = -\partial_t \mathbf{H}_{NS}(\mathbf{w}, \partial_i \mathbf{w}) - \partial_i \mathbf{F}_{NS}^i(\mathbf{w}, \partial_i \mathbf{w}) \quad (38)$$

and our source contains solely first- and second-order *spatial* derivatives. For a derivation of higher order approximations to time derivatives of primitive variables, see Appendix A.

We will dub this new formulation DEIFY (Dissipative Extension to Ideal Fluid dYnamics) so that $\tilde{\mathbf{R}}_{(1)}$ is the first-order DEIFY source

term. Also observe how the source term for DEIFY is proportional to ϵ whereas the MIS formulation source terms scale as $1/\tau \propto \epsilon^{-1}$. This means that the two forms become stiff in opposing limits – near the ideal regime (small ϵ) DEIFY will be stable as a result of a small source term, and will only become stiff, and potentially unstable, as ϵ grows large. The big benefit of this behaviour is that near the ideal regime we can confidently evolve DEIFY with explicit time integrators, knowing that source contributions will remain small.

In contrast, in the event of very slow-acting (large τ) and large-in-magnitude (large $\{\kappa, \zeta, \eta\}$) viscosities and heat fluxes, it will not be sensible or accurate to evolve DEIFY, even using implicit schemes. Instead, we may revert to an implementation of the MIS formulation in this regime, which is likely to be stable with explicit integrators and therefore less costly. Future work will extend the approach of Wright (2020), where an adaptive model of resistive and ideal MHD was implemented. Ours will be able to switch between different dissipative formulations of hydrodynamics during evolution, ensuring stability, efficiency, and accuracy.

In summary, both the ideal and highly non-ideal limits, we should be able to use explicit integration schemes, which have been shown to provide a speed-up of up to an order of magnitude over implicit schemes in comparable models of resistive/ideal MHD (Wright & Hawke 2020). In Section 5.3, we will prove the validity of this claim, and further explore the intermediate region of non-ideal behaviour between these two extremes.

4.1 First-order source

In order to compute the DEIFY source term(s), we will need to calculate matrices and, crucially, their inverses. For instance, we will need to know the inverse matrix that appears on the RHS of equation (37), which represents the Jacobian of the primitive vector with respect to the non-stiff conserved vector.

Here, we have a choice of how to compute the matrices of interest – that is we can invert them numerically, or try to get the form of the inverted matrix symbolically. Inverting matrices numerically, especially when densely populated, can require a large amount of computation, reducing accuracy as well as slowing down simulations. If the algebraic form of the matrices were at hand, this would lead to a far more efficient simulation, and as we are trying to build a source term to extend ideal hydrodynamics with the intention of being faster to evolve than other forms of dissipative hydrodynamics, it is sensible to adopt the performance gains of a purely symbolic source term.

On this note, let us turn to computing (algebraically) the matrices $(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}})$ and, hence, $(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}})^{-1}$. Here, we will make a simplification so that the terms appearing in these matrices are human readable: we take the low-velocity limit, neglecting terms $\mathcal{O}(v^2)$ and hence setting the Lorentz factor, $W = 1$. This assumption is not made for the numerical implementation. We also have a choice to make over which two thermodynamic variables are present in our primitive variable vector \mathbf{w} . The equation of state, which relates p , ρ , and n , gives us this choice, and we opt to work with $\mathbf{w} = \{p, \rho, \mathbf{v}\}$. Thus, our ideal conserved vector is now given by

$$\mathbf{q}_0(\mathbf{w}) = \begin{pmatrix} D \\ S_j \\ E = \tau + D \end{pmatrix} = \begin{pmatrix} \rho + p/(1-\Gamma) \\ (\rho + p)v_j \\ \rho \end{pmatrix}, \quad (39)$$

where we have used our equation of state to replace n in the expression for D and chosen to work with the conserved variable E for now instead of τ as it takes a simpler form.

This gives us the matrix

$$\left(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}} \right) = \begin{pmatrix} \partial_p n & \partial_\rho n & 0 & 0 & 0 \\ v_1 & v_1 & p + \rho & 0 & 0 \\ v_2 & v_2 & 0 & p + \rho & 0 \\ v_3 & v_3 & 0 & 0 & p + \rho \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix} \quad (40)$$

and, hence, $(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}})^{-1}$ is

$$(p + \rho)^{-1} \begin{pmatrix} (p + \rho)/\partial_p n & 0 & 0 & 0 & -(p + \rho)\partial_\rho n/\partial_p n \\ 0 & 0 & 0 & 0 & (p + \rho) \\ -v_1/\partial_p n & 1 & 0 & -v_1(\partial_p n + \partial_\rho n)/\partial_p n \\ -v_2/\partial_p n & 0 & 1 & -v_2(\partial_p n + \partial_\rho n)/\partial_p n \\ -v_3/\partial_p n & 0 & 0 & 1 & -v_3(\partial_p n + \partial_\rho n)/\partial_p n \end{pmatrix}. \quad (41)$$

Next, using equation (37), this gives us

$$\partial_t \begin{pmatrix} p \\ \rho \\ v_j \end{pmatrix} = (\rho + p)^{-1} \begin{pmatrix} (p + \rho)((1/\partial_p n) - (\partial_\rho n/\partial_p n)\partial_t E) \\ (p + \rho)\partial_t E \\ -v_j((1 + \partial_\rho n/\partial_p n)\partial_t E + (1/\partial_p n)\partial_t D) + \partial_t S_j \end{pmatrix} \quad (42)$$

where we can exchange the time-derivatives of the conserved variables for spatial derivatives of the fluxes using equation (30). Doing this, and using our equation of state $p = (\Gamma - 1)(\rho - n)$ to replace the partial derivatives of primitive variables, we arrive at

$$\partial_t \begin{pmatrix} p \\ \rho \\ v_j \end{pmatrix} = \begin{pmatrix} (1 - \Gamma)(1 + \partial_i S^i) \\ -\partial_i S^i \\ (\rho + p)^{-1} [v_j((2 - \Gamma)\partial_i S^i + (1 - \Gamma)\partial_i(Dv^i)) - \partial_i S_j^i] \end{pmatrix} \quad (43)$$

which represents expressions for the partial time derivatives of the primitive variables in terms of purely spatial-derivatives.

Let us now demonstrate what the MISCE sources look like. These are too complex to write in full, so we consider the case of bulk viscosity only, restrict to one spatial dimension and again work in the low-velocity approximation such that $W = 1$ (but not neglecting terms $\mathcal{O}(v^2)$). Then, $\Pi_{NS} = -\zeta \partial_x v^x$. The leading order source term, $\tilde{\mathbf{R}}_{(1)} = -\partial_t \mathbf{H}_{(1)} - \partial_x \mathbf{F}_{(1)}^x$ in full is

$$\tilde{\mathbf{R}}_{(1)} = -\partial_t \begin{pmatrix} 0 \\ 0 \\ -\Pi_{NS} \end{pmatrix} - \partial_x \begin{pmatrix} 0 \\ \Pi_{NS} v_x^2 \\ 0 \end{pmatrix} \quad (44a)$$

$$= \zeta \left[\begin{pmatrix} 0 \\ 0 \\ -\partial_t \partial_x v^x \end{pmatrix} + \partial_x \begin{pmatrix} 0 \\ v_x^2 \partial_{xx} v^x + 2v_x(\partial_x v^x)^2 \\ 0 \end{pmatrix} \right], \quad (44b)$$

where from equation (43) we have

$$\partial_t \partial_x v^x = (\rho + p)^{-1} [((2 - \Gamma)(v_x \partial_{xx} S^x + (\partial_x S^x)(\partial_x v_x)) + (1 - \Gamma)\partial_{xx}(Dv_x)) - \partial_{xx}(S^x v_x + p)] \quad (45a)$$

$$= [((2 - \Gamma)(2v_x \partial_{xx} v_x + 3v_x(\partial_x v_x)\partial_x(\rho + p) + (\partial_x v_x)^2) + (\rho + p)^{-1}(1 - \Gamma)(v_x \partial_{xx} n + n\partial_{xx} v_x) - \partial_{xx} v_x - v_x \partial_x(\rho + p) - (\rho + p)^{-1}\partial_{xx} p)] \quad (45b)$$

and, finally, we have an expression for $\tilde{\mathbf{R}}_{(1)}$ that is expressed purely in (second-order) spatial gradients of the primitive variables. The full expressions (without simplification) are not human-readable, but are attached as additional material with this submission. The code to derive them can be found at <https://www.github.com/MarcusHatton/ComputerAlgebra>, whilst their implementation can be seen at <https://www.github.com/MarcusHatton/METHOD>.

4.2 Second-order source

A similar but more complex calculation can be made to derive the next order (second) of dissipative correction to ideal hydrodynamics. Beginning again with the conservation law

$$\partial_t [\mathbf{q}_0(\mathbf{w}) + \mathbf{H}_{(1)} + \mathbf{H}_{(2)}] + \partial_i [f_0^i(\mathbf{w}) + \mathbf{F}_{(1)}^i + \mathbf{F}_{(2)}^i] = 0 \quad (46)$$

it follows that

$$\tilde{\mathbf{R}}_{(2)} = -\partial_t \mathbf{H}_{(2)} - \partial_i \mathbf{F}_2^{(i)}. \quad (47)$$

At first order, the form of \mathbf{H} , \mathbf{F} and hence $\tilde{\mathbf{R}}_{(1)}$ followed simply from the definition of the Navier–Stokes terms which are of $\mathcal{O}(\epsilon)$. At second order, we use the stiff subsystem

$$\partial_t (nW\bar{\mathbf{w}}) + \partial_i (nWv^i\bar{\mathbf{w}}) = \frac{n}{\tau}(\bar{\mathbf{w}} - \bar{\mathbf{w}}_{NS}) \quad (48a)$$

$$\partial_t \bar{\mathbf{q}} + \partial_i \bar{\mathbf{f}}^i(\mathbf{q}, \bar{\mathbf{q}}) = \frac{1}{\epsilon}(\bar{\mathbf{q}} - \bar{\mathbf{q}}_{NS}) \quad (48b)$$

and make the substitution $\bar{\mathbf{q}} = \bar{\mathbf{q}}_0 + \bar{\mathbf{q}}_1 + \bar{\mathbf{q}}_2 \equiv \mathbf{0} + \bar{\mathbf{q}}_{NS} + \bar{\mathbf{q}}_2$ to obtain, at order $\mathcal{O}(\epsilon^2)$,

$$\epsilon \left[\partial_t \bar{\mathbf{q}}_{NS} + \partial_i \bar{\mathbf{f}}^i(\mathbf{q}, \bar{\mathbf{q}}_{NS}) \right] = \bar{\mathbf{q}}_2. \quad (49)$$

Because the NS forms of the stiff variables can be defined entirely in terms of the non-stiff primitive variables ($\bar{\mathbf{q}}_{NS} \equiv \bar{\mathbf{q}}_{NS}(\mathbf{q})$), so too can $\bar{\mathbf{q}}_2$. The vectors $\mathbf{H}_{(2)}$, $\mathbf{F}_{(2)}$ are given by

$$\mathbf{H}_{(2)}(\mathbf{w}) = \begin{pmatrix} 0 \\ W(q_{0,(2)}v_j + q_{j,(2)}) + \pi_{0j,(2)} \\ 2q_{0,(2)}W - \Pi_{(2)} + \pi_{00,(2)} \end{pmatrix}$$

and

$$\mathbf{F}_{(2)}^i(\mathbf{w}) = \begin{pmatrix} 0 \\ \Pi_{(2)}(v^i v_j + \delta_j^i)W^2 + W(q_{(2)}^i v_j + q_{j,(2)}v^i) + \pi_{j,(2)}^i \\ W(q_{0,(2)}v^i + q_{(2)}^i) + \pi_{0,(2)}^i \end{pmatrix}.$$

Putting these results together and making substitutions wherever we find time-derivatives of the primitive variables (as before) allows us to arrive at a purely spatial form for $\tilde{\mathbf{R}}_{(2)}$.

5 RESULTS

Next, we will perform numerical tests of our implementation of the MISCE model, with comparison primarily to the established MIS model from which it has been derived. A code named Multifluid Electromagneto-HydroDynamics (METHOD) was used to perform these simulations, which may be found at <https://www.github.com/MarcusHatton/METHOD>, having been forked and extended from its creator's repository at <https://www.github.com/AlexJamesWright/METHOD>. Instructions on how to run the simulations and reproduce the results of this chapter are to be found on the 'MISCE Paper' branch.

For explicit time-integration, we mostly use a second-order, operator-split Runge–Kutta (RK) scheme and occasionally a fourth-order RK scheme for comparison (Gottlieb & Shu 1998). For semi-implicit time-integration, we use a Strong Stability Preserving (SSP) scheme from Pareschi & Russo (2005) that uses two explicit and two implicit sub-steps and gives second-order convergence in time-step Δt . To calculate the fluxes, we use a flux-vector splitting approach paired with third- or fifth-order Weighted Essentially Non-Oscillatory (WENO) reconstruction (Shu 1998). Our numerical grid is a static Cartesian one. For further details on the numerical scheme employed, see Wright & Hawke (2019).

These will be standard tests such that we may compare results against the literature and check for agreement. These tests are also

chosen to reflect the physics we are interested in capturing for actual neutron star mergers.

5.1 Shocktubes

Shocktubes are simple, 1D tests useful for closely analysing the behaviour of the fluid model and its numerical implementation. They are designed to produce a set of forward- and backward-travelling waves, and in particular, discontinuities in the fluid's properties. These waves (contact, rarefaction, and shock) are the fundamental propagation modes of the fluid and will certainly be produced at the point of merger, and thereafter whenever a sharp jump in density, pressure, or temperature occurs such as between different phases of matter within the neutron star. They also involve advection of the fluid, which will be important for the inspiral phase of the merger as tidal forces will drag fluid around the star as they orbit their mutual centre of mass.

The initial data for these tests are similar to those of Takamoto & Inutsuka (2011) but with zero initial velocity. We also share the same equation of state, allowing for a favourable quantitative comparison to be made. A domain of one spatial dimension is initially split into bordering left and right states $[L, R]$ where the primitive variables in the two states are

$$L : \begin{pmatrix} p \\ n \\ \rho \end{pmatrix} = \begin{pmatrix} 10 \\ 8 \\ 23 \end{pmatrix} \quad (50)$$

for the left state and

$$R : \begin{pmatrix} p \\ n \\ \rho \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \\ 2.5 \end{pmatrix} \quad (51)$$

for the right state. The pressure, p , and number density, n , are set in the initial data and the equation of state, $p = (\Gamma - 1)(\rho - n)$, determines the energy density, ρ . We set a value of $\Gamma = 5/3$ for the adiabatic index.

In Fig. 4, we see the expected production of the three travelling waves: the left-moving rarefaction; right-moving contact-wave; and (faster) right-moving shockwave. These are most easily seen in the energy density plot in the top-left. The bulk viscosity has a smoothing effect on these waves, particularly on the shockwave, where it also significantly increases the propagation speed of the shock-front – this can also be seen directly in the velocity plot (top-right). The bulk viscosity itself (bottom-right) spikes at the shock where the velocity gradients are highest. Its positivity there indicates a resistance to the rapid compression of the fluid by the shockwave, with the reverse being true for the rarefaction.

In Fig. 5, we take a closer look at a shocktube profile for the fluid's number density with bulk viscosity and heat flux present. Three results from the MIS model are plotted for differing dissipative time-scales τ (the same for both types of dissipation), and one for the MISCE model (at leading order, so the time-scale does not enter into the EoM). In particular, we see convergence of the MIS result to the MISCE result as τ is decreased. This is expected given that for the MIS model, in the $\tau \rightarrow 0$ limit, any off-shell deviations from relativistic Navier–Stokes behaviour are instantaneously quenched. This means that the MIS model's behaviour should match that of the leading-order MISCE model in this limit, where terms $\mathcal{O}(\tau)$ and higher are neglected.

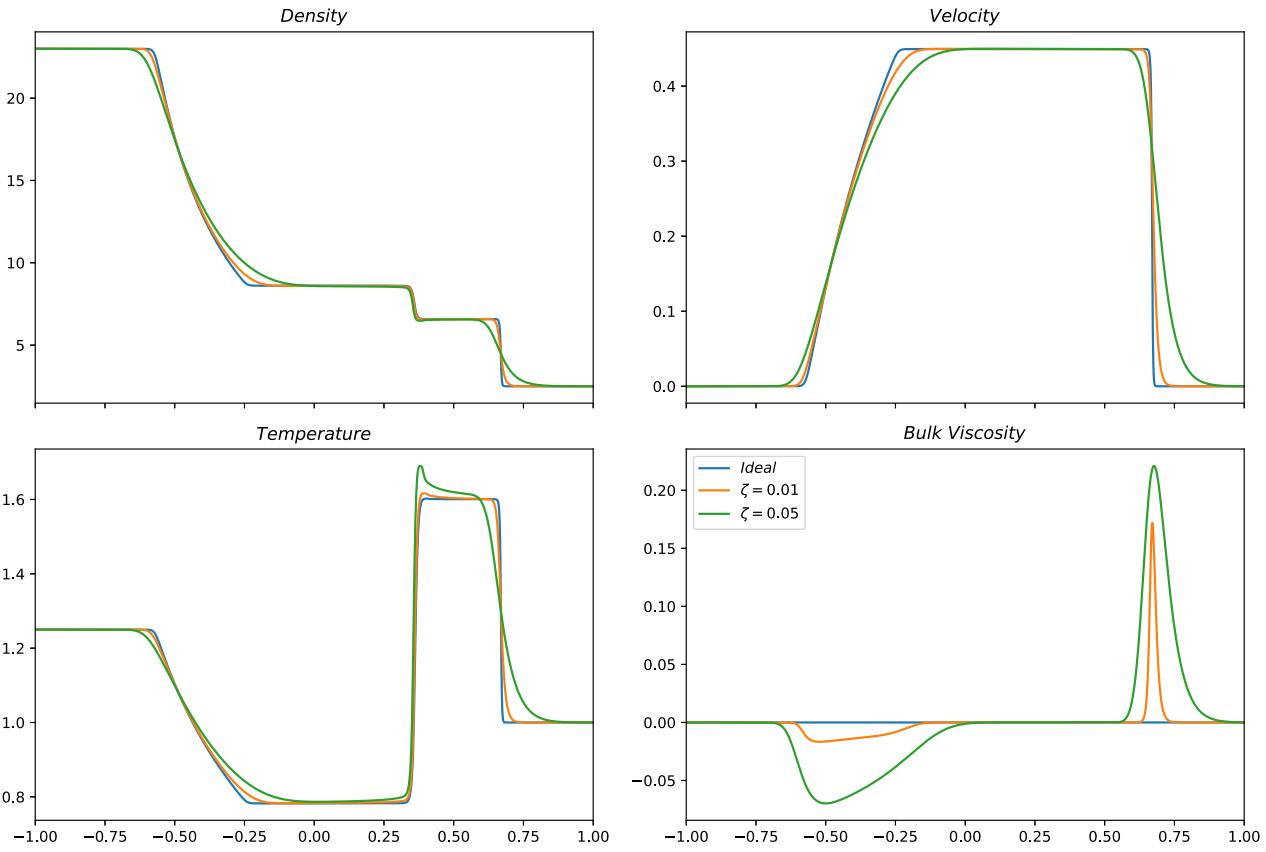


Figure 4. A shocktube simulation comparing results of an ideal fluid with a viscous one using the MISCE formulation at leading order. A bulk viscosity parameter of $\zeta = 1 \times 10^{-2}$ and $\zeta = 5 \times 10^{-2}$ is used in each of the two viscous cases. The energy density, velocity, Navier–Stokes bulk viscosity and temperature are plotted. The increase in shock propagation speed and smearing of discontinuities due to the inclusion of bulk viscosity are both visible physical effects.

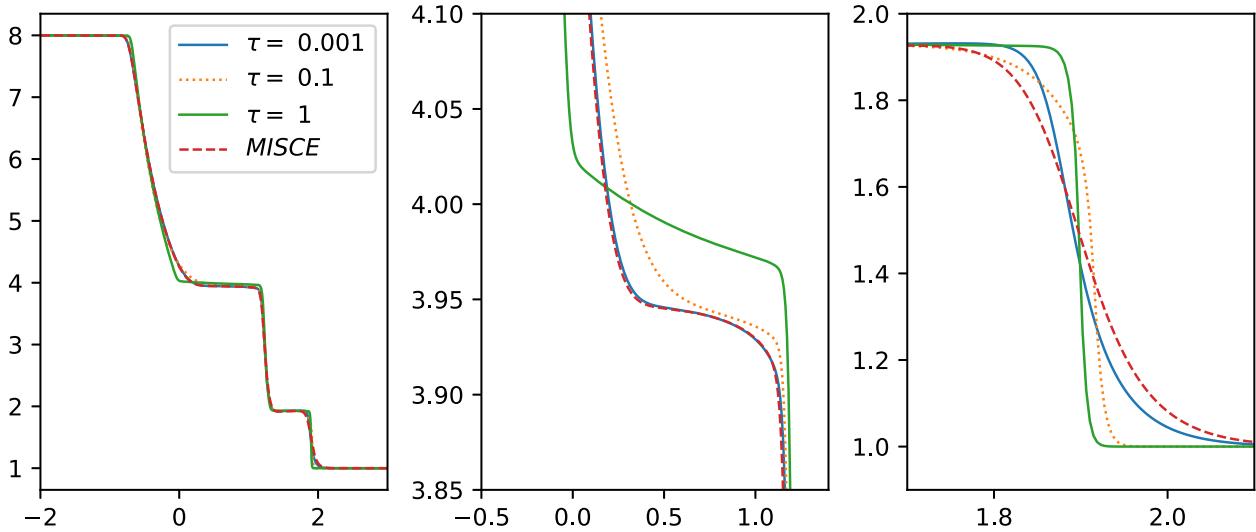


Figure 5. The evolution of the number density for a ‘stillshock’ test – a shocktube with zero initial velocity. There is bulk viscosity and heat flux present with coefficients $\zeta = 5 \times 10^{-2}$ and $\kappa = 5 \times 10^{-3}$. The three panels show the entire domain (left), the rarefaction wave (centre) and the shockwave (right). The two models (MIS, MISCE) are compared in all three panels, with the dissipative time-scale τ varying for the MIS model but held constant at zero for the MISCE model (making it the leading-order version). One can see the approach of the MIS solution towards the MISCE solution as $\tau \rightarrow 0$. For the rarefaction wave they have converged in the fastest case, but for the shock there are still differences. In particular, one can see the increase in speed of the shock as the ratio ζ/τ_π increases for the MIS model. It is catching up to the MISCE solution, which can be thought of as its limiting case.

5.2 Kelvin–Helmholtz instabilities

The Kelvin–Helmholtz instability (KHI) is a shearing instability that results when two (or more) fluid regions flow in opposite directions past each other, each usually differing in density. A wide range of fluid behaviours can be observed depending on the precise initial data, but here we will be varying the shear viscosity only, to focus on its effect. In the most interesting cases, there is an initial linear growth phase of the instability at the interface, followed by a non-linear phase where the creation of vortices and a complex network of shocks typically precedes the onset of smaller-scale turbulence.

Neutron star matter in mergers is likely to be Kelvin–Helmholtz unstable as the two objects collide and shearing flows develop. The KHI is known to play an important role in post-merger dynamics where it moderates the cascade of energy between macroscopic and microscopic scales through the action of shear viscosity in the fluid. This is important in the spin-down of the remnant where the rotational energy of the fluid is converted to small-scale turbulence and then to either magnetic energy through the dynamo effect or dissipated through viscous heating. We will also analyse the integrated power spectrum of kinetic energy resulting from turbulence induced by the KHI. This has famously been shown by Kolmogorov to have a universal scaling relation with wavenumber for at least part of its spectrum, a result which was generalized by Qian (1994).

To investigate this process, we use the initial conditions from Beckwith & Stone (2011), as well as the spectral analysis laid out by them. The data are defined within a 2D domain where $x \in [-1.0, 1.0]$ and $y \in [-0.5, 0.5]$. The domain is then divided into two-fluid regions, with the inner region contained roughly within $x \in [-0.5, 0.5]$ and the outer elsewhere. The two-fluid regions have differing densities and flow past each other with velocities directed in the positive and negative y -directions. There is a narrow transition layer between the two where a small, spatially varying perturbation to the x -directed velocity is also introduced to induce mixing. The primitive variables are

$$\begin{pmatrix} v_y \\ \rho \\ v_x \end{pmatrix} = \begin{pmatrix} v_{sh} \tanh\left(\frac{x-0.5}{a}\right) \\ \rho_0 + \rho_1 \tanh\left(\frac{x-0.5}{a}\right) \\ A_0 v_{sh} \sin(2\pi y) \exp\left(\frac{-(x-0.5)^2}{l^2}\right) \end{pmatrix}; x > 0.0 \quad (52a)$$

and

$$\begin{pmatrix} v_y \\ \rho \\ v_x \end{pmatrix} = \begin{pmatrix} -v_{sh} \tanh\left(\frac{x+0.5}{a}\right) \\ \rho_0 - \rho_1 \tanh\left(\frac{x+0.5}{a}\right) \\ -A_0 v_{sh} \sin(2\pi y) \exp\left(\frac{-(x+0.5)^2}{l^2}\right) \end{pmatrix}; x \leq 0.0, \quad (52b)$$

where the shear velocity is $v_{sh} = 0.5$, the boundary layer thickness is $a = 0.01$, the densities are given by $(\rho_0, \rho_1) = (0.55, 0.45)$, and the perturbation has an amplitude $A_0 = 0.1$ over a characteristic length $l = 0.1$. The initial pressure is uniform, $p = 1.0$, and the adiabatic index is set to $\Gamma = 4/3$. We use periodic boundaries in the both the x and y directions.

Figs 6 to 7 show the development of the KHI for the fluid's number density. Figs 6 and 8 show its development for an ideal (inviscid) fluid. For the former, the early-time behaviour is the focus, with the initial growth of the interface instability visible, followed by large-scale mixing and finally the formation of small-scale structure as energy cascades from longer to shorter scales. In the latter, the asymmetry of the initial perturbation has had time to grow into a macroscopic asymmetry. One can also see vortices forming and the onset of turbulence in the wide mixing layer.

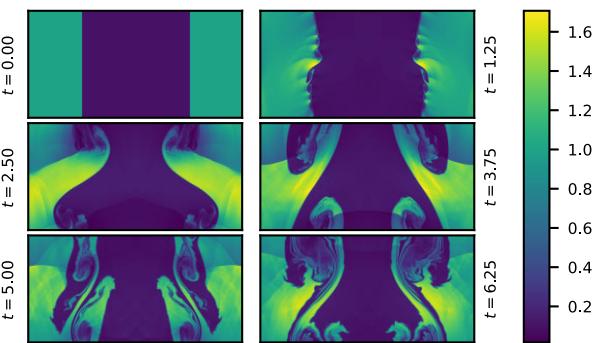


Figure 6. The short-term development of a Kelvin–Helmholtz unstable fluid with negligible viscosity until $t = 6.25$. This uses the MISCE model that reduces to the ideal Euler equations in the inviscid limit. The number density is shown in colour, as is the case for all KHI plots here. The initial perturbation grows rapidly until the interface breaks and large-scaling mixing occurs, followed by the onset of turbulent behaviour which produces shocks and smaller-scale vortices.

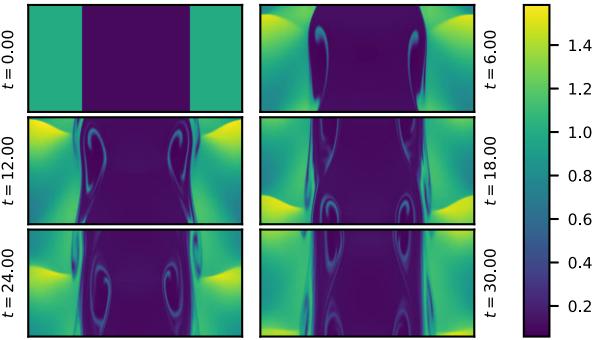


Figure 7. The long-term evolution of the KHI using the MISCE model at leading order with a shear viscosity parameter of $\eta = 1 \times 10^{-3}$. The shear viscosity has an intermediate value here: it suppresses large-scale mixing of the two fluids but vortices still form in a narrow shearing layer that is stable even at late times. The asymmetry is again visible here, but obscured for similar reasons.

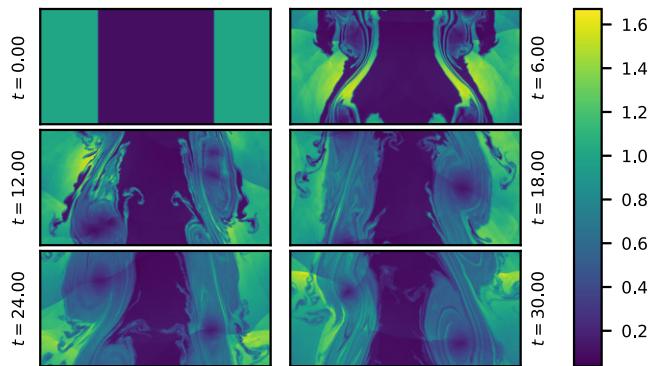


Figure 8. The long-term development of the Kelvin–Helmholtz instability until $t = 30.0$, for an ideal fluid with negligible viscosity, again using the MISCE model. The longer simulation time allows the asymmetric initial perturbation at the interface to give rise to large-scale asymmetric vortex formation.

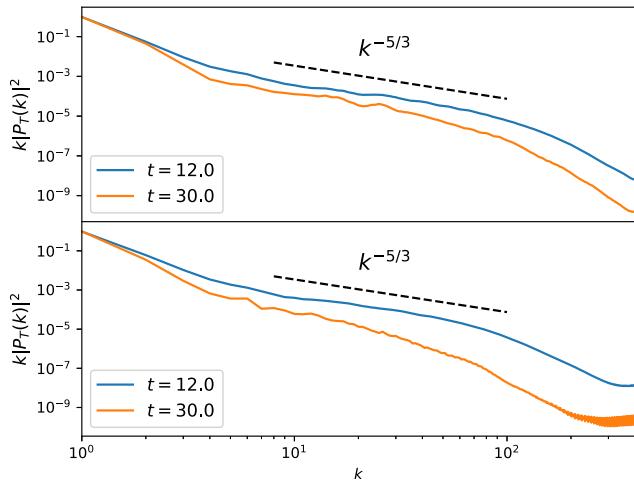


Figure 9. The power spectra for the kinetic energy density in the KHI at medium and long times, for an inviscid fluid (top) and one with weak shear viscosity, $\eta = 1 \times 10^{-4}$ (bottom). This uses the MISCE formulation at leading order with a grid of size $N_x = N_y = 800$. The expected Kolmogorov scaling of the power spectrum is seen in the inertial regime at earlier times. In the inviscid case, the numerical viscosity has a minor damping effect on the power spectrum at late times and high wavenumbers (short length-scales). A greater damping effect is seen in the viscous case, as well as a ‘ringing’ at high wavenumbers due to coupled action of fluid element discretization and local viscosity: these wavenumbers correspond to length-scales of a few, or even a single, cell(s).

In Fig. 7, the long-term behaviour for a viscous fluid is shown. Viscosity suppresses the perturbation’s growth and stabilizes the mixing at the interface. Vortices do form, still, but they are confined to a smaller corridor between the two bulk fluid regions, and in general the behaviour is less chaotic. We observe similar qualitative behaviour to Takamoto & Inutsuka (2011), who performed comparable simulations. Fig. 9 shows the power spectra for the kinetic energy in our KHI simulations. Two comparisons are made: one between early ($t = 12.0$, blue curve) and late ($t = 30.0$, orange curve) times; and one between an inviscid (top panel) and viscous (bottom panel) fluid. In both cases, the system loses energy over time. For the inviscid case, this is due to numerical dissipation. For the viscous case, there is the additional effect of viscous dissipation, which causes the steeper drop-off for the orange versus the blue curve. The expected Kolmogorov scaling for the inertial range is plotted and matches well with the data for all but the late-time viscous case, where dissipation has more efficiently moved energy to the shorter length-scales, giving a steeper dependence on wavenumber. Finally, the ‘ringing’ effect seen for the highest wavenumbers in the late-time, viscous case is, we believe, a numerical artefact. Dissipative behaviour in the MISCE model is captured using a complex mixture of spatial gradients of the primitive variables, generally calculated using simple central differencing rather than, for example, a WENO scheme which is designed to be non-oscillatory. For the highest wavenumbers here, corresponding to a few or even a single cell(s), these derivatives may be causing small-scale oscillations in fluid variables that have no qualitative significance.

5.3 Code performance

There are a few key metrics of code performance we must now consider. First, how the runtime of simulations scales with resolution.

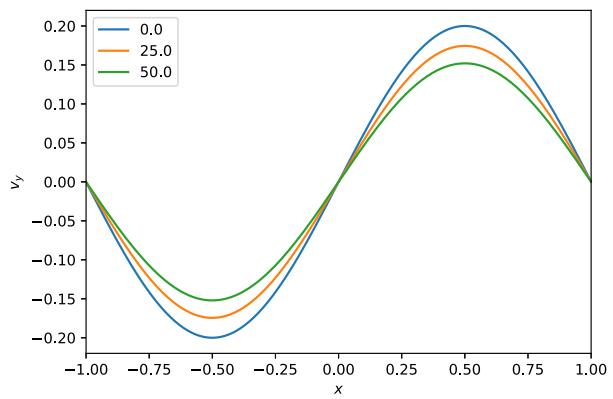


Figure 10. The ‘SineWave’ initial data (and its evolution with the MISCE model at second order) used to assess numerical convergence with resolution. Shear viscosity causes the flattening of features in the y -direction velocity across the x -domain here. This simulation was performed using 3200 cells in one dimension up to a code time of 50.0.

Secondly, the convergence of the simulation output, which is assessed in two ways: the self-convergence of the MISCE results to a very high resolution simulation output; the asymptotic approach of the MISCE results to either leading-order or ideal fluid behaviour as the non-ideal coefficients approach zero. Finally, we present a comparison of runtimes, showing the significant speed-up achieved by the MISCE model.

5.3.1 Scaling and convergence

By evolving smooth ‘SineWave’ initial data (Fig. 10), we are able to assess the convergence of our MIS and MISCE implementations with resolution. Considering the error due to finite resolution, we define it as the difference between ‘true’ solution (the one obtained at infinite resolution) and the finite-resolution solutions our code actually produces: $\mathcal{E} = Q_{true} - Q_{num}$.

Then we make the usual assumption that this error follows a power-law scaling in the grid-size: $\mathcal{E} \propto \Delta x^n$. Different approaches exist for extracting the value of n in this expression, and we choose here to use self-convergence, where each resolution’s solution is compared to its neighbours to produce a set of convergence powers at different resolutions. We do this because different components of our numerical scheme (the time-integrator, cell-interface reconstruction method etcetera) each have individual expected convergence rates that blend together to give an overall convergence. This means that different components can dominate the error at different resolutions, and we are able to assess the transition between them using this approach.

We show in Table 2, a summary of convergence orders for different models and resolutions. In Fig. 11 we plot convergence rate against resolution for our MISCE model, and in Fig. 12 we plot the convergence rate over time for a simulations at different resolutions, again using our MISCE model. In summary, we see a transition from high-order convergence at low resolutions to lower order convergence at higher resolutions. For both models, the error at low resolutions is dominated by the time integrator and reconstruction algorithm, which are high-order schemes and hence their error converges away quickly.

At high resolutions, we see a drop in the overall convergence order. For the MISCE model, this is because there are many spatial deriva-

Table 2. The self-convergence of a smooth sin-wave evolution using different models of non-ideal hydrodynamics and different numerical schemes. The expected orders of convergence are seen. At very high resolution, the first-order central differencing used in the MISCE model source's spatial derivatives causes the convergence order to drop to 2^{nd} . For the MIS model, we use lagged-updates to calculate the required time derivatives. This similarly caps the order of convergence at first when very high resolutions are used and error from other components of the numerical scheme are tiny.

Model	Integrator	Self-convergence		
		Reconstruction	Resolution	Order
MISCE	RK2	WENO3	100	2.45
MISCE	RK4	WENO5	50	4.5
MISCE	RK4	WENO5	100	4.7
MISCE	RK4	WENO5	200	3.9
MISCE	RK4	WENO5	400	2.6
MISCE	RK4	WENO5	800	2.1
MISCE	RK4	WENO5	1600	2.0
MIS	SSP2	WENO5	400	4.0
MIS	SSP2	WENO5	800	5.7
MIS	SSP2	WENO5	1600	1.0

Table 3. A comparison of computational time required for different hydrodynamic models and time-integrators. These results are for KHI simulations using 40 CPU nodes and MPI memory management on the Iridis5 supercomputer. The MISCE model gives about an order of magnitude speed-up compared to the MIS model (when evolved with explicit methods instead of implicit ones). RK2 refers to an operator-split, second-order RK scheme and SSP2(222) refers to a second-order implicit-explicit scheme.

Average runtime				
Model	Integrator	Resolution	Endtime	Runtime (Speed-up)
MIS	SSP2	200x400	6.25	1h22m
MISCE	RK2	200x400	6.25	6m ($\times 14$)
MIS	SSP2	400x800	6.25	3h22m
MISCE	RK2	400x800	6.25	29m ($\times 7$)
MIS	SSP2	800x1600	6.25	26h10m
MISCE	RK2	800x1600	6.25	3h7m ($\times 8.4$)
MIS	RK2	800x1600	18.0	15h45m
MISCE	RK2	800x1600	18.0	9h8m ($\times 1.7$)
MIS	SSP2	800x1600	3.75	22h3m
MISCE	RK2	800x1600	3.75	1h54m ($\times 11.6$)

tives of the primitive variables in the complex source terms, which are evaluated using second-order central differencing. Increasing the order of this central differencing does increase the convergence rate, but makes negligible difference to quantitative results. For the MIS model, we require temporal derivatives of the primitive variables. These are evaluated using backwards-differencing on the primitives' values at the current and previous time-step. This introduces a first-order error due to these lagging updates that does not converge away with resolution, and hence appears as the dominant error at high resolutions.

5.3.2 Model comparison

We show in Table 3, a comparison of runtimes between the MIS and MISCE models for the KHI. We primarily present results comparing the MIS model evolved with the SSP2(222) IMEX time-integrator (Pareschi & Russo 2005) and the MISCE model evolved with an operator-split RK2 time-integrator. Whilst this comparison may seem 'unfair' at first, due to the costly nature of IMEX schemes compared

to explicit ones, it is justified. Whilst for much of a merger simulation the neutron star fluid may be accurately treated as ideal or near-ideal, when dissipation does become significant its parameter space will certainly extend into the region where the MIS model becomes stiff and IMEX schemes are needed to evolve it stably.

In this case, a significant speed-up of about an order of magnitude is achieved using the MISCE model. When the two models are compared using explicit time-integrators for both, a speed-up of nearly a factor of 2 occurs, owing to the reduced system size.

6 FURTHER CONSIDERATIONS

6.1 Rapid evolution of reduced initial data

The source terms in MIS models drive the dissipative variables towards their equilibrium values on time-scales τ . We therefore expect that, when our initial data (or otherwise) put us significantly out-of-equilibrium, at times $t \lesssim \tau$, there will be a systematic error in the dissipative variables that decays roughly as $e^{-t/\tau}$.

However, a CE-expanded model does not possess this type of source term nor indeed any explicit dissipative variables at all. Instead, the primitive variables and their derivatives are used to produce dissipative effects. We therefore expect that we will need to make modifications to the primitive variables' initial values to reflect their out-of-equilibrium status in lieu of having terms that explicitly define our out-of-equilibrium state.

Let us demonstrate the effect of not making appropriate adjustments to the primitive variables to reflect their out-of-equilibrium state. We take the simple heat model presented earlier in equation (18), in one dimension:

$$\partial_t T + \partial_x q = 0, \quad (53a)$$

$$\partial_t q = -\frac{1}{\tau_q}(\kappa \partial_x T + q); \quad (53b)$$

and its CE form, equation (22),

$$\partial_t T = \kappa [\partial_x^{(2)} T - \kappa \tau_q \partial_x^{(4)} T]. \quad (54)$$

By introducing a fast time variable $\mathcal{T} = t/\tau_q$ on the scale of the relaxation rate, we can perform a matched asymptotic expansion valid even at small times. This transforms equation (53) into

$$\partial_{\mathcal{T}} T + \tau_q \partial_x q = 0, \quad (55a)$$

$$\partial_{\mathcal{T}} q = -q - \kappa \partial_x T. \quad (55b)$$

From equation (55), the power series expansion now gives that the temperature T is independent of \mathcal{T} to leading order and

$$\partial_{\mathcal{T}} q_0 = -q_0 \quad (56)$$

which can be integrated directly to give

$$q_0 = C_0 e^{-\mathcal{T}}, \quad (57)$$

where C_0 is a constant of integration. We immediately see that this exponential behaviour in fast time, \mathcal{T} , cannot be captured by a power series expansion in the original time, t .

Noting that $C_0 = q(t=0) + \mathcal{O}(\tau_q)$, we relabel C_0 as Δq_0 because it represents an initial offset of the heat-flux at $\mathcal{T} \rightarrow 0^+$. To compare this early time behaviour between the two models (equations 53 and 54), we can Taylor-expand each in terms of \mathcal{T} to an arbitrary (small) time $\mathcal{T} = 1$ about $\mathcal{T} = 0$. This is equivalent to considering the evolution up to time $t = \tau_q$. From the relaxation model equation (53)

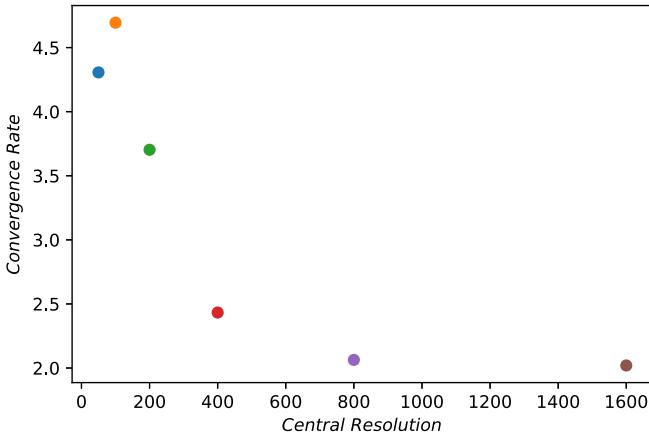


Figure 11. The self-convergence rate of the MISCE model (at second order) for different resolutions. For lower resolutions, the convergence order is between fourth and fifth due to the use of an RK4 time-integrator and a WENO5 reconstruction scheme. At higher resolutions, a transition to second-order convergence is seen due to the presence of first-order central differencing used for spatial derivatives in the MISCE source terms.

we have

$$T(\mathcal{T} = 1) \simeq T(\mathcal{T} = 0) + \left(\frac{\partial T_0}{\partial \mathcal{T}} + \tau_q \frac{\partial T_1}{\partial \mathcal{T}} + \dots \right) (\mathcal{T} = 0) + \dots \quad (58a)$$

$$= T(\mathcal{T} = 0) - \tau_q \frac{\partial \Delta q_0}{\partial x} + \mathcal{O}(\tau_q^2) \quad (58b)$$

whilst from the CE model equation (54), we have

$$T(\mathcal{T} = 1) \simeq T(\mathcal{T} = 0) + \left(\frac{\partial T_0}{\partial \mathcal{T}} + \tau_q \frac{\partial T_1}{\partial \mathcal{T}} + \dots \right) (\mathcal{T} = 0) + \dots \quad (59a)$$

$$= T(\mathcal{T} = 0) + \mathcal{O}(\tau_q^2). \quad (59b)$$

Comparing the two, we see that we can match the two temperatures at small times by making an initial-data adjustment given by

$$T(t = 0, x) \rightarrow T(t = 0, x) - \tau_q \partial_x \Delta q_0. \quad (60)$$

This accounts for the fast-relaxation behaviour and removes the exponentially decaying, leading-order systematic error in the solution. In words, we are taking into account the heat flux (q) that would have produced our initial, out-of-equilibrium temperature (T) distribution. Otherwise, our reduced system does not have access to this knowledge and will not equilibrate accordingly.

This can be seen in Fig. 13, where a 1D ‘top-hat’ temperature profile evolves up to $t = 1 \times 10^{-3}$, using non-ideal parameter values $\kappa = 1 \times 10^{-3}$ and $\tau_q = 1 \times 10^{-4}$. Here, we are not interested in the usual, long-term evolution where heat would slowly diffuse outwards and the temperature profile would adopt a decaying Gaussian shape. Instead, we are interested in the very short-term evolution due to the inclusion of an initial heat flux $\Delta q_0 = \sin(6\pi x)$ into the MIS-derived heat model given by equation (53).

In the left panel, the temperature of the MIS-derived relaxation model is shown with solid lines, whilst the initial temperature with the offset computed in equation (60) is shown with dotted lines. Excellent agreement is seen, indicating that this offset would work when applied to a reduced order model such as the MISCE approach. The right panel shows the heat flux of the MIS-derived relaxation model,

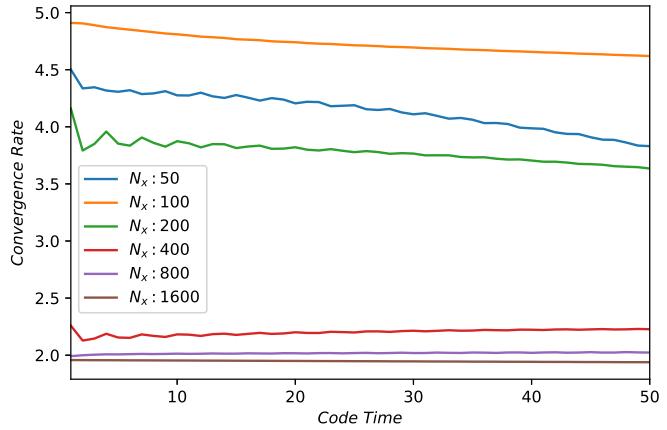


Figure 12. The self-convergence of the (second-order) MISCE model across time for a range of resolutions. The simulation data can be seen in Fig. 10, although this convergence test is carried out on the number density, n . For lower resolutions, the convergence order is between fourth and fifth due to the use of an RK4 time-integrator and a WENO5 reconstruction scheme. At higher resolutions a transition to second-order convergence is seen due to the presence of first-order central differencing used for spatial derivatives in the MISCE source terms.

showing that the system has relaxed to equilibrium, illustrating that this applied offset has the appropriate magnitude.

In summary, injecting an initial heat-flux into a relaxation-type system leads to an exponentially fast adjustment of the corresponding conjugate primitive variable: the temperature, in this case. We are able to derive an analytical expression for this adjustment behaviour that depends on the spatial gradient of the injection and the non-ideal parameter controlling its time-scale, τ_q in this case. Even in the reduced system found using the CE-expansion, we are able to adjust the sole remaining variable (the temperature) to capture the offset that is quickly arrived at by the original relaxation system.

The same effect would be observed when using the full MIS model of non-ideal hydrodynamics, where an initial bulk or shear viscosity would lead to an exponentially fast adjustment of the velocity, albeit likely small in magnitude. If one uses our MISCE model for capturing far out-of-equilibrium dissipation, the initial conditions of the non-ideal variables (viscosity, heat-flux) can and should still be taken account of by adjustment of their conjugate primitive variables (such as velocity and temperature).

6.2 Stability analysis

It is important to consider the numerical stability of the CE systems introduced here. Usually, conservation laws are evolved for hydrodynamic simulations of ideal fluids in special relativity of the form

$$\partial_t q + \partial_x f = 0, \quad (61)$$

where we choose to write it in one spatial dimension for simplicity. The Courant–Friedrichs–Lowy (CFL) condition sets a stability criterion for these strongly hyperbolic systems given by

$$C := \left| \frac{\partial f}{\partial q} \right| \frac{\Delta t}{\Delta x} \leq C_{max} \quad (62)$$

where Δt is the time-step and Δx is the spatial resolution. C is the Courant number and C_{max} is a constant that determines its maximum stable value and depends on the particular numerical scheme employed. Typically, $C_{max} = \mathcal{O}(1)$ for explicit schemes.

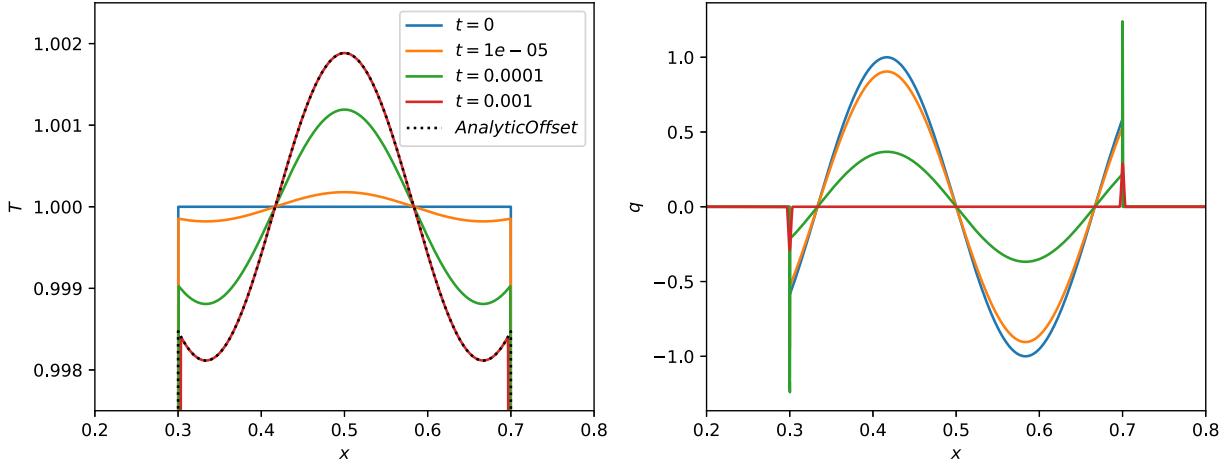


Figure 13. The evolution of the temperature and heat flux for the initial data described in Section 6.1 using the relaxation model given by equation (53). The initial heat flux means the data is initially out-of-equilibrium. The non-ideal parameters are $\kappa = 1 \times 10^{-3}$ and $\tau_q = 1 \times 10^{-4}$, so the system relaxes to equilibrium on the time-scale shown here, as seen by the heat flux relaxing to nearly zero. The analytical result for the appropriate adjustment to the initial data, derived in Section 6.1, is also plotted in the left-hand panel (dotted) and shows excellent agreement with the numerical evolution result.

For linear systems involving higher-order spatial derivatives, such as the heat equation given by

$$\partial_t T = \kappa \partial_x^{(2)} T, \quad (63)$$

a von Neumann stability analysis is usually performed, where the numerical errors are decomposed into a Fourier series. For a forward-time, centre-spaced (explicit) numerical scheme this analysis finds that

$$\kappa \frac{\Delta t}{\Delta x^2} \leq \frac{1}{2} \quad (64)$$

is the condition necessary for stability. For the CE heat system given by equation (22)

$$\partial_t T = \kappa \partial_x^{(2)} T - \kappa^2 \tau_q \partial_x^{(4)} T, \quad (65)$$

the conditions for stability are more complex. See Bevilacqua et al. (2011) for a treatment of higher-order differential terms relevant to our work here. The analysis gives us the inequality for stability of

$$\kappa \frac{\Delta t}{\Delta x^2} + 4\tau_q \kappa^2 \frac{\Delta t}{\Delta x^4} \leq \frac{1}{2}. \quad (66)$$

This clearly yields the previous heat-equation limit when the first term is dominant. In the limit where $\Delta x \rightarrow 0$ and the second term becomes dominant, we instead obtain the condition $\Delta t \leq \frac{1}{8} \frac{\Delta x^4}{\tau_q \kappa^2}$ which is the stricter condition of the two in this limit. We expect a cross-over of stability between the two criteria when $\Delta x^2 = 4\tau_q \kappa$. At this point, the overall stability condition given by equation (66) above yields $\Delta t \leq \tau_q$. See Fig. 14 for a visualisation of these stability criteria.

The stability criteria for the full, non-linear MISCE model we have developed will be more complex still, given the presence of many mixed derivatives. Hence, we primarily investigate its stable parameter space empirically. However, we can first glean some insight analytically, although the usual von Neumann stability analysis is not applicable to the non-linear terms, and we therefore consider the linear terms only here.

We make the ansatz that the solution can be written as $q_l^n = q^n \exp(i\alpha \Delta x)$ where n and l index the time-step and a grid-point, respectively, and α is a spatial frequency present in the data. Using central finite differencing, the MISCE sources will produce a solution

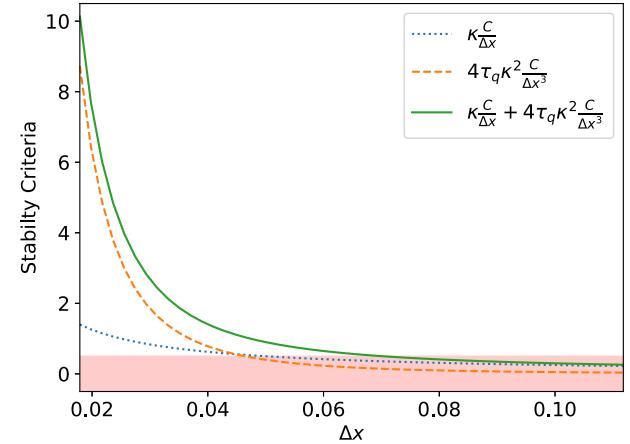


Figure 14. The stability criteria of equation (66) are plotted separately (dashed, dotted) and summed (solid). The red, shaded region shows where the simple heat flux model given by equation (65) should be stable, according to the standard von Neumann analysis technique using a Fourier series of errors. The heat dissipation parameter values here are $\tau_q = 0.01$, $\kappa = 0.05$, and the CFL factor is $C = 0.5$ hence the crossover between the two stability criteria occurs at $\Delta x \simeq 0.045$.

growth rate per step, q , with the following form:

$$q = 1 - \xi \frac{A \Delta t}{\Delta x^2} \sin^2 \left(\frac{\theta}{2} \right) + \xi \tau \frac{B \Delta t}{\Delta x^3} (\sin(2\theta) - 2 \sin(\theta)) + \xi^2 \tau \frac{C \Delta t}{\Delta x^4} \sin^4 \left(\frac{\theta}{2} \right),$$

where $\theta = \alpha \Delta x / 2$, $\xi \equiv \{\zeta, \kappa, \eta\}$ and A, B , and C are functions of the primitive variables. We anticipate a cross-over between the various stability criteria as resolution varies.

First, we note that the validity of our expansion only applies when $\tau \ll 1$ and indeed we find that our simulations are unstable when $\tau \gtrsim 10^{-2}$. In Hiscock & Lindblom (1983), conditions are provided for the stability (and causality) of MIS theory. There, the β coefficients given in equation (12), which represent the ratio of dissipation strengths ($\xi \equiv \{\zeta, \kappa, \eta\}$) to time-scales ($\tau \equiv \{\tau_\Pi, \tau_q, \tau_\pi\}$), are used

to determine the stability of the theory. Unsurprisingly then, we find the same is true here: a lower bound from causality appears on the ratio τ/ξ .

In particular, we find that for shocktube tests, $\tau_q/\kappa = T\beta_1 \gtrsim 0.5$ as well as $\tau_{\Pi}/\xi = \beta_0 \gtrsim 0.1$ grants stability. Similarly, for KHI tests, $\tau_{\pi}/\eta = \beta_2 \gtrsim 0.2$ gives stability. These ratio conditions coupled with the small- τ requirement work together to create a stability region bounded at either end, at least when next-to-leading order terms $\propto \tau$ are included in the MISCE source.

For a Boltzmann gas, the β coefficients have thermodynamic forms that we can calculate analytically. We expect them to usually be $\sim \mathcal{O}(1)$. Indeed, we have implemented these thermodynamic forms such that the timescales used are dynamically adjusted during the simulation—little difference is made to using preset values.

6.3 Extension to GR

Here we briefly consider the form our model would take if we were to extend it to General Relativity, in anticipation of future work. We follow the fairly standard numerical relativity definitions as laid out in Chabanov et al. (2021). In particular, we introduce the unit time-like vector \mathbf{n} that is normal to the spatial hypersurfaces, on which an induced spatial metric $\gamma_{\mu\nu} \doteq g_{\mu\nu} + n_{\mu}n_{\nu}$ is defined, where $g_{\mu\nu}$ is the space-time metric and $g \doteq \det(g_{\mu\nu}) = \alpha\sqrt{\gamma}$ relates the determinants of the two. The components of \mathbf{n} are given by $n_{\mu} = (-\alpha, 0, 0, 0)$ and $n^{\mu} = \alpha^{-1}(1, -\beta^i)^T$ where α is the lapse scalar and $\beta_i = \gamma_{ij}\beta^j$ is the shift-vector.

The hydrodynamic equations in GR then become

$$\partial_t(\sqrt{\gamma}\mathbf{U}) + \partial_i(\sqrt{\gamma}\mathbf{F}^i(\mathbf{U})) = \sqrt{\gamma}\mathbf{S}, \quad (67)$$

where

$$\begin{aligned} \mathbf{U} &= \begin{pmatrix} D \\ S_j \\ \tau \\ X \\ Y^j \\ Z^{jk} \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi)W^2v_j + W(q_0v_j + q_j) + \pi_{0j} \\ (\rho + p + \Pi)W^2 + 2q_0W - (\rho + \Pi - \pi_{00}) - nW \\ nW\Pi \\ nWq^j \\ nW\pi^{jk} \end{pmatrix} \\ \mathbf{F}^i &= \begin{pmatrix} DV^i \\ \alpha S_j^i - \beta^i S_j \\ \alpha S^i - \beta^i E - DV^i \\ XV^i \\ Y^j V^i \\ Z^{jk} V^i \end{pmatrix}, \\ \mathbf{S} &= \begin{pmatrix} 0 \\ \frac{1}{2}\alpha S^{ik}\partial_j\gamma_{ik} + S_i\partial_j\beta^i - E\partial_j\alpha \\ \alpha S^i K_{ij} - S^j\partial_j\alpha \\ \frac{\alpha n}{\tau_{\Pi}}(\Pi_{\text{NS}} - \Pi + \Delta_{\Pi}) \\ \frac{\alpha n}{\tau_q}(\gamma_{\mu}^j q_{\text{NS}}^{\mu} - q^j + \Delta_q^j + \tau_q\mathcal{G}_q^j + \tau_q\mathcal{H}_q^j) \\ \frac{\alpha n}{\tau_{\pi}}(\gamma_{\mu}^j \gamma_{\nu}^k \pi_{\text{NS}}^{\mu\nu} - \pi^{jk} + \Delta_{\pi}^{jk} + \tau_{\pi}\mathcal{G}_{\pi}^{jk} + \tau_{\pi}\mathcal{H}_{\pi}^{jk}) \end{pmatrix} \end{aligned} \quad (68)$$

and we have defined $V^i = \alpha v^i - \beta^i$, $K_{\mu\nu} = -\nabla_{\mu}n_{\nu} - n_{\mu}a_{\nu}^{(n)}$ as the extrinsic curvature and $a_{(n)}^{\mu} = n^{\nu}\nabla_{\nu}n^{\mu}$ as the acceleration of the normal observer. The Δ terms represent the next-to-leading order contributions to the dissipative terms' sources and the \mathcal{G} and \mathcal{H} terms represent couplings between the dissipative variables and the gauge and curvative variables.

Let us consider how the MISCE derivation will work now: more details and an illustration in a different context are given in Wright (2020). We consider the $\tau \rightarrow 0$ limit so that the source terms \mathcal{G} and \mathcal{H} are irrelevant for the expansion about this limit. We also neglect higher-order corrections as usual so that the Δ terms vanish.

The perturbative expansion about the Navier–Stokes equilibrium is conceptually the same, if algebraically more complex. The procedure that replaces time derivatives of the perturbative quantities with spatial derivatives of equilibrium quantities does, however, involve derivatives of more terms that appear in equation (68).

For instance, derivatives of the (square root of) spatial metric's determinant ($\sqrt{\gamma}$) appear in equation (67). The spatial derivatives are fairly straightforward to calculate using finite-differencing. However, we would like to avoid using finite-differencing for the temporal derivatives, which requires additional memory allocation and is generally low-order accurate. To this end, we use the equation, $\nabla_{\mu}\sqrt{-g} = 0$, as well as the relation, $\sqrt{-g} = \alpha\gamma$, to arrive at the identity

$$\alpha\partial_{\mu}\sqrt{\gamma} + \sqrt{\gamma}\partial_{\mu}\alpha = 0 \quad (69)$$

which, with index $\mu = 0$, gives

$$\partial_t\sqrt{\gamma} = \frac{-\sqrt{\gamma}}{\alpha}\partial_t\alpha, \quad (70)$$

where the time-derivative of α is usually calculated anyway as it is needed for the evolution of the gauge.

7 CONCLUSIONS

We have presented a dissipative extension to the relativistic, ideal hydrodynamic equations often used in astrophysical simulations. Motivated by the relaxation form of the MIS sources for the dissipative variables, new source terms are derived by writing the dissipative variables as a series expansion in deviations from their equilibrium, relativistic Navier–Stokes values. The series is parameterized by the dissipation strength and time-scale coefficients, and its terms are found using an order-by-order comparison of the MIS equations of motion. This leads to a rapidly convergent series in the case of fast-acting, weak dissipation, which we term the MISCE formulation.

This new system is numerically non-stiff in the exact limit where the commonly used MIS equations of motion are stiff and inaccurate in the opposing limit. Because much of the matter in a neutron star may be treated as a near-ideal fluid, the MISCE equations of motion may be evolved explicitly, giving accurate results with execution times that are about an order of magnitude reduced. Even when both models are evolved with the same, explicit integrators, the MISCE formulation is nearly twice as fast, owing to its reduced system size. It also converges to the Euler equations in the zero-dissipation limit, allowing for the natural evolution of a fluid, which is mostly ideal with some areas of non-equilibrium behaviour.

Within its domain of validity, we have demonstrated it to produce highly similar results to the MIS formulation for a range of initial data. It is able to capture dissipative effects near discontinuous data without the onset of Gibbs oscillations and shows little error growth (compared to MIS results) for smooth solutions over dissipation strengths and time-scales spanning many orders of magnitude. For more complex simulations of Kelvin–Helmholtz instabilities, the expected scaling laws are reproduced for the kinetic power spectrum across the inertial range of wavenumbers.

The realm of stability for our new model is considered in Section 6.2 and is dependent both upon the dissipation coefficients (in particular the ratio of strength to time-scale) and the simulation's space–time resolution, with a sharper dependence on the spatial resolution. The presence of many mixed-order derivatives in the source terms can lead to instabilities when spatial resolutions are pushed very high, though this effect may be mitigated somewhat by using better numerical-derivative approximations (than simple

finite-differences) such as slope-limiting ones. In Section 5.3.1, we presented results showing the expected convergence for the fourth-order RK and fifth-order WENO schemes we use. One caveat is that at high resolutions, the MISCE formulation, which makes use of second-order central-differencing of the primitive variables, starts converging at second-order in the grid-spacing. Similarly, the MIS formulation starts converging at first-order for high resolutions, when the dominant error contribution is the first-order time derivatives calculated using lagged updates.

In Section 4, to simplify the form of the matrices in the source that we present, we made the assumption that terms of $\mathcal{O}(v^2)$ and higher were negligible, and hence that the Lorentz factor, W , could be set to unity. Whilst for simulations we use the entire, non-simplified expressions that we derived using computer algebra, the differences this made to results were small, and were generally eclipsed by resolution effects. However, the differences may be more significant for fluid velocities approaching the speed of light, such as in the final orbits of a binary neutron star pre-merger, or for the significantly out-of-equilibrium matter created in the merger itself.

Although all simulations have been performed in the special relativistic limit, the techniques we have used are not limited to this alone. A general relativistic extension to MISCE is (in principle, at least) straightforward and already underway – we have sketched the structure of this in Section 6.3. In addition, we have developed an adaptive code prototype that evolves different dissipative fluid formulations in different physical regimes (e.g. MIS and MISCE) to minimize computational work and maximize accuracy and stability. Together, this should allow for more efficient, dissipative simulations of neutron star mergers and accretion on to compact objects.

The authors acknowledge the use of the IRIDIS High Performance Computing Facility, and associated support services at the University of Southampton, in the completion of this work. Open source software used includes SYMPY (Meurer et al. 2017), MATPLOTLIB (Hunter 2007), and cminpack (Devernay 2017).

DATA AVAILABILITY

The data used to perform the analyses and produce the figures in this article were obtained using the METHOD (<https://www.github.com/MarcusHatton/METHOD>) codebase, originally developed by Alex James Wright (Wright & Hawke 2019, 2020; Wright 2020), and will be shared with any interested party upon reasonable request.

REFERENCES

- Abbott B. et al., 2017, *Phys. Rev. Lett.*, 119, 161101
 Abbott B. et al., 2018, *Phys. Rev. Lett.*, 121, 161101
 Abbott B. et al., 2019, *Phys. Rev. X*, 9, 011001
 Ackley K. et al., 2020, *Publ. Astron. Soc. Aust.*, 37, e047
 Alford M. G., Bovard L., Hanuske M., Rezzolla L., Schwenzer K., 2018, *Phys. Rev. Lett.*, 120, 041101
 Beckwith K., Stone J. M., 2011, *ApJS*, 193, 6
 Bemfica F. S., Disconzi M. M., Noronha J., 2019, *Phys. Rev. Lett.*, 122, 221602
 Bemfica F. S., Disconzi M. M., Hoang V., Noronha J., Radosz M., 2021, *Phys. Rev. Lett.*, 126, 222301
 Bemfica F. S., Disconzi M. M., Noronha J., 2022, *Phys. Rev. X*, 12, 021044
 Berti E. et al., 2022, preprint ([arXiv:2203.06240](https://arxiv.org/abs/2203.06240))
 Bevilacqua L., Galeão A. C. N. R., Costa F. P., 2011, *J. Braz. Soc. Mech. Sci. Eng.*, 33, 166
 Biswas R., Dash A., Haque N., Pu S., Roy V., 2020, *J. High Energy Phys.*, 2020, 171
 Carrasco F., Viganò D., Palenzuela C., 2020, *Phys. Rev. D*, 101, 063003
 Cattaneo C., 1948, *Atti Sem. Mat. Fis. Univ. Modena*, 3, 83
 Celora T., Andersson N., Hawke I., Comer G., 2021, *Phys. Rev. D*, 104, 084090
 Cercignani C., Kremer G. M., 2002, *Progress in Mathematical Physics*, The Relativistic Boltzmann Equation: Theory and Applications. Birkhäuser, Basel, available at: <http://link.springer.com/10.1007/978-3-0348-8165-4>
 Chabanov M., Rezzolla L., Rischke D. H., 2021, *MNRAS*, 505, 5910
 Chapman S., Cowling T. G., 1990, *The Mathematical Theory of Non-Uniform Gases: An Account of the Kinetic Theory of Viscosity, Thermal Conduction, and Diffusion in Gases*, 3rd edn., Cambridge Univ. Press, Cambridge
 Chugunov A. I., Yakovlev D. G., 2005, *Astron. Rep.*, 49, 724
 Dash A., Shokri M., Rezzolla L., Rischke D. H., 2023, *Phys. Rev. D*, 107, 056003
 Del Zanna L. et al., 2013, *Eur. Phys. J. C*, 73, 2524
 Dionysopoulou K., Alic D., Palenzuela C., Rezzolla L., Giacomazzo B., 2013, *Phys. Rev. D*, 88, 044020
 Du L., Heinz U., 2020, *Comput. Phys. Commun.*, 251, 107090
 Duez M. D., Liu Y. T., Shapiro S. L., Stephens B. C., 2004, *Phys. Rev. D*, 69, 104030
 Gottlieb S., Shu C.-W., 1998, *Math. Comput.*, 67, 73
 Hammond P., Hawke I., Andersson N., 2021, *Phys. Rev. D*, 104, 103006
 Hammond P., Hawke I., Andersson N., 2023, *Phys. Rev. D*, 107, 043023
 Heller M. P., Serantes A., Spaliński M., Withers B., 2023, *Phys. Rev. Lett.*, 130, 261601
 Hiscock W. A., Lindblom L., 1983, *Ann. Phys.*, 151, 466
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Israel W., 1976, *Ann. Phys.*, 100, 310
 Israel W., Stewart J. M., 1979, *Ann. Phys.*, 118, 341
 Kiuchi K., Kyutoku K., Sekiguchi Y., Shibata M., 2018, *Phys. Rev. D*, 97, 124039
 Kumar K., 1967, *Aust. J. Phys.*, 20, 205
 LeVeque R. J., 2002, *Cambridge Texts in Applied Mathematics: Finite Volume Methods for Hyperbolic Problems*. Cambridge Univ. Press, Cambridge, p. 410
 Manuel C., Tolos L., 2011, *Phys. Rev. D*, 84, 123007
 Meurer A. et al., 2017, *PeerJ Comput. Sci.*, 3, e103
 Miranda-Aranguren S., Aloy M. A., Rembiasz T., 2018, *MNRAS*, 476, 3837
 Molnár E., Niemi H., Rischke D. H., 2010, *Eur. Phys. J. C*, 65, 615
 Most E. R. et al., 2021, *MNRAS*, 509, 1096
 Most E. R., Haber A., Harris S. P., Zhang Z., Alford M. G., Noronha J., 2022, 967, L14
 Palenzuela C., Lehner L., Reula O., Rezzolla L., 2009, *MNRAS*, 394, 1727
 Pandya A., Most E. R., Pretorius F., 2022, *Phys. Rev. D*, 105, 123001
 Pareschi L., Russo G., 2005, *J. Sci. Comput.*, 25, 129
 Punturo M. et al., 2010, *Class. Quantum Gravity*, 27, 194002
 Qian J., 1994, *Phys. Rev. E*, 50, 611
 Radice D., 2017, *ApJ*, 838, L2
 Radice D., Hawke I., 2024, *Liv. Rev. Comput. Astrophys.*, 10, 1
 Reitze D. et al., 2019, *BAAS*, 51, 35
 Rezzolla L., Pizzochero P., Jones D. I., Rea N., Vidaña I., 2018, in Steven N., ed., *Astrophysics and Space Science Library*, Vol. 457, *The Physics and Astrophysics of Neutron Stars*. Springer-Verlag, Berlin
 Ripperda B. et al., 2019, *ApJS*, 244, 10
 Schmitt A., Shternin P., 2018, in Rezzolla L., Pizzochero P., Jones D. I., Rea N., Vidaña I., eds, *Astrophysics and Space Science Library*, Vol. 457, *The Physics and Astrophysics of Neutron Stars*. Springer-Verlag, Berlin, p. 455
 Shibata M., Kiuchi K., Sekiguchi Y.-i., 2017, *Phys. Rev. D*, 95, 083005
 Shu C.-W., 1998, in Cockburn B., Shu C.-W., Johnson C., Tadmor E., Quarteroni A., eds, *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations*. Springer, Berlin, Heidelberg, p. 325
 Smagorinsky J., 1963, *Monthly Weather Rev.*, 91, 99
 Takamoto M., Inutsuka S.-i., 2011, *J. Comput. Phys.*, 230, 7002
 Thompson C., Duncan R. C., 1993, *ApJ*, 408, 194
 Viganò D., Aguilera-Miret R., Carrasco F., Miñano B., Palenzuela C., 2020, *Phys. Rev. D*, 101, 123019
 Wagner D., Gavassino L., 2024, *Phys. Rev. D*, 109, 016019

- Wright A. J., 2020, PhD thesis, University of Southampton, available at: <https://eprints.soton.ac.uk/448034/>
- Wright A. J., Hawke I., 2019, *ApJS*, 240, 8
- Wright A. J., Hawke I., 2020, *MNRAS*, 491, 5510
- Yang Y., Hippert M., Speranza E., Noronha J., 2024, *Phys. Rev. C*, 109, 015805

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

suppl.data

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: APPROXIMATING TIME DERIVATIVES

In Section 4 we showed that by making use of the fundamental conservation-law equation

$$\partial_t \mathbf{q}(\mathbf{w}) + \partial_i \mathbf{f}^i(\mathbf{w}) = \mathbf{0} \quad (\text{A1})$$

and the simple chain-rule for derivatives, we are able to arrive at an approximation to time derivatives of the primitive variables, \mathbf{w} , containing only spatial derivatives of the fluxes (or primitive variables).

$$\frac{\partial \mathbf{w}}{\partial t} = \frac{\partial \mathbf{w}}{\partial \mathbf{q}} \frac{\partial \mathbf{q}}{\partial t} = - \left(\frac{\partial \mathbf{q}}{\partial \mathbf{w}} \right)^{-1} \partial_i \mathbf{f}^i = - \mathbf{A} \mathbf{B} \quad (\text{A2})$$

where

$$\mathbf{A} = \left(\frac{\partial \mathbf{q}}{\partial \mathbf{w}} \right)^{-1}, \quad \mathbf{B} = \partial_i \mathbf{f}^i. \quad (\text{A3})$$

In the case of BDNK models of dissipative fluids, and of our MISCE model presented here (to first-order), we may write both

the conservative and flux vectors in an expanded form that separates (first-order) derivatives in the primitive quantities:

$$\begin{aligned} \partial_t [\mathbf{q}_{(0)}(\mathbf{w}) + \epsilon \mathbf{q}_{(1)}(\mathbf{w}, \partial_t \mathbf{w}, \partial_i \mathbf{w})] \\ + \partial_i [\mathbf{f}_{(0)}^i(\mathbf{w}) + \epsilon \mathbf{f}_{(1)}^i(\mathbf{w}, \partial_t \mathbf{w}, \partial_i \mathbf{w})] = 0 \end{aligned} \quad (\text{A4})$$

where ϵ parametrizes the size of dissipation in the fluid model. Under the assumption that dissipation is small compared to the bulk behaviour of the fluid, ϵ is small. This is the regime when using the MISCE makes sense, anyway.

Now we cannot only consider the contribution of the fluid variables themselves to the time-derivative of the state vector ($\partial_t \mathbf{w}$), but also the contribution of the temporal and spatial derivatives ($\dot{\mathbf{w}}, \mathbf{w}'$). We first rewrite \mathbf{A} and \mathbf{B} as

$$\mathbf{A} = [\partial_u \mathbf{q}_0 + \epsilon \partial_w \mathbf{q}_1]^{-1}, \quad \mathbf{B} = [\partial_i \mathbf{f}_0^i + \epsilon \partial_i \mathbf{f}_1^i]. \quad (\text{A5})$$

After much manipulation, and using the assertion that ϵ is indeed small to expand a sum of matrices to leading order, we arrive at

$$\mathbf{A} = \left[\left(\mathbf{I} - \epsilon \left(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}} \right)^{-1} \frac{\partial \mathbf{q}_1}{\partial \mathbf{w}} \right) \left(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}} \right)^{-1} \right], \quad (\text{A6a})$$

$$\mathbf{B} = \left[\frac{\partial \mathbf{f}_0^i}{\partial \mathbf{w}} \mathbf{w}' + \epsilon \left(\frac{\partial \mathbf{f}_1^i}{\partial \mathbf{w}} \mathbf{w}' + \frac{\partial \mathbf{f}_1^i}{\partial \dot{\mathbf{w}}} \dot{\mathbf{w}}' + \frac{\partial \mathbf{f}_1^i}{\partial \mathbf{w}'} \mathbf{w}'' \right) \right]. \quad (\text{A6b})$$

Note that we can choose to use the form of \mathbf{A} in equation (A3) and invert the sum of matrices, rather than using the approximate small- ϵ trick that leads to \mathbf{A} in equation (A6a). Similarly, we can choose the expression for \mathbf{B} from equation (A3) which makes use of the fluxes themselves directly, or we may use its form in equation (A6a) which requires evaluation of second-order spatial differences of the primitives.

This paper has been typeset from a *TeX/LaTeX* file prepared by the author.