

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]

University of Southampton

Faculty of Social Sciences
School of Mathematical Sciences

**Simulating Relativistic, Dissipative
Hydrodynamics for Neutron Star
Mergers**

by

Marcus John Hatton

BA (Cantab), MSc

ORCID: [0000-0003-3301-2166](https://orcid.org/0000-0003-3301-2166)

*Thesis for the degree of
Doctor of Philosophy*

23 May 2025

University of Southampton

Abstract

Faculty of Social Sciences
School of Mathematical Sciences

Doctor of Philosophy

**Simulating Relativistic, Dissipative Hydrodynamics for Neutron Star
Mergers**

by Marcus John Hatton

Neutron stars are astrophysical compact objects that contain a host of extreme physics including strong gravity, immense magnetic fields and highly dense nuclear matter. Neutron star mergers are violent enough to probe this physics in a meaningful way through our detection of the electromagnetic and gravitational waves we receive from these events.

However, their complexity necessitates the use of numerical simulations to evolve the curved spacetime, as well as the fluid used to describe the neutron stars themselves. Most simulations to date use an ‘ideal fluid’ description - a simplification that reduces computational cost but misses out-of-equilibrium effects like dissipation.

Here, we compare results from a number of both established and emerging models of non-ideal hydrodynamics. This includes results from our novel formulation of dissipative hydrodynamics, which we show is able to capture viscosity and heat conductivity near the ideal fluid limit in a highly efficient computational way.

Even the highest resolution simulations of binary neutron star mergers are unable to resolve all the dynamical scales of these systems, leaving unresolved, subgrid microphysics. This can produce macroscale effects, which our field of research has recently begun to model using explicit large-eddy simulations.

To advance this direction, we also present the first fully-covariant, Lagrangian filtering scheme applied to relativistic turbulence. We show that a dissipative fluid prescription, with statistically-fitted non-ideal parameters, may be used as a closure scheme to describe the residuals introduced by explicitly filtering fine-scale fluid flow.

These pieces of work, together, will enable us to perform higher resolution simulations of neutron star mergers that include more of the relevant physics. This will serve us well when trying to obtain accurate gravitational waveforms and electromagnetic signals to compare against future detections from next-generation detectors.

Publications

Parts of this thesis are available as published work, and can also be viewed on the [arXiv](#) website.

In particular, the work of chapter 5 is [viewable as a preprint online](#) (Hatton and Hawke, 2024) and has [been published](#) in the journal Monthly Notices of the Royal Astronomical Society.

The work of chapter 6 is also [viewable as a preprint online](#) (Celora, Hatton, et al., 2024) and [has recently been published](#) in the journal Physical Review D. Note that T. Celora and myself contributed equally to this work.

Finally, I am involved in one further publication (Celora, Andersson, et al., 2024) which is [viewable as a preprint online](#) and [has also recently been published](#) in the journal Physical Review D. However, little of this work is present in this thesis.

Contents

List of Figures	xi
List of Tables	xix
Declaration of Authorship	xxi
Acknowledgements	xxiii
1 Introduction	1
1.1 Motivation - The Physics of Neutron Stars	1
1.1.1 Formation	2
1.1.2 Structure	3
1.1.3 Equation of State	4
1.1.4 Magnetic Fields	6
1.1.5 Transport	7
1.1.6 Neutron Star Merger Events	9
1.1.6.1 In Nature	9
1.1.6.2 In Silico	11
1.2 Summary and Thesis Structure	13
2 Theoretical Foundation	15
2.1 General Relativity	15
2.1.1 Numerical Relativity Formulation	16
2.2 Fluid Dynamics	18
2.2.1 Gradient Expansion Models	19
2.2.1.1 The Non-Relativistic Navier-Stokes Equations	20
2.2.1.2 Ideal Fluid Equations in Relativity	20
2.2.1.3 Equations of Non-Ideal, Relativistic Fluids	22
2.2.1.4 Relativistic Navier-Stokes Equations	23
2.2.1.5 The Müller-Israel-Stewart Formulation	24
2.2.1.6 BDNK Formulation	25
2.2.2 Valencia Formulation of Balance Law Equations	26
2.2.3 Variational Models	28
2.2.4 Smoothed Particle Hydrodynamics	28
2.3 Turbulence Modelling and Explicit Large Eddy Schemes	29
2.4 Kinetic Theory and the Chapman-Enskog Expansion	33
2.4.1 Viscosity via Chapman-Enskog Expansion	34
2.4.2 Heat Conductivity via Chapman-Enskog Expansion	37

3	Numerical Methods	39
3.0.1	Numerical Schemes for Relativistic Viscous Hydrodynamics in Astrophysics	40
3.1	The Balance-Law Form of Hyperbolic Hydrodynamic Equations	40
3.2	A Simple Numerical Scheme	43
3.3	Riemann Problems	44
3.4	Flux Methods	46
3.5	Time Integrators	47
3.5.1	Explicit Schemes	48
3.5.2	Implicit Schemes	49
3.5.3	IMEX Schemes	50
3.6	Boundary Conditions	52
3.7	Parallel Programming	53
3.8	Primitive Variable Recovery	54
3.9	Summary	55
4	Simulations with the Müller-Israel-Stewart Model	57
4.1	Reduced MIS Models	57
4.1.1	Heat Flux	60
4.1.1.1	Limiting case	61
4.1.1.2	Cattaneo equation	61
4.1.1.3	Chapman-Enskog expansion	62
4.1.1.4	Numerical Results	63
4.1.2	Bulk Viscosity	65
4.1.2.1	Limiting case	65
4.1.2.2	Chapman-Enskog expansion	66
4.1.2.3	Wave Speeds	66
4.1.2.4	Numerical Results	68
4.1.3	Shear Viscosity	68
4.1.3.1	Limiting case	69
4.1.3.2	Chapman-Enskog Expansion in 1D	71
4.1.3.3	Chapman-Enskog Expansion in Pseudo-2D	72
4.1.3.4	Wave Speeds	72
4.1.3.5	Numerical Results	74
4.1.4	Summary	74
4.2	The Full MIS Model	75
4.2.1	An Alternative C2P Scheme	77
5	Simulations with our ‘MISCE’ Model	85
5.1	General Balance-Law Derivation	87
5.1.1	First-Order Source	91
5.1.2	Second-Order Source	94
5.2	Simulation Results	95
5.3	Code Performance	100
5.3.1	Scaling and Convergence	100
5.3.2	Model Comparison	102
5.4	Further Considerations	104

5.4.1	Rapid Evolution of Reduced Initial Data	104
5.4.2	Stability Analysis	107
5.4.3	Extension to GR	110
5.5	Summary	111
6	A Covariant Filtering Implementation	113
6.1	Introduction	113
6.1.1	Logic of the Scheme	114
6.1.2	Box Simulations of the Kelvin-Helmoltz Instability	118
6.2	Stress-Energy Residuals and Closure Ingredients	126
6.3	Discriminating Between Models	129
6.3.1	Scaling with filter width	129
6.3.2	Linear regression in log-space	130
6.3.3	Model interpretation and viability	132
6.4	The Equation of State Residual	133
6.5	Modelling the Isotropic Stresses and Momentum Flux Residuals	135
6.6	Extracting the Shear Viscosity, Component-Wise	138
6.7	Extension to MHD	139
6.8	Summary	141
7	Other Avenues	145
7.1	Simulations with the BDNK Model	145
7.1.1	An Alternative Evolution Scheme	147
7.2	Subgrid Modelling with Dissipative Hydrodynamics	153
7.2.1	Assessing the SubGrid Model	158
8	Conclusions & Outlook	161
	Appendix A Approximating Time Derivatives	165
	Appendix B Extracting Numerical Viscosity	167

List of Figures

- 1.1 Neutron star structure. a: Schematic slice through a neutron star. Letters N , n , p , e and μ refer to the presence of nuclei, fluid neutrons and protons, electrons and muons, respectively. The inner core composition is still uncertain and various exotic possibilities exist, including hyperons and deconfined quark matter. b: A detailed picture of the composition of the inner crust. At lower densities, a lattice of superheavy, neutron-rich nuclei is immersed in a fluid of neutrons (which are likely to be superfluid) and a relativistic electron gas. At high enough densities the nuclei might deform and connect along certain directions to form extended tubes, sheets and bubbles of nuclear matter. These nuclear pasta phases might form a layer at the base of the neutron star crust, sometimes referred to as the mantle; searching for observational signatures of such phases is of great interest. Ranges of density and thickness given for each layer represent current uncertainties in the physics of neutron star crusts. This figure was taken from <https://compstar.uni-frankfurt.de/outreach/short-articles/the-nuclear-pasta-phase/>. 3
- 1.2 A schematic diagram illustrating the various stages in the evolution of an equal-mass binary system of neutron stars as a function of the initial mass of the binary. Depending on the initial total mass of the binary, M , and on how it relates to the maximum mass of a non-rotating neutron star, M_{TOV} , the binary can either collapse promptly to a black hole surrounded by a torus (top row), or give rise to a hypermassive (or supermassive) neutron star that ultimately collapses to a black hole and torus (middle row), or even lead to a supramassive neutron star (first differentially and subsequently uniformly rotating) neutron star that eventually yields a black hole or a non-rotating neutron star (bottom row). Also indicated in red are the typical frequencies at which gravitational waves are expected to be emitted. This figure was taken from (Baiotti and Rezzolla, 2017). 9
- 1.3 A visual summary of the (both current and proposed) sources and detectors of gravitational waves across the spectrum of frequencies/wave-periods. Taken from the [LISA website](#). 12
- 1.4 Example of a BNS merger simulation using General Relativistic Magnetohydrodynamics (from the models presented in Ciolfi, Kastaun, et al., 2017). The temporal sequence shows the bulk of the NS(s) (in white together) with color-coded isodensity surfaces. Deformation of the neutron stars through tidal forces can be seen pre-merger. Post-merger, the remnant neutron star is also deformed by the uneven distribution of angular momentum within it. 13

2.1	Schematic representation of the 3+1 decomposition of spacetime with hypersurfaces of constant time coordinate Σ_t and Σ_{t+dt} foliating the space-time. The four-vector \mathbf{t} represents the direction of evolution of the time coordinate t and can be split into a timelike component $\alpha\mathbf{n}$, where \mathbf{n} is a time-like unit normal to the hypersurface, and into a spacelike component, represented by the spacelike four-vector $\boldsymbol{\beta}$. Taken from Bishop and Rezzolla, 2016	17
2.2	Relativistic filtering of a fine-grain velocity field where black lines indicate streamlines. The red line is the coarse-grained, ‘average’ four-velocity. If the fine-grain field is filtered with respect to that observer, the coarse-grained field is found. This procedure is independent of the coordinates, but its inherent non-linearity makes it impractical for use within a simulation.	32
3.1	Wave pattern for the Riemann problem applied to a fluid with initial left and right states (L, R) taken from Harpole, 2018. The initial discontinuity decays into two non-linear waves $(W_{\leftarrow}, W_{\rightarrow})$ between which there are two new constant states (L^*, R^*) separated by a contact discontinuity \mathcal{C} . This wave pattern will be clearly seen in the shocktube test cases presented in section 4. The initial data for these tests essentially represent large-scale Riemann problems where W_{\leftarrow} is a rarefaction wave and W_{\rightarrow} is a shockwave.	45
4.1	The spreading of a spherically-symmetric hot “blob” by the heat flux, using equations (4.11). The 2D grid is 128×128 cells in size and the simulation runs until $t_{\text{final}} = 5$ with $\kappa = 1 \times 10^{-3}$ and $\tau_q = 1 \times 10^{-3}$	64
4.2	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 Chapman-Enskog expansion. There is excellent visual agreement. This is a 1D-slice view through the 2D domain seen in figure 4.1.	64
4.3	The velocity and bulk viscosity are plotted throughout our spatial domain. The initial velocity profile was a top-hat function with $v_x = 0.7$ for $0.3 \leq x \leq 0.7$ and $v_x = 0$ otherwise. In the ideal case, the top figure shows similar behaviour to the simple Burger’s equation, as expected. The non-ideal behaviour is given by the bulk viscosity in the lower plot, and its impact through smoothing strong gradients is clear in the velocity plot. The simulation ran with parameters $t_{\text{final}} = 0.2$, $\zeta = 1 \times 10^{-2}$ and $\tau_{\Pi} = 1 \times 10^{-3}$	70

- 4.4 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: firstly, the MIS-derived simple shear model given by eq. (4.35) (solid lines); secondly, the model obtained applying the CE-expansion to it, given by eq. (4.36) (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 analytic 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. 74
- 4.5 The development of 4 quantities (velocity, number density, pressure and bulk viscosity) for a shocktube setup with zero initial velocity as the bulk viscosity coefficient ζ is varied over several orders of magnitude. One can see several prominent features, such as the smearing of discontinuities and the acceleration of the shock-front. 78
- 4.6 The development of a Kelvin-Helmholtz unstable fluid with negligible viscosity until $t = 6.25$. 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. 79
- 4.7 The development of the Kelvin-Helmholtz instability until $t = 30.0$, for an ideal fluid with negligible viscosity. The longer simulation time allows the asymmetric initial perturbation at the interface to give rise to large-scale asymmetric vortex formation. 79
- 4.8 The short-term evolution of the Kelvin Helmholtz instability using the MIS model of non-ideal hydrodynamics with a shear viscosity parameter of $\eta = 1 \times 10^{-3}$ and timescale $\tau_\pi = 5 \times 10^{-3}$ 80
- 4.9 The long-term evolution of the Kelvin Helmholtz instability using the MIS model of non-ideal hydrodynamics with a shear viscosity parameter of $\eta = 1 \times 10^{-3}$ and timescale $\tau_\pi = 5 \times 10^{-3}$ 80
- 4.10 The power spectrum, $P_{\hat{u}(k_x)}$, for the kinetic energy density in the Kelvin-Helmholtz instability at $t = 15.0$ and $t = 30.0$, for both an inviscid and a viscous fluid. In the viscous cases, The expected Kolmogorov scaling in the inertial range is also plotted and matches the data well. In the viscous cases, dissipation has reduced the total kinetic energy present. It has also smoothed the minor oscillations present in the ideal curves which stem from the small-scale mixing modes at the interface which are suppressed by viscosity. 81
- 5.1 The evolution of four primitive variables for a 'stillshock' test - a shocktube with zero initial velocity - up until time $t = 1.5$. An ideal fluid is compared with a viscous one using the MISCE formulation at both leading and next-to-leading order, for a range of parameter values. The increase in shock propagation speed and smearing of discontinuities due to the inclusion of bulk viscosity are both visible physical effects. The second-order terms have little impact on the solution as they are $\mathcal{O}(\zeta\tau_\pi \sim 1 \times 10^{-4})$ but do make some visible difference to the temperature and number density at the contact wave discontinuity. 96

- 5.2 The evolution of the number density for the same stillshock test, but now until time $t = 5.0$, with the x -coordinate plotted horizontally. 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 timescale τ varying for the MIS model but held constant at zero for the MISCE model. 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. 97
- 5.3 The development of a Kelvin-Helmholtz unstable fluid with negligible viscosity until $t = 6.25$. 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. 97
- 5.4 The development of the Kelvin-Helmholtz instability until $t = 30.0$, for an ideal fluid with negligible viscosity. The longer simulation time allows the asymmetric initial perturbation at the interface to give rise to large-scale asymmetric vortex formation. 98
- 5.5 The long-term evolution of the Kelvin Helmholtz instability 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. 98
- 5.6 The evolution of the temperature for the same shocktube test with zero initial velocity and varying heat flux parameter, κ , across three orders of magnitude. An effect is only seen for the MISCE model in the highest- κ case, where an additional heating of fluid between the shock and contact discontinuities is seen. We originally considered this to be the breakdown of the MISCE approximation occurring, and MIS to be the ‘true’ solution. However, we argue that the MISCE solution here is correct in its effective limit of $\tau_q \rightarrow 0$, which cannot be replicated in the MIS model due to numerical stability bounds on the ratio κ/τ_q 99
- 5.7 The power spectra for the kinetic energy density in the Kelvin-Helmholtz instability 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 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 lengthscales). 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 lengthscales of a few, or even a single, cell(s). 100

- 5.8 The “SineWave” initial data (and its evolution with the MISCE model) 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. 101
- 5.9 The self-convergence rate of the MISCE model for different resolutions using the number density in the ‘sinewave’ test seen in fig. 5.8. 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. 102
- 5.10 The difference between the MISCE model at $\tau = 0.0$ (leading order) and the MIS model with varying τ , where all three forms of dissipation are present. The L_2 -norm of the difference between number densities (left) and velocities (right) are plotted for shocktube initial data as seen in fig. 5.2. Please see the text for an explanation of the observed behaviour. 104
- 5.11 The evolution of the temperature and heat flux for the ‘top-hat’ initial data described in section 5.4.1 using the relaxation model given by eq. (5.26). 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 timescale shown here, as seen by the heat flux relaxing to nearly zero. The analytic result for the appropriate adjustment to the initial data, derived in section 5.4.1, is also plotted in the left panel (dotted) and shows excellent agreement with the numerical evolution result. 107
- 5.12 The stability criteria of eq. (5.39) are plotted separately (dashed, dotted) and summed (solid). The red, shared region shows where the simple heat flux model given by eq. (5.38) 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 $\Delta x \simeq 0.045$ 109
- 6.1 The development of a Kelvin-Helmholtz instability for an ideal fluid with negligible viscosity. The number density is plotted and the simulation ran until a code time of $t = 16.0$. Breakdown of the interface, seeded by a random initial perturbation, occurs quickly and leads to vortex formation and the onset turbulence across the domain. 119
- 6.2 The adjusted, integrated power spectra of the Lorentz factor for fluid in the Kelvin-Helmholtz simulation seen above in fig. 6.1. The inertial range, where the power scales as $k^{-5/3}$, would correspond to a horizontal line here. On the left, three curves for times $t = 4.0$, $t = 10.0$ and $t = 20.0$ are plotted for the micromodel data (straight from simulation). On the right, four spectra are plotted, all at $t = 10.0$: one for the micromodel data again, and three for the mesomodel with varying parameters. CG is the coarse-graining factor of the mesomodel’s grid compared to the micromodel’s one and FW is the filter width used for the averaging operation. In all cases the inertial range is seen over at least an order of magnitude of wavenumbers. 120

- 6.3 Comparison between the fine-scale velocity u^μ (middle panel) and the filtering observers U^μ (right panel), computed using a box-length of $8\Delta x$ and focusing on the time component. Data from a snapshot taken at around $t = 10$. In the left-most panel we show the relative difference between the two velocities, noting that this is typically of the order of a few percent and peaks where u^μ presents sharp gradients. 123
- 6.4 Comparison between the fine-scale n^μ and filtered $\langle n^\mu \rangle$ baryon current (filter-size $L = 8\Delta x$), focusing on the time component, at a representative time $t = 10$. The left panel shows the relative difference between the two over the full grid. In the middle and right panel we plot the fine-scale and filtered data zooming in the region $x \in (0.15, 0.35)$, $y \in (0.5, 0.7)$ (the box indicated in the left panel). We do so to visually appreciate the effects of filtering: the right panel presents the same ‘bands’ as the middle one but the boundaries between them are smeared out. 124
- 6.5 Comparison between the fine-scale n^μ and filtered $\langle n^\mu \rangle$ baryon current, focusing on the time component, at a representative time $t = 10$. The left-most panel shows the fine-scale data, whilst in the panels to its right we plot the relative difference between this and the filtered data—with filter sizes of $2, 4, 8\Delta x$. As we observed the maximum values of these to roughly double as we double the filter-size, we plot here the re-scaled relative differences, that is we divide by the filter-size L . This allows us to use a single colour map for the three panels to the right and provides a better comparison among them. 124
- 6.6 Comparison between fine-scale baryon current n^μ and filtered data using various filter-widths. The left-most panel shows a snapshot of the fine-scale data at around $t = 10$. In the panels to its right we plot the filtered baryon current with filter-sizes $2, 4, 8\Delta x$, zooming in the region $x \in (0.1, 0.4)$, $y \in (0.4, 0.7)$. The corresponding patch is indicated in the left panel. Data for this figure has been produced using an observer box-length equal to the filter size and explicitly coarse-graining the meso grid. Coarse-graining is particularly evident in the third and fourth panel, which makes the images appear pixelated. 125
- 6.7 Plotting the magnitude of the stress-energy tensor residuals. From left to right: the isotropic residual stresses $\tilde{\Pi}$, the anisotropic residual stresses $\tilde{\pi}^{\mu\nu}$ and residual momentum flux \tilde{q}^μ . Data underlying this figure is obtained from filtering a snapshot at around $t = 10$ using a filter-size of $8\Delta x$. 127
- 6.8 Results for various gradients relevant to the modelling of the stress-energy tensor residuals. From left to right we show: the expansion rate $\tilde{\theta}$, the square-root of the trace of the squared shear tensor $\sqrt{\sigma_{\mu\nu}\sigma^{\mu\nu}}$, and the temperature gradients corrected by the heat inertia $\sqrt{\tilde{\Theta}_\mu\tilde{\Theta}^\mu}$, where $\tilde{\Theta}^\mu = \tilde{\mathcal{I}}^{\mu\nu}(\nabla_b\tilde{T} + \tilde{T}\tilde{a}_b)$. Data for this figure was obtained from filtering a snapshot at around $t = 10$ using a filter size of $8\Delta x$ 128
- 6.9 Comparing distributions of stress-energy tensor residuals with their values from a first-order gradient model that assumes constant coefficients. From left to right we show: the isotropic residual stresses $\tilde{\Pi}$, the anisotropic residual stresses $\tilde{\pi}^{\mu\nu}$ and the residual momentum flux \tilde{q}^μ 129

6.10	Comparing distributions at different filter sizes. From left to right we show: i) the second invariant of the shear tensor, ii) the ‘magnitude’ of the anisotropic stress residuals and iii) the re-scaled, extracted shear viscosity.	130
6.11	Best regression model in log space for η . From left to right we show: i) scatter-plot of the extracted data vs. model predictions ii) probability distributions of extracted η vs its modelling iii) probability distributions of the extracted residual vs its modelling.	133
6.12	Visualizing the different residuals contributing to the trace of the stress-energy tensor. In the left-most panel we show distributions corresponding to the EoS residual, the isotropic stresses residual and the sum of the two. In the panels to its right we plot the same quantities individually. Data used in this figure has been filtered with filter-size $L = 8 \Delta x$ using snapshots around $t = 10$	134
6.13	Comparing the 1 st adiabatic index Γ_1 , computed locally using eq. (6.19), to the expected value of $4/3$. The left panel corresponds to the fine-scale data, with a largest relative difference on the order of 10^{-11} . The right panel corresponds to the filtered data, explicitly ignoring the EoS residual. For this case we observe differences up to percent level. The data underlying this figure has been filtered with filter size $L = 8 \Delta x$ using snapshots around $t = 10$	136
6.14	Comparing distributions of ζ (left) and κ (right) at different filter sizes. The almost perfect overlap of the re-scaled distributions demonstrate that the coefficients (and the residuals) follow the same scaling as discussed for η	137
6.15	Equivalent to fig. 6.11 but focusing on the heat conductivity. From left to right we show: i) scatter-plot of the extracted data vs. model predictions ii) probability distributions of extracted κ vs its modelling iii) probability distributions of the extracted residual vs its modelling.	137
6.16	Same as fig. 6.11 but focusing on the bulk viscosity. From left to right we show: i) scatter-plot of the extracted data vs. model predictions ii) probability distributions of extracted ζ vs its modelling iii) probability distributions of the extracted residual vs its modelling.	137
6.17	Extracting η component-wise: comparing the magnitude distributions obtained in a component-wise fashion (one per independent component) vs that via squaring. Data underlying this figure has been filtered with filter-size $L = 8 \Delta x$ using snapshots around $t = 10$	140
6.18	Comparing the histograms for positive and negative values of η , component-wise and via squaring. For each panel we report the total counts of positive vs negative values, so to be able to appreciate whether there are more positive than negative values or vice versa. Data underlying this figure has been filtered with filter size $L = 8 \Delta x$ and snapshots around $t = 10$	140
7.1	The evolution of discontinuous initial data (a shocktube) for the ideal and bulk-viscous ($\zeta = 1 \times 10^{-2}$) cases. The BDNK formulation is compared to the MIS model, showing excellent visual agreement in both cases.	148
7.2	The evolution of a spinning ‘rotor’ with initial data similar to that in Pandya, Most, et al., 2022b using the BDNK formulation with a shear viscosity coefficient of $\eta = 1 \times 10^{-3}$. The weak shear viscosity allows small-scale vortex structure to develop but eventually causes breaking of the fluid.	149

7.3	The long-term evolution of the Kelvin Helmholtz instability using the BDNK formulation of non-ideal hydrodynamics with no non-ideal effects and 800×800 cells.	150
7.4	The long-term evolution of the Kelvin Helmholtz instability using the BDNK formulation of non-ideal hydrodynamics with a viscosity parameter of $\eta = 1 \times 10^{-4}$ and 800×800 cells. The small but finite shear viscosity here initially suppresses the breaking of the interface but then causes large-scale features to develop (roll-up, vortices) with fine-scale details within them.	151
7.5	The evolution of the Kelvin-Helmholtz instability (KHI) and onset of turbulence for an ideal fluid. The number density, n , is plotted over a low-resolution domain of 200×200 cells. No subgrid model is active but this simulation is the one used to calculate the model's coefficients in the plots (figs. 7.6 to 7.9) that follow.	155
7.6	The subgrid model coefficient κ plotted over the domain of the turbulent KHI simulation seen above. The spatial correlation with 'sharp' features (strong gradients) in the fluid flow is clear, in particular at the shocks. . .	155
7.7	The same plot but for the coefficient ζ . Some spatial structure is seen but much of the domain is 'dark' (zero) because of the clipping we perform when the expansion $\tilde{\theta}$ diverges, which has a negative power fitted in our subgrid model.	156
7.8	For the coefficient η , spikes in its value are seen along the discontinuities, and there is clearly spatial correlation with the vortices.	156
7.9	The distribution of values of κ , ζ and η in the subgrid model for a 200×200 resolution box simulation with initial data creating a Kelvin-Helmholtz instability.	157
7.10	Low resolution (200×200) simulation of the KHI for an ideal fluid with our subgrid model active. Compared to the case with no subgrid model, seen in fig. 7.5, there are small but significant visual differences. In particular, the shocks are noticeably shifted in position due to the inclusion of (effective) dissipative effects that modify their propagation speeds. Also, the maximum number density is slightly suppressed, which is unsurprising given that dissipation smooths gradients in fluid properties.	159
Appendix B.1	The smearing of an initial discontinuity due to numerical viscosity for a number of resolutions. The simulation ran until time $t = 20.0$ with an ideal fluid description. This behaviour matches the analytic solution of equation eq. (B.2) which is described by a broadening error function.	168
Appendix B.2	The power-law scaling of the effective viscosity η_{eff} with grid spacing Δx for the smearing of an initial discontinuity as seen in fig. B.1. As expected, the effective viscosity is proportion to the grid spacing and hence vanishes in the continuum limit of infinite resolution.	169

List of Tables

3.1	Butcher tableau table used to concisely summarise Runge-Kutta methods.	48
3.2	Butcher tableau for the explicit forward-Euler method.	49
3.3	Butcher tableau for the implicit backward-Euler method.	49
3.4	Butcher tableau for the ‘classic’, explicit RK4 method.	49
3.5	Butcher tableau of the explicit coefficients a_{ij}, b_j for the SSP2(2,2,2) IMEX scheme.	51
3.6	Butcher tableau of the implicit coefficients \bar{a}_{ij}, \bar{b}_j for the SSP2(2,2,2) IMEX scheme. $c = 1 - \frac{1}{\sqrt{2}}$.	51
4.1	Code run-times for our toy bulk viscous model in one dimension, simulating the data seen in figure 4.3. The expected scaling of the run-time with resolution, $t_{run} \propto N_x^2$, can be seen for both the full model and the Chapman-Enskog expansion (the latter in parentheses). Crucially, a nearly three-fold speed-up is achieved by using the Chapman-Enskog expansion, even with identical numerical methods.	69
5.1	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.	102
5.2	A comparison of computational time required for different hydrodynamic models and time-integrators. These results are for Kelvin-Helmholtz instability 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, 2^{nd} -order Runge-Kutta scheme and SSP2(222) refers to a 2^{nd} -order implicit-explicit scheme.	103

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as: "A dissipative extension to ideal hydrodynamics" (Hatton and Hawke, 2024), and "Covariant approach to relativistic large-eddy simulations: Lagrangian filtering" (Celora, Hatton, et al., 2024).

Acknowledgements

My sincerest gratitude must go firstly to Ian for your constant guidance and wisdom in all things physics, neutron stars and PhDs! I consider myself very lucky in having had you as my supervisor and I hope that we can continue to work together in the future. All the best.

Thank you also to Nils for your mentorship and leading our Gravity group so enigmatically here at Southampton. Thanks again to Ian & Nils, as well as my other very clever collaborators, Thomas & Greg, for bringing me on board with your vision for the field and sharing so much knowledge through our many insightful & entertaining discussions together.

To the rest of the Gravity group - Ian J., Andreas, Carsten, Shanshan, Tom, Yashaswi, Grant, Rhys, Rahime, Geraint, Savvas, Orestis, Pete, Fabian and others I am no doubt forgetting - thank you for gracefully enduring my many talks on the fine details of numerical methods (and still asking insightful questions at the end!), for your companionship on conference trips and for making the office and hence my PhD a much nicer place to be than it otherwise would have.

Amina, thank you for being the most wonderful partner through it all and making my life outside of work so bright. I look forward to our future together as a (hopefully soon) pair of Drs! I love you lots x

To my family - my Mum, my Dad & my Aunty - thank you for all your support over the years in everything I've done. We may be few in number but that just makes us stronger! I would not have made it here without you, of course. Thank you also to Seynosh for being so welcoming and caring - I hope you are as happy being 'adopted' as I am.

Lastly but certainly not least, thank you to all the friends I've made in Southampton. In particular, "The Bests" (you know who you are) for making my time so fun and *entertaining*. It is a great joy to make friends with whom you are able to share almost anything, and I hope these friendships last our lifetimes.

Marcus

Chapter 1

Introduction

1.1 Motivation - The Physics of Neutron Stars

Neutron Stars seem more likely to be an invention of science fiction than a genuine part of our universe. Their existence was first “tentatively” proposed in 1933 by Walter Baade and Fritz Zwicky (Baade and Zwicky, 1934a; Baade and Zwicky, 1934b) and the Nobel prize in physics was awarded for their discovery (Hewish, Bell, et al., 1968) over four decades later.

Almost five decades after that, the merger of two neutron stars was simultaneously detected in gravitational and electromagnetic waves in late 2017 (Abbott, Abbott, et al., 2017b), giving unprecedented information on matter and gravity in extreme cases. Prior to this, gravitational waves with no electromagnetic counterpart were seen from a merging binary black hole (BBH) system for the first time in 2015 (Abbott, Abbott, et al., 2016). And in 2020, the final combination of merging compact objects was seen when GWs from a black hole merging with a neutron star (BHNS) were detected (Abbott, Abbott, et al., 2021). To date, a total of seven signals have been received from mergers involving neutron stars as at least one of the progenitors. These events are currently detectable within a range from Earth of a few hundred megaparsecs, whilst black hole mergers are detectable at distances of up to several gigaparsecs.

Here we will give a brief summary of what is known about neutron stars. This knowledge has been built from both theoretical grounds and observation, and has been constrained considerably by analysis of the recent multi-messenger (combined gravitational wave and electromagnetic) signals. For a more historic reference on compact objects in general, see Shapiro and Teukolsky, 1983, for example.

1.1.1 Formation

Poetically, neutron stars begin their lives following the death of other, massive stars. When the fusion engine of a main-sequence star runs low on its light-elements fuel, it is no longer supported against its own immense gravitational pull by thermal & radiation pressure. Its fate is then determined by its mass. Lighter stars, like our sun, expand in their red giant phase before shrinking again to leave a dense white dwarf, supported primarily by electron degeneracy pressure, a quantum mechanical effect deriving from the Pauli exclusion principle. Stars in this category have masses $\lesssim 8 M_{\odot}$, where M_{\odot} is the solar mass unit equivalent to $\approx 2 \times 10^{30}$ kg, roughly the mass of our Sun.

Heavier stars have a brief (a few millions of years...) red supergiant phase during which they fuse heavier elements until they produce Iron-56, which sits atop the stability ranking of nuclei. After those are burnt through, its core collapses inwards, with a re-bouncing shock that travels back outwards but typically stalls due to in-falling matter from the outer layers of the progenitor. The immense pressure in the core causes protons and electrons to combine, producing both neutrons and a flood of neutrinos, the latter of which may escape the transparent outer layers of the core. They are then captured by the hot, dense matter behind the shock, re-powering it and triggering a supernova explosion that blows off the outer layers, leaving behind a remnant.

If the main-sequence progenitor had a mass in the range $8 - 25 M_{\odot}$, a neutron star is now typically formed, with a mass loosely in the range of $1 - 2 M_{\odot}$ and a radius of $\mathcal{O}(10 \text{ km})$. It is supported in part by neutron degeneracy pressure, and also by repulsive strong nuclear forces, which together prevent further collapse to a black hole - the usual fate of stars with masses $\gtrsim 25 M_{\odot}$.

Binary neutron star systems are of particular interest to us. Unfortunately, only about 5% of known neutron stars exist in binary systems (Tauris and Heuvel, 2006). These are thought to form in two main scenarios. The primary source is from a binary system of main-sequence stars. Once one of the pair has collapsed to form a neutron star, a ‘common envelope’ is formed that encompasses both stars. The neutron star then briefly orbits in its partner’s extended outer atmosphere, with drag causing the radius of the orbit to shrink rapidly until a second supernova explosion of the remaining main-sequence star occurs. If the momentum imparted by the ejecta of this second supernova does not unbind the pair, a binary neutron star (BNS) system is formed. The second, narrower formation channel involves the ‘dynamical capture’ of one isolated neutron star by another in dense stellar regions such as globular clusters (Lee, Ramirez-Ruiz, et al., 2010).

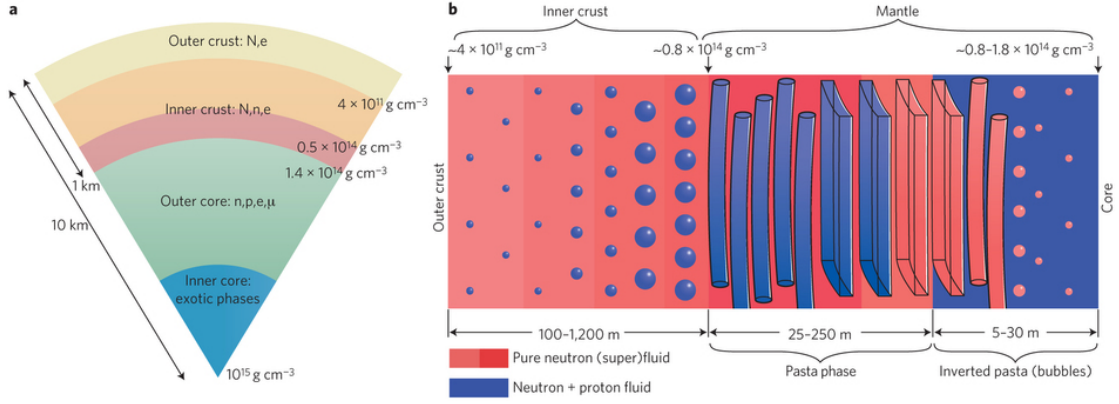


FIGURE 1.1: Neutron star structure. a: Schematic slice through a neutron star. Letters N , n , p , e and μ refer to the presence of nuclei, fluid neutrons and protons, electrons and muons, respectively. The inner core composition is still uncertain and various exotic possibilities exist, including hyperons and deconfined quark matter. b: A detailed picture of the composition of the inner crust. At lower densities, a lattice of superheavy, neutron-rich nuclei is immersed in a fluid of neutrons (which are likely to be superfluid) and a relativistic electron gas. At high enough densities the nuclei might deform and connect along certain directions to form extended tubes, sheets and bubbles of nuclear matter. These nuclear pasta phases might form a layer at the base of the neutron star crust, sometimes referred to as the mantle; searching for observational signatures of such phases is of great interest. Ranges of density and thickness given for each layer represent current uncertainties in the physics of neutron star crusts. This figure was taken from <https://compstar.uni-frankfurt.de/outreach/short-articles/the-nuclear-pasta-phase/>.

1.1.2 Structure

Whilst many properties of a neutron star vary smoothly upon its traversal, one can identify five important regions in a neutron star, as shown in the left-hand image of figure 1.1: the atmosphere, the outer & inner crusts and the outer & inner cores. Let's follow the direction of increasing mystery and work our way inwards. The atmosphere is very thin compared to main-sequence stars (due to the immense gravity) with a thickness of only a few centimeters and a maximum density of about $\rho_{atm} \simeq 1 \text{ g cm}^{-3}$. The surface of a NS is almost perfectly smooth (Abbott, Abbott, et al., 2017b): if a NS were scaled to the size of an apple, deviations in surface height would be less than the width of a hair. Note that this is not to say that NSs are spherical - their shape is distorted to an oblate spheroid by the enormous angular momentum they possess, thanks to the 'ice skater effect'¹. Next, the outer crust is about 10 – 100 m metres thick, composed of nuclei and a degenerate electron gas and reaching a density of $\rho \sim 10^{10} - 10^{11} \text{ g cm}^{-3}$; it is important because its composition controls heat flow into the atmosphere and, hence, the cooling of the neutron star. Cooling is an important topic in neutron star study but one which we will not devote much attention to in this thesis - for a review of cooling and thermal emission, see Page, Lattimer, et al., 2004.

¹The 'ice-skater effect' refers to the spin-up of a rotating body when its moment of inertia is reduced. For an ice-skater, this is done by retracting their limbs; for a neutron star this happens as it collapses, shrinking its radius.

Whilst neutron stars are hot in astrophysical terms, with a surface temperature of around 6×10^5 K, a nuclear physicist would describe them as cold because their matter is below its Fermi temperature. For a further discussion of temperature effects in neutron star mergers see Hammond, Hawke, et al., 2021.

Once the density exceeds $\approx 4 \times 10^{11} \text{ g cm}^{-3}$ (entering the inner crust), the chemical potential of neutrons becomes positive² and some neutrons become unbound from nuclei in a process known as ‘neutron drip’. The mantle, seen in the right-hand side of figure 1.1 (and the red region on the left-hand side), is the transition region between inner crust and outer core, with a density between about a quarter and a half of nuclear saturation density, $\rho_{sat} \simeq 2.7 \times 10^{14} \text{ g cm}^{-3}$. It is here that clusters of unusual, non-spherically shaped nuclear matter are expected to form due to the competing (strong) nuclear force and Coulomb interaction (Ravenhall, Pethick, et al., 1983). A phase transition occurs upon entering the outer core, where nuclei have dissolved almost completely to form a superconducting fluid and the density $\rho \approx \rho_{sat}/2$ whilst the proton fraction $Y_p \approx 0.02 - 0.04$.

It is not known whether another phase transition occurs at the ‘boundary’ between the outer & inner cores. The proton fraction follows an opposite trend to that within the crust, generally increasing with density in the core. Many exotic states of matter have been predicted (see e.g. Kovensky and Schmitt, 2020 and Spinella, 2017 for a review) to exist in the inner core of a neutron star such as hyperons, condensed kaons, or pions in addition to nucleons, or deconfined quark matter in a pure phase or a mixed phase with hadrons.

Quantum chromodynamics (QCD) is the standard model’s gauge theory used to describe the strong-force interactions between quarks mediated by gluons. Based on perturbative quantum chromodynamics (pQCD) calculations, it is expected that matter at very high densities consists of asymptotically free quark matter. However, since these calculations are valid only at densities larger than about $40\rho_{sat}$, and the maximum density in a neutron star is less than about $10\rho_{sat}$, it is not yet clear whether such a phase exists in the interiors of neutron stars. At lower densities, lattice QCD may be used to investigate non-perturbative phenomena such as confinement and quark–gluon plasma formation.

1.1.3 Equation of State

The first pair of questions that one might ask when thinking about any star’s physical properties are: What is its mass? And what is its radius? So it is no surprise that these questions have been continually asked about neutron stars since their inception.

²Strictly, in the relativistic convention, it exceeds the neutron mass: $\mu_n > m_n$.

As early as 1939, the exact, spherically-symmetric, general-relativistic equations of hydrostatic equilibrium were derived (Oppenheimer and Volkoff, 1939)

$$\frac{dp}{dr} = -\frac{G}{r^2} \left(\rho + \frac{p}{c^2} \right) \left(m(r) + 4\pi r^3 \frac{p}{c^2} \right) \left(1 - \frac{Gm(r)}{c^2 r} \right)^{-1} \quad (1.1)$$

where $m(r)$ is the total mass enclosed inside radius r . When paired with an equation of state (EoS) relating the pressure p to the density ρ , the Tolman-Oppenheimer-Volkov (TOV) equation (1.1) completely determines the structure of a spherically symmetric body of isotropic material in equilibrium.

Unsurprisingly then, much work both theoretical (e.g. Potekhin, Fantina, et al., 2013; Gandolfi, Carlson, et al., 2014; Raduta, Nacu, et al., 2021; Burgio, Schulze, et al., 2021; Kovensky, Poole, et al., 2022) and observational (e.g. Lattimer, 2004; Lattimer, 2012; Özel and Freire, 2016; Hanauske, Steinheimer, et al., 2019; Baiotti, 2019; Lattimer, 2021) has gone into providing increasingly reliable equations of state for the dense nuclear matter found in neutron stars. These fall into a number of classes, with some being purely hadronic and hence the only particles present are nucleons (protons, neutrons) and leptons (electrons, muons) with the interactions described by relativistic mean-field theory, or chiral effective field theory. Others include hyperons as additional degrees of freedom, which typically softens the equation of state and lowers the maximum neutron star mass. At ultra-high densities, there may be a transition to a deconfined quark phase that can be described as a color-superconducting quark-gluon plasma. If these exotic states exist, they will alter the transport properties of the neutron star and may even be detectable in gravitational wave signals.

The maximum mass of neutron stars lies somewhere in the range of $2.0 - 2.5 M_\odot$ with a radius at that mass of $10 - 14$ km. Often, the radius of a neutron star with a ‘typical’ mass of $1.4 M_\odot$ is quoted for comparison and this generally falls in the range $11 - 12.5$ km, depending on the EoS. The stiffness (pressure: density gradient) of a given nuclear EoS affects the predicted mass–radius relationship of the neutron star. Therefore, terrestrial experiments aimed at constraining the variation of the pressure with respect to other parameters of nuclear matter are being carried out (Baldo and Burgio, 2016; Mammei, Horowitz, et al., 2024).

Some equations of state include finite-temperature (non-zero) effects. These are particularly relevant when describing the hot nuclear matter found in proto-neutron stars formed in core-collapse scenarios, or for remnant neutron stars left by merger events. For these, the pressure is often given as a function of temperature, T , baryon number density, n_B , and electron fraction, Y_e . In this thesis, investigating equation-of-state effects is not our focus. We make use of a simple, two-parameter EoS such that there is a function $F(\rho, p, n) = 0$ which defines the relationship between the (equilibrium) energy density, pressure and baryon number density at any given point in the fluid.

For drawing quantitative inference about the EoS from simulations, one may turn to [CompOSE](#), the CompStar Online Supernovae Equations of State, which is “a database of equations of state with detailed information on the thermodynamic, compositional and microscopic properties of dense matter that can be used in astrophysical simulations” (Typel, 2021). From CompOSE, we may obtain EoSs, both in tabulated and fitted form. This would enable us in the future to perform quantitative assessment of varying the equation-of-state.

1.1.4 Magnetic Fields

Neutron stars are endowed with incredibly strong magnetic fields, the strongest of which are known as ‘magnetars’ and have field strengths of 10^{14} - 10^{15} G (Baiotti and Rezzolla, 2017). This class of neutron star likely represents at least 10% of the total population of young neutron stars and various theoretical explanations for how these immense fields develop have been presented. Amplification of the progenitor star’s magnetic field through magnetic compression alone during its collapse is not sufficient to explain how it grows so large. Instead, other non-linear growth mechanisms must be present. For the central compact object, these typically involve dynamo effects - see Raynaud, Guilet, et al., 2020, Thompson and Duncan, 1993, and Barrère, Guilet, et al., 2022. Similarly, non-linear magnetic field amplification may occur in the matter surrounding the proto-neutron star through the action of magnetohydrodynamic instabilities such as the Kelvin-Helmholtz instability (KHI) and the magneto-rotational instability (MRI) (Reboul-Salze, Guilet, et al., 2021; Miravet-Tenés and Pessah, 2024).

Magnetars exhibit a plethora of phenomena related to their fields. Unsurprisingly, they emit strongly in the electromagnetic spectrum, and were initially classified as soft gamma repeaters or anomalous X-ray pulsars due to the regular, periodic nature of our detections of them. Magnetars that display this ‘lighthouse’ effect³ are known as pulsars and pulse with time periods ranging from seconds to milliseconds. Emissions are often coupled with interesting timing behavior including enhanced spin-down, glitches, anti-glitches and glitch/anti-glitch candidates. Only around 30 magnetars have been detected as yet but their study has lead to over 1000 publications due to the exotic signals we receive from them and their importance in understanding the isolated neutron star population. See, for instance, Kaspi and Beloborodov, 2017 for a review of magnetars.

A neutron star’s field decays throughout its lifetime such that those of older, merging neutron stars are commonly in the range 10^8 - 10^{10} G. Still, electromagnetism is expected to play a powerful role in the dynamics of neutron star mergers. Similarly

³A spinning magnetar whose magnetic field axis is misaligned with its spin axis will emit a constant flux of electromagnetic radiation in a beam. This beam will rotate with the magnetar, causing us to receive regular pulses of this radiation whenever it is incident on the Earth, creating a ‘lighthouse’ effect.

instabilities and dynamo effects (as in proto-neutron stars) can create very strong magnetic field in the merger's neutron star remnant (Shibata, Fujibayashi, et al., 2021; Aguilera-Miret, Viganò, et al., 2021; Most, 2023; Kiuchi, Reboul-Salze, et al., 2024). These fields can power electromagnetic transients signals such as short gamma-ray bursts (Rezzolla, Giacomazzo, et al., 2011). Much work has been done on simulating magnetic fields in neutron star mergers - for a modern review, see Cioffi, 2020.

1.1.5 Transport

Transport describes how conserved quantities such as energy, momentum, particle number, or electric charge are transferred from one region to another. Such a transfer occurs if the system is out of equilibrium, for instance through a temperature gradient or a non-uniform chemical composition. Different theoretical methods are used to understand transport, depending on how far the system is away from its equilibrium state. If the system is close to equilibrium locally and perturbations are on large scales in space and time, hydrodynamics is a powerful technique. When considering transport due to microphysical particle interactions, kinetic theory is most useful, and can be used to provide forms for the transport coefficients included in the hydrodynamic equations to model dissipation.

In this thesis, we will be particularly concerned with transport properties because it is their addition to, and modification of, ideal hydrodynamics that we will be modelling numerically. Whilst we do not directly model the underlying microphysics that gives rise to these non-ideal transport properties (it occurs on both timescales and lengthscales below those we can resolve in simulations), we take results and expressions from fundamental descriptions of transport in neutron stars to constrain our simulation parameters to sensible values.

Theoretical work (Chugunov and Yakovlev, 2005; Manuel and Tolos, 2011; Potekhin, Pons, et al., 2015; Schmitt and Shternin, 2018) and numerical investigations (Hammond, Hawke, et al., 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 reactions⁴ operate at an atomic scale and may give rise to an effective bulk viscosity at the fluid scale (Celora, Hawke, et al., 2022) that quantitatively affects the gravitational wave signal we obtain from the merger and its remnant's ringdown (Alford, Bovard, et al., 2018; Most, Harris, et al., 2021; Most, Haber, et al., 2022; Hammond, Hawke, et al., 2023). Similarly, work has been done to investigate the possible effects of both shear viscosity (Duez, Liu, et al., 2004) and heat transport (Alford, Bovard, et al., 2018) in binary neutron star mergers, particularly for modulating the turbulence that ensues post-merger, both in the remnant itself, and its

⁴These are weak nuclear reactions that alter the proton fraction to restore beta-equilibrium in the matter.

associated accretion disk. Viscous braking redistributes momentum in a differentially-rotating remnant, removing centrifugal support which can aid in the collapse of the core into a black hole. It is also able to provide thermal support from viscous heating, negating this effect. This viscous heating can enhance mass ejection, particularly in the outer regions of the disk.

The strength and timescale of these dissipative effects depends strongly on the thermodynamic conditions of the matter they act in. At low temperatures, direct Urca processes are slow and their timescales exceed the millisecond ones that oscillations and gravitational waves act on ($\tau_{Urca} \gg \tau_{GW}, \tau_{Hydro}$). Additionally, when the density of the nuclear matter is below the direct-Urca threshold, kinematic constraints on the particles' Fermi momenta prevents these reactions from taking place. Instead, modified-Urca reactions with a bystander nucleon, N , occur:

$$p + e^- + N \leftrightarrow n + \nu_e + N \quad (1.2)$$

which are far less efficient at restoring beta-equilibrium. The chemical composition of the matter is effectively ‘frozen-in’ during the merger and ringdown such that viscous damping has little effect.

On the other hand, when the nuclear matter is hot and above this density threshold, direct-Urca reactions happen fast compared to the other timescales of the system ($\tau_{Urca} \ll \tau_{GW}, \tau_{Hydro}$). This leads again to a weak viscous contribution to the merger dynamics. However, in the intermediate regime, where $\tau_{Urca} \sim \tau_{GW}, \tau_{Hydro}$, a resonance between the oscillations and reactions leads to the greatest impact of viscous damping. Even here, there is only a minor weakening of the gravitational wave signal, and a slight shift in its phase and frequency.

On the millisecond timescale, strong interactions ensure that the neutrons, protons, and electrons are always in thermal equilibrium, described by Fermi-Dirac momentum distributions. However, neutrino interactions can also contribute to viscosity, especially in the hotter, denser regions near the core, but they generally remain subdominant.

In the disk formed in the merger around the remnant, turbulent viscosity, primarily driven by the magnetorotational instability (MRI), is expected to dominate. This mechanism also operates with characteristic timescales of approximately 10 to 100 milliseconds, during which it efficiently redistributes angular momentum. The strength of this viscous effect is often quantified using an effective α -viscosity parameter, with typical values estimated in the range $\alpha \sim 0.01$ to 0.1 , depending on factors such as magnetic field strength and numerical resolution in simulations. These viscous processes play a key role in driving accretion onto the central remnant and launching neutrino-cooled winds that contribute to observable electromagnetic counterparts like kilonovae. Overall, viscous transport significantly affects both the dynamical and thermal evolution of the post-merger system.

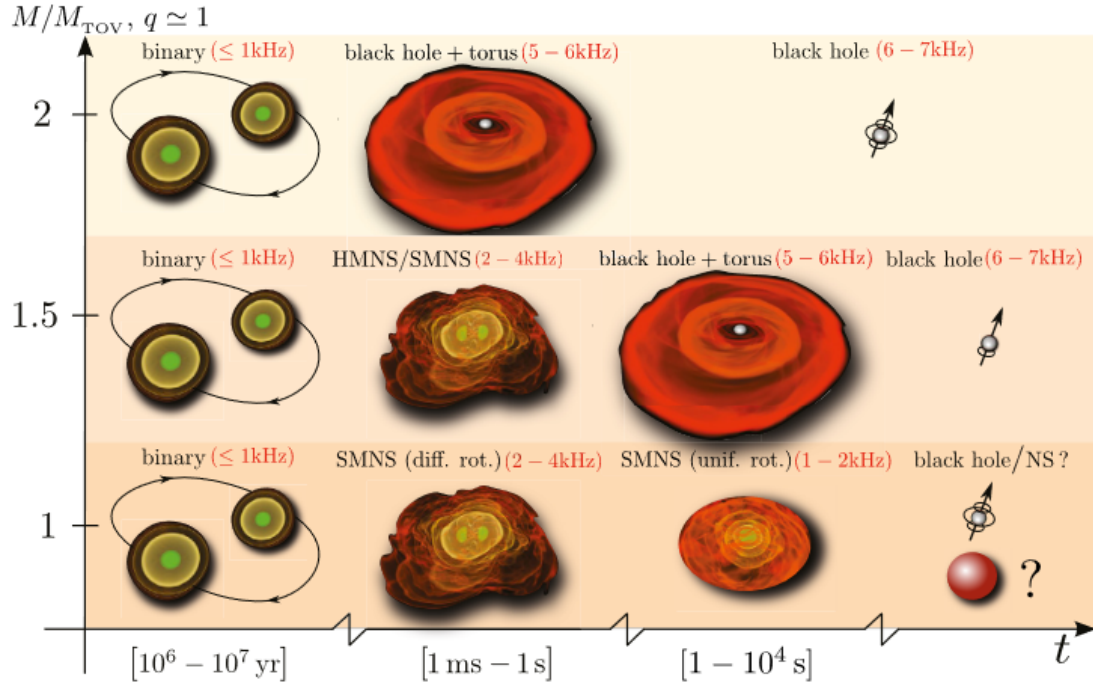


FIGURE 1.2: A schematic diagram illustrating the various stages in the evolution of an equal-mass binary system of neutron stars as a function of the initial mass of the binary, M , and on how it relates to the maximum mass of a non-rotating neutron star, M_{TOV} , the binary can either collapse promptly to a black hole surrounded by a torus (top row), or give rise to a hypermassive (or supermassive) neutron star that ultimately collapses to a black hole and torus (middle row), or even lead to a supramassive neutron star (first differentially and subsequently uniformly rotating) neutron star that eventually yields a black hole or a non-rotating neutron star (bottom row). Also indicated in red are the typical frequencies at which gravitational waves are expected to be emitted. This figure was taken from (Baiotti and Rezzolla, 2017).

1.1.6 Neutron Star Merger Events

1.1.6.1 In Nature

Binary neutron star (BNS) mergers are currently a rare treat for us to detect, with only two confirmed detections to date (Abbott, Abbott, et al., 2017c; Abbott, Abbott, et al., 2020) and a (rather unconstrained) predicted local occurrence rate of $250 - 2810 \text{ Gpc}^{-3} \text{ yr}^{-1}$. In the first of these events, GW170817, the component masses are constrained to lie in the range $1.16 - 1.60 M$ with a total mass for the system of about $2.74M$ (Abbott, Abbott, et al., 2017c). Constraints on the radii, spins and tidal deformability of the two neutron stars were also able to be placed (De, Finstad, et al., 2018; Abbott, Abbott, et al., 2018), ruling out several equation of state models with high confidence. The remnant left by the merger was likely a hypermassive neutron star with a mass exceeding the maximum stable mass of a static neutron star, as predicted by most equations of state. There is some debate over whether the collapse to a black hole was prompt (Abbott, Abbott, et al., 2019) or occurred more than a

second after merger (Murguia-Berthier, Ramirez-Ruiz, et al., 2021). There was an associated detection of a γ -ray burst (GRB 170817A) 1.7 s after merger by the Fermi-GMB satellite (Goldstein, Veres, et al., 2017), aiding in sky-localization of the source and providing the first direct evidence of a link between these mergers and short γ -ray bursts. Subsequent identification of transient counterparts across the electromagnetic spectrum in the same location further supports the interpretation of this event as a neutron star merger (Abbott, Abbott, et al., 2017c).

A summary of the stages of a BNS merger is presented in figure 1.2 for a pair of equal mass binaries, taken from Baiotti and Rezzolla, 2017. This equal-mass assumption is fairly robust as the difference in the masses of NSs in binary pairs are typically $\approx 10\%$ and are at most $\approx 30\%$ (Baiotti and Rezzolla, 2017). Gravitational radiation produced by the time-varying mass-quadrupole moment will shrink the initial orbit of a bound pair of compact objects such that a merger will occur within the Hubble time (roughly the age of the universe) if the initial orbital period $P_{orbit} \lesssim 1$ d (Rosswog, 2015). At the point of contact, the stars typically have an orbital period of milliseconds, radiating $\approx 10^{53}$ erg s $^{-1}$ in the form of gravitational waves (Faber and Rasio, 2012). This radiation, detected already from two mergers, GW170817 (Abbott, Abbott, et al., 2017b) and GW190415 (Abbott, Abbott, et al., 2020), travelled to Earth at the speed of light with a precision exceeding 10^{-14} . This observation alone ruled out many alternative theories of gravity and provided Einstein’s theory of general relativity with yet more evidence.

The multi-messenger nature of some BNS mergers allows them to be used as ‘standard sirens’, meaning that the absolute distance to the source can be determined directly from the gravitational-wave measurements and used as a distance calibration tool for analysing the electromagnetic signal received alongside them. This then gives a third, independent measurement of the Hubble constant which parametrizes the local, current rate of expansion of our universe (Abbott, Abbott, et al., 2017a). One hopes that this will (eventually) resolve some of the ‘tension’ that has arisen as the uncertainties on the two other measures (from calibrated cosmic distance ladder techniques and cosmic microwave background (CMB) data) have shrunk to the point that they are now in disagreement (Di Valentino, Mena, et al., 2021).

Following the merger, a small delay, $\mathcal{O}(\text{ms} - \text{s})$, is expected to precede a short-duration gamma-ray burst (sGRB): an emission type of particular interest (see, for example, Lee, Ramirez-Ruiz, et al., 2010; Rezzolla, Giacomazzo, et al., 2011). The exact mechanism powering the sGRB is still an area of active research. It may be that, in the prompt collapse scenario, the turbulent, hyper-accreting torus surrounding the black hole extracts energy from it via the Blandford-Znajek process, leading to a relativistic outflow collimated by disk winds and a highly poloidal magnetic field (Janiuk_black_2017 ; Janiuk and Yuan, 2010). If instead a magnetar is left by the merger, longer-lived emission including an X-ray afterglow powered by pulsar-like

spin-down and large-scale magnetic reconnection events will be seen (Metzger, Giannios, et al., 2011; Lü, Zhang, et al., 2015; Bucciantini, Metzger, et al., 2012).

There is also now strong evidence that BNS mergers are responsible for the production of much of the heavy elements (heavier than Iron-56) in our universe. This occurs in the hot ejecta where rapid neutron capture, referred to as the *r*-process, should take place. It is even possible to infer the masses and relative abundances (Freiburghaus, Rosswog, et al., 1999) of nuclear elements synthesized in this way and, hence, the nuclear make-up of the universe as whole. There was also observation of an optical transient following the merger seen in 2017: a *kilonova* spanning the UV, optical & IR parts of the EM spectrum, powered by the radioactive decay of by-products created in the aforementioned *r*-process.

1.1.6.2 In Silico

There are still open questions, of course, in the study of BNS mergers. Numerical modeling of the merger process via general relativistic magnetohydrodynamics (GRMHD) simulations offers the best chance to tackle the open questions and to establish a reliable connection between the merger & post-merger dynamics and the observable GW and EM emission. For this, see Cioffi and Kalinani, 2020, Rezzolla, Giacomazzo, et al., 2011, Aguilera-Miret, Viganò, et al., 2021. In addition, experimental work is being done to establish, more precisely, the neutron star equation of state using simulations (Prakash, Radice, et al., 2021).

Already, people are looking ahead to the next generation of ground-based gravitational wave detectors such as LIGO-voyager (Berti, Cardoso, et al., 2022), NEMO (Ackley, Adya, et al., 2020), Cosmic Explorer (Reitze, Adhikari, et al., 2019) and the Einstein Telescope (Punturo, Abernathy, et al., 2010; Maggiore, Broeck, et al., 2020). These instruments will bring an order of magnitude improvement in gravitational wave sensitivity, leading to a significant increase in the merger detection rate. Improved sensitivity at higher (kHz) frequencies should also allow detection of the post-merger signal. It is in this phase that the subtleties of modelling the remnant's behaviour become important for discerning the precise physics involved. With more precise data comes the need for more accurate numerical simulations to draw meaningful inference from the observations. In particular, the post-merger gravitational-wave emission is sensitive to the small-scale dynamics of nuclear reactions (Alford, Bovard, et al., 2018; Most, Haber, et al., 2022; Hammond, Hawke, et al., 2023; Chabanov and Rezzolla, 2023), electromagnetism (Aguilera-Miret, Palenzuela, et al., 2023; Most, 2023; Kiuchi, Reboul-Salze, et al., 2024) and turbulence (Radice, 2020; Palenzuela, Aguilera-Miret, et al., 2022).




FIGURE 1.3: A visual summary of the (both current and proposed) sources and detectors of gravitational waves across the spectrum of frequencies/wave-periods. Taken from the [LISA website](#).

Meanwhile, the [Laser Interferometer Space Antenna \(LISA\)](#) mission is a proposed space-based gravitational-wave detector with very low noise and very long (5×10^6 km) interferometer arms, giving it the ability to detect low-frequency (0.1 mHz – 1 Hz) gravitational waves from other sources entirely. In particular, it will allow detection of gravitational waves from the capture of compact objects by supermassive black holes at the centres of galaxies, from the merger of these supermassive black holes themselves, and from ultra-compact and eccentric binaries (Cornelisse, 1996; Danzmann, 2000; Danzmann and R Diger, 2003; Bender, Begelman, et al., 2013). It will even search for seed black holes, found at the cosmic dawn (Amaro Seoane, Arca Sedda, et al., 2022). Figure 1.3 summarises the (both current and proposed) sources and detectors of gravitational waves across the spectrum of frequencies.

The disparate array of dynamical time and length scales present in binary neutron star mergers presents a major challenge for any numerical code looking to directly resolve them all. This is, in fact, an impossibility for now and the foreseeable future. To illustrate this quantitatively, consider that the highest simulation resolutions achieved so far have grids with side-length $\mathcal{O}(10\text{ m})$ (Kiuchi, Cerdá-Durán, et al., 2015) whilst the molecular dissipation scale in the neutron star fluid is either $\mathcal{O}(1\text{ cm})$ when set by neutrinos, or $\mathcal{O}(1\text{ nm})$ when set by electron scattering (Thompson and Duncan, 1993). Alternatively, consider that the Reynolds number of the flows in neutron stars may range as widely as 10^6 to 10^{15} during a merger event (Radice and Hawke, 2024). Nuclear reactions, turbulence and neutrino processes all lead to variations in fluid properties below the resolved length-scale. This violates the fundamental assumption of fluid dynamics: uniformity within a fluid element. Whether we view this discrepancy as stochastic fluctuations or as genuine microphysics may affect our motivations and approach to the problem, but in any case it needs attention.

Turbulence, seeded by Kelvin-Helmholtz instabilities that grow in the shearing layer as




FIGURE 1.4: Example of a BNS merger simulation using General Relativistic Magnetohydrodynamics (from the models presented in Cioffi, Kastaun, et al., 2017). The temporal sequence shows the bulk of the NS(s) (in white together) with color-coded isodensity surfaces. Deformation of the neutron stars through tidal forces can be seen pre-merger. Post-merger, the remnant neutron star is also deformed by the uneven distribution of angular momentum within it.

the two stars come into contact, causes a cascade of kinetic energy from the largest length scales down to the smallest, at which point viscous dissipation converts it to heat. Unfortunately, the non-linear coupling of scales in turbulent flows means that uncontrolled small-scale effects may spoil the large-scale behaviour. This issue is perhaps epitomized by the non-convergence of the maximum magnetic field strength seen in magnetohydrodynamic (MHD) simulations as resolution is increased. As a solution to this problem, large eddy schemes (LES) which include subgrid closures have been introduced as extensions to the usual general-relativistic MHD prescription used in simulations. These general relativistic LES (GRLES) models were able to solve this particular issue of magnetic field amplification and led to convergence of its saturated strength. Following this, LES schemes are now seeing more widespread usage (Radice, 2017; Radice, 2020; Viganò, Aguilera-Miret, et al., 2020a; Carrasco, Viganò, et al., 2020; Palenzuela, Liebling, et al., 2022; Wang, Yuan, et al., 2022; Izquierdo, Bezares, et al., 2024; Aguilera-Miret, Viganò, et al., 2020; Miravet-Tenés, Cerdá-Durán, et al., 2022; Miravet-Tenés, Cerdá-Durán, et al., 2023). For an extensive discussion of the current status of this kind of modelling, see Radice and Hawke, 2024.

1.2 Summary and Thesis Structure

In this chapter we have introduced neutron stars, their key features and their mergers with each other. The extreme physics they exhibit clearly warrants study, and mergers involving neutron stars are our best way of doing this. The rest of this thesis will focus on improving the current state of numerical modelling of NS mergers, by focusing in particular on the merits of using non-ideal, dissipative fluid models instead of the simpler, ideal ones that are often currently employed. The motivation for this is twofold: to simulate genuine dissipate effects; and to use a dissipative fluid description as a subgrid closure model. The former application is particularly important in

understanding the impact of non-ideal effects on the products of mergers: gravitational waves, electromagnetic signals and mass ejecta. The latter is crucial for modelling the turbulence and magnetic field evolution of a neutron star remnant.

In chapter 2 we will lay out the theoretical understanding of general relativity and relativistic hydrodynamics required for describing neutron stars and their strong-gravity environments. In chapter 3 we will introduce the foundations of the computational methods used to solve the equations provided by the aforementioned theories. In chapters 4 and 5 we will present results from simulations we have performed of relativistic, dissipative hydrodynamics using the well-established Müller-Israel-Stewart model and our novel ‘MISCE’ model. In chapter 6 we will introduce our covariant filtering codebase and apply it to turbulence to derive a subgrid closure model based on a dissipative fluid description. Finally, chapter 7 covers preliminary results from outstanding work using simulations of dissipative fluids. Specifically, results from simulating the newly-popularised BDNK model, and from performing large-eddy simulations using the subgrid closure model derived in chapter 6. Chapter 8 makes concluding remarks before discussing the outlook for using dissipative hydrodynamics in neutron star merger simulations.

Chapter 2

Theoretical Foundation

2.1 General Relativity

Infamous thought experiments involving elevators and free-fall, as well as a great deal of mathematical machinery drawn from the field of differential geometry, led Einstein to a remarkably simple-looking set of equations that govern the evolution of spacetime and the matter that resides in it:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu}. \quad (2.1)$$

Here, $G_{\mu\nu}$ is known as the Einstein tensor and describes the geometry of spacetime whilst $T_{\mu\nu}$ is the stress-energy tensor describing the distribution of mass and energy within the spacetime. In addition, $R_{\mu\nu}$ is the Ricci tensor whilst R is the Ricci scalar, both descriptions of the curvature of spacetime. The metric tensor $g_{\mu\nu}$ defines proper distances between points on the 4-dimensional manifold describing spacetime, \mathcal{M} . G is Newton's gravitational constant and c is the speed of light. Note that we are working in a (3+1) spacetime convention where Greek indices run over four values (0, 1, 2, 3) and Latin indices run over the three spatial values (1, 2, 3) only. We use Einstein summation convention such that repeated indices must appear once 'up' (contravariant) and once 'down' (covariant) and represent a sum over their possible values. Our spacetime metric, $g_{\mu\nu}$, has a signature of $(-+++)$. This set of ten coupled differential equations represent the mathematical formulation of Einstein's theory of General Relativity (GR).

We know that general relativity is not the 'final' theory of gravity. It is not quantum in nature and admits non-physical singularities where the curvature of spacetime becomes infinite. There remain many outstanding questions regarding the nature of dark matter, dark energy and their relation to gravity. Whether an alternative theory would provide as yet elusive answers to these questions is still to be seen. However, GR

has proved remarkably powerful over the last century. It has withstood stringent tests within our solar-system (Turyshev, 2008) and galaxy (Liu, Li, et al., 2022). Its predictive power has led to the advent of technologies such as the Global Positioning System (GPS). So far, its predictions match observations of more extreme, distant systems such as pulsars and black holes (Kramer, Backer, et al., 2004; Kramer, Stairs, et al., 2006).

Neutron stars are dense enough to create a strong-field regime of gravity that will distort spacetime sufficiently such that one must couple their description of neutron star matter to GR. One way that this can be seen theoretically is by considering that the Schwarzschild radius of a neutron star, $r_s = 2GM/c^2 \sim a \text{ few km}$, is of the same order of magnitude as its physical radius ($\sim 10 \text{ km}$). Hence, realistic neutron star merger simulations require evolution of the background spacetime as well as the matter itself. This is not something we have done in our simulations as yet (our formulations are special relativistic but not general). Instead, we have focused our study on models of fluid dynamics that describe the neutron star matter itself. Extending our simulations to GR is one possible avenue of further work that will be discussed in more detail in chapter 8. However, we will now give a brief introduction to the field of numerical relativity, which is used to computationally model systems exhibiting strong gravity.

2.1.1 Numerical Relativity Formulation

Equation 2.1 describes the interaction between matter and spacetime. It is manifestly covariant and hence space and time are placed on an equal footing. Whilst this is logical from a theoretical perspective, it does not immediately lend itself well to numerical evolution in this form, where space and time are traditionally decoupled. Instead, alternative but equivalent ways of formulating Einstein’s equations have been developed over the last 70 or so years—ones that retain covariance and can naturally describe non-Cartesian coordinate systems, but which are more computationally tractable. These are typically based on the pioneering work by Arnowitt, Deser and Misner (Arnowitt, Deser, et al., 1959) which led to the ADM formalism of general relativity. Popular modern extensions such as that of Baumgarte, Shapiro, Shibata and Nakamura (BSSN) (Shibata and Nakamura, 1995; Baumgarte and Shapiro, 1998) are based on this work.

Their approach was to create a “(3+1)” split of spacetime whereby space and time are decoupled. This approach relies on slicing the four-dimensional spacetime into a series of three-dimensional, spacelike hypersurfaces so that the Lorentzian spacetime metric (signature $(-+++)$) induces a Riemannian metric on them (signature $(+,+,+)$). This allows one to formulate Einstein’s equations into first-order, balance-law forms such that finding their solution represents an initial value problem with additional, elliptic




FIGURE 2.1: Schematic representation of the 3+1 decomposition of spacetime with hypersurfaces of constant time coordinate Σ_t and Σ_{t+dt} foliating the spacetime. The four-vector \mathbf{t} represents the direction of evolution of the time coordinate t and can be split into a timelike component $\alpha\mathbf{n}$, where \mathbf{n} is a time-like unit normal to the hypersurface, and into a spacelike component, represented by the spacelike four-vector $\boldsymbol{\beta}$. Taken from Bishop and Rezzolla, 2016

constraint equations that must remain satisfied throughout the evolution of the system. Figure 2.1 shows a schematic representation of the formalism. The hypersurfaces Σ are separated by time dt . In that time, the point $x^i(t)$ has been transported along a vector $\mathbf{t} = \alpha\mathbf{n} + \boldsymbol{\beta}$. Here, α is the lapse scalar which measures the proper time between successive hypersurfaces, \mathbf{n} is the time-like normal to the hypersurface, and $\boldsymbol{\beta}$ is the space-like shift 4-vector whose spatial components measure the change in coordinates from one hypersurface to the next. The spacetime distance ds between two points is therefore given by

$$ds^2 = -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt) \quad (2.2)$$

where γ_{ij} are the spatial components of the induced metric $\gamma_{\mu\nu} = g_{\mu\nu} + n_\mu n_\nu$ on each hypersurface and the covariant & contravariant components of the normal vector are given respectively by

$$n_\mu = (-\alpha, 0, 0, 0); \quad n^\mu = \frac{1}{\alpha}(1, \beta^i). \quad (2.3)$$

Finally, we may relate proper time $d\tau$ intervals to spacetime separation ds and therefore to coordinate timesteps dt by

$$d\tau^2 = -ds^2 = \alpha^2 dt^2. \quad (2.4)$$

Whilst numerical relativity is necessary for performing binary neutron star merger simulations, it is not the focus of this thesis. Instead, we choose to focus on hydrodynamics, which describes the matter in neutron stars. As such, our results come from evolving a fluid in flat, stationary spacetime for which we do not need to employ

numerical relativity.

2.2 Fluid Dynamics

Fluid dynamics is the branch of physics that originally sought to understand and predict the behaviour of liquids, gases and plasmas in motion. Early examples of fluid dynamics include the well-known principles of Archimedes¹ and Bernoulli². These principles explain, for instance, how an enormous aircraft carrier, though appearing precariously balanced, remains buoyant and stable in water; how an aircraft wing creates lift and a sail generates forward-thrust; or why your shower curtain is sucked inwards when you turn on the water.

The modern view of fluid dynamics is more general: it is the long-wavelength, slow-acting limit of an underlying microphysical theory. This generality allows fluid dynamics to be used widely in describing, for instance, terrestrial atmospheric & ocean modelling, astrophysical compact objects & their accretion disks, and even high-energy heavy-ion collision products at particle colliders (Del Zanna, Chandra, et al., 2013).

Their success lies in their ability to obfuscate the complex and ostensibly insignificant details of the underlying physics occurring at scales below those of the coarse fluid ‘elements’ that comprise the fluid. The whole of the fluid within an element is then described by physical & thermodynamic properties (such as velocity, pressure and energy density) which together form scalar and vector fields over the fluid’s domain. An inherent assumption, then, is that one may assign the fluid’s state variables each a single value for all fluid within one element. This is an assumption that is not expected to always hold true for current neutron star merger simulations because the physical size of the domains used is of order metres to tens of metres and, hence, the fluid elements contain significant velocity, temperature and compositional gradients. Whilst finer-scale simulations are currently out of reach due to their high computational cost, Large Eddy Simulations (LES) are a proposed method of capturing subgrid scale physics, one which we will discuss in chapter 6.

In the remainder of this chapter, we will lay out the theoretical framework currently used to describe (special) relativistic fluids. These are those whose velocity is an appreciable fraction of the speed of light, such that their kinetic energy may be comparable to or exceed their rest-mass energy. We will start with a brief introduction using the typical non-relativistic, dissipative description before moving to the ideal, relativistic case and finally covering a number of models of non-ideal, relativistic fluids

¹Archimedes’ principle states that the upward buoyant force that is exerted on a body immersed in a fluid, whether fully or partially, is equal to the weight of the fluid that the body displaces.

²Bernoulli’s principle states that an increase in the speed of a fluid occurs simultaneously with a decrease in static pressure or a decrease in the fluid’s potential energy.

which are those most relevant for us. We will primarily focus on gradient expansion models of hydrodynamics, which are by far the most popular in our field and those of primary use to us here. However, we will also briefly touch upon variational models and smoothed particle hydrodynamics, which give alternative ways of deriving a fluid description. The particular guiding literature for this chapter is Romatschke, 2010 and Rocha, Wagner, et al., 2024, which provide a comprehensive overview of the field. For an excellent summary of special relativistic hydrodynamic theory for the purpose of modeling heavy-ion collisions, see Jaiswal and Roy, 2016.

2.2.1 Gradient Expansion Models

The first and simplest models, known as ‘zero’th-order’, describe fluids in local thermodynamic equilibrium. Differences in thermodynamic properties between neighbouring fluid elements equilibrate on timescales that are effectively instantaneous compared to those that the fluid acts on, $\tau_{diss} \ll \tau_{fluid} \sim L_{fluid}/v_{fluid}$. Although appropriate for some simple systems (such as stable, plane flows), higher-order gradient expansions in these quantities, which represent a departure from equilibrium, are necessarily included to capture more complex behaviour such as that seen in heavy-ion collisions, black hole ejecta and certainly neutron star mergers. First-order theories add dissipative effects expressed as first-order gradients in the ‘perfect’ (ideal) fluid properties. These were first formulated by Eckart (Eckart, 1940) and Landau & Lifshitz (Landau and Lifshitz, 1987) by making simple extensions to the non-relativistic Navier-Stokes (NS) equations, which result from applying Newton’s second law to fluids. Initially, first-order theories seemed promising until Hiscock and Lindblom (Hiscock and Lindblom, 1985) showed that they admitted instabilities that would allow runaway from equilibrium, as well as giving acausal behaviour due to their parabolic nature. These pathologies were wrongly attributed to the first-order nature of these theories, leading to the development of a number of second-order theories, for instance by Baier & colleagues (Baier, Romatschke, et al., 2008) and Carter (Carter, 1985; Carter, 1988).

Today, we know that some first-order theories are stable, causal and do not disobey the second law of thermodynamics. This class of models was first introduced in 2018 by Bemfica, Disconzi and Noronha (Bemfica, Disconzi, et al., 2018) and “First-order relativistic hydrodynamics is stable” is the triumphant title of Kovtun’s paper in 2019 (Kovtun, 2019). In section 7.1 of this thesis, we will use one of these first-order formulations of recent popularity, named BDNK after the four aforementioned authors. See also Bemfica, Disconzi, et al., 2019b; Bemfica, Disconzi, et al., 2021; Pandya and Pretorius, 2021 as references. First, let us lay down some fundamental principles, starting with fluid dynamics in the non-relativistic setting.

2.2.1.1 The Non-Relativistic Navier-Stokes Equations

The starting point for fluid dynamics in Newtonian gravity can be obtained through application of Newton's second law to a fluid element, which results in the Navier-Stokes equations for a viscous fluid:

$$\partial_t v^i + v^k \partial_k v^i = \rho^{-1} (\partial_i p + \partial_k \Pi^{ki}) \quad (2.5a)$$

$$\Pi^{ki} = -\eta(\partial^k v^i + \partial^i v^k - \frac{2}{3}\delta^{ki}\partial_l v^l) - \zeta\delta^{ik}\partial_l v^l. \quad (2.5b)$$

Here, $v^i(t, \mathbf{x})$ is the spatial velocity of the fluid, $\rho(t, \mathbf{x})$ is its mass-density, $p(t, \mathbf{x})$ is the pressure, and Π^{ki} is the viscous stress tensor which contains contributions from both shear and bulk viscosity and accounts for non-ideal behaviour. The strength of the shear and bulk viscosities are modulated via their coefficients, η & ζ . Additionally, the conservation of mass is enforced through the continuity equation

$$\partial_t \rho + \partial_i(\rho v^i) = 0. \quad (2.6)$$

The system is then typically closed with an equation of state which relates the density to the pressure: $p = p(\rho)$. The widespread success of this equation means that any relativistic version we derive should reduce to it in the low-velocity limit.

2.2.1.2 Ideal Fluid Equations in Relativity

In relativity we describe matter through the stress energy tensor $T^{\mu\nu}$. After demanding Lorentz invariance, and putting sensible physical requirements on its components (see Romatschke, 2010), we arrive at the most general form allowed, for an ideal fluid at least, which is

$$T_{(0)}^{\mu\nu} = (\rho + p)u^\mu u^\nu - pg^{\mu\nu} = \rho u^\mu u^\nu - p\Delta^{\mu\nu} \quad (2.7)$$

where ρ has now been promoted to a total energy density, $\Delta^{\mu\nu} \doteq g^{\mu\nu} - u^\mu u^\nu$ is a projection operator into the space orthogonal to the 4-velocity, u^μ , which is defined as

$$u^\mu = \frac{dx^\mu}{d\mathcal{T}} = W(1, \mathbf{v}) \quad (2.8)$$

where $W = (1 - \mathbf{v} \cdot \mathbf{v})^{-1/2}$ is the Lorentz factor and \mathcal{T} is the proper time. For the type of non-conformal fluids that we deal with, the SET is typically paired with a statement about the number current of the particle species that comprises the fluid. This is written as

$$N^\mu = nu^\mu + n^\mu \quad (2.9)$$

where n is the number density of the species and n^μ is the ‘drift-current’ which captures the part of the number current misaligned with the fluid’s 4-velocity. This presents one possible choice of hydrodynamic frame: to define the fluid’s 4-velocity as that parallel to the total number current such that $n^\mu = 0$ and $N^\mu = nu^\mu$. This is known as the Eckart frame.

To obtain the fluid equations of motion, we invoke the fundamental principles of conservation of mass, energy and momentum. These are captured by the following pair of equations:

$$\nabla_\mu N^\mu = 0, \quad (2.10a)$$

$$\nabla_\mu T^{\mu\nu} = 0. \quad (2.10b)$$

Here, ∇_μ is the covariant derivative which contains contributions from two sources related to the object it acts on: intrinsic variations in the object itself across the spacetime; and changes arising due to variations in the spacetime coordinates by which the object is described. The former is known as the continuity equation and enforces the local conservation of the rest mass of particles. The latter equation, in words, says that the covariant derivative of the stress-energy tensor (sometimes, the energy-momentum tensor) vanishes for all matter (non-gravitational) fields. This is a consequence of the stress-energy tensor, $T^{\mu\nu}$, being the conserved Noether current for spacetime translations. For this thesis, we concern ourselves almost entirely with modelling fluids in flat spacetime. Hence, the covariant derivatives become partial ones, $\nabla^\mu \rightarrow \partial^\mu$.

To obtain the equations of motion for the fluid, it is useful to project eq. (2.10) both parallel and orthogonal to u^μ :

$$u_\nu \partial_\mu T^{\mu\nu} = 0 = \Delta^\alpha_\nu \partial_\mu T^{\mu\nu} \quad (2.11)$$

which for an ideal fluid leads to

$$D\rho + (\rho + p)\partial_\mu u^\mu = 0, \quad (2.12a)$$

$$(\rho + p)Du^\mu - D^\mu p = 0. \quad (2.12b)$$

We have defined the derivative operators $D \doteq u^\mu \partial_\mu$ and $D^\nu \doteq \Delta^{\mu\nu} \partial_\mu$ which reduce, in the small-velocity limit $|\mathbf{v}| \ll 1$, to

$$D \simeq \partial_t + v^i \partial_i + \mathcal{O}(|\mathbf{v}|^2); \quad D^i = \Delta^{i\mu} \partial_\mu \simeq \partial^i + \mathcal{O}(|\mathbf{v}|) \quad (2.13)$$

and are hence relativistic versions of the (convective) time and space derivatives.

Together with the assumption that the energy density is dominated by the mass density for a non-relativistic fluid ($\rho \simeq p$) and that a non-relativistic equation of state

obeys $p \ll \rho$, equations 2.13 reduce to the Navier-Stokes equations 2.5 in the non-relativistic limit.

2.2.1.3 Equations of Non-Ideal, Relativistic Fluids

We may extend the relativistic SET to include dissipative effects by writing

$$T^{\mu\nu} = T_{Ideal}^{\mu\nu} + q^{(\mu} u^{\nu)} + \Pi^{\mu\nu} \quad (2.14)$$

where q^μ is the energy-momentum diffusion current relative to u^μ (now possible in relativity) and $\Pi^{\mu\nu}$ is the viscous stress tensor that includes the contributions to $T^{\mu\nu}$ stemming from viscous dissipation. For non-perfect fluids, the two vector & tensor N^μ & $T^{\mu\nu}$ now together contain 14 independent variables, instead of the 5 for perfect fluids.

Another choice of hydrodynamic frame is now possible, one which we will make here. In the Landau frame³, the 4-velocity u^μ is defined as being parallel to the energy-momentum flow, and $q^\mu = 0$. The 3 degrees-of-freedom that this seemingly removes are shifted into the drift term within the rest-mass current such that $N^\mu = nu^\mu + n^\mu$ with $n^\mu \neq 0$.

The equations of motion are then given by

$$D\rho + (\rho + p)\partial_\mu u^\mu - \Pi^{\mu\nu}\nabla_{(\mu}u_{\nu)} = 0, \quad (2.15a)$$

$$(\rho + p)Du^\mu - \nabla^\mu p + \Delta_\nu^\mu \partial_\alpha \Pi^{\alpha\nu} = 0. \quad (2.15b)$$

The total viscous stress tensor is typically split into two contributions - one traceless part, $\pi^{\mu\nu}$, and one with non-vanishing trace, Π :

$$\Pi^{\mu\nu} = \pi^{\mu\nu} + \Delta^{\mu\nu}\Pi. \quad (2.16)$$

Π is the bulk viscosity (a scalar which also contributes to the total isotropic pressure) and $\pi^{\mu\nu}$ is the shear-stress tensor with the following properties: it is symmetric ($\pi^{[\mu\nu]} = 0$), purely spatial ($\pi^{\mu\nu}u_\mu = 0$), and trace-free ($\pi_\mu^\mu \equiv \pi^{\mu\nu}g_{\mu\nu} = 0$).

Essentially, models of dissipative, relativistic fluid dynamics are built by setting the hydrodynamic frame, and prescribing some form for the non-ideal contribution $\Pi^{\mu\nu}$. The latter of these choices is typically made either by providing a constitutive relation or by promoting the additional degrees of freedom to independent variables with their

³Note that the Eckart and Landau frames, whilst ostensibly sensible and physically-motivated choices, are not the only ones one can make. Indeed, the BDNK formalism discussed in section 2.2.1.6 makes use of this degree of freedom to design a stable and causal first-order theory of dissipative hydrodynamics in relativity.

own evolution equations. We will give an example of each of these approaches below, specifically for the MIS and BDNK models that we will go on to simulate in chapter 4.

2.2.1.4 Relativistic Navier-Stokes Equations

The most straightforward starting point for a relativistic description of a dissipative fluid is made by extending the Newtonian Navier-Stokes equations 2.5. Because dissipation is an irreversible, entropy-generating process, we consider the thermodynamics of the fluid. The Gibbs relation describes the free-energy contributions to the system as

$$\rho + p = Ts + \mu n, \quad (2.17a)$$

$$d\rho = Tds + \mu dn, \quad (2.17b)$$

where T is the temperature, s is the entropy density and μ is the chemical potential of the species with number density n . To enforce the second law of thermodynamics, that entropy may never decrease, we stipulate in a covariant way that in local-equilibrium

$$\nabla_\mu s^\mu \geq 0; \quad s^\mu = su^\mu \quad (2.18)$$

where s^μ is the entropy current carried with the fluid. Using equations 2.17 and 2.18, as well as the fluid evolution equations 2.15 we are able to rewrite the second law as

$$\nabla_\mu s^\mu = \frac{1}{2T} \pi^{\mu\nu} \nabla_{<\mu} u_{\nu>} + \frac{1}{T} \Pi \nabla_\mu u^\mu \geq 0 \quad (2.19)$$

where we have defined $\nabla_{<\mu} u_{\nu>} \doteq 2\nabla_{(\mu} u_{\nu)} - \frac{2}{3} \Delta_{\mu\nu} \nabla_\alpha u^\alpha$ as the traceless part of $\nabla_{(\mu} u_{\nu)}$.

The inequality is then guaranteed to be fulfilled if we choose the shear stress tensor and bulk viscous pressure⁴ to be

$$\pi^{\mu\nu} = \eta \nabla^{<\mu} u^{\nu>}, \quad \Pi = \zeta \nabla_\mu u^\mu \quad (2.20)$$

where their coefficients, the bulk & shear viscosities $\zeta, \eta \geq 0$. Such a choice guarantees that s^μ is a sum of positive squares. These forms also arise naturally out of kinetic theory if systematic expansions of the fluid theory (in gradients) are considered. We discuss this briefly in section 2.4.

This description reassuringly reduces to the regular Navier-Stokes equations in the non-relativistic limit, but is unfortunately plagued with acausal and unstable

⁴Note that we have now used both the term ‘bulk viscosity’ and also the term ‘bulk viscous pressure’. Viscosity is a non-ideal, dissipative effect that resists differentials in fluid motion and is not strictly a pressure in the same way that the hydrostatic pressure is, for instance. However, bulk viscosity often acts as an effective pressure due to its resistance to compression or expansion of a fluid, and appears as a pressure contribution in the equations of motion for a viscous fluid as will be seen in chapter 4. Therefore, we often use the terms interchangeably throughout this thesis.

pathologies Hiscock and Lindblom, 1983. The source of these is the fact that we have truncated the entropy current at first order in gradients. meaning the total entropy of these fluids, restricted to the dynamically accessible states, has no upper bound (Gavassino, Antonelli, et al., 2020). This motivated the development of new, higher-order formulations of dissipative hydrodynamics, which we will continue onto now.

2.2.1.5 The Müller-Israel-Stewart Formulation

The Müller-Israel-Stewart formalism was developed in the late 1970s Israel, 1976; Stewart, 1977; Israel and Stewart, 1979 and inspired by the relaxation-type form of the Maxwell-Cattaneo equation (Cattaneo, 1948) usually used for regularizing the acausal signal propagation of the standard heat equations.

Now, the dissipative terms obey dynamical equations of a relaxation-type form with relaxation timescales τ_Π , τ_π and τ_q given by

$$\tau_\Pi \dot{\Pi} + \Pi = \Pi_{NS} + \mathcal{O}_2, \quad (2.21a)$$

$$\tau_\pi \dot{\pi}^{<\mu\nu>} + \pi^{\mu\nu} = \pi_{NS} + \mathcal{O}_2, \quad (2.21b)$$

$$\tau_q \dot{q}^{<\mu>} + q^\mu = q_{NS} + \mathcal{O}_2, \quad (2.21c)$$

where $\dot{\Pi} = D\Pi$, $\dot{q}^{<\mu>} = \Delta_\alpha^\mu Dq^\mu$ and $\dot{\pi}^{\mu\nu} = \Delta_\alpha^\mu \Delta_\beta^\nu D\pi^{\alpha\beta}$ and we have defined the relativistic Navier-Stokes terms to which the dissipative variables relax as

$$\Pi_{NS} = -\zeta\theta, \quad (2.22a)$$

$$\pi_{NS} = -2\eta\sigma^{\mu\nu}, \quad (2.22b)$$

$$q_{NS} = -\kappa T(\nabla^{<\mu>} \ln T + a^\mu), \quad (2.22c)$$

where $\sigma^{\mu\nu} \equiv \frac{1}{2}\nabla^{<\mu}u^{>\nu}$, $\theta = \nabla_\mu u^\mu$ is the expansion and $a_\mu = u^\nu \nabla_\nu u_\mu$ is the acceleration. We have neglected terms of higher-order than first (in gradients, Knudsen number or inverse Reynolds number) in equations 2.21.

The relaxation timescales and dissipation strengths may be related through the expressions

$$\tau_\Pi = \beta_0 \zeta, \quad (2.23a)$$

$$\tau_q = \beta_1 \kappa T, \quad (2.23b)$$

$$\tau_\pi = 2\beta_2 \eta \quad (2.23c)$$

where the β -functions have analytic expressions for certain types of gases - see section 4.1.

The theoretical properties of the MIS model have received thorough investigation (Molnár, Niemi, et al., 2010; Biswas, Dash, et al., 2020; Bemfica, Disconzi, et al., 2021; Bemfica, Disconzi, et al., 2022; Wagner and 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, Chandra, et al., 2013; Du and Heinz, 2020), as well as in the astrophysical community for modelling viscous black-hole accretion (Chabanov, Rezzolla, et al., 2021), for example. The stability of these models is directly related to the fact that the entropy now has a well-defined maximum, rather than having a saddle-point as in the case of naive first-order models discussed previously. Still, entropy conservation is only enforced perturbatively and is therefore only valid within the regime of validity of the model.

2.2.1.6 BDNK Formulation

The BDNK formulation (Bemfica, Disconzi, et al., 2018; Kovtun, 2019) prescribes constitutive relations for the dissipative variables at first-order in gradients of the primitive fluid variables. The primitive fluid variables describe the fluid's fundamental physical properties such as its pressure, velocity, temperature and so on. Despite the instability of other first-order models, the BDNK model stabilises the system in small departures from equilibrium by retaining the time-like derivatives that are systematically replaced by space-like ones in other approaches. This gives a remedy to the problem of unbound entropy growth in other first-order theories, not by providing a maximum, but by allowing for small violations of the second law of thermodynamics outside its regime of strict applicability.

Whilst the causal nature and well-posedness of the MIS formalism (coupled to gravity) has been formally shown in the presence of bulk viscosity only (Bemfica, Disconzi, et al., 2019a), the BDNK formalism is able to describe shear & bulk viscous effects, as well as particle and energy diffusion in a stable and causal way, when coupled to Einstein's equations in curved spacetime. It produces equations of motion with highly desirable properties like strong hyperbolicity and local well-posedness of the initial value problem (Rocha, Wagner, et al., 2024). No other formalism has been shown to achieve all this.

The designation of hydrodynamic frame from being either Landau or Eckart is relaxed, and an additional degree of freedom represented by the scalar A is introduced as a result. The stress-energy tensor now reads

$$T^{\mu\nu} = \rho u^\mu u^\nu + (p + \Pi + A)\Delta^{\mu\nu} + 2q^{(\mu}u^{\nu)} + \pi^{\mu\nu} \quad (2.24)$$

whilst the dissipative variables are given by the expressions

$$A = \tau_\epsilon [u^\mu \nabla_\mu \rho + (\rho + p) \nabla_\mu u^\mu], \quad (2.25a)$$

$$\Pi = -\zeta \Theta + \tau_\Pi / \tau_\epsilon A, \quad (2.25b)$$

$$q^\mu = \tau_q (\rho + p) u^\nu \nabla_\nu u^\mu + \beta_\epsilon \Delta^{\mu\nu} \nabla_\nu \rho + \beta_n \Delta^{\mu\nu} \nabla_\nu n, \quad (2.25c)$$

$$\pi^{\mu\nu} = -\eta \left[\Delta^{\mu\alpha} \Delta^{\nu\beta} \nabla_\alpha u_\beta + \Delta^{\mu\alpha} \Delta^{\nu\beta} \nabla_\beta u_\alpha - \frac{2}{3} \Delta^{\mu\nu} \Delta^{\alpha\beta} \nabla_\alpha u_\beta \right]. \quad (2.25d)$$

The BDNK formulation has already been used in earnest to simulate a fluid with underlying conformal symmetry (Pandya, Most, et al., 2022b) and ideal gas microphysics (Pandya, Most, et al., 2022a). However, we are unaware of simulations using the BDNK model with a realistic equation of state for neutron star matter. We perform simulations of this model in section 7.1 with our usual ideal (relativistic) gas equation of state and compare the results to that of MIS. Whilst not all non-ideal effects in BNS mergers take the BDNK form⁵, one avenue for future work is to perform merger simulations using it.

2.2.2 Valencia Formulation of Balance Law Equations

The ‘Valencia formulation’ of general-relativistic hydrodynamics (GRHD) allows one to rewrite the hydrodynamic equations of motion introduced above in a ‘flux-conservative’ form Banyuls, Font, et al., 1997; Antón, Zanotti, et al., 2006. This is done by building upon the spacetime description laid out in section 2.1.1 for numerical relativity.

An Eulerian observer \mathcal{O}_E with 4-velocity \mathbf{n} is at rest in slice Σ_t . We define an adapted coordinate basis for this observer, using the orthonormal tetrad

$$e_\mu^{(A)} = \{\partial_t, \partial_i\} \quad (2.26)$$

where (A) is a label and

$$\partial_t = \alpha \mathbf{n} + \beta^i \partial_i \quad (2.27)$$

We also make use of the following identity for the covariant derivative of an arbitrary vector Y^μ

$$\nabla_\mu Y^\mu = \frac{1}{\sqrt{-g}} \partial_\mu (\sqrt{-g} Y^\mu) \quad (2.28)$$

where $\sqrt{-g} = \alpha \sqrt{\gamma}$ is the square-root of (minus) the determinant of the 4-metric $g^{\mu\nu}$ whilst $\sqrt{\gamma}$ is the square-root of the determinant of the induced 3-metric γ_{ij} . This allows one to swap the total derivatives for partial derivatives in the hydrodynamic equations of motion (conservation of energy-momentum and baryon number) such that

⁵Bulk viscosity sourced from Urca reactions, as discussed in section 1.1.5, is equivalent to a resumed MIS system, for example.

in full general relativity they become

$$\partial_\mu \left(\sqrt{-g} T_\nu^\mu e_{(A)}^\nu \right) = \sqrt{-g} T_\nu^\mu \nabla_\nu e_{(A)}^\nu, \quad (2.29a)$$

$$\partial_\mu (\sqrt{-g} N^\mu) = 0 \quad (2.29b)$$

where

$$\nabla_\nu e_{(A)}^\nu = \partial_\nu e_{(A)}^\nu + \Gamma_{\mu\lambda}^\nu e_{(A)}^\lambda = \Gamma_{\mu\lambda}^\nu e_{(A)}^\lambda \quad (2.30)$$

and $\Gamma_{\nu\lambda}^\mu$ are the Christoffel symbols given explicitly by

$$\Gamma_{\nu\lambda}^\mu = \frac{1}{2} g^{\mu\kappa} (\partial_\lambda g_{\kappa\nu} + \partial_\nu g_{\kappa\lambda} - \partial_\kappa g_{\nu\lambda}) \quad (2.31)$$

whilst the orthonormal tetrad is

$$e_{(0)\nu} = -\alpha \delta_{0\nu}, \quad (2.32a)$$

$$e_{(k)\nu} = g_{k\nu} = (\beta_k, \gamma_{kj}) \quad (2.32b)$$

with δ_μ^ν as the Kronecker-delta function.

The quantities directly measured by the Eulerian observer \mathcal{O}_E are the rest-mass density, D , the momentum-density in all 3 spatial directions, S_j , and the total energy density, E . We can relate these to the primitive variables $\mathbf{w} = (n, v_j, \rho, p)^T$ as

$$D \doteq -N^\mu n_\mu = nW, \quad (2.33a)$$

$$S_j \doteq -T^{\mu\nu} n_\mu e_{\nu(j)} = (\rho + p)W^2 v_j, \quad (2.33b)$$

$$E \doteq T^{\mu\nu} n_\mu n_\nu = (\rho + p)W^2 - p. \quad (2.33c)$$

This results in the balance-law form (a conservation law in flat spacetime)

$$\frac{1}{\sqrt{-g}} \left[\frac{\partial \sqrt{\gamma} \mathbf{U}(\mathbf{w})}{\partial t} + \frac{\partial \sqrt{-g} \mathbf{F}^i(\mathbf{w})}{\partial x^i} \right] = \mathbf{S}(\mathbf{w}). \quad (2.34)$$

where

$$\mathbf{U}(\mathbf{w}) = (D, S_j, \tau)^T, \quad (2.35a)$$

$$\mathbf{F}^i(\mathbf{w}) = [D\mathcal{V}^i, S_j \mathcal{V}^i + p\delta_j^i, \tau \mathcal{V}^i + p v^i]^T, \quad (2.35b)$$

$$\mathbf{S}(\mathbf{w}) = \left[0, T^{\mu\nu} \left(\frac{\partial g_{\nu j}}{\partial x^\mu} - \Gamma_{\nu\mu}^\delta g_{\delta j} \right), \alpha \left(T^{\mu 0} \frac{\partial \ln \alpha}{\partial x^\mu} - T^{\mu\nu} \Gamma_{\nu\mu}^0 \right) \right]^T \quad (2.35c)$$

with $\mathcal{V}^i \equiv v^i - \frac{\beta^i}{\alpha}$ and $\tau \equiv E - D$.

2.2.3 Variational Models

Variational models are ubiquitous in physics and may indeed be used to model dissipative fluids in general relativity. They begin with an extension to the Einstein-Hilbert action to include matter contributions such that with $G = c = 1$ the action reads

$$S = \frac{1}{16\pi} \int (R + \Lambda) \sqrt{-g} d^4x \quad (2.36)$$

where R is the Ricci scalar of general relativity and Λ is the ‘master-function’ governing matter’s contribution to the action. It is this latter piece, of course, that describes the fluid or, often, multi-fluid system. These multiple fluids may be genuine physically-separable fluids or in-fact thermodynamic currents within a single, continuous fluid that describe the direction and magnitude of the flow of heat, entropy and of course particle number. These currents are not at all necessarily aligned for dissipative fluids. Despite the success of variation models in their ability to produce multi-fluid models of dissipative fluids (see Andersson and Comer, 2015, for example) they do not currently form the basis of numerical simulations of BNS merger simulations. This is because one does not generally obtain balance-law equations from them - see Andersson, Comer, et al., 2017; Andersson, Dionysopoulou, et al., 2017; Andersson, Hawke, et al., 2017, for example. Instead, multi-fluid models are useful for describing neutron star cores (e.g. Haskell, Andersson, et al., 2012; Rau and Wasserman, 2020) where the various exotic particle species and superfluidity present must be modelled using multiple fluids. Similarly, capturing the magnetic field evolution of neutron stars (e.g. Hoyos, Reisenegger, et al., 2008; Glampedakis, Andersson, et al., 2012) may be done by treating the charged species as separate fluids.

2.2.4 Smoothed Particle Hydrodynamics

The smoothed particle hydrodynamics (SPH) method is not so much a theory of fluid dynamics but rather a way of representing a fluid, particularly in numerical simulations. It will be briefly introduced here due to its relevance to simulations of compact objects. Using SPH was first suggested in 1977 by Gingold and Monaghan, 1977 and Lucy, 1977. Instead of existing on a Eulerian grid, fluids are modelled as being comprised of a number of discrete Lagrangian ‘particles’⁶. These are not microphysical particles but rather more like fluid elements whose properties are (spatially) averaged over to obtain bulk fluid properties. Typically, particles have an interaction ‘sphere’ with a radial size given by a small multiple of a ‘smoothing length’ used to set the dimensions of the averaging domain.

⁶An Eulerian view of a fluid focuses on a point in space and analyses the flow past that point over time, whereas a Lagrangian view tracks with the flow of a fluid element.

SPH has both advantages and disadvantages when compared to the traditional, Eulerian fluid models that are more commonly used to describe compact objects and their mergers. In particular, SPH codes exactly conserve mass, energy, momentum and angular momentum (by construction) and can more easily handle vacuum regions of space (which are simply the absence of particles). For merger simulations, this means that ejecta are tracked with ease and that the neutron star surface remains well-behaved. SPH simulations are less restrictive in terms of spatial domain as particles are individually tracked, which also makes tracing back the history and properties of a given particle far easier.

However, there are drawbacks to the use of SPH models: magnetohydrodynamics is notoriously hard to include in simulations, in particular maintaining the vanishing divergence of the magnetic field. Whilst this is not a great concern as yet for us, they can also struggle to handle shocks and fluid instabilities due to surface tension forces that emerge near contact discontinuities. Finally, SPH simulations can become prohibitively costly when performing 3-dimensional simulations which require the inclusion of many particles to accurately capture fine-scale, thermodynamic behaviour - which is our aim! For these reasons, we will devote no further attention to SPH methods in this thesis but direct the interested reader to Rosswog, 2015 for a thorough review of the approach and to their recent work performing BNS merger simulations (Rosswog, Diener, et al., 2022; Rosswog, Torsello, et al., 2023) using the SPHINCS_BSSN code they have developed (Rosswog and Diener, 2021; Rosswog and Diener, 2024).

2.3 Turbulence Modelling and Explicit Large Eddy Schemes

As mentioned at numerous points already in this thesis, we are not currently able to resolve the finest scales of physics within simulations of neutron star mergers. This has proved problematic, for instance when trying to simulate the local, small-scale magnetic turbulence that may explain how remnant neutron star field strengths grow so large. Such an effect is expected to be initially driven by the Kelvin-Helmholtz instability in the shear layer between the two neutron stars during the first few milliseconds post-merger. The highest resolutions achieved so far have grid spacing $\Delta \sim \mathcal{O}(10m)$ (Kiuchi, Kyutoku, et al., 2018) and magnetic field amplifications by a factor of $\sim 10^3$ due to the KHI have been seen (Kiuchi, Cerdá-Durán, et al., 2015). Yet, even when approaching this scale, numerical convergence is not reached, meaning the dynamically relevant scales are unresolved. Whilst the primary magnetohydrodynamic processes during and post-merger are qualitatively understood, quantitative modelling and predictions are out of reach due to the wide range of spatial and temporal scales involved, as well as the non-linearity of the equations. This

is similarly true for the modelling of viscous effects and heat fluxes in BNS mergers and the evolution of the remnant they leave behind, where the cascade of energy from long to short length-scales will inevitably cause unresolved fluid motion.

Explicit Large Eddy Simulations (LES) have been introduced recently to our field of research (Radice, 2017) as a computationally inexpensive way of capturing subgrid scale (SGS) turbulence in BNS mergers. They have been developed over the last 40 or so years and used extensively in engineering applications (see, for instance, Lu and Rutland, 2016) where they are able to resolve complex flows in greater detail and with more predictive power than traditional Reynolds-Averaged Navier-Stokes models that, roughly speaking, time-average the fluid equations. Technically, all computational models of fluid dynamics include (at least) an *implicit* LES scheme. This is because the finite temporal and spatial resolution of any hydrodynamic simulation means there are always unresolved physical scales whose contribution to the solution arises intrinsically (and in an unphysical, uncontrolled way) through the dissipation & dispersion of the numerical scheme employed. Fluid dynamics itself involves an implicit spatial filtering over a local ensemble of individual particles, from which a fluid element emerges, often with a size much greater than that the microphysical dissipation scale but smaller than the turbulent scale. Using implicit LES models to model turbulent flows is surprisingly effective and has been subject to a considerable amount of development & analysis, albeit not for the purposes of neutron star merger simulations. See, for example Margolin and Rider, 2005; Margolin, Rider, et al., 2006; Aspden, Nikiforakis, et al., 2008.

Instead, in explicit LES, the resolvable fluid fields are (spatially) averaged or (temporally) filtered and hence necessarily augmented with closure terms which should account for the additional degrees of freedom introduced in this process. In reality, whether these additional terms genuinely represent unresolved, non-ideal micro-scale physics like particle viscosities, or are simply the result of coarse-graining an ideal fluid is an open question for any given problem. Answering this question is the focus of Berselli, Iliescu, et al., 2006, for instance.

In Celora, Andersson, et al., 2021, a covariant approach to filtering & averaging is developed from first principles using a fibration of spacetime. There, it is found that one may formulate the closure terms in an algebraic form that mimics a relativistic analogue of heat flux and (bulk & shear) viscosity.

Let us demonstrate the basic mathematical principles for implementing an explicit LES closure model for hydrodynamic turbulence. Following Radice, 2017, we start from the Euler equations for conservation of baryon number, momentum and energy in

flat spacetime,

$$\partial_t D + \partial^i D v_i = 0, \quad (2.37a)$$

$$\partial_t S_j + \partial^i (S_j v_i + p \eta_{ij}) = 0, \quad (2.37b)$$

$$\partial_t E + \partial^i S_i = 0 \quad (2.37c)$$

where

$$\begin{pmatrix} D \\ S_j \\ E \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p)W^2 v_j \\ (\rho + p)W^2 - p \end{pmatrix}. \quad (2.38)$$

We apply an arbitrary linear filter to these equations such that a quantity A becomes a filtered quantity \bar{A} , removing features below a particular (unspecified here) scale. The above equations transform simply into

$$\partial_t \bar{D} + \partial^i (\bar{D} v_i) = 0, \quad (2.39a)$$

$$\partial_t \bar{S}_j + \partial^i (\bar{S}_j v_i + \bar{p} \eta_{ij}) = 0, \quad (2.39b)$$

$$\partial_t \bar{E} + \partial^i \bar{S}_i = 0 \quad (2.39c)$$

which are exact but no longer closed as the term $\bar{S}_j v_i$ can not be written solely in terms of other, coarse-grained quantities. A closure relation is needed of the form

$$\bar{S}_j v_i = \bar{S}_j \bar{v}_i + \tau_{ij} \quad (2.40)$$

where τ_{ij} represents SGS turbulence. Relying on numerical viscosity and setting $\tau_{ij} = 0$ is the usual, implicit LES approach. Alternatively, we may choose, for example, a relativistic analogue of that introduced by Smagorinsky, 1963 such that

$$\tau_{ij} = -2\nu_T(\rho + p)W^2 \left[\frac{1}{2} (\nabla_i \bar{v}_j + \nabla_j \bar{v}_i) - \frac{1}{3} \nabla_k \bar{v}^k \eta_{ij} \right]. \quad (2.41)$$

This expression has a familiar form! That of the viscous stress tensor in the Navier-Stokes equations eq. (2.5) or its relativistic counterpart eq. (2.20) introduced into the stress-energy tensor (and hence equations of motion) for non-ideal fluids at first-order in a gradient expansion. Note that it is not covariant as it is written purely in terms of spatial components and is defined at the equation of motion level rather than a higher one - see Duez, Knight, et al., 2020 for this, and for a wider comparison of momentum transport models used in numerical relativity.

However, there are differences: for example, the ‘viscosity’ factor ν_T is no longer a physically-determinable quantity but rather dependent on both the particular filtering operation chosen and the Eulerian observer ‘doing’ the filtering⁷. On dimensional

⁷In relativity, averaging and filtering have no global definition and instead become observer dependent operations.




FIGURE 2.2: Relativistic filtering of a fine-grain velocity field where black lines indicate streamlines. The red line is the coarse-grained, ‘average’ four-velocity. If the fine-grain field is filtered with respect to that observer, the coarse-grained field is found. This procedure is independent of the coordinates, but its inherent non-linearity makes it impractical for use within a simulation.

grounds, $\nu_T \sim \lambda_{mix} c_s$ where c_s is the sound speed in the fluid and λ_{mix} is a characteristic mixing length on which turbulence occurs. Because of this difference, ν_T should be calibrated using comparison of filtered, high-resolution simulation results to unfiltered, lower-resolution ones as in Carrasco, Viganò, et al., 2020; Viganò, Aguilera-Miret, et al., 2020a. Alternatively, its value may be set using expressions for the wavelength of instabilities such as the magneto-rotational instability and sound speed in a neutron star as in Duez, Liu, et al., 2006. Interestingly, the introduction of this viscous-tensor-like closure term into the filtered, relativistic Euler equations does not induce the same pathologies that the relativistic Navier-Stokes equations suffer from (Radice, 2017). The stability seen in practice may indicate special properties of the filtering terms, or may only show that the numerical resolutions used introduce an

effective frequency cut-off that artificially removes the instabilities. See Duez, Knight, et al., 2020 for a highly pertinent discussion of instabilities arising whilst using SGS turbulence modelling for BNS mergers.

For our interests, the good news is that we already have models coded to include terms that may serve as a closure piece for explicit LES simulations, even if their inclusion were initially motivated on physical grounds. However, a word of caution on applying the sort of closure scheme introduced above, and indeed all current published implementations of explicit LES closures for capturing unresolved dynamics in BNS mergers. It has been applied after the standard (3+1) division of spacetime, solely to the spatial equations of motion. It is therefore not covariant and the filtered quantities will depend, as previously mentioned, on the observer ‘doing’ the filtering. In Celora, Hawke, et al., 2022, a covariant formulation of filtering is developed by finding a local observer at each point in spacetime which is aligned with the number current at the coarse, filtered level. In chapter 6, we implement the scheme laid out therein and model the residuals it introduces.

2.4 Kinetic Theory and the Chapman-Enskog Expansion

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 in its relativistic form as

$$p^\mu \partial_\mu f(\mathbf{x}, \mathbf{p}) = \mathcal{C}[f] \quad (2.42)$$

where p^μ is the 4-momentum of the particle and $f(\mathbf{x}, \mathbf{p})$ is the 8-dimensional 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 $n(\mathbf{x})$ is given by

$$n(\mathbf{x}) = \int \frac{d^3p}{(2\pi)^3} f(\mathbf{x}, \mathbf{p}) \quad (2.43)$$

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

$$N^\mu = \int_p p^\mu f(\mathbf{x}, \mathbf{p}), \quad (2.44a)$$

$$T^{\mu\nu} = \int_p p^\mu p^\nu f(\mathbf{x}, \mathbf{p}) \quad (2.44b)$$

where $\int_p \cdots \equiv \int d^3p / [(2\pi)^3 p^0]$ 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 and Cowling, 1990 and Kumar, 1967. Their approach allows one to derive macroscopic descriptions of fluids, including the familiar Euler and Navier-Stokes equations (Cercignani and 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 Chapman-Enskog (C-E) analysis is that the particle collision duration time is not only far less than the time between collisions, but also far smaller than *extrinsic* timescales: the timescales associated with macroscopic fluid behaviour. The solution is in fact obtained as an asymptotic expansion in the closely-related Knudsen number⁸, where at zero'th order we have the equilibrium solution.

2.4.1 Viscosity via Chapman-Enskog Expansion

Here we derive both the bulk and shear viscosity coefficients, ζ & η , by applying this method to the relativistic Boltzmann equation.

The Chapman-Enskog method expands f in powers of a small parameter ϵ (e.g. the Knudsen number) representing deviations from equilibrium:

$$f = f^{(0)} + \epsilon f^{(1)} + \epsilon^2 f^{(2)} + \dots, \quad (2.45)$$

where $f^{(0)}$ is the local equilibrium distribution. Inserting into the Boltzmann equation and collecting first-order terms yields

$$p^\mu \partial_\mu f^{(0)} = -\mathcal{C}[f^{(1)}], \quad (2.46)$$

where \mathcal{C} is the linearized collision operator.

Shear Viscosity

To solve for $f^{(1)}$, we postulate the following ansatz appropriate for shear viscous corrections:

$$f^{(1)} = -f^{(0)} \phi(p), \quad \text{with} \quad \phi(p) = A(p) p^{\langle\mu} p^{\nu\rangle} \sigma_{\mu\nu}. \quad (2.47)$$

Here, $p^{\langle\mu} p^{\nu\rangle}$ is the symmetric, traceless part of the momentum tensor:

$$p^{\langle\mu} p^{\nu\rangle} = p^\mu p^\nu - \frac{1}{3} \Delta^{\mu\nu} (p \cdot u)^2 \quad (2.48)$$

⁸The 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.

where $p \cdot u \equiv p^\mu u_\mu$. The function $A(p)$ depends on the particle energy and encodes the strength of the deviation.

The shear part of the energy-momentum tensor is

$$\pi^{\mu\nu} = \int \frac{d^3p}{p^0} p^\mu p^\nu \delta f = - \int \frac{d^3p}{p^0} p^\mu p^\nu f^{(0)} \phi(p). \quad (2.49)$$

Substituting the ansatz for $\phi(p)$ and using the symmetry of $\sigma^{\mu\nu}$, we extract η via

$$\pi^{\mu\nu} = -\sigma_{\alpha\beta} \int \frac{d^3p}{p^0} f^{(0)} A(p) p^\mu p^\nu p^{\langle\alpha} p^{\beta\rangle}, \quad (2.50)$$

$$= -\eta \sigma^{\mu\nu}. \quad (2.51)$$

Hence,

$$\eta = \frac{1}{10T} \int \frac{d^3p}{(2\pi)^3} \frac{1}{E} A(p) \left(p^\mu p^\nu - \frac{1}{3} \Delta^{\mu\nu} p^2 \right)^2 f^{(0)}. \quad (2.52)$$

To find $A(p)$, one solves the linearized Boltzmann equation

$$\mathcal{C}[f^{(1)}] = -p^\mu \partial_\mu f^{(0)}, \quad (2.53)$$

which becomes a linear integral equation for $A(p)$ involving the collision kernel. In practice, this is often solved numerically or approximated via variational methods or model collision terms (via the Anderson–Witting model, for example).

In the relaxation time approximation (RTA), we write the collision operator as

$$\mathcal{C}[f^{(1)}] = -\frac{p^\mu u_\mu}{\tau_R} f^{(1)}, \quad (2.54)$$

yielding

$$A(p) = \frac{\tau_R}{T(p^\mu u_\mu)}. \quad (2.55)$$

Substituting back, we recover the RTA result for shear viscosity:

$$\eta = \frac{\tau_R}{10T} \int \frac{d^3p}{(2\pi)^3} \frac{1}{E} \left(p^\mu p^\nu - \frac{1}{3} \Delta^{\mu\nu} p^2 \right)^2 f^{(0)} \quad (2.56)$$

where $E = p \cdot u$. This provides the connection between the Chapman–Enskog formalism and the RTA derivation discussed previously.

Bulk Viscosity

Bulk viscosity ζ quantifies the fluid’s resistance to isotropic expansion or compression and arises from the scalar deviation of the energy-momentum tensor away from its

equilibrium form. We derive ζ from the relativistic Boltzmann equation using the Chapman–Enskog method in the relaxation time approximation.

For the bulk viscosity, we assume the deviation from equilibrium is proportional to scalar hydrodynamic gradients, particularly the expansion scalar $\theta = \partial_\mu u^\mu$. The appropriate ansatz for $f^{(1)}$ is now

$$f^{(1)} = -f^{(0)} \phi(p), \quad \phi(p) = \mathcal{A}(p) \theta. \quad (2.57)$$

To determine $\mathcal{A}(p)$, we match the kinetic theory energy-momentum tensor to its hydrodynamic decomposition.

The first-order correction to the energy-momentum tensor is

$$\delta T^{\mu\nu} = \int \frac{d^3p}{p^0} p^\mu p^\nu \delta f = -\theta \int \frac{d^3p}{p^0} p^\mu p^\nu f^{(0)} \mathcal{A}(p). \quad (2.58)$$

In hydrodynamics, the bulk viscous stress is given by

$$\Pi = -\zeta \theta, \quad T^{\mu\nu} = T_{\text{ideal}}^{\mu\nu} + \Pi \Delta^{\mu\nu}. \quad (2.59)$$

So we identify the bulk part by taking the trace:

$$\Pi = \frac{1}{3} \Delta_{\mu\nu} \delta T^{\mu\nu} = -\theta \int \frac{d^3p}{p^0} \left(\frac{1}{3} \Delta_{\mu\nu} p^\mu p^\nu \right) f^{(0)} \mathcal{A}(p). \quad (2.60)$$

Thus, the bulk viscosity becomes

$$\zeta = \int \frac{d^3p}{(2\pi)^3} \frac{1}{E} \left(\frac{1}{3} \Delta_{\mu\nu} p^\mu p^\nu \right) f^{(0)} \mathcal{A}(p). \quad (2.61)$$

To find $\mathcal{A}(p)$, we return to the first-order Boltzmann equation

$$p^\mu \partial_\mu f^{(0)} = \frac{p^\mu u_\mu}{\tau_R} f^{(0)} \mathcal{A}(p) \theta. \quad (2.62)$$

where for a local equilibrium distribution $f^{(0)} = \exp[-(p \cdot u - \mu)/T]$, we compute

$$p^\mu \partial_\mu f^{(0)} = -f^{(0)} \left[\frac{p^\mu p^\nu}{T} \nabla_\mu u_\nu - \left(\frac{(p \cdot u)^2}{T^2} - \frac{m^2}{T^2} \right) \nabla \cdot u \right] + \dots. \quad (2.63)$$

Isolating the scalar (bulk) contribution, we match the coefficients of θ to identify

$$\mathcal{A}(p) = \frac{\tau_R}{T(p \cdot u)} \left[\left(\frac{1}{3} - c_s^2 \right) (p \cdot u)^2 - c_s^2 m^2 \right], \quad (2.64)$$

where $c_s^2 = \frac{dp}{d\rho}$ is the speed of sound squared.

Substituting into the expression for ζ , we find

$$\zeta = \frac{\tau_R}{T} \int \frac{d^3p}{(2\pi)^3} \frac{1}{E^2} \left[\left(\frac{1}{3} - c_s^2 \right) E^2 - c_s^2 m^2 \right]^2 f^{(0)}. \quad (2.65)$$

This result vanishes in the conformal limit $c_s^2 = 1/3$ and $m \rightarrow 0$, as expected. The nonzero value of ζ therefore reflects the breaking of conformal symmetry due to finite mass and interaction effects.

2.4.2 Heat Conductivity via Chapman-Enskog Expansion

Heat conductivity κ quantifies the response of a fluid to temperature gradients, corresponding to energy transport in the presence of a nonzero heat flow. In relativistic kinetic theory, it is extracted from the first-order deviation of the distribution function around equilibrium. We use the Chapman-Enskog expansion of the relativistic Boltzmann equation in the relaxation time approximation (RTA).

In the Landau frame, the energy-momentum tensor satisfies $T^{\mu\nu}u_\nu = \epsilon u^\mu$, and the particle current decomposes as

$$N^\mu = n u^\mu + q^\mu, \quad (2.66)$$

where q^μ is the heat flow, orthogonal to u^μ , i.e., $u_\mu q^\mu = 0$.

The first-order correction to the particle current is

$$\delta N^\mu = \int \frac{d^3p}{(2\pi)^3 E} p^\mu \delta f. \quad (2.67)$$

To isolate the heat conductivity, we consider the linear response of the heat flow to temperature and chemical potential gradients:

$$q^\mu = -\kappa \left(\frac{\nabla^\mu T}{T} - \frac{T}{h} \nabla^\mu \left(\frac{\mu}{T} \right) \right), \quad (2.68)$$

where $h = \epsilon + P$ is the enthalpy density, and $\nabla^\mu = \Delta^{\mu\nu} \partial_\nu$ projects gradients orthogonal to u^μ .

We take the ansatz

$$f^{(1)} = -f^{(0)} \phi(p), \quad \phi(p) = \mathcal{B}^\mu(p) \nabla_\mu \left(\frac{\mu}{T} \right). \quad (2.69)$$

From the Boltzmann equation, using

$$p^\mu \partial_\mu f^{(0)} = f^{(0)} \left[- \left(\frac{p^\mu u_\mu - \mu}{T} \right) p^\nu \nabla_\nu \left(\frac{\mu}{T} \right) + \dots \right], \quad (2.70)$$

we equate terms and obtain

$$\mathcal{B}^\mu(p) = \frac{\tau_R}{T(p \cdot u)} (p^\mu - hu^\mu) \left(\frac{p \cdot u - \mu}{T} \right). \quad (2.71)$$

The heat flow is then

$$q^\mu = \Delta^\mu{}_\nu \int \frac{d^3p}{(2\pi)^3 E} p^\nu \delta f = -\nabla^\alpha \left(\frac{\mu}{T} \right) \int \frac{d^3p}{(2\pi)^3 E} (p^\mu - hu^\mu) (p_\alpha - hu_\alpha) f^{(0)} \frac{\tau_R}{T(p \cdot u)} \left(\frac{p \cdot u - \mu}{T} \right) \quad (2.72)$$

such that after evaluating the tensor structure and taking into account isotropy and orthogonality to u^μ , the heat conductivity becomes

$$\kappa = \frac{\tau_R}{3T^2} \int \frac{d^3p}{(2\pi)^3 E^2} (p \cdot u - h)^2 (p \cdot u - \mu)^2 f^{(0)}. \quad (2.73)$$

This expression vanishes in equilibrium and grows with relaxation time and particle energy, as expected.

Our Analogous Expansion

Inspired by this method's success, work has been done (Wright and Hawke, 2019; Wright and Hawke, 2020; Hatton and Hawke, 2024) in applying this method to describe a system where a dissipative source acts rapidly to relax the solution towards equilibrium. For the MIS formulation of dissipative hydrodynamics, we may indeed have a source of this form, provided that the relaxation timescales $\{\tau\}$ are small. In this case, we consider the forms of the source terms to motivate an ansatz for the dissipative variables which is a series expansion in the timescales on which they act. Up to first order, this is:

$$q = q_{\text{NS}} + \tau_q q_1 + \mathcal{O}(\tau_q^2) \quad (2.74a)$$

$$\Pi = \Pi_{\text{NS}} + \tau_\Pi \Pi_1 + \mathcal{O}(\tau_\Pi^2) \quad (2.74b)$$

$$\pi = \pi_{\text{NS}} + \tau_\pi \pi_1 + \mathcal{O}(\tau_\pi^2) \quad (2.74c)$$

where we have chosen to omit the indices for the heat flux vector and shear viscosity tensor and at zero'th order in the timescales we have the Navier-Stokes form for the dissipative variables. We must solve for the form of the dissipative variables at first-order in the timescales. Making this substitution eliminates the reciprocal timescales that appear in the coefficients of our source terms in the MIS equations, removing their stiffness - see section 3.5 for a discussion of why this is numerically helpful. This substitution is only accurate when $\{\tau\} \ll 1$ and we expect deviations from the 'true' solution when this is not the case.

Chapter 3

Numerical Methods

The complexity of the world we live in makes analytic solutions to even simple-looking systems intractable. Even when all the underlying physics is well-understood, accurately modelling a real-world scenario is challenging with pen and paper. Numerical methods take on many forms but all seek to approach the ‘true’ solution to a problem with the precision & accuracy needed to make the answer meaningful and the method predictive. They do this whilst circumventing the parts of the problem that make an analytical, closed-form solution impossible.

One of the oldest and most tedious examples of a numerical method is that of Archimedes, who drew circles and regular polygons inscribed inside them with increasingly many sides in an effort to arrive at an accurate and precise value of π . Incredibly, mathematicians could not rid themselves of this painstaking process until the genius of Isaac Newton produced an analytic approach using an infinite series to arrive at a value for π - still a numerical method.

In modern science, the complexity of systems we wish to model numerically mandates use of computers, often “supercomputers”. We are grateful for access to Southampton University’s High Performance Computing (HPC) clusters, Iridis 4 & 5. They have allowed us to run the simulations that form the basis of the results presented in this thesis. For the numerical evolution of our fluid models, we make heavy use of the METHOD (<https://github.com/AlexJamesWright/METHOD>) codebase originally developed by Alex James Wright as part of his PhD thesis work and extended for our purposes here (<https://github.com/MarcusJHatton/METHOD>).

This codebase implements modern numerical methods in order to solve a hyperbolic, balance-law form of the relativistic hydrodynamic equations. In the following section, we will give a fairly pedagogical introduction to tackling these types of problems, and then discuss in further detail their specific implementation in METHOD. For a comprehensive review of numerical hydrodynamics in special relativity, see Marti and Mueller, 1999.

3.0.1 Numerical Schemes for Relativistic Viscous Hydrodynamics in Astrophysics

Relativistic viscous hydrodynamics has become an essential framework for modeling high-energy astrophysical phenomena, including neutron star mergers, accretion disks around black holes, and core-collapse supernovae. In such environments, relativistic speeds, strong gravitational fields, and dissipative processes like shear and bulk viscosity, as well as heat conduction, must be accurately captured. Solving the relativistic viscous equations numerically is challenging due to their hyperbolic-parabolic structure and the stiff relaxation dynamics introduced by causal dissipative theories. High-resolution shock-capturing (HRSC) techniques are standard, using Riemann solvers (e.g., HLL, HLLE, HLLC) and spatial reconstruction schemes such as MUSCL, WENO, or PPM to handle shocks and steep gradients without generating spurious oscillations (Marti and Mueller, 1999; Rezzolla and Zanotti, 2013).

To preserve causality and stability in relativistic settings, viscous effects are often implemented via Müller–Israel–Stewart theory or its modern generalizations, such as the Denicol–Niemi–Molnar–Rischke (DNMR) (Denicol, Niemi, et al., 2012) formulation. As already seen, these theories may introduce relaxation equations for the shear stress tensor $\pi^{\mu\nu}$, bulk viscous pressure Π , and heat flux q^μ which can become stiff and often require implicit-explicit (IMEX) Runge–Kutta schemes or operator splitting methods to evolve them stably.

3.1 The Balance-Law Form of Hyperbolic Hydrodynamic Equations

As discussed in section 2.2.2, the pair of equations governing conservation of mass and stress-energy may be manipulated into a hyperbolic, flux-conservative form - commonly, the “Valencia” formulation (Banyuls, Font, et al., 1997; Antón, Zanotti, et al., 2006). This allows for the application of high resolution shock capturing methods (Ma, 2001; Marti and Mueller, 1999) whose fundamental property is the ability to accurately solve Riemann problems when discontinuous fluid data arises in the numerical simulation¹. Following the format of LeVeque, 2002, this balance-law form reads

$$\partial_t \mathbf{U}(\mathbf{w}) + \partial_i \mathbf{F}^i(\mathbf{U}) = \mathbf{S}(\mathbf{U}) \quad (3.1)$$

¹Using a flux-conservative form also has important numerical implications due to a pair of powerful theorems: Lax-Wendroff (1960) and Hou-LeFlock (1994). The first essentially guarantees that the function to which the numerical method converges is indeed a solution to the original conservation law. The second says that non-conservative formulations to not converge do the correct solution in the presence of shock waves.

where \mathbf{U} is the conserved vector, a function of the primitive fluid variables, \mathbf{w} , and \mathbf{F} and \mathbf{S} are the flux and source vectors, respectively, which may be thought of as functions of either the conserved variables, or directly as functions of the primitive variables. Without loss of generality, we may assume (or rather manipulate the source and flux vectors such that) all spatial derivatives are contained within the flux vector. Then, when we write the balance-law as

$$\partial_t \mathbf{U}(\mathbf{w}) + \left(\frac{\partial \mathbf{F}^i}{\partial \mathbf{U}} \right) \partial_i \mathbf{U} = \mathbf{S}(\mathbf{U}), \quad (3.2)$$

where the Jacobian, $\underline{J}^i(\mathbf{U}) := \partial_{\mathbf{U}} \mathbf{F}^i(\mathbf{U})$, governs the propagation speeds of information within the fluid. More technically, a perturbation to the j^{th} component of \mathbf{U} travels at speed λ_j in the i^{th} direction within the fluid, where $\{\lambda_j\}^i$ is the set of eigenvalues of the Jacobian, \underline{J}^i . If all these eigenvalues are real and distinct, the system is said to be strictly hyperbolic, and the perturbation propagation or ‘sound’ speeds within the fluid are physical. If the eigenvalues are such that $|\lambda_{min}| \ll |\lambda_{max}|$, the system is said to be stiff as dynamics will occur across a wide range of time and length-scales, making it difficult to capture the full spectrum of physics involved in the problem. This issue will come to bear in our implementations of dissipative hydrodynamics, as seen most clearly in section 4.2.

Note that in this thesis we will work with a number of non-zero source vectors which contain temporal & spatial derivatives. These are generally small in value when compared to the flux vector so will make minor alterations to the characteristic speeds of information travel from their non-sourced values. For the models of relativistic, dissipative hydrodynamics presented in sections 2.2.1.5 and 2.2.1.6, analytic expressions for the eigenvalues of the Jacobian cannot reasonably be found. Instead, the effects of altered propagation speeds will be seen in the primitive values. However, for the reduced models presented in section 4.1, analytic expressions are obtained for the characteristic speeds.

To solve a balance-law system of Dissipative Hydrodynamic (DHD) equations like (3.1), we apply the following series of steps:

1. Discretize our spacetime using foliations: space-like hypersurfaces with perpendicular time-like vectors that describe the direction of evolution of time.
2. These three-dimensional hypersurfaces are then discretized themselves when living in the memory of a computer, by a spatial grid spanning three dimensions. In this thesis we work only with Cartesian coordinates. The grid contains cuboid-shaped with sides of lengths Δx , Δy and Δz and we take time-steps forward of size Δt .

3. Across the first hypersurface at $t = 0$, specify initial data for the primitive variables. From these, calculate the initial state, flux and source vectors $(\mathbf{U}, \mathbf{F}, \mathbf{S})$.
4. Calculate the fluxes through the cell interfaces (spatial derivatives of the flux vector, $\partial_i \mathbf{F}^i(\mathbf{U})$). This typically requires reconstruction of the flux vector at the cell edges.
5. Use a time integrator to progress the system (specifically, the state vector) to a later time.
6. Apply boundary conditions to cells within the ‘ghost-zones’. Ghost cells are additional cells added in the exterior of our regular (interior) domain. They allow both boundary conditions to be applied, and derivatives to be calculated, to the edges of the interior region.
7. Perform a conserved-to-primitive transformation to obtain updated values of the primitive variables, allowing this iterative process to start again and continue until an end-time is reached.

We write the series of state vectors that we obtain in this process as \mathbf{U}_n , where the subscript is an integer label for the current iteration, i.e. $\mathbf{U}_n \equiv \mathbf{U}(t_n)$ and $t_{n+1} = t_n + \Delta t$. Similarly, $\mathbf{F}_i \equiv \mathbf{F}(x_i)$, where the subscript now labels a particular cell in space and all cells together form a grid over the whole spatial domain of our problem. How exactly the time and spatial derivatives are calculated in a numerical way varies greatly and much work has been devoted to doing this in an efficient and accurate manner, as one would expect when it is the crux of the numerical method. Next, we will expand upon the series of steps described above and cover the pertinent points for our purposes.

Below, we will introduce a finite-difference method for solving a PDE, where the spatial derivatives are calculated by taking differences between the values of discrete variables. An alternative is to use finite element methods, or the closely related spectral methods. These represent the solution as a sum of basis functions (trigonometric functions or Legendre polynomials, for instance) which is substituted into the PDE and the coefficients of the basis functions are calculated over time to solve the system. Whilst finite element methods typically involve compact support, meaning that the basis functions are non-zero (and hence the variables connected) only over local sub-domains, spectral methods have global support. Because of this, spectral methods can give exponentially-convergent error with resolution for smooth solutions, giving highly accurate results. However, their inability to capture shocks renders them impractical for our desired applications here. They have been used in codes written to solve the spacetime evolution in black hole mergers, for instance by

the SpEC² code employed by the Simulating Extreme Spacetimes collaboration. For a review, see Canuto, Hussaini, et al., 2007.

3.2 A Simple Numerical Scheme

One simple finite-difference scheme is constructed using first-order-in-time forward-differencing for the time differential and second-order-in-space central-differencing for the spatial derivatives. Then, in one spatial dimension,

$$\frac{\partial \mathbf{U}_n}{\partial t} \approx \frac{\mathbf{U}_{n+1} - \mathbf{U}_n}{\Delta t} \quad (3.3)$$

and

$$\frac{\partial \mathbf{F}_i}{\partial x} \approx \frac{\mathbf{F}_{i+1} - \mathbf{F}_{i-1}}{2\Delta x} \quad (3.4)$$

so that the numerical equivalent of (3.1) using this forward-time, centered-space approximation (“FTCS”) is

$$\mathbf{U}_{n+1,i} \approx \mathbf{U}_{n,i} + \frac{\Delta t}{2\Delta x} (\mathbf{F}_{n,i+1} - \mathbf{F}_{n,i-1}) + \Delta t \mathbf{S}_{n,i}. \quad (3.5)$$

This gives an **explicit** scheme for obtaining the solution ($\mathbf{U}_{n+1,i}$) for the i ’th cell at the $(n+1)$ ’th timestep because it may be performed given information wholly contained within the n ’th timestep. An example of an **implicit** scheme is the backward-time centered-space or “BTCS” method which has the form

$$\mathbf{U}_{n+1,i} \approx \mathbf{U}_{n,i} + \frac{\Delta t}{2\Delta x} (\mathbf{F}_{n+1,i+1} - \mathbf{F}_{n+1,i-1}) + \Delta t \mathbf{S}_{n+1,i}. \quad (3.6)$$

Because this scheme requires knowing $\mathbf{F}_{n+1}(\mathbf{U}_{n+1})$ to calculate \mathbf{U}_{n+1} , one must solve a set of nonlinear equations at each timestep to obtain the updated solution. This is computationally expensive. However, it is guaranteed to be stable and convergent. In contrast, the explicit scheme is only guaranteed to be stable and convergent if the Courant-Friedrichs-Lewy (CFL) condition is met. This can be written in a general form as

$$C \equiv \left| \frac{\partial \mathbf{F}}{\partial \mathbf{U}} \right| \frac{\Delta t}{\Delta x} \leq C_{max} \quad (3.7)$$

where C is the Courant number and the partial differential represents the advection velocity of the fluid. The maximum allowed Courant number, C_{max} , is dependent on the numerical method used to solve the system, with implicit schemes typically permitting larger values than explicit ones (≈ 100 compared to ≈ 1). Meeting this condition can force us to take very small timesteps, especially for a high-resolution grid of cells. For example, to double the resolution of a simulation in three spatial dimensions, the computational cost is increased by a factor of 2^4 (three factors of 2

²www.black-holes.org

come from doubling the number of cells along each axis and the fourth comes from needing to halve the timestep's size to continue to meet the CFL condition)!

3.3 Riemann Problems

The simple schemes introduced above already present a number of questions, some more immediate than others. Firstly, can we make better numerical approximations for the temporal and spatial derivatives than simple first-order ones that converge more sharply with increasing resolution. In mergers (and heavy ion collisions, supernovae, jets etcetera) we expect sharp gradients in primitive quantities to form that may present problems for the simple low-order, finite-difference approximations to spatial derivatives introduced above. Even without the presence of physical discontinuities, our discretization of space means that the cell-averaged values of variables will jump at the interfaces between cells. This represents a type of Riemann problem, a class of problems that involves two piecewise constant states separated by a sharp interface. Given the inapplicability of the strong-form of the balance-law equation (3.1) when discontinuities are present, we consider instead the weak-form where we integrate over the cell's domain from left (x_L) to right (x_R) interface. We continue to work in one dimension and drop our source so that, without loss of generality,

$$\frac{d}{dt} \int_{x_L}^{x_R} \mathbf{U}(t, x) dx = \mathbf{F}(\mathbf{U}(t, x_L)) - \mathbf{F}(\mathbf{U}(t, x_R)) \quad (3.8)$$

Across a shock with position $X(t)$ located between x_L and x_R , we then have

$$\frac{dX(t)}{dt} (\mathbf{U}_L - \mathbf{U}_R) = (\mathbf{F}_L - \mathbf{F}_R) \quad (3.9)$$

where dX/dt is the speed of propagation of the discontinuity. Riemann solvers are usually part of modern High Resolution Shock Capturing (HRSC) methods which are used in relativistic hydrodynamic codes - see Ma, 2001. They maintain high (2^{nd} to 4^{th}) order accuracy in smooth flows and preserve stability (with first order accuracy) at discontinuities. First, the left ($\mathbf{U}(t, x_L)$) and right ($\mathbf{U}(t, x_R)$) states are *reconstructed* - see section 3.4. This is often done using a Total Variation Diminishing (TVD) scheme. The total variation (TV) of a scheme is defined as

$$TV(\mathbf{Q}^n) = \sum_j \|\mathbf{Q}_{j+1}(t^n) - \mathbf{Q}_j(t^n)\| \quad (3.10)$$

and a TVD scheme then satisfies

$$TV(\mathbf{Q}^{n+1}) \leq TV(\mathbf{Q}^n) \quad (3.11)$$




FIGURE 3.1: Wave pattern for the Riemann problem applied to a fluid with initial left and right states (L, R) taken from Harpole, 2018. The initial discontinuity decays into two non-linear waves $(W_{\leftarrow}, W_{\rightarrow})$ between which there are two new constant states (L^*, R^*) separated by a contact discontinuity \mathcal{C} . This wave pattern will be clearly seen in the shocktube test cases presented in section 4. The initial data for these tests essentially represent large-scale Riemann problems where W_{\leftarrow} is a rarefaction wave and W_{\rightarrow} is a shockwave.

where

$$\mathcal{Q}_i^n = \frac{1}{\Delta x_i} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \mathbf{u}(t^n, x) dx \quad (3.12)$$

is the cell-averaged value of the conserved state vector. Schemes that are TVD minimize the onset of spurious or unphysical oscillations at discontinuities Harten (1983). Methods in which the stricter, local condition $\|\mathcal{Q}^{n+1}\| \leq \|\mathcal{Q}^n\|$ is also satisfied are known as strong stability preserving (SSP).

Solving Riemann problems exactly is time-consuming, particularly in the case of multidimensional special-relativistic hydrodynamics where the Lorentz factor couples the set of equations together. Typically, as a far faster yet still highly accurate alternative, approximate Riemann solvers are used. Early developments were made by Roe, 1981, where an eigen-decomposition of the Jacobian is used to track the travelling waves present in the system. This is prohibitively costly for large systems of equations such as those used in relativistic, dissipative hydrodynamics. Many alternatives exist, and in section 3.4 we outline the Lax-Friedrichs approach. For a review see, for example, Marti and Mueller, 1999.

3.4 Flux Methods

The numerical fluxes $(\partial_i \mathbf{F}^i(\mathbf{U}))$, which are derivatives of the flux vector, are required to evolve the system given by eq. (3.1). We can calculate the net flux through the j^{th} cell, $\mathbf{F}_j^i(\mathbf{U})$, by considering the fluxes that are both entering and leaving the cell through its left and right edges. Dropping the spatial-direction index superscript, we write these as $\{\mathbf{F}_L(\mathbf{U}), \mathbf{F}_R(\mathbf{U})\} \equiv \{\mathbf{F}_{j-\frac{1}{2}}(\mathbf{U}), \mathbf{F}_{j+\frac{1}{2}}(\mathbf{U})\}$.

Because the values of the state and primitive variables are defined at cell centres, we must perform reconstruction to be able to calculate the flux vector at the cell interfaces where mass/energy/momentum flows between cells. This may be done either by first reconstructing the state variables at the cell faces, $\mathbf{U}_{j-\frac{1}{2}}, \mathbf{U}_{j+\frac{1}{2}}$, and then calculating the fluxes from them: $\mathbf{F}_{j-\frac{1}{2}} = \mathbf{F}(\mathbf{U}_{j-\frac{1}{2}}), \mathbf{F}_{j+\frac{1}{2}} = \mathbf{F}(\mathbf{U}_{j+\frac{1}{2}})$; or, by directly reconstructing the fluxes themselves at the cell faces:

$\mathbf{F}_{j-\frac{1}{2}} = \mathbf{F}(\dots, \mathbf{F}_{j-1}, \mathbf{F}_j, \dots), \mathbf{F}_{j+\frac{1}{2}} = \mathbf{F}(\dots, \mathbf{F}_j, \mathbf{F}_{j+1}, \dots)$ using interpolation.

The flux vector splitting (FVS) technique, originally proposed by Shu, 1998, opts for the latter approach. FVS is relatively inexpensive, simple to implement and robust. It is constructed to minimise numerical oscillations at shocks when applied with weighted, essentially non-oscillatory (WENO) reconstruction methods which were first introduced by Liu, Osher, et al., 1994. The specific implementation in METHOD will be outlined here.

To calculate the inter-cell fluxes, we split the flux at the cell boundaries into upwind (+) and downwind (−) components, which are opposite in sign and considered to be entirely right- and left-moving, respectively. To do this, we use the Lax-Friedrichs flux splitting approximation, which mixes the overall contribution to the flux between the flux vector and the state vector. We can ensure the positivity/negativity of each contribution by writing $\mathbf{F}_j^\pm = \frac{1}{2}(\mathbf{F}_j \pm \alpha \mathbf{U}_j)$ such that $\mathbf{F}_j = \mathbf{F}_j^+ + \mathbf{F}_j^-$. The coefficient α introduces an artificial viscosity that scales with the size of the jump in the conserved variables—this smooths shocks and damps discontinuities. One normally defines α as being the maximum wave-speed of the system, $\alpha = \max_p |\lambda_p|$ with λ_p as the p^{th} -eigenvalue of the Jacobian of the system. In practice, it is therefore the fastest wave-speed contained within the stencil used to apply the particular order of WENO method chosen. However, for simplicity, we set $\alpha = 1$, corresponding to light propagation speed, admittedly adding some unnecessary additional diffusivity but ensuring sufficient damping at discontinuities.

Applying this at the cell boundaries, we get

$$\mathbf{F}_{j+\frac{1}{2}} = \mathbf{F}_{j+\frac{1}{2}}^+ + \mathbf{F}_{j+\frac{1}{2}}^-, \quad (3.13)$$

$$\mathbf{F}_{j-\frac{1}{2}} = \mathbf{F}_{j-\frac{1}{2}}^+ + \mathbf{F}_{j-\frac{1}{2}}^-, \quad (3.14)$$

where we must construct both the upwind and downwind fluxes at each face using a WENO scheme. For a second-order WENO scheme this means that

$$\mathbf{F}_{j+\frac{1}{2}}^+ = \text{WENO}(\mathbf{F}_{j-2}^+, \mathbf{F}_{j-1}^+, \mathbf{F}_j^+), \quad (3.15)$$

$$\mathbf{F}_{j-\frac{1}{2}}^- = \text{WENO}(\mathbf{F}_{j+1}^-, \mathbf{F}_j^-, \mathbf{F}_{j-1}^-). \quad (3.16)$$

Finally, we can write the total net flux through the j^{th} cell as the difference between the fluxes at each of its cell faces:

$$\mathbf{F}_j = \mathbf{F}_{j-\frac{1}{2}} - \mathbf{F}_{j+\frac{1}{2}}. \quad (3.17)$$

3.5 Time Integrators

Time integrators advance the state of a numerical system forward in time. They are applied to an ordinary differential equation (ODE) or partial differential equation (PDE) of the form

$$\frac{d\mathbf{Q}(t, \mathbf{x})}{dt} = \mathcal{L}(\mathbf{Q}(t, \mathbf{x})) \quad \text{or} \quad \frac{\partial \mathbf{Q}(t, \mathbf{x})}{\partial t} = \mathcal{L}(\mathbf{Q})(t, \mathbf{x}), \quad (3.18)$$

where $\mathcal{L}(\mathbf{Q})$ is the sum of the flux and source terms, known colloquially as the ‘right-hand side’ (RHS) of the system. There are different classes of scheme used to solve this system, useful in different regimes.

The most common distinction is that of explicit vs. implicit. If $\mathbf{Q}(t, \mathbf{x})$ is the current state of the system, then at a small timestep Δt later, the solution is simply $\mathbf{Q}(t + \Delta t, \mathbf{x})$. An explicit method solves the equation

$$\mathbf{Q}(t + \Delta t, \mathbf{x}) = F(\mathbf{Q}(t, \mathbf{x})) \quad (3.19)$$

whilst an implicit method solves the equation

$$G(\mathbf{Q}(t, \mathbf{x}), \mathbf{Q}(t + \Delta t, \mathbf{x})) = 0. \quad (3.20)$$

In words, an explicit method performs an update to the next time using an algorithm that relies on information wholly known at the current time. Whereas an implicit method performs an update by using a root-finding algorithm involving information at the current *and* the next time. The root is then the solution at the next time.

3.5.1 Explicit Schemes

Some of the most popular explicit schemes for discretized time integration are Runge-Kutta (RK) methods, which involve a series of Euler substeps that together amount to a whole timestep. Many have have proven to be strong stability preserving (Gottlieb and Shu, 1998; Gottlieb, Ketcheson, et al., 2009) and can be written in the general form

$$\mathbf{Q}^{(0)} = \mathbf{Q}^n, \quad (3.21a)$$

$$\mathbf{Q}^{(i)} = \sum_{k=0}^{i-1} \left[\alpha_{ik} \mathbf{Q}^{(k)} + \Delta t \beta_{ik} \mathcal{L}(\mathbf{Q}^{(k)}) \right], \quad (3.21b)$$

$$\mathbf{Q}^{n+1} = \mathbf{Q}^{(m)} \quad (3.21c)$$

where $i = 1, 2, \dots, m$ and m is the number of sub-timesteps taken by the RK scheme, whilst α_{ik}, β_{ik} are positive constants. One example of an “RK2” scheme (second-order, Runge-Kutta) is given by

$$\mathbf{Q}^{(0)} = \mathbf{Q}^n, \quad (3.22a)$$

$$\mathbf{Q}^{(1)} = \mathbf{Q}^{(0)} + \Delta t \mathcal{L}(\mathbf{Q}^{(0)}), \quad (3.22b)$$

$$\mathbf{Q}^{n+1} = \mathbf{Q}^{(2)} = \frac{1}{2} \left[\mathbf{Q}^{(0)} + \mathbf{Q}^{(1)} + \Delta t \mathcal{L}(\mathbf{Q}^{(1)}) \right] \quad (3.22c)$$

such that the non-zero coefficients are $\alpha_{10} = 1, \alpha_{20} = \alpha_{21} = \frac{1}{2}$ and $\beta_{10} = 1, \beta_{21} = \frac{1}{2}$. METHOD has a number of higher-order integration techniques implemented into it, such as third-, fourth- and fifth-order RK schemes. For non-stiff sources, we may simply sum the numerical flux approximation and source terms to form the ‘right-hand side’ ($\mathcal{L}(\mathbf{Q})$).

Butcher tableaus can be used to summarise Runge-Kutta methods, both explicit and implicit. They provide a table for the coefficients of the form where for explicit

c_1	a_{11}	a_{12}	\dots	a_{1s}
c_2	a_{21}	a_{22}	\dots	a_{2s}
\vdots	\vdots	\vdots	\ddots	\vdots
c_s	a_{s1}	a_{s2}	\dots	a_{ss}
	b_1	b_2	\dots	b_s

TABLE 3.1: Butcher tableau table used to concisely summarise Runge-Kutta methods.

methods the coefficient matrix $[a_{ik}]$ is lower triangular. For example, the explicit forward-Euler method used in section 3.2 has the Butcher tableau shown in table 3.2 whilst the implicit backward-Euler method has the tableau shown in table 3.3. A more complex case is the ‘classic’ RK4 method which has the tableau shown in table 3.4.

$$\begin{array}{c|c} 0 & 0 \\ \hline & 1 \end{array}$$

TABLE 3.2: Butcher tableau for the explicit forward-Euler method.

$$\begin{array}{c|c} 1 & 1 \\ \hline & 1 \end{array}$$

TABLE 3.3: Butcher tableau for the implicit backward-Euler method.

$$\begin{array}{c|ccc} 0 & & & \\ 1/2 & 1/2 & & \\ 1/2 & 0 & 1/2 & \\ 1 & 0 & 0 & 1 \\ \hline & 1/6 & 1/3 & 1/3 & 1/6 \end{array}$$

TABLE 3.4: Butcher tableau for the ‘classic’, explicit RK4 method.

3.5.2 Implicit Schemes

Firstly, let us demonstrate the power of implicit methods when the timescale that the ‘right-hand side’ (RHS) acts upon is very small. Taking our simple ODE of the form

$$\frac{d\mathbf{Q}}{dt} = \mathcal{L}(\mathbf{Q}) \quad (3.23)$$

we may rewrite it as

$$\frac{d\mathbf{Q}}{dt} \sim -\frac{1}{\tau}\mathbf{Q} \quad (3.24)$$

where τ is the timescale upon which the RHS acts. A single *explicit* Euler step gives a recursive update equation of the form

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n - \frac{\Delta t}{\tau}\mathbf{Q}^n \quad (3.25)$$

such that the final solution after N timesteps is

$$\mathbf{Q}^N = \left(1 - \frac{\Delta t}{\tau}\right)^N \mathbf{Q}^0. \quad (3.26)$$

From this expression we can see that the solution diverges when $\tau \ll \Delta t$ and for stability we require $\tau \gtrsim \Delta t$. Sometimes this necessitates us to take such small timesteps that the computational cost of the simulation makes it unfeasible. Let us contrast this result with that where, instead, we take *implicit* Euler steps. Now,

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n - \frac{\Delta t}{\tau}\mathbf{Q}^{n+1} \quad (3.27)$$

which gives the final form

$$\mathbf{Q}^N = \frac{\mathbf{Q}^0}{\left(1 + \frac{\Delta t}{\tau}\right)^N}. \quad (3.28)$$

The solution is now bounded and the scheme is stable, unconditionally. However, large timesteps will still lead to large numerical errors. Many implicit schemes require a root-find for the solution at each timestep because they produce a non-linear equation for the update step. This means they can become very costly but are still generally superior to explicit schemes for fast-acting fluxes and/or sources.

3.5.3 IMEX Schemes

In many systems, physical processes are occurring on a wide range of timescales. For instance, the reactions taking place in the hot nuclear matter created in a BNS merger are occurring many orders of magnitude faster than the dynamical timescales of the bulk fluid motion. It can be difficult for numerical methods to simultaneously capture both the fast- and slow-acting processes in a stable and accurate way. For problems like these, a number of semi-implicit time-integrators known as IMEX (Implicit-Explicit) schemes have been developed (Pareschi and Russo, 2005). These evaluate the minimal required set of ‘fast terms’ using an implicit scheme and the rest using an explicit scheme. For our purposes, this typically means evaluating the fluxes explicitly and source terms implicitly. The large, fast-acting source terms we will encounter are due to dissipative fluid effects. However, similar issues arise when modelling electromagnetism or radiation paired with hydrodynamics.

Consider again the general ODE update form

$$\frac{d\mathbf{Q}}{dt} = \mathcal{L}(\mathbf{Q}). \quad (3.29)$$

We can split the right-hand-side update term $\mathcal{L}(\mathbf{Q})$ into a slow-acting flux piece $\mathcal{F}(\mathbf{Q})$ and a fast-acting source term $\mathcal{S}(\mathbf{Q})$, which acts on a timescale τ , such that

$$\frac{d\mathbf{Q}}{dt} = \mathcal{F}(\mathbf{Q}) + \mathcal{S}(\mathbf{Q}). \quad (3.30)$$

Schematically, the IMEX update then looks like

$$\mathbf{Q}^{(i)} = \mathbf{Q}^n + \Delta t \left[\sum_{j=1}^{i-1} a_{ij} \mathcal{F}(\mathbf{Q}^{(j)}) + \sum_{j=1}^s \bar{a}_{ij} \mathcal{S}(\mathbf{Q}^{(j)}) \right], \quad (3.31)$$

$$\mathbf{Q}^{(i)} = \mathbf{Q}^n + \Delta t \left[\sum_{j=1}^s b_j \mathcal{F}(\mathbf{Q}^{(j)}) + \sum_{j=1}^s \bar{b}_j \mathcal{S}(\mathbf{Q}^{(j)}) \right] \quad (3.32)$$

where a_{ij}, \bar{a}_{ij} are the elements of an $s \times s$ matrix and b_j, \bar{b}_j are the elements of a vector of length s , with s being the total number of implicit stages used in the scheme. In general, we may classify a strong-stability-preserving IMEX scheme using the notation $\text{SSP}k(s, \sigma, p)$ where k is the order of the SSP scheme and s and σ are (respectively) the number of implicit and explicit stages of the IMEX scheme, whilst p is its order of

convergence in Δt . The coefficients are derived by comparing the solution with the Taylor expansion of the exact solution, as in other RK methods.

METHOD implements a number of higher-order IMEX schemes such as the SSP2(2,2,2) implicit scheme which has the following tableaux for the explicit coefficients a_{ij}, b_j and the implicit coefficients \bar{a}_{ij}, \bar{b}_j , respectively. These lead to the

0	0	0
1	1	0
<hr/>		
	1/2	1/2

TABLE 3.5: Butcher tableau of the explicit coefficients a_{ij}, b_j for the SSP2(2,2,2) IMEX scheme.

c	c	0
$1-c$	$2-c$	c
<hr/>		
	1/2	1/2

TABLE 3.6: Butcher tableau of the implicit coefficients \bar{a}_{ij}, \bar{b}_j for the SSP2(2,2,2) IMEX scheme. $c = 1 - \frac{1}{\sqrt{2}}$.

evolution equation

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n + \Delta t [\mathcal{F}(\mathbf{Q}^{(1)}) + \mathcal{F}(\mathbf{Q}^{(2)})] + \Delta t [\mathcal{S}(\mathbf{Q}^{(1)}) + \mathcal{S}(\mathbf{Q}^{(2)})] \quad (3.33)$$

in which

$$\mathbf{Q}^{(1)} = \mathbf{Q}^n + \gamma \Delta t \mathcal{S}(\mathbf{Q}^{(1)}), \quad (3.34a)$$

$$\mathbf{Q}^{(2)} = \mathbf{Q}^n + \Delta t \mathcal{F}(\mathbf{Q}^{(1)}) + \Delta t [(1 - 2\gamma) \mathcal{S}(\mathbf{Q}^{(1)}) + \gamma \mathcal{S}(\mathbf{Q}^{(2)})]. \quad (3.34b)$$

Determining the solution to equation 3.33 requires a non-linear root-find due to the construction of the implicit scheme. The most common tool for this purpose is an N -dimensional secant method (Balsara, 1994; Palenzuela, Lehner, et al., 2009), where an estimate of the stage one and two solutions are given, $\mathbf{Q}^{(1,2)*}$, and the Jacobian is approximated via a finite-difference method. A satisfactory estimate for the interstage solutions is the solution at the previous timestep, \mathbf{Q}^n .

It is important to notice that the implicit source term is a function of the guess $\mathbf{Q}^{(1,2)*}$, and that the source is computed using both the values of the conserved and primitive variables. As a result, not only must one determine the primitive variables at each time step for the source and flux contribution of equation 3.33, but also for every new estimate of the solution used in the root-find presented in equations 3.34 until a root has been found. The recovery of the primitive variables is outlined in section 3.8 and model specific, but for dissipative, relativistic fluids will involve an additional four-dimensional Newton method. As a result of this nested root-find, the IMEX

scheme naturally requires a significant amount of additional computation compared to fully explicit schemes.

We can see how the need for a semi-implicit integration scheme has dramatically increased the required amount of computation for each time-step. It is also worth noting that the size of the conserved vector is problem dependent. Models of an ideal, single-fluid have size $N = 5$, whereas the non-ideal Müller-Israel-Stewart model we consider here has size $N = 15$. One advantage of the other non-ideal models we consider is that they have size $N = 5$ as in the ideal case, but with increased complexity. Multi-fluid models or the inclusion of GR will increase the system size further. For these larger systems, the scheme described above will require substantially more computation per timestep, especially for larger systems of equations, but in stiff regions should relax the CFL condition to such a degree that the total execution time is still less than for explicit schemes.

See Takamoto and Inutsuka, 2011 for an example of a closely related but somewhat different approach to evolving the stiff equations of the MIS model. There, they use a Strang-splitting (Strang, 1968) technique to separate the equations of motion into their inviscid and dissipative parts. The assumption that the dissipative relaxation timescales are *very* short allows one to approximately solve the relaxation part of the dissipative evolution equations analytically. This means that they can avoid using an implicit solver for the fast-acting dissipative part of the system, and the inviscid part of the system may also be evolved using explicit time integration as usual.

3.6 Boundary Conditions

Any computational domain is finite in size and often comes to a rather abrupt end at the edges of the space it covers. Whilst in reality the physical system would define what happens at its boundaries, in simulations we must try to capture these ‘boundary conditions’ (BCs) numerically.

This typically involves padding the *actual* domain with a layer of several ‘ghost cells’ that are used to impose BCs in one of a number of ways. With METHOD, we make use of two types of BCs: outflow and periodic.

Outflow BCs simply copy the final edge cell’s value within the domain into all of the ghost cells outside that edge. Mathematically, this means that with a rectangular domain, the ghost cells padded in the x -direction obey

$$U_{-N_g+i,j,k} = U_{1,j,k} \quad (3.35)$$

at the left edge and

$$U_{N_x+i,j,k} = U_{N_x,j,k} \quad (3.36)$$

at the right edge, for $i = 1, 2, \dots, N_g$ and $1 \leq j \leq N_y$, $1 \leq k \leq N_z$ where N_g is the number of ghost zones. The same copying is performed in the y, z directions.

Periodic BCs mean that the domain effectively wraps-round onto itself, and the boundaries of opposite sides of the domain are connected. The ghost cells of one side are filled with the values of variables from the cells inside the domain on the opposite side, and vice versa. Similarly, this gives

$$U_{-i,j,k} = U_{N_x-i,j,k} \quad (3.37)$$

at the left edge and

$$U_{N_x+i,j,k} = U_{i,j,k} \quad (3.38)$$

at the right edge. The same mirroring is performed in the y, z directions.

For a one-dimensional domain, the number of ghost cells is simply $2N_g$. For a two-dimensional domain, it is $2(N_x + N_y)N_g$ and for a three-dimensional domain it is $2(N_xN_yN_g + N_xN_zN_g + N_yN_zN_g)$. One can see that the amount of memory allocated to ghost cells scales rapidly with dimensionality and domain size. Still, the total memory allocated for the ghost cells is generally small compared to the memory required for the bulk of the domain ($\sim N_g/N_X$).

However, if the domain is split up into sections that are worked on separately and concurrently by processors, the ghost zones and their memory usage present a significant complication and problem. This is because every section needs to have access to its neighbours' adjacent cells, effectively increasing the number of ghost cells dramatically and, crucially, requiring their cell values to be frequently communicated between the processors.

3.7 Parallel Programming

The immense computational cost of BNS merger simulations necessitates the splitting up of the total workload between multiple computer processor units (CPUs) that work simultaneously, typically on separate patches of the spatial domain, which are then stitched together again to produce the final output. Indeed, even the flat spacetime, hydrodynamic-only simulations we present the results of in this thesis are run using parallel computation such that they take hours rather than days to run.

METHOD is built with both [Open Multi-Processing \(OpenMP\)](#) and [Message Passing Interface \(MPI\)](#) software. This means that on the Iridis5 Supercomputer here at the University of Southampton, we may take advantage of the compute cluster architecture to run on many multi-core nodes at once. OpenMP handles parallelism between cores within a node, whilst MPI handles parallelism between nodes.

We typically perform parallel runs on 40 physical cores. Using 40 cores doesn't quite give a 40-fold speed-up, mainly because of the additional overheads associated with communication between cores. In fact, for very large scale (such as exascale) simulation, this becomes the limiting factor in achieving speed-up with parallel computing.

METHOD also implements the [Compute Unified Device Architecture \(CUDA\)](#) software to utilize the power of Graphical Processing Units (GPUs). GPUs allow one to very efficiently perform an operation on many individual pieces of data simultaneously - an example of the Single Instruction Multiple Data (SIMD) parallel processing paradigm. Not all parts of a code are suitable for parallelization in this way but METHOD implements routines to use GPUs for the conserved-to-primitive transformation, time-integration and flux reconstruction/calculation.

3.8 Primitive Variable Recovery

Another crucial aspect of any numerical evolution scheme is known as primitive variable recovery or, colloquially, the "C2P" (conserved-to-primitive) scheme. This is how one recovers the primitive fluid variables, in our case $\mathbf{w} = \{p, \rho, n, v_1, v_2, v_3, \Pi, q^\mu, \pi_\nu^\mu\}$, from the state vector, $\mathbf{U}(\mathbf{w})$, so that we may then calculate the succeeding flux vector, $\mathbf{F}(\mathbf{w})$. In some cases, this may be done using a closed-form analytic expression - see Pandya and Pretorius, [2021](#) or Pandya, Most, et al., [2022b](#) for instance where, within the BDNK formulation of first-order hydrodynamics, use of a conformal fluid with a high degree of symmetry makes this possible. In most cases, however, it requires a numerical iteration scheme: one guesses the value of one of the primitive variables, in our case often pressure or a pressure-like term, \tilde{p} ; next, one uses the guess and a scheme involving the conserved vector's elements to solve for the suite of primitive variables; this includes a calculated value for the guessed quantity, \bar{p} , against which one can compare the initial guess. Typically, a gradient descent root-finding method is used that solves for the zero of $f(\tilde{p}) = \tilde{p} - \bar{p}$. In METHOD, a Newton-secant method is used.

Unsurprisingly, much work has gone into producing primitive recovery algorithms that are robust (stable for a wide parameter space) and computationally inexpensive (converge in a reasonable number of iterations). Primitive recovery may be avoided entirely by using the approach taken in Aguayo-Ortiz, Mendoza, et al., [2018](#), where the C2P scheme is swapped for algebraic manipulation of the balance-law form, but this is accurate only to first-order in the primitive variables and non-convergent at shocks. There has also been (perhaps unsurprisingly) recent work where machine learning methods are used to perform primitive variable recovery (Dieselhorst, Cook, et al., [2021](#)). This shows promising results, for the exact equation of state we use here in fact!

The form of our conserved vectors for the models studied in this thesis suggested a simple extension to the common recovery scheme used for the Euler equations (where the pressure is guessed) but we have also implemented (see section 4.2.1) a modified version of an alternative scheme from Kastaun, Kalinani, et al., 2021 which is further explored in Kalinani, Ciolfi, et al., 2022. These schemes were developed in response to recovery schemes failing for fast flows (and/or with strong magnetic fields present) and have the advantages of proven single-limit convergence and the existence of a root within a known interval. We have a testing setup for our recovery schemes that checks convergence for parameter values covering the realm of physically-plausible ones by beginning with known conserved quantities and ensuring we arrive back at them after primitive recovery.

3.9 Summary

More details about the precise numerical schemes we employ will be given in chapter 4 once the models of hydrodynamics that we work with have been introduced. Our main difficulty will be in resolving the physics occurring over a wide range of spatial and temporal scales. Macroscopic fluid motion occurs at acoustic speeds governed by the Jacobian of the flux vector ($|\partial \mathbf{F} / \partial \mathbf{U}|$) and our timestep limitations here will be set by meeting the CFL condition necessary for stable evolution.

Microscopic dissipation typically occurs on much shorter timescales. From a physical standpoint, these timescales should be set by the specific dissipation mechanism occurring (or at least the one being modelled) which in turn sets the parameter values used in the dissipative fluid description. For example, in the case of shear viscosity, causality will naturally enforce a lower bound on the dissipation timescale of $\tau \sim (\rho + p) / \eta$.

Here, stiffness becomes the issue. The Müller-Israel-Stewart model introduced in section 2.2.1.5 (with simulation results in section 4.2) has evolution equations for the dissipative variables with source terms proportion to $1/\tau$, the reciprocal of small timescales. As seen earlier in section 3.5.2, numerical instability is introduced whenever $\Delta t < \tau$. Using implicit time-integration methods can provide stability, but overstepping the dissipation timescale still leads to causality violation and inaccurate results. The Chapman-Enskog expansion, applied to the MIS model, does remedy this problem but is only appropriate in the small-timescale limit and is only ever an approximate model, albeit an accurate one.

Finally, the BDNK formulation introduced in section 7.1 ostensibly avoids this latter constraint by not evolving the dissipative variables and instead provides constitutive expressions for them. However, the time derivatives of primitive variables are solved for in the conserved-to-primitive procedure and are proportional to either $\chi \equiv \{\zeta, \kappa, \eta\}$

(the dissipation strengths) or reciprocal timescales, $1/\tau_X$, multiplied by changes in the conserved quantities. Because of this, the stiffness is actually moved into the conserved-to-primitive transformation, making it far more sensitive and unreliable. This may still force us to take shorter timesteps in order to reduce the rate of evolution of the conserved variables, coupling the two disparate systems once again.

Chapter 4

Simulations with the Müller-Israel-Stewart Model

This chapter presents results from our simulations of special-relativistic, dissipative fluids using the formalism of Müller-Israel-Stewart Israel, 1976; Israel and Stewart, 1979.

4.1 Reduced MIS Models

In section 4.2, we will simulate the ‘full’ Müller-Israel-Stewart equations of motion. However, it is instructive to develop a series of ‘reduced’ models beforehand. These give us stepping stones to jump to along the way and let us gain insight into the effects of each physical mode of dissipation in turn. In order to derive these reduced models, we must first present the full MIS equations, and then apply simplifications to them.

We follow the considerable amount of previous work, such as that of Chabanov, Rezzolla, et al., 2021; Takamoto and Inutsuka, 2011; Del Zanna, Chandra, et al., 2013; Du and Heinz, 2020; Peitz and Appl, 1998; Peitz and Appl, 1999, and write the equations of motion out here in part to fix notation. A balance-law equation is of the form

$$\partial_t \mathbf{U} + \partial_i \mathbf{F}^{(i)} = \mathbf{S}, \quad (4.1)$$

and in particular for the MIS model by

$$\mathbf{U} = \begin{pmatrix} D \\ S_j \\ \tau \\ U \\ Y_j \\ Z_{jk} \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi)W^2 v_j + W(q_0 v_j + q_j) + \pi_{0j} \\ (\rho + p + \Pi)W^2 + 2q_0 W - (p + \Pi - \pi_{00}) - nW \\ nW\Pi \\ nWq_j \\ nW\pi_{jk} \end{pmatrix}, \quad (4.2a)$$

$$\mathbf{F}^{(i)} = \begin{pmatrix} Dv^i \\ S_j^i \\ S^i - Dv^i \\ Uv^i \\ Y_j v^i \\ Z_{jk} v^i \end{pmatrix}, \quad (4.2b)$$

$$\mathbf{S} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{n}{\tau_\Pi}(\Pi_{\text{NS}} - \Pi) \\ \frac{n}{\tau_q}(q_{j,\text{NS}} - q_j) \\ \frac{n}{\tau_\pi}(\pi_{jk,\text{NS}} - \pi_{jk}) \end{pmatrix} \quad (4.2c)$$

where we have followed the convention of the majority of the referenced papers and ignored second-order (in derivatives) terms in the sources. Here we have $\rho = mn(1 + \epsilon)$ being the energy density, m the mass per baryon (set equal to unity from now onwards as we consider only single-species fluids), n the number density, and ϵ the specific internal energy. The hydrodynamic pressure p is given by an equation of state to close the system - here it is most convenient to calculate it as $p \equiv p(n, \rho)$. For simplicity, we fix a gamma-law equation of state such that the pressure is given by

$$\begin{aligned} p &= (\Gamma - 1)mne \\ &= (\Gamma - 1)(\rho - mn). \end{aligned} \quad (4.3a)$$

We also have the spatial three-velocity v_j and the Lorentz factor $W = (1 - v_j v^j)^{-1/2}$, which make up the four-velocity $u_\mu = W(1, v_j)$. The non-ideal terms are the bulk viscous pressure Π , the heat flux q_j , and the shear viscosity π_{jk} . The heat flux and shear viscosity are orthogonal to the four velocity on all indices, and the shear viscosity is trace free, implying

$$q_0 = v^k q_k, \quad (4.4a)$$

$$\pi_{0j} = v^k \pi_{kj}, \quad (4.4b)$$

$$\pi_{j0} = v^k \pi_{jk}, \quad (4.4c)$$

$$\pi_{00} = -\pi_k^k. \quad (4.4d)$$

We need the spatial projections of the stress-energy tensor to get the fluxes, which gives us the remaining flux term

$$S_j^i = (\rho + p + \Pi)W^2 v^i v_j + W(q^i v_j + q_j v^i) + (p + \Pi)\delta_j^i + \pi_j^i. \quad (4.5)$$

Two more important thermodynamic quantities of interest are the specific enthalpy $h = 1 + \epsilon + p/(mn) = (\rho + p)/(mn)$ and the temperature which is given simply by $T = p/n$ for our chosen EoS. We also define the thermodynamic inverse temperature $\beta = m/T$. From these we can link the non-ideal timescales, $\tau_X = \{\tau_\Pi, \tau_q, \tau_\pi\}$, to the non-ideal coefficients $\chi = \{\zeta, \kappa, \eta\}$ through the relations

$$\tau_\Pi = \zeta\beta_0, \quad (4.6a)$$

$$\tau_q = \kappa T\beta_1, \quad (4.6b)$$

$$\tau_\pi = 2\eta\beta_2 \quad (4.6c)$$

where (see Hiscock and Lindblom, 1983, for example)

$$\Omega = 3\Gamma - 5 + \frac{3\Gamma}{h\beta}, \quad (4.7a)$$

$$\Omega^* = 5 - 3\Gamma + 3(10 - 7\Gamma)\frac{h}{\beta}, \quad (4.7b)$$

$$\Omega^{**} = 5 - 3\Gamma + \frac{3\Gamma^2 h^2 \beta^2}{\Gamma - 1}, \quad (4.7c)$$

$$\beta_0 = \frac{3\Omega^*}{h^2 \Omega^2 p}, \quad (4.7d)$$

$$\beta_1 = \left(\frac{\Gamma - 1}{\Gamma}\right)^2 \frac{\beta}{hp} \left(5h^2 - \frac{\Gamma}{\Gamma - 1}\right), \quad (4.7e)$$

$$\beta_2 = \frac{1 + 6h/\beta}{2h^2 p}. \quad (4.7f)$$

These parameters set the scaling relations between the dissipation strengths and timescales. Whilst we have implemented these thermodynamic expressions and used them in earnest, we find that for the purposes of numerical exploration and a clearer control of our simulation parameters, we set the values of the timescales and strengths independently, by-hand.

We need to specify the first-order ‘‘Navier-Stokes’’ terms to which the dissipative system relaxes. These are

$$\Pi_{\text{NS}} = -\zeta\Theta, \quad (4.8a)$$

$$q_{j,\text{NS}} = -\kappa T(\partial_j \log T + a_j), \quad (4.8b)$$

$$\pi_{jk,\text{NS}} = -2\eta\sigma_{jk}. \quad (4.8c)$$

These are defined in terms of the non-ideal coefficients, the temperature, and first order derivatives of existing variables. Specifically,

$$\Theta = \partial_\mu u^\mu, \quad (4.9a)$$

$$a_\mu = u^\nu \partial_\nu u_\mu, \quad (4.9b)$$

$$\sigma_{\mu\nu} = \left(\partial_\mu u_\nu + \partial_\nu u_\mu - \frac{2}{3} \eta_{\mu\nu} \Theta \right) \quad (4.9c)$$

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

We have studied three reduced models: one for the heat flux; another for the bulk viscosity; and a third for the shear viscosity. For each, we make a number of simplifying assumptions before arriving at a two-equation model with a stiff source term present in one of the equations. We then perform the Chapman-Enskog expansion which gives a single equation that governs the evolution of the variable of interest for that particular reduced model. In our case, this is either the temperature or the velocity. Lastly, we make a comparison between numerical results obtained from simulating the toy model's original pair of equations with its Chapman-Enskog expansion equation.

4.1.1 Heat Flux

For our first reduced model, we want to arrive at a system of equations that govern two intrinsically linked quantities: the temperature, T , of our fluid and the heat flux, q , within our fluid. 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 $u_\mu q^\mu = 0$ 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

$$\partial_t T + \partial_i q^i = 0, \quad (4.10a)$$

$$\partial_t q_j = \frac{1}{\tau_q} (q_{j,\text{NS}} - q_j). \quad (4.10b)$$

We note that the acceleration term usually present in the heat-flux's source will vanish so that $q_{j,\text{NS}} \rightarrow -\kappa \partial_j T$ 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:

$$\partial_t T + \partial_i q^i = 0, \quad (4.11a)$$

$$\partial_t q_j = -\frac{1}{\tau_q} (\kappa \partial_j T + q_j). \quad (4.11b)$$

4.1.1.1 Limiting case

We pause briefly to analyse the equilibrium behaviour of our toy model equations. In the limit of infinitely fast dissipation, $\tau_q \rightarrow 0$, the heat flux also adopts its equilibrium value so $q_j \rightarrow q_{j,\text{NS}} = -\kappa \partial_j T$. This means the toy model reduces to the single equation

$$\partial_t T = \kappa \partial_i \partial^i T. \quad (4.12)$$

This is the heat equation, which is the standard linear parabolic equation, and the limit that we expect to find. Some key features of interest are that the solutions are continuous for $t > 0$ irrespective of the initial data, and that the speed of propagation of information is infinite. This latter point is problematic for a theory that is meant to represent something relativistic.

4.1.1.2 Cattaneo equation

Let us examine the speeds of wave propagation more closely. If we take the original toy model (4.11) then we can rearrange it into a single second order equation as

$$\partial_{tt} T = -\partial_{ti} q^i = \frac{1}{\tau_q} \left(\kappa \partial_i \partial^i T - \partial_t T \right). \quad (4.13)$$

This is a hyperbolic equation, the telegraph or Cattaneo equation (Cattaneo, 1948).

We can compute the wave speeds of this equation by looking for travelling wave solutions of the form $T \propto \exp\{i(\omega t + k^l x_l)\}$. The resulting dispersion relation is

$$\tau_q \omega^2 = \kappa k^2 + i\omega. \quad (4.14)$$

The solutions are

$$\omega = \frac{1}{2\tau_q} \left(i \pm \sqrt{-1 + 4\kappa\tau_q k^2} \right). \quad (4.15)$$

As $\kappa, \tau_q > 0$ this should always be stable (the imaginary part is always positive). The wavespeed (phase velocity) is given by the real part divided by the wave number. There are two cases. When $k^2 < 1/(4\kappa\tau_q)$ the real part vanishes and the mode is pure imaginary, and damped. For larger k the wavespeed is

$$\frac{|\text{Re}\{\omega\}|}{k} = \sqrt{\frac{\kappa}{\tau_q} - \frac{1}{4k^2\tau_q^2}}. \quad (4.16)$$

We see that this is bounded by $\sqrt{\kappa/\tau_q}$, so the wavespeed is finite as expected. At the cross-over point ($4k^2\kappa\tau_q = 1$) the wavespeed vanishes.

The outcome of this is that we should expect $\kappa \lesssim \tau_q$ in order to get a *physical* (subluminal) wavespeed. This looks reasonable when we consider the thermodynamic

relation between the dissipation size coefficient, κ , and the dissipation timescale, τ_q , given in equation (4.6). This finding is analogous to that found for the full MIS system with bulk viscosity by Bemfica, Disconzi, et al., 2019a, where the existence of a minimum (and resolved) scale for τ_Π is necessary in order to have well-posed evolution equations.

Consider the wavespeeds for the whole (T, q) system, which we find by eigenvalue analysis of equations (4.11) – see section 4.1.2 for more detail on this method. They are zero when only the LHS is considered and $\sqrt{\frac{\kappa}{\tau_q}}$ when the derivative on the RHS in the definition of $q_{j,NS}$ is included. Again, this shows why $\frac{\kappa}{\tau_q}$ must be sensibly bounded. One possible numerical approach is to pretend that the timestep limit for the hyperbolic step is set by the left-hand-side alone, and that some semi-implicit update scheme will allow us to bypass the limit from the source term (even though it contains a derivative). The question, for practical cases, is whether that is true: for the tests considered here, it does seem to be.

4.1.1.3 Chapman-Enskog expansion

Restricting to one dimension for now, we will perform the Chapman-Enskog expansion to reduce the two-equation system given by eq. (4.10a) & eq. (4.10b) to a single equation.

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 correction term of order τ_q , so

$$q_j = q_{j,NS} + \tau_q q_j^{(1)}. \quad (4.17)$$

We can then write the pair of equations, to first order in τ_q , as

$$\partial_t T + \partial_i q_{NS}^i = -\partial_i \left(\tau_q q^{(1),i} \right), \quad (4.18a)$$

$$\partial_t q_{j,NS} = q_j^{(1)}. \quad (4.18b)$$

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

$$\partial_t T = \partial_i \left(\kappa \partial^i T \right) - \partial_i \left(\tau_q q^{(1),i} \right), \quad (4.19a)$$

$$-\partial_t (\kappa \partial_j T) = q_j^{(1)} \quad (4.19b)$$

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

For simplicity we will first assume that κ and τ_q are constants in time and space. By commuting the temporal and spatial derivatives in eq. (4.19b), we can now substitute the leading order form (zero'th order in τ_q) of the equation of motion for T ,

eq. (4.19a), into the relaxation equation eq. (4.19b), to determine the correction $q_j^{(1)} = \kappa^2 \partial_j^{(3)} T$ as purely spatial derivatives. Inserting this result back into eq. (4.19a), and writing the result in one spatial dimension, we finally have the CE form

$$\partial_t T = \kappa \left[\partial_x^{(2)} T + \kappa^2 \beta_q \partial_x^{(4)} T \right] + \mathcal{O}(\tau_q^2). \quad (4.20)$$

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 eq. (4.11). This reduced model is non-stiff (as the source term is multiplied by the small timescale τ_q , not by its reciprocal), but may lead to other numerical problems due to the higher derivatives. In particular, like the heat equation, this form we obtain is also non-hyperbolic and is not Lorentz-boost invariant. 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.

Equation (4.20) is essentially a power series expansion in $\{\kappa, \tau\}$, at least while κ and τ_q are treated as constant. When we (more realistically) let the coefficients depend on T , we get

$$\partial_t T = \partial_x (\kappa \partial_x T) + \partial_x \{ \tau_q [\kappa \partial_{xx} (\kappa \partial_x T) + \partial_T \kappa \partial_x T \partial_x (\kappa \partial_x T)] \}. \quad (4.21)$$

The constant coefficient form is linear and in a diffusion-retention equation form. 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, et al., 2011).

4.1.1.4 Numerical Results

We now look at numerical results for the systems developed above. We use both an implicit and an explicit integration scheme within the METHOD codebase to generate these results. An explicit scheme without the Chapman-Enskog expansion would not be stable for small values of κ and τ_q , so an implicit scheme is used. See section 3.5.1 and section 3.5.2 for a discussion of time-integration schemes.

The generic behaviour of the full system is shown in a two dimensional simulation in figure 4.1. To directly compare the results of the Chapman-Enskog approximation and the full scheme, we look at a one dimensional slice through these two dimensional results. This is shown in figure 4.2. The results are visually indistinguishable, which is encouraging. As the Chapman-Enskog expansion has eliminated the stiff source term, the computation is significantly sped up. This will be quantified further in the toy model for the bulk viscous pressure, section 4.1.2.4.




FIGURE 4.1: The spreading of a spherically-symmetric hot “blob” by the heat flux, using equations (4.11). The 2D grid is 128×128 cells in size and the simulation runs until $t_{\text{final}} = 5$ with $\kappa = 1 \times 10^{-3}$ and $\tau_q = 1 \times 10^{-3}$.




FIGURE 4.2: 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 Chapman-Enskog expansion. There is excellent visual agreement. This is a 1D-slice view through the 2D domain seen in figure 4.1.

4.1.2 Bulk Viscosity

Next, we will derive a toy model for a fluid that exhibits bulk viscosity – that is one which resists the shearless compression or expansion of a fluid. One may gain a better visual image of the action of bulk viscosity by considering its form in the Navier-Stokes equations: $\Pi_{\text{NS}} = -\zeta \partial_\mu u^\mu$; where the negative sign indicates its resistance to divergences in the fluid velocity, scaled by a coefficient.

We set the shear and heat terms to vanish, but rather than working with a static fluid, we work in the slow motion limit so that $W \rightarrow 1$, but the 3-velocity does not vanish – a necessary condition because the bulk viscosity depends on gradients in the fluid velocity. In addition, we assume that $\Pi \ll \rho + p$ and n, ρ, p are all constant, giving great simplification. In this case the continuity and energy equations become trivial, and the momentum terms reduce as: $S_j \rightarrow v_j$; $U \rightarrow \Pi$; $S_j^i \rightarrow v^i v_j + \Pi \delta_j^i$. We are left with

$$\partial_t v_j + \partial_i (v^i v_j + \Pi \delta_j^i) = 0, \quad (4.22a)$$

$$\partial_t \Pi + \partial_i (\Pi v^i) = \frac{1}{\tau_\Pi} (\Pi_{\text{NS}} - \Pi). \quad (4.22b)$$

Finally we note that $\Theta \rightarrow \partial_k v^k$, giving

$$\partial_t v_j + \partial_i (v^i v_j + \Pi \delta_j^i) = 0, \quad (4.23a)$$

$$\partial_t \Pi + \partial_i (\Pi v^i) = -\frac{1}{\tau_\Pi} (\zeta \partial_k v^k + \Pi). \quad (4.23b)$$

This is again a closed system with the same size as the heat flux case. It is slightly more complex as the fluxes are nonlinear.

4.1.2.1 Limiting case

Again we look at the limiting case as $\tau_\Pi \rightarrow 0$, giving $\Pi \rightarrow -\zeta \partial_k v^k$. This reduces the toy model to

$$\partial_t v_j + \partial_i (v^i v_j) = \zeta \partial_{ik} (\delta_j^i v^k). \quad (4.24)$$

This is essentially the viscous Burger's equation, only missing a factor of half in the flux compared to its canonical form. In the limit as $\zeta \rightarrow 0$ we expect the standard behaviour of Burger's equation – shock formation being the key feature here. The dissipative term modifies this, reducing any gradients and leading to the standard travelling wave solutions. Our imposed simplifications seem to be consistent.

4.1.2.2 Chapman-Enskog expansion

We again summarize the expansion about the equilibrium $\Pi \rightarrow -\zeta \partial_x v$, which gives (to leading order in τ_Π and reducing to one dimension)

$$\partial_t v + \partial_x v^2 = \partial_x (\zeta \partial_x v) + \tau_\Pi \partial_x \left\{ \partial_x (\zeta v \partial_x v) - \zeta \partial_{xx} v^2 \right\}. \quad (4.25)$$

This simplifies slightly with constant ζ :

$$\partial_t v + 2v \partial_x v = \zeta \partial_{xx} v + \tau_\Pi \zeta (\partial_x v \partial_{xx} v - v \partial_{xxx} v). \quad (4.26)$$

The first term on the right hand side is a standard dissipation whilst the others are more complex mixed terms that contain elements of propagation ($\partial_x v$), dissipation ($\partial_{xx} v$) and dispersion ($\partial_{xxx} v$).

4.1.2.3 Wave Speeds

Let us determine the wave speeds of our bulk viscous toy model. We cannot take the same approach as that for the heat flux toy model because the governing equations are now non-linear in v and so determining a phase velocity from a dispersion relation will not work.

Instead, start by considering the limit $\zeta \rightarrow 0$ in equations (4.23), in which case the equilibrium value of Π , Π_{NS} , is automatically zero. We can see that equation (4.23b), which governs the evolution of Π , has a characteristic wavespeed of v and constantly drives Π towards its equilibrium value of zero. Then, from (4.23a), the characteristic wave speed for the velocity becomes $2v$. Given this, we expect the additional degree of freedom provided by the re-introduction of Π to give us two wave speeds, both functions of (Π, v) , where each speed is a deviation from the “equilibrium wave speed” values of $2v$ and v .

For a system of equations written in the vector form

$$\partial_t \mathbf{U} + \underline{\underline{\Lambda}}^{(i)} \partial_i \mathbf{U} = \mathbf{S}. \quad (4.27)$$

where Λ is in general a matrix, the characteristic speeds in the i 'th direction are given by the eigenvalues of $\underline{\underline{\Lambda}}^{(i)}$. With this in mind, let us take equations (4.23) and work in just one spatial dimension so that we have

$$\partial_t \begin{pmatrix} v \\ \Pi \end{pmatrix} + \begin{pmatrix} 2v & 1 \\ \Pi & v \end{pmatrix} \partial_x \begin{pmatrix} v \\ \Pi \end{pmatrix} = \begin{pmatrix} 0 \\ -\frac{1}{\tau_\Pi} (\zeta \partial_x v + \Pi) \end{pmatrix} \quad (4.28)$$

where $v \equiv v_x$. Notice that the source vector \mathbf{S} contains the derivative $\partial_x v$ but has been initially left on the RHS of (4.28) because it is the one-dimensional version of the

divergence term, Θ , which appears as a stiff source term that happens to contain derivatives in the velocity. Now, moving it across to be incorporated into the Λ matrix, we have

$$\partial_t \begin{pmatrix} v \\ \Pi \end{pmatrix} + \begin{pmatrix} 2v & 1 \\ \Pi + \frac{\zeta}{\tau_\Pi} & v \end{pmatrix} \partial_x \begin{pmatrix} v \\ \Pi \end{pmatrix} = \begin{pmatrix} 0 \\ -\frac{1}{\tau_\Pi} \Pi \end{pmatrix}. \quad (4.29)$$

Solving for the eigenvalues of Λ gives

$$\lambda_{\pm} = \frac{3v}{2} \pm \sqrt{\frac{v^2}{4} + \Pi + \frac{\zeta}{\tau_\Pi}}. \quad (4.30)$$

which are our characteristic wave speeds. Taking the ideal equilibrium limits $\zeta \rightarrow 0$, $\Pi \rightarrow 0$ gives

$$\lambda_{\pm} \rightarrow 2v, v \quad (4.31)$$

so we do indeed recover our original wave speeds of $2v$ and v . Now consider the effect of the additional terms inside the square root of (4.30). The action of bulk viscosity is to resist motion in the fluid, specifically the shear-less compression or expansion of a fluid, so it seems logical that it would reduce the magnitude of fluid motion, but not necessarily sound speed(s) within the fluid. In a full non-ideal hydrodynamic theory we would expect the dissipation timescales and magnitude coefficients (τ_Π , ζ) to be of the same sign and the same order of magnitude. This modifies the wavespeed by a positive quantity that may not be small in size compared to the dissipation variable itself (Π , here).

Next, we note that Π may be both positive and negative. Consider its equilibrium value of $\Pi_{NS} = -\zeta \partial_k v^k$: for uniform expansion it is negative; for compression it is positive. For our 1D wave speeds to be purely real we require

$$\left(\Pi + \frac{\zeta}{\tau_\Pi} \right) > -\frac{v^2}{4}. \quad (4.32)$$

We can see now why some initial conditions (those with steep, positive velocity gradients) are prone to instability: they give Π_{NS} a negative value which drives Π to be negative, too. A steep, positive velocity gradient represents a fluid where the downstream elements are moving faster than the upstream ones. Essentially, all elements within the fluid are moving apart and at an increasing rate, not unlike the currently accelerating expansion of the universe. In this scenario, one does not expect the bulk viscosity to be able to act when its speed of propagation is less than that of the fluid motion, which is perhaps represented by the relation above.

In conclusion, we note that the additional terms (Π , $\frac{\zeta}{\tau_\Pi}$) in (4.30) are capable of both diverging and focusing the characteristic speeds, depending on the sign of $(\Pi + \frac{\zeta}{\tau_\Pi})$.

Having calculated the eigenvalues of the system, let us consider their eigenvectors. Up to a constant factor, they are given by

$$\mathbf{v}_{\pm} = \begin{pmatrix} -2 \\ v \mp \sqrt{v^2 + 4(\Pi + \frac{\zeta}{\tau_1})} \end{pmatrix} \quad (4.33)$$

where the \pm subscript on \mathbf{v} corresponds to the eigenvalues defined above. We see that the characteristic information associated with non-zero non-ideal parameter values will appear as a modification in the contribution of the non-ideal Π term.

4.1.2.4 Numerical Results

Figure 4.3 shows the effect of including bulk viscosity on the evolution of fluid which is initially moving with uniform velocity. The moving fluid is surrounded by still fluid, creating a ‘top-hat’ profile for the velocity analogous to the one for the temperature seen in figure 4.2. In particular we see the behaviour expected from a (viscous) Burger’s model: the formation of shocks and rarefactions, with the steep gradients being attenuated in the non-ideal case.

We also include table 4.1 which contains results concerning run-times of the code for the simulation above. We compare the relaxation model given by eq. (4.23b) (in one dimension) with that produced using the CE expansion, given by eq. (4.26).

This shows a few different effects. Firstly, doubling the number of grid cells demands twice as many calculations per timestep, but also halves the timestep size due to the CFL condition, leading to a factor four increase in run-time. This is the case for both models. Secondly, it is seen that the Chapman-Enskog expansion leads to significantly faster execution times, by a factor between 2 and 3, in general. This is purely because of the reduced system size - the numerical methods used are the same. Finally, this speed-up between models is less when the run-time of the simulations is shorter. This is because there is some computational cost in setting up the initial data and starting the simulation, which is similar for both models.

4.1.3 Shear Viscosity

Finally, we will model a fluid that resists shear motions within itself which arise when neighbouring fluid elements with differing velocities move past each other. As for the bulk viscosity, looking at the Navier-Stokes form of the shear viscous tensor, $\pi_{\mu\nu,NS} = -2\eta(\partial_{\mu}u_{\nu} + \partial_{\nu}u_{\mu} - \frac{2}{3}h_{\mu\nu}\partial_{\gamma}u^{\gamma})$, is somewhat illuminating. When there is a gradient in a component of the fluid’s 4-velocity, a resistive force will arise, even in directions perpendicular to that velocity component.

Grid size	t_{final}	Runtime [s]		Scaling	Speed-up
		Full	Chapman-Enskog		
2048	0.2	1.20	0.56	– (–)	2.1
4096	0.2	5.19	2.07	4.3 (3.7)	2.5
8192	0.2	19.04	7.95	3.7 (3.8)	2.4
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

TABLE 4.1: Code run-times for our toy bulk viscous model in one dimension, simulating the data seen in figure 4.3. The expected scaling of the run-time with resolution, $t_{\text{run}} \propto N_x^2$, can be seen for both the full model and the Chapman-Enskog expansion (the latter in parentheses). Crucially, a nearly three-fold speed-up is achieved by using the Chapman-Enskog expansion, even with identical numerical methods.

We run through the same approximations as in the bulk pressure case in section 4.1.2. Note that the slow motion approximation means that the components $\pi_{0a} \rightarrow 0$ in the state and flux vectors (as we assume them small compared to the constant $\rho + p$), leading to

$$\partial_t v_j + \partial_i (v^i v_j + \pi_j^i) = 0, \quad (4.34a)$$

$$\partial_t \pi_{jk} + \partial_i (\pi_{jk} v^i) = \frac{1}{\tau_\pi} (\pi_{jk, \text{NS}} - \pi_{jk}). \quad (4.34b)$$

It is clear that the model is very similar to the bulk viscous pressure case, and the only times this will give something distinct is when we consider pure shear flows. As an example, let us work with x derivatives only, but set $v_x \equiv 0$ whilst allowing for non-zero flow in other directions; for simplicity, only consider $v_y \neq 0$ (which means that only $\pi_y^x \neq 0$). Then we have

$$\partial_t v_y + \partial_x \pi_y^x = 0, \quad (4.35a)$$

$$\partial_t \pi_{xy} = -\frac{1}{\tau_\pi} (2\eta \sigma_{xy} + \pi_{xy}). \quad (4.35b)$$

The slow motion approximation and our consideration of only x derivatives means $\Theta \rightarrow 0$ and $h_{jk} \rightarrow \eta_{jk}$ in the definition of $\pi_{jk, \text{NS}}$, hence $\sigma_{xy} \rightarrow \partial_x v_y$. This is very close to the shear flow example of Takamoto and Inutsuka, 2011, and can be rewritten as

$$\partial_t v_y = 2\eta \partial_{xx} v_y + \tau_\pi \partial_{tx} \pi_{xy}, \quad (4.36)$$

giving the heat equation behaviour if $\tau_\pi \ll \eta$.

4.1.3.1 Limiting case

For completeness we give the limiting case here, although the pure shear case and bulk viscous toy model have given us a clear idea of what to expect. We look at the limiting




FIGURE 4.3: The velocity and bulk viscosity are plotted throughout our spatial domain. The initial velocity profile was a top-hat function with $v_x = 0.7$ for $0.3 \leq x \leq 0.7$ and $v_x = 0$ otherwise. In the ideal case, the top figure shows similar behaviour to the simple Burger's equation, as expected. The non-ideal behaviour is given by the bulk viscosity in the lower plot, and its impact through smoothing strong gradients is clear in the velocity plot. The simulation ran with parameters $t_{\text{final}} = 0.2$, $\zeta = 1 \times 10^{-2}$ and $\tau_{\Pi} = 1 \times 10^{-3}$.

case as $\tau_{\pi} \rightarrow 0$, giving $\pi_{jk} \rightarrow -\eta \partial_{<j} v_{k>}$. This reduces the toy model to

$$\partial_t v_j + \partial_i (v^i v_j) = \eta \partial^i \partial_{<i} v_{j>}. \quad (4.37)$$

Once again this is Burger's equation, but now the dissipation only acts on cross-terms, not in the isotropic directions.

4.1.3.2 Chapman-Enskog Expansion in 1D

As in the previous toy models, we perform the Chapman-Enskog expansion by setting π to the sum of its equilibrium value and a term representing departure from equilibrium, which is multiplied by the relevant timescale, τ_π in this case.

$\pi_{jk} \rightarrow \pi_{NS} + \tau_\pi \pi_1 = -2\eta\sigma_{jk} + \tau_\pi \pi_1$ where $\sigma_{jk} = \left(\partial_j v_k + \partial_j v_k - \frac{2}{3}\delta_{jk}\Theta\right)$. Restricting to one dimension gives a great simplification: our shear tensor becomes a scalar; $\sigma_{ij} \rightarrow \sigma_{xx} = \frac{4}{3}\partial_x v$.

Making the substitution discussed previously in equations (4.35), we get for the first (retaining terms of order τ_π^0)

$$\pi_1 = \frac{8}{3}\eta \left[\partial_{tx} v + (\partial_x v)^2 + v \partial_{xx} v \right]. \quad (4.38)$$

For the second, retaining all terms yields

$$\partial_t v + 2v \partial_x v - \frac{8}{3}\eta \partial_x v + \tau_\pi \partial_x \pi_1 = 0. \quad (4.39)$$

Note that throughout this derivation we are treating τ_π and η as constants, rather than the thermodynamic functions which they more accurately are. We need to remove the mixed $t - x$ derivative in equation (4.38) so we differentiate equation (4.39), neglecting the final term of order τ_π

$$\partial_{tx} v = \frac{8}{3}\eta \partial_{xx} v - 2(\partial_x v)^2 - 2v \partial_{xx} v \quad (4.40)$$

and substitute into (4.38) to get

$$\begin{aligned} \pi_1 &= \frac{8}{3}\eta \left[\frac{8}{3}\eta \partial_{xx} v - 2(\partial_x v)^2 - 2v \partial_{xx} v + (\partial_x v)^2 + v \partial_{xx} v \right] \\ &= \frac{8}{3}\eta \left[\frac{8}{3}\eta \partial_{xx} v - (\partial_x v)^2 - v \partial_{xx} v \right]. \end{aligned} \quad (4.41)$$

Similarly, we can substitute the spatial derivative of π_1

$$\partial_x \pi_1 = \frac{8}{3}\eta \left[\frac{8}{3}\eta \partial_{xxx} v - 3(\partial_x v)(\partial_{xx} v) - v \partial_{xxx} v \right], \quad (4.42)$$

and use the relation $\tau_\pi = 2\eta\beta_2$ to link our timescale and dissipation coefficient and arrive at

$$\partial_t v + 2v \partial_x v = \frac{8}{3}\eta \left[\partial_x v + 2\eta\beta_2 \left(3(\partial_x v)(\partial_{xx} v) + v \partial_{xxx} v - \frac{8}{3}\eta \partial_{xxx} v \right) \right]. \quad (4.43)$$

4.1.3.3 Chapman-Enskog Expansion in Pseudo-2D

Let us look at the 2-dimensional case with the Chapman-Enskog expansion. We will enforce the same assumptions as in section 4.1.3 and so start with equations (4.35):

$$\partial_t v_y + \partial_x \pi_{xy} = 0, \quad (4.35a)$$

$$\partial_t \pi_{xy} = -\frac{1}{\tau_\pi} (2\eta \sigma_{xy} + \pi_{xy}). \quad (4.35b)$$

Working through the same steps as before, we arrive at

$$\partial_t v_y = 2\eta \partial_x^{(2)} v_y - 4\eta^2 \tau_\pi \partial_x^{(4)} v_y. \quad (4.45)$$

This is the analogous equation to (4.36), but with the Chapman-Enskog expansion. Again, we recover the heat equation in the limit of small τ_π , in which case the wave speed is expected to diverge nonphysically. However, the notion of a wavespeed for a parabolic equation is tenuous. We can attempt to establish another dispersion relation here by writing $v_y \propto \exp(i(\omega t \pm kx))$ which gives

$$\omega = i(2\eta k^2 + 4\eta^2 \tau_\pi k^4). \quad (4.46)$$

Given the pure-imaginary and positive-definiteness of the above wave frequency, only decaying velocity-waves are permitted, with a wavespeed of zero.

4.1.3.4 Wave Speeds

As in the case of the bulk viscosity, our governing equations for the shear viscosity and the fluid velocity in this toy model are non-linear. We therefore look again at a vector form of the equations with a Jacobian matrix pre-factor for the spatial derivatives. In one spatial dimension this looks like

$$\partial_t \begin{pmatrix} v \\ \pi_{xx} \end{pmatrix} + \begin{pmatrix} 2v & 1 \\ \pi_{xx} & v \end{pmatrix} \partial_x \begin{pmatrix} v \\ \pi_{xx} \end{pmatrix} = \begin{pmatrix} 0 \\ -\frac{1}{\tau_\pi} (\frac{8}{3}\eta \partial_x v + \pi_{xx}) \end{pmatrix}. \quad (4.47)$$

We see that we have a very similar form to that for the bulk viscous pressure. Perhaps the only difference worth noting is that in the expression linking τ_π and η ($\tau_\pi = 2\eta\beta_2$), β_2 is strictly positive, which will have implications for the speed(s) of sound in this model.

Moving to two spatial dimensions may introduce more interesting behaviour, but at the cost of much more complexity. We will now have 5 independent variables (2

velocities and 3 components of π_{ij}). The equivalent of (4.47) is

$$\begin{aligned} \partial_t \begin{pmatrix} v_x \\ v_y \\ \pi_{xx} \\ \pi_{xy} \\ \pi_{yy} \end{pmatrix} + \begin{pmatrix} 2v_x & 0 & 1 & 0 & 0 \\ v_y & v_x & 0 & 1 & 0 \\ \pi_{xx} & 0 & v_x & 0 & 0 \\ \pi_{xy} & 0 & 0 & v_x & 0 \\ \pi_{yy} & 0 & 0 & 0 & v_x \end{pmatrix} \partial_x \begin{pmatrix} v_x \\ v_y \\ \pi_{xx} \\ \pi_{xy} \\ \pi_{yy} \end{pmatrix} \\ + \begin{pmatrix} v_y & v_x & 0 & 1 & 0 \\ 0 & 2v_y & 0 & 0 & 1 \\ 0 & \pi_{xx} & v_y & 0 & 0 \\ 0 & \pi_{xy} & 0 & v_y & 0 \\ 0 & \pi_{yy} & 0 & 0 & v_y \end{pmatrix} \partial_y \begin{pmatrix} v_x \\ v_y \\ \pi_{xx} \\ \pi_{xy} \\ \pi_{yy} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ -\frac{1}{\tau_\pi}(2\eta\sigma_{xx} + \pi_{xx}) \\ -\frac{1}{\tau_\pi}(2\eta\sigma_{xy} + \pi_{xy}) \\ -\frac{1}{\tau_\pi}(2\eta\sigma_{yy} + \pi_{yy}) \end{pmatrix}. \end{aligned} \quad (4.48)$$

Solving for the eigenvalues of the first matrix gives

$$\lambda = v_x, \frac{3v_x}{2} \pm \sqrt{\frac{v_x^2}{4} + \pi_{xx}} \quad (4.49)$$

and for the second gives an equivalent result but with $x \rightarrow y$. The first eigenvalue, v_x , has multiplicity 3 and the corresponding eigenvectors represent linear waves that advect the v_y and π_{yy} terms. These values look sensible; they are essentially the same as those obtained for the bulk viscosity but with the shear viscosity in its place and split into the two orthogonal directions. Let us now include the source terms' contributions to the Jacobians of the partial derivatives. Remembering that $\sigma_{ij} = \partial_i v_j + \partial_j v_i - \frac{2}{3}\eta_{ij}\Theta$ and working in flat spacetime we have

$$\sigma_{xx} = \frac{4}{3}\partial_x v_x - \frac{2}{3}\partial_y v_y, \quad (4.50a)$$

$$\sigma_{xy} = \frac{2}{3}(\partial_x v_y + \partial_y v_x), \quad (4.50b)$$

$$\sigma_{yy} = \frac{4}{3}\partial_y v_y - \frac{2}{3}\partial_x v_x. \quad (4.50c)$$

The Jacobian for the flux in the x -direction therefore becomes

$$\begin{pmatrix} 2v_x & 0 & 1 & 0 & 0 \\ v_y & v_x & 0 & 1 & 0 \\ \pi_{xx} + \frac{8\eta}{3\tau_\pi} & 0 & v_x & 0 & 0 \\ \pi_{xy} & \frac{4\eta}{3\tau_\pi} & 0 & v_x & 0 \\ \pi_{yy} - \frac{4\eta}{3\tau_\pi} & 0 & 0 & 0 & v_x \end{pmatrix} \quad (4.51)$$

which has eigenvalues

$$\lambda = \left\{ v_x, v_x \pm \sqrt{\frac{4\eta}{3\tau_\pi}}, \frac{3v_x}{2} \pm \sqrt{\frac{v_x^2}{4} + \pi_{xx} + \frac{8\eta}{3\tau_\pi}} \right\}. \quad (4.52)$$




FIGURE 4.4: 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: firstly, the MIS-derived simple shear model given by eq. (4.35) (solid lines); secondly, the model obtained applying the CE-expansion to it, given by eq. (4.36) (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 analytic 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.

The degeneracy of the $\lambda = v_x$ eigenvalue has been broken and two new wavespeeds have been created by the splitting of these linear waves.

4.1.3.5 Numerical Results

In figure 4.4, the numerical solutions to the systems given in equations (4.35) and (4.36) are shown. The initial data form a step-function in the y -directed velocity across the x -domain, with the discontinuity occurring at $x = 0.0$. The effect of shear viscosity highlighted here is its damping of velocities in directions perpendicular to the velocity itself. The analytic solution in the limit $\tau_\pi \rightarrow 0$ is of the form $v_y(t) \sim v_y(0) \operatorname{erf}\left(\frac{|x|}{\eta t}\right)$ whilst $\tau_\pi > 0$ introduces small deviations from this - see appendix B for the exact solution. Excellent agreement is seen between the numerical results with and without the C-E expansion, with the differences being $\mathcal{O}(\tau_\pi) \sim 1 \times 10^{-4}$ or smaller, as expected.

4.1.4 Summary

From the Müller-Israel-Stewart equations of motion, we have derived simple models describing each dissipation mechanism in turn. By applying the Chapman-Enskog

expansion to them, we have shown its potential as an approach for efficiently simulating non-ideal fluids. In particular, for small non-ideal parameters, the numerical results of the approximation schemes are indistinguishable from the full relaxation-type models, but are produced much faster.

However, the use of a truncated expansion at the equation of motion level (rather than at the stress-energy tensor level) breaks the covariant nature of the original models, meaning the resulting expressions are no longer Lorentz-boost invariant and are not applicable in all frames of reference. This limits their physical applicability in simulations.

We also made significant approximations, such as taking the low-velocity limit and removing the coupling between many fluid variables. We improve upon this now by simulating a fully relativistic non-ideal model - the well-studied Müller-Israel-Stewart formalism introduced earlier.

4.2 The Full MIS Model

This section will look at results from simulations using the MIS formulation of non-ideal hydrodynamics. Writing the system out again explicitly and aligning closely with Chabanov, Rezzolla, et al., 2021 specifically, we have

$$\partial_t \mathbf{U} + \partial_i \mathbf{F}^{(i)} = \mathbf{S}, \quad (4.53)$$

where

$$\mathbf{U} = \begin{pmatrix} D \\ S_j \\ E \\ Y_j \\ U \\ Z_j^l \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi)W^2 v_j + (q_j + q_k v^k v_j)W + \pi_j^0 \\ (\rho + p + \Pi)W^2 - (p + \Pi) + 2q_k v^k W + \pi_0^0 \\ nW q_j \\ nW \Pi \\ nW \pi_j^l \end{pmatrix}, \quad (4.54a)$$

$$\mathbf{F}^{(i)} = \begin{pmatrix} Dv^i \\ S_j v^i + (p + \Pi)\delta_j^i + W(q^i v_j - q_k v^k v_j v^i) + \pi_j^i \\ (E + p)v^i + W(q^i - q_k v^k v^i) + \pi_0^i \\ Yv^i \\ Uv^i \\ Z_j^l v^i \end{pmatrix}, \quad (4.54b)$$

$$\mathbf{S} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{n}{\tau_q}(\tilde{q}_j - q_j) \\ \frac{n}{\tau_\Pi}(\tilde{\Pi} - \Pi) \\ \frac{n}{\tau_\pi}(\tilde{\pi}_j^l - \pi_j^l) \end{pmatrix}. \quad (4.54c)$$

All variables are consistent with their earlier definitions.

Once our numerical scheme has updated the conserved state vector

$(D, S_j, \tau \equiv E - D, U, Y_j)^T$, we need to recompute the primitive variables including velocity, temperature, number density, and all other quantities derived from them (such as pressure, acceleration and so on). The conserved-to-primitive (C2P) scheme that we use here is inspired by the one commonly used for ideal fluid formulations, whereby the pressure is the guessed quantity. Here, we simply extend our guesses (there are now four of them) to include the additional dissipative terms that must be subtracted or divided by in order to obtain an expression for the square of the velocity. It reads

1. Calculate q_j as Y_j/D , Π as U/D and π_j^l as Z_j^l/D .
2. Guess $\chi = p + \Pi - 2q_k v^k W - \pi_k^k$, $\sigma_j = (q_j + q_k v^k v_j)W + \pi_j^0$.
3. Compute $v^2 = \gamma^{ij}(S_i - \sigma_i)(S_j - \sigma_j)/(E + \chi)^2$
4. Compute $W = (1 - v^2)^{-1/2}$
5. Compute $n = D/W$
6. Compute $\rho + p = (E + \chi)/W^2 - \Pi$
7. (In general, may need to now get v_j from $(S_j - \sigma_j)/[(\rho + p + \Pi)W^2]$).

8. Invert the EoS to get p from n , $(\rho + p)$. In the ideal case: $\rho = n(1 + \epsilon)$, $p = (\Gamma - 1)n\epsilon$ and so, after some manipulation, $p = \frac{\Gamma-1}{\Gamma}(\rho + p - n)$.
9. Compute v_j from $(S_j - \sigma_j)/[(\rho + p + \Pi)W^2]$.
10. Compute $\chi = p - 2q_k v^k W$, $\sigma_j = (q_j + q_k v^k v_j)W$ from our quantities: compare to guesses (and root-find to minimize the difference until it is zero to within a small tolerance, typically $\mathcal{O}(10^{-8})$).

4.2.1 An Alternative C2P Scheme

In Kastaun, Kalinani, et al., 2021, concerns are raised over the stability and accuracy of common C2P schemes, in particular for fast-moving, highly magnetized flows (we have no magnetic field). A new scheme is presented and analysed with desirable features. Although rarely facing stability issues arising from our C2P scheme (for the MIS formulation, at least, the stability is controlled strongly by the strengths and timescales) we choose to implement a close-adaptation of this new scheme, again adjusting the guessed quantities to account for our dissipative terms. First, we define $\tilde{h} = (\rho + p + \Pi)/n = 1 + \epsilon + (p + \Pi)/n$ as a modified specific enthalpy and a new, guessed variable $\mu = 1/(W\tilde{h})$ in addition to, as before, $\sigma_j = (q_j + q_k v^k v_j)W + \pi_j^0$. Our conserved quantities may now be written as

$$D = nW, \tag{4.55a}$$

$$S_j = DW\tilde{h}v_j + (q_j + q_k v^k v_j)W + \pi_j^0, \tag{4.55b}$$

$$\tau = D(\tilde{h}W - 1) - p - \Pi + 2q_k v^k W + \pi_0^0. \tag{4.55c}$$

And our C2P scheme becomes

1. Compute Π, q_j, π_j^l as usual.
2. Guess μ, σ_j .
3. Compute $v^2 = \mu^2 \gamma^{ij}(S_i - \sigma_i)(S_j - \sigma_j)/D^2$.
4. Compute W, n as usual.
5. Compute $v_j = \mu(S_j - \sigma_j)/D$.
6. Compute $\pi_0^0 = v_j v_k \pi^{jk}$ using orthogonality.
7. Compute p from τ , using

$$p = D(1/\mu - 1) - \Pi + 2q_k v^k W + \pi_0^0 - \tau = D(1/\mu - 1) - \Pi + 2q_k v^k W + \pi_k^k - \tau.$$
8. Invert our EOS to get ρ from n, p : $\rho = n + p/(\Gamma - 1)$.

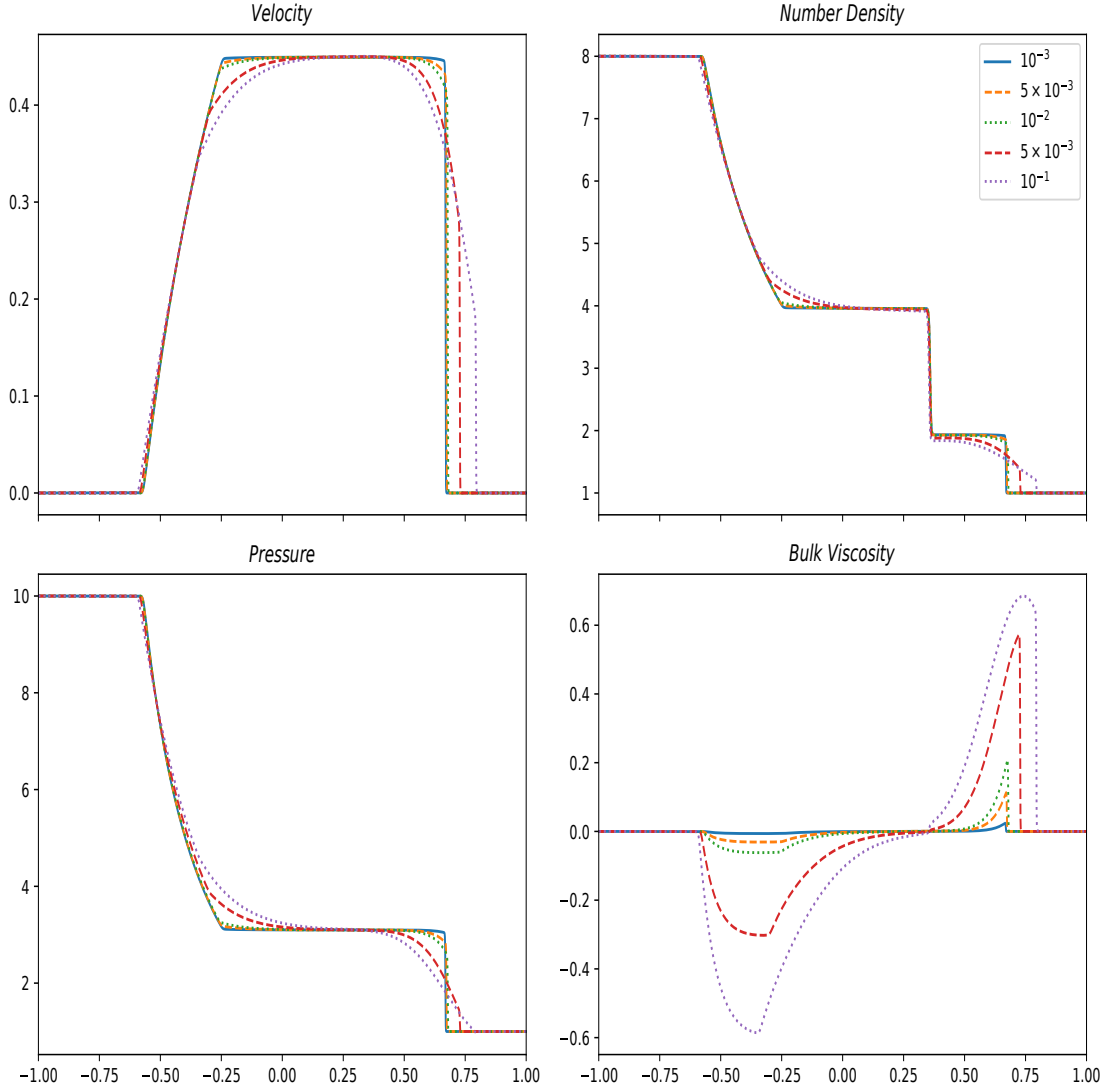


FIGURE 4.5: The development of 4 quantities (velocity, number density, pressure and bulk viscosity) for a shocktube setup with zero initial velocity as the bulk viscosity coefficient ζ is varied over several orders of magnitude. One can see several prominent features, such as the smearing of discontinuities and the acceleration of the shock-front.

9. Compute \tilde{h} and hence $\mu = 1/(W\tilde{h})$ to compare to its guessed value. This is given by $\tilde{h} = 1 + (p\Gamma/(\Gamma - 1) + \Pi)/n$.

Next, we perform numerical tests of our implementation of the MIS model. These are 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. In particular,

- Shocktubes are simple, one-dimensional tests useful for closely analysing the behaviour of the model & its numerical implementation when there are discontinuities present. These will occur at the point of merger, and also if there exists a sharp phase (and hence density & pressure) transition inside the neutron




FIGURE 4.6: The development of a Kelvin-Helmholtz unstable fluid with negligible viscosity until $t = 6.25$. 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 4.7: The development of the Kelvin-Helmholtz instability until $t = 30.0$, for an ideal fluid with negligible viscosity. The longer simulation time allows the asymmetric initial perturbation at the interface to give rise to large-scale asymmetric vortex formation.




FIGURE 4.8: The short-term evolution of the Kelvin Helmholtz instability using the MIS model of non-ideal hydrodynamics with a shear viscosity parameter of $\eta = 1 \times 10^{-3}$ and timescale $\tau_\pi = 5 \times 10^{-3}$.

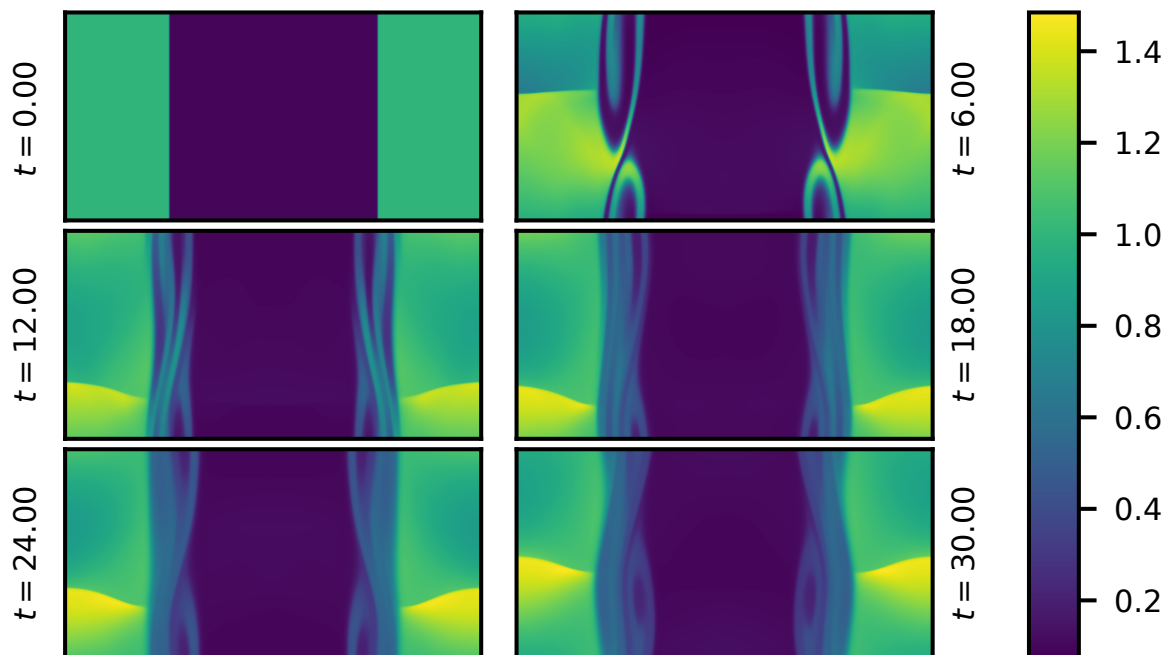


FIGURE 4.9: The long-term evolution of the Kelvin Helmholtz instability using the MIS model of non-ideal hydrodynamics with a shear viscosity parameter of $\eta = 1 \times 10^{-3}$ and timescale $\tau_\pi = 5 \times 10^{-3}$.




FIGURE 4.10: The power spectrum, $P_{\hat{u}(k_x)}$, for the kinetic energy density in the Kelvin-Helmholtz instability at $t = 15.0$ and $t = 30.0$, for both an inviscid and a viscous fluid. In the viscous cases, The expected Kolmogorov scaling in the inertial range is also plotted and matches the data well. In the viscous cases, dissipation has reduced the total kinetic energy present. It has also smoothed the minor oscillations present in the ideal curves which stem from the small-scale mixing modes at the interface which are suppressed by viscosity.

star core. 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 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 of a differing density. A wide range of behaviours can be observed depending on the parameter values - we will be varying the shear viscosity only. 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. 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 (essentially inverse wavelength) for at least part of its spectrum (the ‘inertial sub-range’) and generalized by Qian, 1994.

See figure 4.5 for results from a shocktube setup where an ideal fluid is compared to a viscous one with several different strengths of dissipation. The initial data for this test is similar to those of Takamoto and 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

$$L : \begin{pmatrix} p \\ n \\ \rho \end{pmatrix} = \begin{pmatrix} 10 \\ 8 \\ 23 \end{pmatrix} \quad (4.56)$$

for the left state and

$$R : \begin{pmatrix} p \\ n \\ \rho \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 2.5 \end{pmatrix} \quad (4.57)$$

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.

The three degrees of freedom in the model, represented by the three conserved quantities (D, S_x, τ) , produce three waves: a left-travelling rarefaction wave, a right-travelling contact wave and a right-travelling (and faster) shockwave. Note the smearing of the rarefaction & shockwave, and the increase in shock speed for the non-ideal case.

In figures fig. 4.6, fig. 4.7, fig. 4.8 and fig. 4.9 we show results of Kelvin-Helmholtz instability simulations, both with and without shear viscosity. To investigate this process, we use the initial conditions from Beckwith and 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

$$\begin{pmatrix} v_y \\ \rho \\ v_x \end{pmatrix} = \begin{pmatrix} v_{sh} \tanh\left(\frac{x-0.5}{a}\right) \\ \rho_0 + \rho_1 \tanh\left(\frac{x-0.5}{a}\right) \\ A_0 v_{sh} \sin(2\pi y) \exp\left(\frac{-(x-0.5)^2}{l^2}\right) \end{pmatrix} ; \quad x > 0.0 \quad (4.58a)$$

and

$$\begin{pmatrix} v_y \\ \rho \\ v_x \end{pmatrix} = \begin{pmatrix} -v_{sh} \tanh\left(\frac{x+0.5}{a}\right) \\ \rho_0 - \rho_1 \tanh\left(\frac{x+0.5}{a}\right) \\ -A_0 v_{sh} \sin(2\pi y) \exp\left(\frac{-(x+0.5)^2}{l^2}\right) \end{pmatrix} ; \quad x \leq 0.0 \quad (4.58b)$$

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.

One can see that, in general, adding viscosity suppresses the formation of small-scale features and the mixing of the two fluids. Also, in figure 4.10 we plot the normalized, integrated power spectrum, $P_{\hat{u}}(k_x)$, of the kinetic energy density. This quantity is defined as

$$P_{\hat{u}}(k_x) = \sum_{n_y=0}^{N_y} |\hat{u}(k_x, n_y)|^2 / \sum_{k_x=0}^{N_{Ny}} \sum_{n_y=0}^{N_y} |\hat{u}(k_x, n_y)|^2 \quad (4.59)$$

where $N_{Ny} = N_x/2$ is the Nyquist frequency and $\hat{u}(k_x, n_y)$ is the 1D discrete Fourier transform of the discretized variable $u(n_x, n_y)$ in the x -direction, defined as

$$\hat{u}(k_x, n_y) = \sum_{n_x=0}^{N_x} \frac{1}{N_x} u(n_x, n_y) \exp^{-2\pi i k_x \frac{n_x}{N_x}} \quad (4.60)$$

where in our case $u(n_x, n_y)$ is specifically given by $u = nW(W-1)$ and is hence a “kinetic energy density”. Note that we could have performed the Fourier transform from position space to wavenumber space ($n \rightarrow k$) in the y -direction, or even performed a 2D discrete Fourier transform in both directions. In all cases the resulting spectrum is similar and demonstrates that there is an intermediate range of wavenumbers where the expected Kolmogorov scaling is obeyed and $P_{\hat{u}}(k_x) \propto k_x^{-5/3}$. Below this range (longer wavelengths, lower wavenumber) there is a greater power due to the kinetic energy in the large-scale features and above this range (shorter wavelengths, higher wavenumber) there is a fall-off in kinetic energy as viscosity, either genuine or numerical, extracts it from the system.

Chapter 5

Simulations with our ‘MISCE’ Model

In this chapter, we presents results from our own ‘Müller-Israel-Stewart-Chapman-Enskog’ (MISCE) model, derived from the Müller-Israel-Stewart formalism already introduced in section 2.2.1.5. Large sections of this chapter are adapted directly from our publication this year (Hatton and Hawke, 2024).

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 timescales become small and the sources become ‘stiff’. The relaxation timescales 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 timestep of the simulation drastically or adopt implicit time-integrator methods to ensure accurate and stable numerical evolution. See Palenzuela, Lehner, et al. (2009), Dionysopoulou, Alic, et al. (2013), Miranda-Aranguren, Aloy, et al. (2018), Ripperda, Bacchini, et al. (2019), Wright and Hawke (2020), Wright (2020), and Dash, Shokri, 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 and Russo, 2005) may be used but in any case, the computational cost increases, potentially by an order of magnitude or two, when source terms become stiff near the ideal limit (Most and Noronha, 2021).

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 lengthscales well below that of the grid cells’ size. Estimates of the dissipation lengthscale above which structure can form

through turbulence suggest that simulations may need to resolve scales below the cm level (Radice and Hawke, 2024; Thompson and Duncan, 1993). However the current highest-resolution simulations have fluid elements with sizes $\approx 10\text{m}$ (Kiuchi, Kyutoku, et al., 2018). To bridge this gap computationally is impractical for the foreseeable future.

Instead, to address this ‘subgrid’ 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 subgrid microphysics or mathematical artefacts resulting from the implicit filtering process introduced by coarse simulations.

Subgrid 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 subgrid 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 subgrid 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 subgrid scale turbulence, the collapse of the hyper-massive neutron star remnant is altered. Other work by Viganò, Aguilera-Miret, et al. (2020b) and Carrasco, Viganò, et al. (2020) uses a gradient expansion approach to prescribe the unresolved fields in the MHD equations. See Radice and Hawke, 2024 for a modern review of the field.

One might ask why these subgrid models are relevant to the non-ideal hydrodynamic formulation presented here. In Celora, Andersson, 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 which 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 subgrid 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. We dedicate chapter 6 to our work in this direction.

In this chapter, 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 chapter is laid out as follows. Section 5.1 presents the full ‘MISCE’ model in balance-law form, with its source derived from applying the CE expansion to the MIS model. In section 5.2 we show results of simulations that use the MISCE formulation of dissipative hydrodynamics and then assess code performance in section 5.3. Particularly, we quantitatively compare results and performance with the MIS model. Section 5.4 covers additional considerations for this model, specifically those related to initial data choices and stability. Finally, in section 5.5, we summarise our findings from simulating this new model and discuss how they fit into the state of current astrophysical simulations.

5.1 General Balance-Law Derivation

In this section we will derive a source term that captures the behaviour due to dissipation introduced in the MIS formulation, but that avoids the numerical difficulties in the first-order limit of this second-order system. We will do this using a Chapman-Enskog-type analysis.

We can re-write system (4.54) in the following, more compact way:

$$\partial_t \mathbf{q}(\mathbf{w}, \overline{\mathbf{w}}) + \partial_i \mathbf{f}^i(\mathbf{w}, \overline{\mathbf{w}}) = \mathbf{s} = \mathbf{0}, \quad (5.1a)$$

$$\partial_t \overline{\mathbf{q}}(\mathbf{w}, \overline{\mathbf{w}}) + \partial_i \overline{\mathbf{f}}^i(\mathbf{w}, \overline{\mathbf{w}}) = \frac{\overline{\mathbf{s}}(\overline{\mathbf{w}})}{\epsilon}, \quad (5.1b)$$

where we indicate terms which become stiff as $\epsilon \rightarrow 0$ with an over-bar. This means that the stiff conserved variables are $\overline{\mathbf{q}} = \{U, Y_j, Z_{jk}\}$ with the corresponding fluxes, $\overline{\mathbf{f}}^i(\mathbf{w}, \overline{\mathbf{w}})$, and sources, $\overline{\mathbf{s}}(\overline{\mathbf{w}})$, as in equations (4.54). These form the dissipative subsystem of equations. The remaining conserved variables describe the fluid evolution and are non-stiff in the ideal limit, denoted $\mathbf{q} = \{D, S_j, \tau\}$. We will also denote the vector of primitive variables present in ideal hydrodynamics as $\mathbf{w} = \{p, \rho, n, v_1, v_2, v_3\}$ and the dissipative variables as $\overline{\mathbf{w}}(\mathbf{w}) = \{\Pi, q_j, \pi_{jk}\}$.

We begin with the full MIS model in balance law form, eq. (5.1), recalling that the non-stiff and stiff conserved variables are labelled \mathbf{q} and $\overline{\mathbf{q}}$ respectively. In order to

maintain finite solutions in the ideal limit, we require that $\lim_{\epsilon \rightarrow 0} \bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}) = \mathbf{0}$. This motivates an expansion of the stiff variables in powers of ϵ , with each increasing order providing a further deviation from the ideal limit.

In fact, because of the mathematical and physical links between the dissipation timescales (τ) and strengths (ξ) discussed earlier, we choose to perform the expansion in powers of ϵ where both τ and ξ are $\mathcal{O}(\epsilon)$. Whilst in theory this means that these parameters should take on similar (small) values, this is not always the case in practice¹. 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

$$\bar{\mathbf{q}} = \bar{\mathbf{q}}_0 + \bar{\mathbf{q}}_1 + \bar{\mathbf{q}}_2 + \mathcal{O}(\epsilon^3) \quad (5.2)$$

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:

$$\bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}) = \frac{1}{\epsilon}(\bar{\mathbf{q}} - \bar{\mathbf{q}}_{NS}) \quad (5.3)$$

and simply rewrite it as

$$\epsilon \bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}) = \bar{\mathbf{q}} - \bar{\mathbf{q}}_{NS}. \quad (5.4)$$

Noting that \mathbf{q}_{NS} is $\mathcal{O}(\epsilon)$, we have that at zero'th 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

$$\partial_t \mathbf{q}_0(\mathbf{w}) + \partial_i \mathbf{f}_0^i(\mathbf{w}) = \mathbf{0} \quad (5.5)$$

where

$$\mathbf{q}_0(\mathbf{w}) = \begin{pmatrix} D \\ S_j \\ \tau \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p)W^2 v_j \\ (\rho + p)W^2 - p - nW \end{pmatrix} \quad (5.6)$$

¹In many numerical simulations of relativistic astrophysical fluids, physical values of transport coefficients (e.g., viscosity, conductivity) correspond to scales that are well below the resolution limit. As a result, simulations often employ artificially large parameter values to retain numerical stability or convergence. While this allows one to qualitatively capture dissipative effects, it does not resolve the intrinsic kinetic or turbulent scales. To bridge this gap, effective mean corrections, such as turbulent viscosities or subgrid closures, are introduced. For example, in neutron star merger simulations, an effective shear viscosity may be used to model angular momentum transport from magnetorotational turbulence, while in accretion disks, α -viscosity prescriptions are commonly adopted to mimic MRI-induced transport.

and

$$\mathbf{f}_0^i(\mathbf{w}) = \begin{pmatrix} nWv^i \\ (\rho + p)W^2v^iv_j + p\delta_j^i \\ (\rho + p)W^2v^i - nWv^i \end{pmatrix}. \quad (5.7)$$

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

$$\begin{aligned} & \partial_t [\mathbf{q}_0(\mathbf{w}) + \mathbf{H}_{(1)}(\mathbf{w}, \partial_t \mathbf{w}, \partial_i \mathbf{w})] + \\ & \partial_i [\mathbf{f}_0^i(\mathbf{w}) + \mathbf{F}_{(1)}^i(\mathbf{w}, \partial_t \mathbf{w}, \partial_i \mathbf{w})] = 0 \end{aligned} \quad (5.8)$$

where

$$\mathbf{H}_{(1)}(\mathbf{w}) = \begin{pmatrix} 0 \\ W(q_{0,NS}v_j + q_{j,NS}) + \pi_{0j,NS} \\ 2q_{0,NS}W - \Pi_{NS} + \pi_{00,NS} \end{pmatrix} \quad (5.9)$$

and

$$\mathbf{F}_{(1)}^i(\mathbf{w}) = \begin{pmatrix} 0 \\ \Pi_{NS}(v^iv_j + \delta_j^i)W^2 + W(q_{NS}^iv_j + q_{j,NS}v^i) + \pi_{j,NS}^i \\ W(q_{0,NS}v^i + q_{NS}^i) + \pi_{0,NS}^i \end{pmatrix}. \quad (5.10)$$

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

$$\partial_t \mathbf{q}_0 + \partial_i \mathbf{f}_0^i = \sum_{p=0} \tilde{\mathbf{R}}_{(p)} \equiv \sum_{p=0} (-\partial_t \mathbf{H}_{(p)} - \partial_i \mathbf{F}_{(p)}^i) \quad (5.11)$$

where each additional term in the series on the RHS of eq. (5.11) 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 eq. (4.20) and eq. (4.26). 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 eq. (4.20) and eq. (4.26) at leading-order.

We choose to perform the series expansion and truncation such that terms $\mathcal{O}(\epsilon)$ contain no timescales and are first-order in the dissipation strengths $\{\zeta, \kappa, \eta\}$. Terms considered to be $\mathcal{O}(\epsilon^2)$ are first-order in the timescales $\{\tau_\Pi, \tau_q, \tau_\pi\}$ 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 timescales 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 and 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. Firstly, 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 eq. (5.5), which contains the time derivative of the ideal state vector and hence the primitive variables that constitute it, we have

$$\frac{\partial \mathbf{w}}{\partial t} = \frac{\partial \mathbf{w}}{\partial q_0} \frac{\partial q_0}{\partial t} + \frac{\partial \mathbf{w}}{\partial q_1} \frac{\partial q_1}{\partial t} + \dots = - \left(\frac{\partial q_0}{\partial \mathbf{w}} \right)^{-1} \partial_i \mathbf{f}_0^i + \mathcal{O}(\epsilon), \quad (5.12)$$

where we again note that \mathbf{w} is the vector of primitives. This means the term $\frac{\partial \mathbf{w}}{\partial q_0}$ has a matrix form that is far more easily obtained through an inversion of the matrix $\frac{\partial q_0}{\partial \mathbf{w}}$. We can use this result to substitute wherever a time-derivative appears in our source such that we then have

$$\tilde{\mathbf{R}}_{(1)} = -\partial_t \mathbf{H}_{\text{NS}}(\mathbf{w}, \partial_i \mathbf{w}) - \partial_i \mathbf{F}_{\text{NS}}^i(\mathbf{w}, \partial_i \mathbf{w}) \quad (5.13)$$

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 Müller-Israel-Stewart 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, in 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 and 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.

5.1.1 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 in equation 5.12 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 $\left(\frac{\partial q_0}{\partial \mathbf{w}}\right)$ and, hence, $\left(\frac{\partial q_0}{\partial \mathbf{w}}\right)^{-1}$. We have some choice over which five non-stiff primitive variables (corresponding to the five conserved quantities we evolve) we choose to include in our differencing vector, \mathbf{w} , when calculating the Jacobian. The equations of state we use in this work relate p , ρ and n , which gives us this freedom, and we opt to work with $\mathbf{w} = \{p, \rho, v\}$.

Our ideal conserved vector is now

$$\mathbf{q}_0(\mathbf{w}) = \begin{pmatrix} D \\ S_j \\ E \equiv \tau + D \end{pmatrix} = \begin{pmatrix} (\rho - p/(\Gamma - 1))W \\ (\rho + p)v_j W^2 \\ (\rho + p)W^2 - p \end{pmatrix} \quad (5.14)$$

where we have chosen to work with the conserved variable E for now instead of τ as it takes an even simpler form. Now, we have

$$\left(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}}\right) = \begin{pmatrix} W/(1-\Gamma) & W & (\rho - p/(\Gamma - 1))W^3 v_1 & (\rho - p/(\Gamma - 1))W^3 v_2 & (\rho - p/(\Gamma - 1))W^3 v_3 \\ W^2 v_1 & W^2 v_1 & (\rho + p)(W^2 + 2W^4 v_1^2) & 2(\rho + p)W^4 v_1 v_2 & 2(\rho + p)W^4 v_1 v_3 \\ W^2 v_2 & W^2 v_2 & 2(\rho + p)W^4 v_2 v_1 & (\rho + p)(W^2 + 2W^4 v_2^2) & 2(\rho + p)W^4 v_2 v_3 \\ W^2 v_3 & W^2 v_3 & 2(\rho + p)W^4 v_3 v_1 & 2(\rho + p)W^4 v_3 v_2 & (\rho + p)(W^2 + 2W^4 v_3^2) \\ W^2 - 1 & W^2 & 2(\rho + p)W^4 v_1 & 2(\rho + p)W^4 v_2 & 2(\rho + p)W^4 v_3 \end{pmatrix}. \quad (5.15)$$

In order to present the following matrix-inversion results in a human-readable fashion, we take the low-velocity limit, neglecting terms $\mathcal{O}(v^2)$ and hence setting the Lorentz factor, $W = 1$. However, this simplification allows us to generalise our expressions to be EoS-agnostic and express certain terms as partial derivatives of the primitive variables, rather than using our specific Gamma-law EoS. This means that we now have

$$\left(\frac{\partial \mathbf{q}_0}{\partial \mathbf{w}}\right) = \begin{pmatrix} \partial_p n & \partial_\rho n & 0 & 0 & 0 \\ v_1 & v_1 & p + \rho & 0 & 0 \\ v_2 & v_2 & 0 & p + \rho & 0 \\ v_3 & v_3 & 0 & 0 & p + \rho \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix} \quad (5.16)$$

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

$$(p + \rho)^{-1} \begin{pmatrix} (p + \rho)/\partial_p n & 0 & 0 & 0 & -(p + \rho)\partial_\rho n/\partial_p n \\ 0 & 0 & 0 & 0 & (p + \rho) \\ -v_1/\partial_p n & 1 & 0 & 0 & -v_1(\partial_p n + \partial_\rho n)/\partial_p n \\ -v_2/\partial_p n & 0 & 1 & 0 & -v_2(\partial_p n + \partial_\rho n)/\partial_p n \\ -v_3/\partial_p n & 0 & 0 & 1 & -v_3(\partial_p n + \partial_\rho n)/\partial_p n \end{pmatrix}. \quad (5.17)$$

Next, using equation 5.12, we have

$$\partial_t \begin{pmatrix} p \\ \rho \\ v_j \end{pmatrix} = (\rho + p)^{-1} \begin{pmatrix} (p + \rho)((1/\partial_p n) - (\partial_\rho n / \partial_p n) \partial_t E) \\ (p + \rho) \partial_t E \\ -v_j((1 + \partial_\rho n / \partial_p n) \partial_t E + (1/\partial_p n) \partial_t D) + \partial_t S_j \end{pmatrix} \quad (5.18)$$

where we can exchange the time-derivatives of the conserved variables for spatial derivatives of the fluxes using eq. (5.5). Doing this, and using our equation of state $p = (\Gamma - 1)(\rho - n)$ to replace the partial derivatives of primitive variables, we arrive at

$$\partial_t \begin{pmatrix} p \\ \rho \\ v_j \end{pmatrix} = \begin{pmatrix} (1 - \Gamma)(1 + \partial_i S^i) \\ -\partial_i S^i \\ (\rho + p)^{-1} [v_j((2 - \Gamma) \partial_i S^i + (1 - \Gamma) \partial_i (Dv^i)) - \partial_i S_j^i] \end{pmatrix} \quad (5.19)$$

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

$$\tilde{\mathbf{R}}_{(1)} = -\partial_t \begin{pmatrix} 0 \\ 0 \\ -\Pi_{NS} \end{pmatrix} - \partial_x \begin{pmatrix} 0 \\ \Pi_{NS} v_x^2 \\ 0 \end{pmatrix} \quad (5.20a)$$

$$= \zeta \left[\begin{pmatrix} 0 \\ 0 \\ -\partial_t \partial_x v^x \end{pmatrix} + \partial_x \begin{pmatrix} 0 \\ v_x^2 \partial_{xx} v^x + 2v_x (\partial_x v^x)^2 \\ 0 \end{pmatrix} \right] \quad (5.20b)$$

where from eq. (5.19) we have

$$\begin{aligned} \partial_x \partial_x v^t &= (\rho + p)^{-1} [(2 - \Gamma)(v_x \partial_{xx} S^x + (\partial_x S^x)(\partial_x v_x)) \\ &\quad + (1 - \Gamma) \partial_{xx} (Dv_x) - \partial_{xx} (S^x v_x + p)] \end{aligned} \quad (5.21a)$$

$$\begin{aligned} &= [(2 - \Gamma)(2v_x \partial_{xx} v_x + 3v_x (\partial_x v_x) \partial_x (\rho + p) + (\partial_x v_x)^2) \\ &\quad + (\rho + p)^{-1} (1 - \Gamma)(v_x \partial_{xx} n + n \partial_{xx} v_x) \\ &\quad - \partial_{xx} v_x - v_x \partial_x (\rho + p) - (\rho + p)^{-1} \partial_{xx} p] \end{aligned} \quad (5.21b)$$

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 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>.

5.1.2 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

$$\partial_t [\mathbf{q}_0(\mathbf{w}) + \mathbf{H}_{(1)} + \mathbf{H}_{(2)}] + \partial_i [\mathbf{f}_0^{(i)}(\mathbf{w}) + \mathbf{F}_{(1)}^{(i)} + \mathbf{F}_{(2)}^{(i)}] = 0 \quad (5.22)$$

it follows that

$$\tilde{\mathbf{R}}^{(2)} = -\partial_t \mathbf{H}_{(2)} - \partial_i \mathbf{F}_{(2)}^{(i)}. \quad (5.23)$$

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:

$$\partial_t (nW\bar{\mathbf{w}}) + \partial_i (nWv^i\bar{\mathbf{w}}) = \frac{n}{\tau}(\bar{\mathbf{w}} - \bar{\mathbf{w}}_{NS}) \quad (5.24a)$$

$$\partial_t \bar{\mathbf{q}} + \partial_i \bar{\mathbf{f}}^i(\mathbf{q}, \bar{\mathbf{q}}) = \frac{1}{\epsilon}(\bar{\mathbf{q}} - \bar{\mathbf{q}}_{NS}) \quad (5.24b)$$

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 $\mathcal{O}(\epsilon^2)$,

$$\epsilon [\partial_t \bar{\mathbf{q}}_{NS} + \partial_i \bar{\mathbf{f}}^i(\mathbf{q}, \bar{\mathbf{q}}_{NS})] = \bar{\mathbf{q}}_2. \quad (5.25)$$

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 $\tilde{\mathbf{R}}^{(2)}$.

5.2 Simulation Results

In fig. 5.1, 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. Increasing the bulk viscosity parameter ζ leads to a strengthening of these features, as expected. Switching on the second-order MISCE source term by setting $\tau_{\Pi} \neq 0$ has a negligible effect on the final state, as expected given the second-order term is $\mathcal{O}(\zeta\tau_{\Pi}) \sim 10^{-4}$ here. Note that, without viscosity, the solution to this test problem would be self-similar, meaning that evolution in time results in a simple re-scaling of the final state.

In fig. 5.2, 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 timescales τ (the same for both types of dissipation), and one for the MISCE model (at leading order, so the timescale 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.

Figures 5.3 to 5.5 show the development of the KHI for the fluid’s number density. Figures 5.3 and 5.4 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.

In fig. 5.5, the long-term behaviour for a viscous fluid is shown. Viscosity suppresses the perturbation’s growth and stabilises 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 and Inutsuka, 2011, who performed comparable simulations.

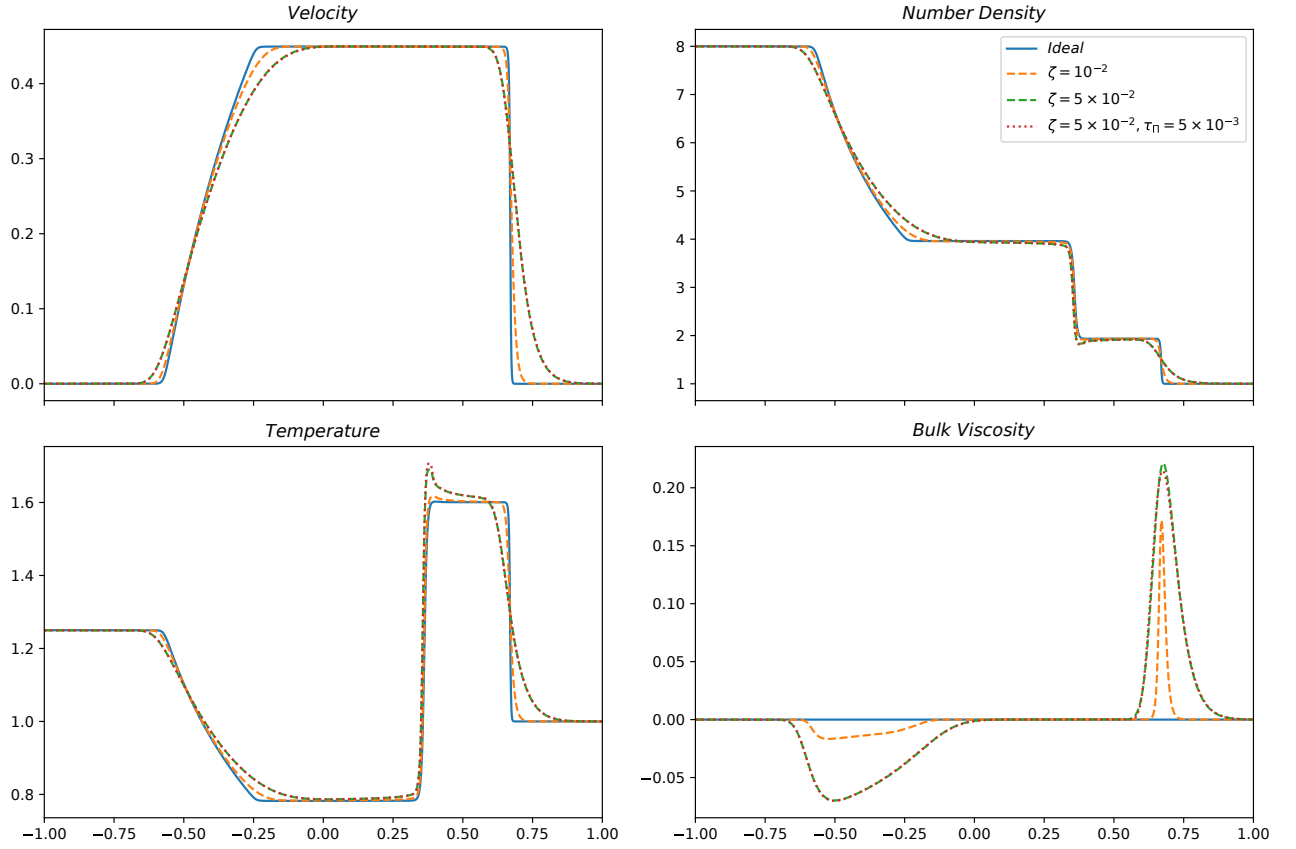


FIGURE 5.1: The evolution of four primitive variables for a ‘stillshock’ test - a shocktube with zero initial velocity - up until time $t = 1.5$. An ideal fluid is compared with a viscous one using the MISCE formulation at both leading and next-to-leading order, for a range of parameter values. The increase in shock propagation speed and smearing of discontinuities due to the inclusion of bulk viscosity are both visible physical effects. The second-order terms have little impact on the solution as they are $\mathcal{O}(\zeta\tau_{II} \sim 1 \times 10^{-4})$ but do make some visible difference to the temperature and number density at the contact wave discontinuity.

In fig. 5.6 we isolate the effect of heat conductivity on the shocktube solution, plotting the temperature for both the MIS and MISCE models. For the MIS model, no effect is seen at any dissipation strength. This is because numerical stability sharply bounds the value of the relaxation timescale τ_q from below, meaning heat conduction happens too slowly to be seen on the dynamical timescales present in the system². This is not the case for the MISCE model, where the effective relaxation timescale is zero ($\tau_q = 0$), and so the effect of a heat flux is seen as an additional heating between the shock and contact waves. We consider the ability of the MISCE model to efficiently capture this physical effect an advantage here. However, an effect is still only seen in the strongest dissipation case, when $\kappa = 0.001$. The other values of κ are simply too small to create a noticeable effect.

²We find, empirically, that the lower bound on the ratio τ_q/κ is much higher than the equivalent bound in the viscous cases. This means that by the time κ is large enough to have an effect, τ_q is also so large that the dissipation acts too slowly to affect the solution.




FIGURE 5.2: The evolution of the number density for the same stillshock test, but now until time $t = 5.0$, with the x -coordinate plotted horizontally. 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 timescale τ varying for the MIS model but held constant at zero for the MISCE model. 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.




FIGURE 5.3: The development of a Kelvin-Helmholtz unstable fluid with negligible viscosity until $t = 6.25$. 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 5.4: The development of the Kelvin-Helmholtz instability until $t = 30.0$, for an ideal fluid with negligible viscosity. The longer simulation time allows the asymmetric initial perturbation at the interface to give rise to large-scale asymmetric vortex formation.




FIGURE 5.5: The long-term evolution of the Kelvin-Helmholtz instability 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 5.6: The evolution of the temperature for the same shocktube test with zero initial velocity and varying heat flux parameter, κ , across three orders of magnitude. An effect is only seen for the MISCE model in the highest- κ case, where an additional heating of fluid between the shock and contact discontinuities is seen. We originally considered this to be the breakdown of the MISCE approximation occurring, and MIS to be the ‘true’ solution. However, we argue that the MISCE solution here is correct in its effective limit of $\tau_q \rightarrow 0$, which cannot be replicated in the MIS model due to numerical stability bounds on the ratio κ/τ_q .

Figure 5.7 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 vs. the blue curve. The expected Kolomogorov 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 lengthscales, 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.




FIGURE 5.7: The power spectra for the kinetic energy density in the Kelvin-Helmholtz instability 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 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 lengthscales). 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 lengthscales of a few, or even a single, cell(s).

5.3 Code Performance

There are a few key metrics of code performance we must now consider. Firstly, how the runtime of simulations scales with resolution. 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. 5.8), we are able to assess the convergence of our MIS and MISCE implementations with resolution. Specifically, the initial data for the y -directed velocity are $v_y(0, x) = 0.2 \sin(\pi(x + 1))$ whilst all other variables are constant across the domain.




FIGURE 5.8: The “SineWave” initial data (and its evolution with the MISCE model) 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.

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 report in table 5.1, and visualize in fig. 5.9, a summary of convergence orders for different models and resolutions. 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 derivatives 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 timestep.




FIGURE 5.9: The self-convergence rate of the MISCE model for different resolutions using the number density in the ‘sinewave’ test seen in fig. 5.8. 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.

Self-Convergence				
Model	Integrator	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 5.1: 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.

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 5.2 a comparison of runtimes between the MIS and MISCE models for the KHI. We primarily present results comparing the MIS model evolved with the

Model	Integrator	Average Runtime		
		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 ($\times 11.6$)

TABLE 5.2: A comparison of computational time required for different hydrodynamic models and time-integrators. These results are for Kelvin-Helmholtz instability 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, 2^{nd} -order Runge-Kutta scheme and SSP2(222) refers to a 2^{nd} -order implicit-explicit scheme.

SSP2(222) IMEX time-integrator (Pareschi and 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 schemes³ compared to explicit ones, it is justified. 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.

We have seen the first-order-in- τ convergence of the MISCE model to its leading-order behaviour as τ is decreased. Now, we also investigate the convergence rate of the MIS model to the leading-order MISCE model (independent of τ) as τ is reduced for the MIS model. This also serves as an indirect check on the behaviour of the MISCE model, as we see how large the effect of being off-shell is for the MIS model, relative to it.

In fig. 5.10, we plot the L_2 -norm of the difference between the MIS and MISCE models, summed across the entire domain, whilst varying τ for the MIS model. The data used are those seen in fig. 5.2. We see that for small values of $\tau \approx 1 \times 10^{-5} - 1 \times 10^{-2}$, where the viscosities and heat flux quickly relax to their Navier-Stokes forms, the difference is small. This is expected given that at leading order, the dissipative

³The implicit part of the solver requires either a complex algebraic calculation involving the inversion of a large matrix, or a numerical root-find to arrive at a self-consistent solution.




FIGURE 5.10: The difference between the MISCE model at $\tau = 0.0$ (leading order) and the MIS model with varying τ , where all three forms of dissipation are present. The L_2 -norm of the difference between number densities (left) and velocities (right) are plotted for shocktube initial data as seen in fig. 5.2. Please see the text for an explanation of the observed behaviour.

variables in the MISCE model are given precisely by the same Navier-Stokes form. For intermediate values when $\tau \approx 1 \times 10^{-2} - 1.0$, the difference rapidly increases as the off-shell behaviour of the MIS model kicks in. Finally, for the largest values of $\tau \approx 1 \times 10^{-2} - 1.0$, the difference starts to plateau as the impact of dissipation becomes negligible—it is too slow compared to the dynamical timescales of the fluid.

5.4 Further Considerations

5.4.1 Rapid Evolution of Reduced Initial Data

The source terms in MIS models drive the dissipative variables towards their equilibrium values on timescales τ . We therefore expect that, when our initial data (or otherwise) puts 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 eq. (4.11), in one dimension:

$$\partial_t T + \partial_x q = 0, \quad (5.26a)$$

$$\partial_t q = -\frac{1}{\tau_q}(\kappa \partial_x T + q); \quad (5.26b)$$

and its CE form, eq. (4.20),

$$\partial_t T = \kappa \left[\partial_x^{(2)} T - \kappa \tau_q \partial_x^{(4)} T \right]. \quad (5.27)$$

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 eq. (5.26) into

$$\partial_{\mathcal{T}} T + \tau_q \partial_x q = 0, \quad (5.28a)$$

$$\partial_{\mathcal{T}} q = -q - \kappa \partial_x T. \quad (5.28b)$$

From eq. (5.28), the power series expansion now gives that the temperature T is independent of \mathcal{T} to leading order and

$$\partial_{\mathcal{T}} q_0 = -q_0 \quad (5.29)$$

which can be integrated directly to give

$$q_0 = C_0 e^{-\mathcal{T}} \quad (5.30)$$

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 (eq. (5.26) and eq. (5.27)), 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 eq. (5.26) we have

$$T(\mathcal{T} = 1) \simeq T(\mathcal{T} = 0) + \left(\frac{\partial T_0}{\partial \mathcal{T}} + \tau_q \frac{\partial T_1}{\partial \mathcal{T}} + \dots \right) (\mathcal{T} = 0) + \dots \quad (5.31a)$$

$$= T(\mathcal{T} = 0) - \tau_q \frac{\partial \Delta q_0}{\partial x} + \mathcal{O}(\tau_q^2) \quad (5.31b)$$

whilst from the CE model eq. (5.27) we have

$$T(\mathcal{T} = 1) \simeq T(\mathcal{T} = 0) + \left(\frac{\partial T_0}{\partial \mathcal{T}} + \tau_q \frac{\partial T_1}{\partial \mathcal{T}} + \dots \right) (\mathcal{T} = 0) + \dots \quad (5.32a)$$

$$= T(\mathcal{T} = 0) + \mathcal{O}(\tau_q^2). \quad (5.32b)$$

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

$$T(t = 0, x) \rightarrow T(t = 0, x) - \tau_q \partial_x \Delta q_0. \quad (5.33)$$

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. 5.11, where a one-dimensional ‘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 eq. (5.26).

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 eq. (5.33) 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, 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 analytic expression for this adjustment behaviour that depends on the spatial gradient of the injection and the non-ideal parameter controlling its timescale, 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).




FIGURE 5.11: The evolution of the temperature and heat flux for the ‘top-hat’ initial data described in section 5.4.1 using the relaxation model given by eq. (5.26). 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 timescale shown here, as seen by the heat flux relaxing to nearly zero. The analytic result for the appropriate adjustment to the initial data, derived in section 5.4.1, is also plotted in the left panel (dotted) and shows excellent agreement with the numerical evolution result.

5.4.2 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

$$\partial_t q + \partial_x f = 0, \quad (5.34)$$

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

$$C := \left| \frac{\partial f}{\partial q} \right| \frac{\Delta t}{\Delta x} \leq C_{max} \quad (5.35)$$

where Δt is the timestep 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.

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

$$\partial_t T = \kappa \partial_x^{(2)} T, \quad (5.36)$$

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

$$\kappa \frac{\Delta t}{\Delta x^2} \leq \frac{1}{2} \quad (5.37)$$

is the condition necessary for stability. This puts a limitation on the maximum timestep allowed for a given spatial resolution, often increasing the computational cost as more timesteps are then required. Super-timestepping is one approach to tackling this issue, often employed in Newtonian hydrodynamics simulations. There, the CFL condition is (somewhat) bypassed by using sub-timesteps within a larger timestep. Only over the larger, super-timestep is the stability condition required to be satisfied. Although this requires a greater number of total (sub-)timesteps, the cost of the simulation is still reduced by allowing greater effective timesteps to be taken. See Alexiades, Amiez, et al., 1996, for instance, where a super-timestepping procedure is applied to an explicit time-integrator (forward Euler method) for the parabolic heat equation (the same one we have here).

For the CE heat system given by eq. (4.20),

$$\partial_t T = \kappa \partial_x^{(2)} T - \kappa^2 \tau_q \partial_x^{(4)} T, \quad (5.38)$$

the conditions for stability are slightly more complex, and are visualized in fig. 5.12.

The same analysis gives us the inequality for stability of

$$\kappa \frac{\Delta t}{\Delta x^2} + 4\tau_q \kappa^2 \frac{\Delta t}{\Delta x^4} \leq \frac{1}{2}. \quad (5.39)$$

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 eq. (5.39) above yields $\Delta t \leq \tau_q$.

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 general solution, Q , can be written as

$Q_l^n = Q^n \exp(il\alpha\Delta x)$ 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




FIGURE 5.12: The stability criteria of eq. (5.39) are plotted separately (dashed, dotted) and summed (solid). The red, shared region shows where the simple heat flux model given by eq. (5.38) 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 $\Delta x \simeq 0.045$.

MISCE sources will produce a solution growth rate per step, R , with the following form

$$R = 1 - \xi \frac{A\Delta t}{\Delta x^2} \sin^2\left(\frac{\theta}{2}\right) + \xi\tau \frac{B\Delta t}{\Delta x^3} (\sin(2\theta) - 2\sin(\theta)) + \xi^2\tau \frac{C\Delta t}{\Delta x^4} \sin^4\left(\frac{\theta}{2}\right) \quad (5.40)$$

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 crossover between the various stability criteria as resolution varies.

Firstly, 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 and Lindblom, 1983, conditions are provided for the stability (and causality) of MIS theory. There, the β -coefficients (written out in section 4.1 of this thesis), which represent the ratio of dissipation strengths ($\xi \equiv \{\zeta, \kappa, \eta\}$) to timescales ($\tau \equiv \{\tau_\Pi, \tau_q, \tau_\pi\}$), 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 $\tau_\Pi/\zeta = \beta_0 \gtrsim 0.1$ 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.

5.4.3 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, Rezzolla, 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 spacetime 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

$$\partial_t (\sqrt{\gamma}\mathbf{U}) + \partial_i (\sqrt{\gamma}\mathbf{F}^i(\mathbf{U})) = \sqrt{\gamma}\mathbf{S} \quad (5.41)$$

where

$$\mathbf{U} = \begin{pmatrix} D \\ S_j \\ \tau \\ X \\ Y^j \\ Z^{jk} \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi)W^2 v_j + W(q_0 v_j + q_j) + \pi_{0j} \\ (\rho + p + \Pi)W^2 + 2q_0 W - (p + \Pi - \pi_{00}) - nW \\ nW\Pi \\ nWq^j \\ nW\pi^{jk} \end{pmatrix},$$

$$\mathbf{F}^i = \begin{pmatrix} DV^i \\ \alpha S_j^i - \beta^i S_j \\ \alpha S^i - \beta^i E - DV^i \\ XV^i \\ Y^j V^i \\ Z^{jk} V^i \end{pmatrix}, \quad \mathbf{S} = \begin{pmatrix} 0 \\ \frac{1}{2}\alpha S^{ik}\partial_j \gamma_{ik} + S_i \partial_j \beta^i - E \partial_j \alpha \\ \alpha S^{ij} K_{ij} - S^j \partial_j \alpha \\ \frac{\alpha n}{\pi_\Pi} (\Pi_{\text{NS}} - \Pi + \Delta_\Pi) \\ \frac{\alpha n}{\tau_q} (\gamma_\mu^j q_{\text{NS}}^\mu - q^j + \Delta_q^j + \tau_q \mathcal{G}_q^j + \tau_q \mathcal{H}_q^j) \\ \frac{\alpha n}{\tau_\pi} (\gamma_\mu^j \gamma_\nu^k \pi_{\text{NS}}^{\mu\nu} - \pi^{jk} + \Delta_\pi^{jk} + \tau_\pi \mathcal{G}_\pi^{jk} + \tau_\pi \mathcal{H}_\pi^{jk}) \end{pmatrix}$$

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} & \mathcal{H} terms represent couplings between the dissipative variables and the gauge & 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} & \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 eq. (5.41).

For instance, derivatives of the (square root of the) spatial metric's determinant ($\sqrt{\gamma}$) appear in eq. (5.41). The spatial derivatives are fairly straightforward to calculate using finite-differencing. However, we would like 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

$$\alpha \partial_\mu \sqrt{\gamma} + \sqrt{\gamma} \partial_\mu \alpha = 0 \quad (5.43)$$

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

$$\partial_t \sqrt{\gamma} = \frac{-\sqrt{\gamma}}{\alpha} \partial_t \alpha \quad (5.44)$$

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

5.5 Summary

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 timescale 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 vice versa. Because much of the matter in a neutron star may be treated as a near-ideal fluid, the MISCE equations of motion are often able to 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 timescales 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 5.4.2, and is dependent both upon the dissipation coefficients (in particular the ratio of strength to timescale) and the simulation’s spacetime 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 Runge-Kutta 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 5.1, 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. 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 onto compact objects.

Chapter 6

A Covariant Filtering Implementation

6.1 Introduction

So far, we have seen how non-ideal fluids are physically important for neutron stars in chapters 1 and 2, and how they can be practically modelled in chapters 4 and 5. However, as discussed in section 2.3, the subgrid-scale (SGS) models used to capture small-scale behaviour may be interpreted as non-ideal fluid models, too. This is particularly important for large-eddy scheme (LES) approaches, which we introduce and then present our work on in this chapter.

The LES strategy consists of two key steps. First, one introduces a filtering operation with an associated filter kernel and filter length-scale. This is done in order to separate a field (or quantity) into a ‘resolved’ part that varies on large scales (that is, scales larger than the filter length-scale) and small-scale fluctuations. Second, the filtering operation is applied to the equations of motion. The net result is that, whenever the equations of motion contain non-linearities, filtering introduces residual terms that capture the impact of the fluctuation dynamics on the resolved scales. By modelling these residuals—e.g. calibrating them on high-resolution simulations of turbulence—one can try to inform a large-scale simulation with relatively low resolution about dynamics happening at smaller scales. These steps form the basis of the LES strategy, in the Newtonian context as well as in relativity.

A significant drawback of all relativistic LES implementations to date is that none of them are covariant, as highlighted in, for instance, Eyink and Drivas, [2018](#); Celora, Andersson, et al., [2021](#). This is because the filtering procedure, which takes one from the ‘micro’-level (where everything is assumed to be known and calculable) to the ‘meso’-level (where simulations are typically carried out) is performed at the level of

the (foliation-based, or 3+1) equations of motion and not on, for example, the stress-energy tensor itself. This makes the procedure frame dependent and so must be performed with respect to a (perhaps implicitly) chosen observer. If we instead choose to set the filtering observer using the fluid’s physical state, we can define a unique observer that will give us the same filtering result independently of the frame of reference or coordinates used. These issues were discussed in Duez, Knight, et al., 2020 and subsequently addressed in Celora, Andersson, et al., 2021, where an alternative, covariant framework using a fibration was put forward. This, together with the higher-level discussion we provide in Celora, Andersson, et al., 2024, serves as the theoretical precursor to the present more practical work.

In this chapter we implement the fibration-based, fully covariant, LES strategy outlined in Celora, Andersson, et al., 2021. We lay out the practical logic in section 6.1.1 and apply the scheme to real numerical data in section 6.1.2. We then build a closure model in section 6.3 and consider the impact on the interpretation of the micro-physical parameters in section 6.4. The main purpose of this chapter is to illustrate the viability of our approach and explore the possible pitfalls of alternative strategies.

6.1.1 Logic of the Scheme

We now discuss a practical implementation of the framework presented in Celora, Andersson, et al., 2021, focusing on the logic of the scheme and deferring to section 6.1.2 a detailed proof of principle analysis. Nonetheless, we anticipate that here the scheme is intended to be applied to special relativistic box simulations of hydrodynamic turbulence. The covariance of the Lagrangian filtering and the fact that this leaves the metric unaffected, allow us to apply the scheme to special relativistic numerical data and lift the results of our analysis to any spacetime. In principle, this allows us to perform explicit large-eddy scheme simulations in curved spacetimes—as required for merger or post-merger—using a subgrid model fitted within our framework. Of course, re-calibration to the specific dynamics of other systems beyond those of our box simulations here may be required for this to be accurate.

The first cornerstone of the scheme is the construction of the observer U^μ , which will be used for filtering and post-processing the numerical data. In essence, the idea is to work around the fact that any filtering operation breaks covariance by linking the operation to a physically meaningful observer—not to a gauge-dependent foliation observer. Then, introducing Fermi-coordinates (see, e.g. Misner, Thorne, et al., 1973) associated with such an observer, one can show that the metric is unaffected by the filtering operation. These two features combined mean that the framework retains compatibility with the tenets of general relativity, while adopting the main ideas of the more traditional LES strategy.

In practice, as for turbulent flows we expect to have changes in the flow properties over small scales, we can introduce the filtering observer as the one that moves with the ‘bulk’ of the fluid. That is to say the observer is aligned with the (baryon) number current of the flow, at the coarse level. This means that, if we consider a ‘box’ in spacetime that is adapted to such an observer—that is, the observer identifies the time-like direction of the box—there should be no (or at least minimized) baryon flux through opposite sides of the box. To ‘extract’ the filtering observer from the numerical data, then, we start by constructing such a box around a grid-point using the fine-scale fluid velocity at the point, u^μ , as initial guess for the observer: $U^\mu = u^\mu$. Starting from an orthonormal basis for the foliation of the spacetime $\{e_{(t)}^\mu, e_{(x)}^\mu, e_{(y)}^\mu, e_{(z)}^\mu\} = \{(1, 0, 0, 0), (0, 1, 0, 0), (0, 0, 1, 0), (0, 0, 0, 1)\}$, we project first spatial leg of the tetrad as

$$E_{(1)}^\mu = e_{(x)}^\mu + U^\mu U_\nu e_{(x)}^\nu, \quad E_{(1)}^\mu E_\mu^{(1)} = 1. \quad (6.1)$$

The remaining two spatial legs can be projected in a similar fashion, namely

$$E_{(2)}^\mu = e_{(y)}^\mu + U^\mu U_\nu e_{(y)}^\nu - E_{(1)}^\mu E_\nu^{(1)} e_{(y)}^\nu, \quad E_{(2)}^\mu E_\mu^{(2)} = 1, \quad (6.2)$$

and similarly for $E_{(3)}^\mu$. Given the tetrad legs, we then consider a space-time volume \mathcal{V}_L aligned with the tetrad $\{U^\mu, E_{(1)}^\mu, E_{(2)}^\mu, E_{(3)}^\mu\}$ with representative lengthscale L about our current point and compute three residual terms

$$r_{(I)} = \int_{\mathcal{V}_L} E_{(I)}^\mu n_\mu d\mathcal{V}_L, \quad I = 1, 2, 3 \quad (6.3)$$

where $n^\mu = nu^\mu$ is the micro-scale baryon current. These residuals measure the average particle drift in the direction I over the space-time volume \mathcal{V}_L . We can then feed these residuals into a root-finding algorithm and reconstruct (from the roots) the observer U^μ that minimizes the particle drift on average¹.

Note that the observer-finding volume’s lengthscale, L , is a free parameter here and in theory represents a separate choice to the filtering kernel size introduced and discussed later. Perhaps it should be made inversely proportional to the Lorentz factor of the observer, to reflect relativistic length contraction for high-velocity flows. Here, however, we decide not to overcomplicate the matter and set the two to be identical. This means that for a turbulent fluid flow, where there are rapid spatial variations in fluid quantities, we at least know that the filtering is being done with respect to an observer (and within a volume) that minimizes the local baryon current drift with respect to itself.

¹We note that we have also implemented and tested a minimization algorithm based on a similar logic, but found this was slower and more sensitive to the initial guess. The root-finding algorithm, however, appears to be robust even for high fluid/observer velocities.

Having discussed the construction of the filtering observer, let us turn to the meso-model and filtering. Given discrete data from a box simulation, the first step is to set up a meso-grid that is aligned with the ‘micro one’ from the simulations². We do so by introducing a coarse-graining factor (CG) in order to be able to reduce explicitly (or not) the resolution of the meso-grid with respect to the micro-grid. Once the grid is set up, we compute the filtering observer at each point on the meso-grid using the procedure discussed above. Then, we perform a Lagrangian spatial filtering of all the relevant variable of the micro-model, where we stress that we filter in the spatial directions relative to the observer. Here we do this using a sharp box filter kernel, so that for a given quantity X we compute its filtered counterpart as

$$\langle X \rangle = \int_{V_L} X dV_L , \quad (6.4)$$

where V_L is a spatial volume adapted to $\{E_{(1)}^\mu, E_{(2)}^\mu, E_{(3)}^\mu\}$ with side L and the integral is performed using a Gauss-Legendre quadrature scheme. We note that, as the spatial directions relative to the filtering observer will in general be tilted with respect to the foliation, the procedure requires introducing abstract coordinates adapted to the tilted box at each point. This operation populates the meso-grid with the computed values for the Lagrangian-filtered meso-model variables. For the simplest case of an ideal, single fluid, which we will be focusing on in the following sections, this involves computing the baryon current $\langle n^\mu \rangle$, stress-energy tensor $\langle T^{\mu\nu} \rangle$ and thermodynamic pressure $\langle p \rangle$.

As extensively discussed in Celora, Andersson, et al., 2021, even if the micro-scale model is that of an ideal fluid with stress-energy tensor given by

$$T^{\mu\nu} = \rho u^\mu u^\nu + p(g^{\mu\nu} + u^\mu u^\nu) , \quad (6.5)$$

where ρ, p are the micro-scale energy density and pressure, filtering will introduce additional ‘non-diagonal’ terms into the stress-energy tensor akin to those that enter dissipative fluid models. In practice, we introduce a ‘Favre’ observer (FO) \tilde{u}^μ at the meso scale as the one associated with the filtered baryon flux $\langle n^\mu \rangle \equiv \tilde{n} \tilde{u}^\mu$ and write the filtered stress-energy tensor as

$$\langle T^{\mu\nu} \rangle = (\tilde{\rho} + \langle p \rangle) \tilde{u}^\mu \tilde{u}^\nu + \langle p \rangle g^{\mu\nu} + 2\tilde{q}^{(a} \tilde{u}^{b)} + \tilde{s}^{\mu\nu} . \quad (6.6)$$

The energy density as measured by the Favre observer is denoted as $\tilde{\rho}$, while the additional residual terms \tilde{q}^μ and $\tilde{s}^{\mu\nu}$ measure effective momentum flux and stresses introduced by the filtering. In essence, these quantities are intended to capture energy and momentum transfer to/from the smaller scales that have been filtered out explicitly. Specifying a closure scheme in this picture amounts to introducing a

²We note that it is not strictly necessary to have the two grids aligned, but this clearly is a sensible choice, particularly for comparison at each grid point.

suitable representation of these residuals in terms of resolved/filtered quantities. Given the evident analogy with non-dissipative fluids and based on the classical work by Smagorinsky, 1963 (see also Radice, 2017; Radice, 2020), in Celora, Andersson, et al., 2021 a closure scheme that builds precisely on this formal connection was proposed. Here we choose to model the residuals as

$$\tilde{s}^{\mu\nu} = \frac{1}{3}\tilde{\Pi}\tilde{\perp}^{\mu\nu} + \tilde{\pi}^{\mu\nu}, \quad \tilde{\Pi} = \zeta\tilde{\theta}, \quad \tilde{\pi}^{\mu\nu} = -\eta\tilde{\sigma}^{\mu\nu}, \quad \tilde{q}^\mu = -\kappa\tilde{T}\tilde{\perp}^{\mu\nu}\left(\frac{1}{\tilde{T}}\nabla_b\tilde{T} + \tilde{a}_b\right), \quad (6.7)$$

where \tilde{T} is the meso temperature, $\tilde{\perp}^{\mu\nu} = g^{\mu\nu} + \tilde{u}^\mu\tilde{u}^\nu$ is the orthogonal projector relative to the Favre observer, while $\tilde{\theta}$, $\tilde{\sigma}^{\mu\nu}$ and \tilde{a}^μ are, respectively, the expansion rate scalar, the shear rate tensor and the acceleration relative to the Favre observer:

$$\tilde{a}_\mu = \tilde{u}^\nu\nabla_b\tilde{u}_\mu, \quad (6.8a)$$

$$\tilde{\theta} = \tilde{\perp}^{\mu\nu}\nabla_\mu\tilde{u}_\nu, \quad (6.8b)$$

$$\tilde{\sigma}_{\mu\nu} = \frac{1}{2}\left(\tilde{\perp}_a^c\tilde{\perp}_b^d + \tilde{\perp}_b^c\tilde{\perp}_a^d\right)\nabla_c\tilde{u}_d - \frac{1}{3}\tilde{\theta}\tilde{\perp}_{\mu\nu}. \quad (6.8c)$$

As extensively discussed in Celora, Andersson, et al., 2021 and Radice and Hawke, 2024, filtering may also impact on the thermodynamics. As the equation of state can be thought of as a non-linear closure, there is no guarantee that its functional form will be preserved by filtering. This implies that there is some freedom in specifying the EoS and related thermodynamic quantities at the filtered level. We do not wish to expand further on this issue here; it will be discussed later in section 6.4. Nonetheless, we note at this point that this may impact on the closure scheme in eq. (6.7) as it requires the introduction and choice of a filtered EoS from which we can compute \tilde{T} given filtered data.

At this point, we have covered the salient points of the scheme, so we can turn our attention to a proof-of-concept implementation of the strategy. In fact, the aim of the next sections is to discuss how the effective dissipative coefficients can be extracted from box simulations and how we can use the results to calibrate a meso-model. Once this is done, the Lagrangian subgrid model is complete. Before we move on, however, we need to comment on the choice of the closure scheme that will be explored in the following sections. At first glance, and with the closure scheme above, the final model resembles the Eckart model for a dissipative fluid, which is well-known to be acausal and suffer from instabilities Hiscock and Lindblom, 1983; Hiscock and Lindblom, 1985.

As far as we are aware, however, all discussions of the stability and causality properties of the Eckart model assume constant transport coefficients. It is not clear what would happen to the ‘classical’ results if we relax this assumption. In fact, as we will see below, it appears that modelling the effective dissipative coefficients only in terms of thermodynamic quantities like temperature and density does not suffice. This casts doubts on the naïve extension of the classical instability results to the problems we are

interested in. Moreover, if the unstable wavenumbers are not resolved in the coarse-grained simulation, then instability would not be triggered at all. We note that a similar perspective has recently been discussed in the context of reaction-sourced bulk-viscosity in mergers (Celora, Hawke, et al., 2022), although bulk viscosity in mergers is mathematically equivalent to resummed MIS (Gavassino, Disconzi, et al., 2024) as mentioned previously in section 1.1.5.

On top of all these considerations, it is also important to note that the various models designed to overcome these issues are all built by extending or generalizing the first-order models of the Landau-Eckart family. Given the novelty of the strategy that is being explored and implemented in this work, and as a first stab at the problem, we will focus on extracting and modelling the effective transport coefficients as in eq. (6.7), which we may also consider as leading order results.

6.1.2 Box Simulations of the Kelvin-Helmholtz Instability

We now turn our attention to applying the practical filtering scheme laid out in Celora, Andersson, et al., 2021 to real simulation data. These data are produced using the `METHOD` codebase; a (3+1D) finite-difference codebase for simulating hydrodynamics in special relativistic spacetimes. It contains a range of explicit and implicit-explicit (IMEX) time-integrators, spatial reconstruction and flux schemes. The `METHOD` codebase has been used previously to simulate resistive magnetohydrodynamics (Wright and Hawke, 2019; Wright, 2020; Wright and Hawke, 2020) and dissipative, relativistic fluids (Hatton and Hawke, 2024). For our purposes here, we perform high-resolution simulations of Kelvin-Helmholtz instabilities using a relativistic, ideal fluid description prescribed by the Euler equations. The data are defined within a 2D ‘box’ domain where $x \in [0.0, 1.0]$; $y \in [0.0, 1.0]$. The domain is filled with two fluid regions of differing densities that flow past each other with velocities directed in the positive and negative y -directions. The inner region of lighter fluid exists roughly in $x \in [0.25, 0.75]$, whilst the outer region of heavier fluid exists elsewhere. The primitive variables are the density and the y -component of the velocity:

$$\begin{pmatrix} \rho \\ v_y \end{pmatrix} = \begin{cases} \begin{pmatrix} 0.1 \\ 0.5 \end{pmatrix} & x_L < x < x_R, \\ \begin{pmatrix} 1.0 \\ -0.5 \end{pmatrix} & \text{otherwise.} \end{cases} \quad (6.9)$$

The interfaces between the two regions—located at $x = x_L \simeq 0.25$ and $x = x_R \simeq 0.75$ —are perturbed in a random fashion to induce mixing of the two fluid regions and encourage the instability to grow. We do so by superimposing 10 different modes that are out of phase and of different amplitudes. Specifically, given

$N = 1, \dots, 10$, we independently perturb the x -position of both interfaces as

$$x \rightarrow x + 0.01 R_{1,N} \cos(R_{2,N} + 2N\pi x) \quad (6.10)$$

where $R_{1,N}$ are ten random numbers between 0 and 1, normalized by their sum, and $R_{2,N}$ are ten random numbers between $-\pi$ and π . For the simulations we use a ‘ Γ -law’ EoS of the form:

$$p = (\Gamma - 1)n\epsilon = (\Gamma - 1)(\rho - n), \quad \rho = mn(1 + \epsilon) \quad (6.11)$$

where p is the pressure, ϵ and ρ are (respectively) the specific internal energy and the energy density, while n is the baryon number density. The initial pressure is uniform, $p = 1.0$, and the adiabatic index is set to $\Gamma = 4/3$ as for an ultra-relativistic ideal fluid in Rezzolla and Zanotti, 2013. We use periodic boundaries in all directions. Whilst the simulations performed here, and the data analysed from them, are spatially two-dimensional, the filtering codebase is written to work in arbitrary spatial dimensions. The standard simulation resolution used below is $N_x = N_y = 800$.




FIGURE 6.1: The development of a Kelvin-Helmholtz instability for an ideal fluid with negligible viscosity. The number density is plotted and the simulation ran until a code time of $t = 16.0$. Breakdown of the interface, seeded by a random initial perturbation, occurs quickly and leads to vortex formation and the onset turbulence across the domain.

When storing METHOD output data for our filtering pipeline, we typically store a series of snapshots (up to ≈ 40) in rapid succession around a central time-slice. These foliations cover (discretely) a portion of spacetime which we are then able to work




FIGURE 6.2: The adjusted, integrated power spectra of the Lorentz factor for fluid in the Kelvin-Helmholtz simulation seen above in fig. 6.1. The inertial range, where the power scales as $k^{-5/3}$, would correspond to a horizontal line here. On the left, three curves for times $t = 4.0$, $t = 10.0$ and $t = 20.0$ are plotted for the micromodel data (straight from simulation). On the right, four spectra are plotted, all at $t = 10.0$: one for the micromodel data again, and three for the mesomodel with varying parameters. CG is the coarse-graining factor of the mesomodel's grid compared to the micromodel's one and FW is the filter width used for the averaging operation. In all cases the inertial range is seen over at least an order of magnitude of wavenumbers.

within. For example, when we calculate the filtering observer at a particular spacetime point, it will typically not be aligned with the normal vector of the foliation at that point. This means we will need to access data outside that snapshot to calculate it. Similarly, when we perform filtering of the microdata on a particular slice to obtain the mesodata, we filter over an $(N + 1)D$ box aligned with this filtering observer. This box will project beyond that particular slice and will also be tilted. Hence, we need additional snapshots around those we strictly filter on. The number we need depends on the filter size L which typically falls in the range $\{2 - 8\}\Delta x$.

Once we have obtained these mesodata, we need to calculate the non-ideal terms in our subgrid closure scheme from them. This will require taking derivatives of the filtered quantities. Spatial derivatives are straightforward to calculate across the meso-grid. However, to calculate temporal derivatives, we again require multiple timeslices of meso-data. Because we are performing covariant filtering, we want to keep time and space on an equal footing. We also want to take derivatives at nearby points in spacetime on fluid with similar global properties. For these reasons, we choose the spatial and temporal capture resolutions to match³ so that $\Delta t = \Delta x = \Delta y$. This property is maintained when we filter, so that if the resolution of the mesodata is half that of the microdata, this is true in the temporal direction, too. For example, when performing a high-resolution simulation we have $\Delta t = \Delta x = \Delta y = 0.00125$ for the simulation microdata. If we then filter with a coarse-graining factor of $CG = 4$, we will choose every fourth timeslice of microdata to work with, such that,

³It is not strictly required to have the time-gap in between timeslices exactly equal to the grid-spacing. This choice, however, is practical as it means that derivatives in space and time evaluated using the same differencing scheme will have the same level of accuracy.

$\Delta t_{meso} = \Delta x_{meso} = \Delta y_{meso} = CG \Delta t_{micro} = 0.005$ for the mesomodel grid. This demonstrates the need for many initial snapshots, so that even after successive trimmings, we have sufficient data at the end to do statistics on. Using simple arguments we can estimate how much data we should need, finding that these should cover a total time-interval of about⁴ $5L$. Given that the maximum value for this is $L = 8\Delta x$ and that we store snapshots separated by a time-gap equal to Δx , this gives a maximum total number of snapshots of ≈ 40 as reported above.

Figure 6.1 shows the evolution of the number density across time for the Kelvin-Helmholtz initial data described above. An asymmetric initial perturbation in the x -directed velocity across the interface leads to mixing of the two fluid regions, which differ in density and velocity. This causes large vortices to form which quickly destroy the distinction between the two fluid regions. In time, small-scale structure develops as kinetic energy cascades down the length scales within the simulation.

Figure 6.2 summarises results of the integrated power spectra across scales, both for the micromodel (simulation) data in the left panel and for filtered (also sometimes coarse-grained) data in the right panel. For this Fourier analysis, we follow the approach laid out in Beckwith and Stone, 2011 which is also seen in Wright and Hawke, 2019; Wright and Hawke, 2020. The results for the micromodel are plotted for three different times so that one may see the features develop. Particularly, a wide inertial range already exists at $t = 4.0$, indicated by the flatness of the blue curve. The inertial range shrinks in time as numerical viscosity plays an increasingly important role in dissipating kinetic energy, which is also shown by the overall lowering of the curves.

The results for the mesomodel are plotted together with the micromodel data in the right panel. A number of effects shown by the three mesomodel curves are worth highlighting. Firstly, their general shape is preserved with respect to the micromodel, at least at low and medium wavenumbers. This is reassuring as we do not expect a large statistical impact on scales well above that which we are filtering at. Secondly, the coarse-graining reduces the maximum wavenumber of modes present, which in turn reduces the extent of the spectrum. Finally, the filtering removes the steep drop-off in kinetic energy seen for the highest wavenumbers in the micromodel data. This is logically what we expect—averaging over small-scale behaviour will remove energy from these scales. However, we also see a long rising tail in the spectrum (red curve) where this filtering is not combined with matched coarse-graining. In this case, it appears energy has actually been shifted into the smallest scales. This is an aliasing affect, where the filtering operation maps some power at low frequencies to the highest resolvable frequencies and vice versa. As there is much more power at low frequencies, this leads to higher power at grid size frequencies. This is not problematic for our

⁴Given a central slice, one needs to move in time about $CG\Delta x/2$ to reach the next useful meso-grid slice, and then up to $2L$ during the observer root-finding step. Because we will work with $CG = L/\Delta x$ and evaluate derivatives using a second order centred finite differencing scheme, we need up to three slices and to span a total time-interval of $\approx 5L$.

analysis as our focus is on capturing the behaviour in the inertial range. However, to minimise this effect we will below make the filtering and coarse-graining widths equal.

The Observers and Filtered Data

Having discussed METHOD and the initial data used for the box simulations, let us turn our attention to the observers and filtering. The first step is to reconstruct the baryon current and stress-energy tensor as these quantities are typically not directly evolved in numerical simulations. Starting from the primitive/auxiliary variables and for the ideal fluid case under consideration, these can be obtained as

$$n^\mu = nu^\mu, \quad u^\mu = (W, Wv_x, Wv_y), \quad (6.12a)$$

$$T^{\mu\nu} = nh u^\mu u^\nu + p \eta^{\mu\nu} \quad (6.12b)$$

where W is the Lorentz factor, v_x, v_y the primitive velocities in the x or y direction, n is the baryon density, h the specific enthalpy, p the fluid pressure and $\eta^{\mu\nu}$ denotes the Minkowski metric. We stress that it is crucial to reconstruct these tensorial quantities from ones that are more conveniently evolved in a hydrodynamic code. The reason for this is that covariance of our Lagrangian scheme follows from filtering independently each component of a geometrically well-defined object, i.e. a tensor, and use these to reconstruct the tensor at the filtered scale. As such, for example, we cannot directly apply the filtering operation to the primitive velocities (v_x, v_y) as these are not the components of a tensor.

Assuming we have computed the baryon current n^μ using the available data from the box simulations, we have all the ingredients we need to find the filtering observers following the logic detailed in section 6.1.1. We demonstrate this by comparing, in fig. 6.3, the fine-scale velocity u^μ against the observers U^μ obtained with $L = 8\Delta x$, focusing on the time component. Data for this figure come from a snapshot around $t = 10.0$. In the middle and right panels we plot data over the entire grid, hence the differences between the two are mainly visible where the fine-scale data shows sharp gradients. As such, and in order to better compare them, we also plot their relative difference in the left-most panel. We can then confirm that the relative difference peaks where there are sharp gradients in the fine-scale velocity, and that typical values are of the order of a few percent. Let us also note that in constructing similar plots for the $x - y$ components, we tend to observe slightly higher relative differences of the order of 10%. Nonetheless, we decided to focus here on the Lorentz factor as this gives us a measure for how tilted we may expect the filtering box to be with respect to the foliation⁵. Finally, we tested the robustness of the algorithm used to find the observers against the choice of initial guess. To do so, we ran the algorithm with 200 different

⁵Since the filtering axes are tilted with respect to the foliation, an interval of $\delta x'$ according to the filtering observer would correspond roughly to a time interval in the foliation of $\delta t = \sqrt{1 - 1/W^2} \delta x'$.




FIGURE 6.3: Comparison between the fine-scale velocity u^μ (middle panel) and the filtering observers U^μ (right panel), computed using a box-length of $8\Delta x$ and focusing on the time component. Data from a snapshot taken at around $t = 10$. In the left-most panel we show the relative difference between the two velocities, noting that this is typically of the order of a few percent and peaks where u^μ presents sharp gradients.

initial velocities, obtained by rotating the fine-scale velocity at the point, and looked at the distributions of the observers found. We repeated the process at several hundred points randomly selected on the grid. In each of these cases, the distributions of the observers found is very well described by a delta-function⁶ thus showing insensitivity to the initial guess. This fact, together with the clear physical interpretation of the algorithm implemented, is of crucial importance for the covariance of the scheme. We now focus our attention on the impact of the Lagrangian filtering. As an example, we show in fig. 6.4 the time component of the micro-scale and filtered baryon current. In order to appreciate visually the impact of filtering, in the right and middle panel we zoom in on the region $x \in (0.5, 0.7)$, $y \in (0.15, 0.35)$. The effect of filtering is visible as it smooths out the features present in the fine-scale data. As the zoomed in patch present several ‘bands’ of different densities, the effects of filtering are particularly evident in that it smears out the ‘boundaries’ between them. In addition, the left-most panel of the figure shows the relative difference between fine-scale and filtered data over the full grid. We then observe that typical values range from a few percent up to 10%.

Larger filter-size means interpolating (or combining more generally) data coming from larger regions around a point. As such, when comparing filtered meso-scale data to the micro-scale, we expect larger differences for larger filter sizes. This would follow automatically had we filtered in the spatial directions identified by the foliation, but such a statement may not be that trivial given the Lagrangian filtering explored in this work. In a sense, this provides a useful sanity check of the strategy and the pipeline implemented. To verify this, we compared the relative differences between filtered and fine-scale data at various filter widths and observed that the maximum relative difference appears to (roughly) double as we double the filter size. In fig. 6.5 we show the (time-component of the) fine-scale baryon current n^μ and its relative difference for

⁶More precisely, we find standard deviations of 11 to 12 orders of magnitude smaller than the mean. Consistently with this, we used a Gaussian kernel density estimator to estimate the observers probability density function, finding this to be clearly limiting to a delta function.




FIGURE 6.4: Comparison between the fine-scale n^μ and filtered $\langle n^\mu \rangle$ baryon current (filter-size $L = 8\Delta x$), focusing on the time component, at a representative time $t = 10$. The left panel shows the relative difference between the two over the full grid. In the middle and right panel we plot the fine-scale and filtered data zooming in the region $x \in (0.15, 0.35)$, $y \in (0.5, 0.7)$ (the box indicated in the left panel). We do so to visually appreciate the effects of filtering: the right panel presents the same ‘bands’ as the middle one but the boundaries between them are smeared out.




FIGURE 6.5: Comparison between the fine-scale n^μ and filtered $\langle n^\mu \rangle$ baryon current, focusing on the time component, at a representative time $t = 10$. The left-most panel shows the fine-scale data, whilst in the panels to its right we plot the relative difference between this and the filtered data—with filter sizes of $2, 4, 8\Delta x$. As we observed the maximum values of these to roughly double as we double the filter-size, we plot here the re-scaled relative differences, that is we divide by the filter-size L . This allows us to use a single colour map for the three panels to the right and provides a better comparison among them.

data filtered with filter lengths of $2, 4, 8\Delta x$. Given the observed scaling in the maximum value—this is about 4 to 5% for $L = 2\Delta x$, 8 to 10% for $L = 4\Delta x$ and up to 15 to 20% for $L = 8\Delta x$ —we plot the relative difference re-scaled by the filter size L . Even though the impact of filtering is less visible in this figure due to the linear scaling of the colour-map, we confirm the anticipated increase with the filter size. We will not try to make a more precise statement at this point given that i) we will get back to this issue in section 6.3 ii) the aim here is merely to demonstrate an increasing impact of filtering in a loose sense. Moreover, we note that here we have varied the filter size while keeping the observers fixed. For these tests, we have also constructed the grid for the filtered data as having the *same* number of points as the fine-scale data. We stress that we chose to do that in order to focus on the impact of the explicit filtering operation implemented. Up until now, we have focused on individual aspects of the




FIGURE 6.6: Comparison between fine-scale baryon current n^μ and filtered data using various filter-widths. The left-most panel shows a snapshot of the fine-scale data at around $t = 10$. In the panels to its right we plot the filtered baryon current with filter-sizes $2, 4, 8 \Delta x$, zooming in the region $x \in (0.1, 0.4)$, $y \in (0.4, 0.7)$. The corresponding patch is indicated in the left panel. Data for this figure has been produced using an observer box-length equal to the filter size and explicitly coarse-graining the meso grid. Coarse-graining is particularly evident in the third and fourth panel, which makes the images appear pixelated.

scheme one at a time, starting from the construction of the observers and then moving on to showing some examples of the filtered data we obtain. We chose to do so for two reasons: i) this is intended as a demonstration of the key points of the logic briefly presented in section 6.1.1; ii) given the novelty of the strategy adopted and of the codebase written, these tests allows us to gain confidence in the validity and robustness of the analysis. In particular, we discussed some properties of the filtering observers computed using $L = 8 \Delta x$ in fig. 6.3. Then we considered the impact of filtering at a fixed filter-size and also varied this (while keeping the observers fixed) to show how the impact of filtering scales with the filter-size (cf. fig. 6.5). In all these previous cases, the filtered data has been produced simply by interpolating data coming directly from simulations, without reducing the number of points of the meso-grid with respect to the fine-scale one. As such, it makes sense to conclude this section by considering the above-mentioned aspects all at once. In fig. 6.6 we show a comparison of the fine-scale and filtered baryon current varying the filter-size and, at the same time, recomputing the observers for each value of this. Similarly, we have also increased the grid-spacing for the filtered data accordingly: if the filter-size is $2 \Delta x$ (as is the observers' box length) the number of points in the meso-grid where such data is stored is reduced by a factor of 2 (in each direction), and so on. We do so because our ultimate goal is to inform 'low-resolution' simulations and make them (somehow) sensitive to physics happening on small scales without increasing the resolution. In essence, we need to keep track of the information that is lost in reducing the grid-resolution.

6.2 Stress-Energy Residuals and Closure Ingredients

Let us now turn to the modelling of the filtered/meso-scale stress-energy tensor $\langle T^{\mu\nu} \rangle$. As discussed in section 6.1.1, filtering introduces additional off-diagonal components in the stress-energy tensor—the residuals—which we need to model. We demonstrate this by first computing the Favre-observers from the filtered baryon current as

$$\tilde{n} = \sqrt{-\langle n^\mu \rangle \langle n_\mu \rangle}, \quad \tilde{u}^\mu = \langle n^\mu \rangle / \tilde{n} \longrightarrow \langle n^\mu \rangle = \tilde{n} \tilde{u}^\mu, \quad (6.13)$$

and decomposing the stress-energy tensors with respect to these. In essence, we compute the residual stresses and momentum flux as

$$\tilde{\Pi} = \frac{1}{3} \tilde{\perp}_{\mu\nu} \langle T^{\mu\nu} \rangle - \langle p \rangle, \quad (6.14a)$$

$$\tilde{\pi}^{\mu\nu} = \tilde{\perp}_c^\mu \tilde{\perp}_d^\nu \langle T^{cd} \rangle - (\langle p \rangle + \tilde{\Pi}) \tilde{\perp}^{\mu\nu}, \quad (6.14b)$$

$$\tilde{q}^\mu = -\tilde{u}_c \tilde{\perp}_b^\mu \langle T^{bc} \rangle, \quad (6.14c)$$

where we recall that $\tilde{\perp}^{\mu\nu} = \eta^{\mu\nu} + \tilde{u}^\mu \tilde{u}^\nu$. Note that the equation of state of the meso-model pressure $\langle p \rangle$ need not be that of the micro-model. For now we make the simplifying assumption that they match and investigate this point further in section 6.4. An example of such residuals is given in fig. 6.7, where once again the data is coming from a snapshot at around $t = 10.0$ and the simulation data has been filtered (and coarse-grained) using a filter size of $8 \Delta x$. As \tilde{q}^μ , $\tilde{\pi}^{\mu\nu}$ are tensors, in the middle and right panel of the figure we plot $\sqrt{\tilde{\pi}_{\mu\nu} \tilde{\pi}^{\mu\nu}}$ and $\sqrt{\tilde{q}_\mu \tilde{q}^\mu}$ as these give us a measure of the overall magnitude. We then confirm these residuals terms are generically non-negligible, as expected based on theoretical grounds.

Even though we have performed the decomposition in terms of the Favre observers \tilde{u}^μ , this is not a unique choice, certainly not from a theoretical point of view. We could have, for example, used the filtering observers U^μ for the decomposition instead, noting that this would not break covariance as the filtering observers have a clear physical meaning in our framework. We choose to work with the Favre observers for two reasons: first, there would be additional drift terms in the filtered baryon current had we decomposed it with respect to U^μ . These additional drift terms are essentially re-absorbed into the definition of \tilde{u}^μ (see Celora, Andersson, et al., 2021) and would need to be modelled otherwise. Second, a similarly defined Favre velocity is used in most of the Newtonian LES work on turbulence (Schmidt, 2015). In any case, given the way we construct the filtering observers, we have checked that there are no large differences between U^μ and \tilde{u}^μ as expected. Having discussed the residuals in the stress-energy tensor, let us turn to their modelling. To do so we first need to compute various gradients of the Favre velocity \tilde{u}^μ and of the temperature \tilde{T} (cf. eq. (6.7)), noting that our closure scheme will require us to combine spatial and temporal derivatives. This is evident if we look at the definition of the shear tensor $\tilde{\sigma}_{\mu\nu}$ in




FIGURE 6.7: Plotting the magnitude of the stress-energy tensor residuals. From left to right: the isotropic residual stresses $\tilde{\Pi}$, the anisotropic residual stresses $\tilde{\pi}^{\mu\nu}$ and residual momentum flux \tilde{q}^μ . Data underlying this figure is obtained from filtering a snapshot at around $t = 10$ using a filter-size of $8\Delta x$.

eq. (6.8): the shear matrix is, by definition, spatial with respect to the Favre observer, but the observer’s spatial directions are tilted with respect to the foliation. We stress that this is not a peculiar feature of the closure scheme we are exploring here, rather it is reasonable to expect that this will happen with *any* covariant closure scheme. For example, if we want the closure scheme to account for the fact that we expect turbulent transport only if there is non-zero shear in a local Lorentz frame moving with the fluid, then this will inevitably require it to mix time and space derivatives (cf. Duez, Knight, et al., 2020). Having said that, we compute space and time derivatives with respect to the foliation—using standard second-order centered finite-differencing—and then from these we reconstruct, say, the velocity gradient decomposition as in eq. (6.8).⁷

Computing spatial derivatives is fairly straightforward since it requires using filtered data coming from a single snapshot. Time derivatives, on the other hand, require some extra work as we also need filtered data coming from different time slices: for a centred first-order finite-differencing scheme we need 3. We also note here that the time gap between the slices play a role at this point. Since space and time derivatives are here computed using the same differencing scheme we choose the time gap between the stored simulation snapshots to be *equal* to the spatial grid spacing. As a first pass, we decided in this work not to interpolate the gridded data. Therefore, having the slices separated by a time gap equal to the grid spacing implies that all derivatives are computed with the same accuracy: at the same order and the same increment.

In fig. 6.8 we show an example of some gradients relevant to the modelling of the residuals plotted in fig. 6.7. In particular, in the left we show the Favre-observers’ expansion rate $\tilde{\theta}$, while in the middle and right panel we plot $\sqrt{\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}}$ and $\sqrt{\tilde{\Theta}_\mu\tilde{\Theta}^\mu}$, where $\Theta^\mu = \tilde{\mathcal{I}}^{\mu\nu}(\nabla_\nu\tilde{T} + \tilde{T}\tilde{a}_\nu)$. Comparing figs. 6.7 and 6.8, we immediately observe a nice degree of spatial correlation between quantities plotted in the middle and

⁷A convenient way of doing this is to store in memory the various derivatives as components of a tensor. This way we can directly apply the algebraic decomposition as in eq. (6.8). We also note that such decomposition makes explicit use of some algebraic constraints which may be violated due to




FIGURE 6.8: Results for various gradients relevant to the modelling of the stress-energy tensor residuals. From left to right we show: the expansion rate $\tilde{\theta}$, the square-root of the trace of the squared shear tensor $\sqrt{\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}}$, and the temperature gradients corrected by the heat inertia $\sqrt{\tilde{\Theta}_\mu\tilde{\Theta}^\mu}$, where $\tilde{\Theta}^\mu = \tilde{\mathbb{T}}^{\mu\nu}(\nabla_b\tilde{T} + \tilde{T}\tilde{a}_b)$. Data for this figure was obtained from filtering a snapshot at around $t = 10$ using a filter size of $8\Delta x$.

right-most panels. Given this, it makes sense to conclude this section by considering the simplest model for the residuals we can think of. Let us first consider the anisotropic residual stresses $\tilde{\pi}^{\mu\nu}$ and residual momentum flux \tilde{q}^μ , as well as the corresponding relevant gradients $\tilde{\sigma}^{\mu\nu}$, $\tilde{\Theta}^\mu$. At each point on the meso-grid, we ‘square’ these and compute their ratio as

$$r_1 = \frac{\sqrt{\tilde{\pi}_{\mu\nu}\tilde{\pi}^{\mu\nu}}}{\sqrt{\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}}}, \quad \text{and } r_2 = \frac{\sqrt{\tilde{q}_\mu\tilde{q}^\mu}}{\sqrt{\tilde{\Theta}_\mu\tilde{\Theta}^\mu}}. \quad (6.15)$$

We then average over the values computed at each grid point—meaning that we average the ratios to obtain \bar{r}_1 , \bar{r}_2 —and construct the models for the residuals as $\tilde{\pi}_{mod}^{\mu\nu} = \bar{r}_1\tilde{\sigma}^{\mu\nu}$ and $\tilde{q}_{mod}^\mu = \bar{r}_2\tilde{\Theta}^\mu$. We do the same for the isotropic residual stresses $\tilde{\Pi}$, noting that we construct the model for this using the absolute value of $\tilde{\theta}$. This is because, while $\tilde{\Pi}$ is positive at all gridpoints, the expansion rate is not (cf. figs. 6.7 and 6.8).

We compare the results to the residuals by plotting the corresponding distributions in fig. 6.9. The first thing we note is that, while we do not expect this model to be particularly good or accurate, the model for the residual momentum flux (right-most panel of the figure) is actually better than expected: while the mean values of the distributions are slightly different, there is a nice overlap between the two. This is suggestive of the fact that a constant heat conductivity model might be not too far off. Having said that, let us turn to the isotropic residual stresses (left-most panel of the figure), which is visibly the worst case of the three. That this was going to be the worst case can be easily explained: when comparing the left-most panels of fig. 6.8 with that of fig. 6.7, we see that $\tilde{\theta}$ shows complex dynamics that are not perfectly represented by $\tilde{\Pi}$. In this sense, it is evident that a constant coefficient model for the

numerical errors. Large deviations from these constraints should be corrected for, although we checked explicitly and this appears not to be required in our case.




FIGURE 6.9: Comparing distributions of stress-energy tensor residuals with their values from a first-order gradient model that assumes constant coefficients. From left to right we show: the isotropic residual stresses $\tilde{\Pi}$, the anisotropic residual stresses $\tilde{\pi}^{\mu\nu}$ and the residual momentum flux \tilde{q}^μ .

isotropic residual stresses will not be particularly good. Also consider the anisotropic residual stresses (middle-panel of the figure). We observe that, while the model almost correctly captures the mean value, it is not able to reproduce either the spread nor the skewness of the distribution. The next section will be devoted to discussing how we can improve on this crude model.

6.3 Discriminating Between Models

The aim of this section is to discuss some initial ideas on how to improve the modelling of the extracted transport coefficients and their corresponding residuals. We will focus on the anisotropic residual stresses, given that these are expected to play the largest role in determining the dynamics of a shearing flow like the one in the Kelvin-Helmholtz instability. We also conducted a similar analysis for the remaining residuals but as the qualitative discussion and key messages are essentially the same, we report on them in section 6.5.

6.3.1 Scaling with filter width

Perhaps the most fundamental thing we need to discuss is how the residuals and the corresponding effective dissipative coefficients change as we vary the filter size. In order to discuss this, we need to compare the distributions obtained with different filter sizes. For the case of the shear viscosity, the relevant distributions to look at are those of $\tilde{\pi}_{\mu\nu}\tilde{\pi}^{\mu\nu}$ and $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}$. These are shown in fig. 6.10 for the three filter sizes considered in this work, $L = 2, 4, 8\Delta x$. Quite notably, while the distributions of $\tilde{\pi}_{\mu\nu}\tilde{\pi}^{\mu\nu}$ change by a constant shift in log space, those of $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}$ do not. In particular, the shift between the $L = 2\Delta x$ and $L = 4\Delta x$ appear to be the same as that between the $L = 4\Delta x$ and $L = 8\Delta x$ distributions—the shift being roughly 1.2 in log-space. As for the shear viscosity coefficient, we compute this (as before) at each grid point as




FIGURE 6.10: Comparing distributions at different filter sizes. From left to right we show: i) the second invariant of the shear tensor, ii) the ‘magnitude’ of the anisotropic stress residuals and iii) the re-scaled, extracted shear viscosity.

$\eta = \sqrt{\tilde{\pi}_{\mu\nu}\tilde{\pi}^{\mu\nu}} / \sqrt{\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}}$. Given that the distributions for $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}$ do not change with the filter size—hence corroborating the idea that this quantity is measuring an intrinsic property of the flow—we then expect the η -distributions to scale as⁸ $(L\Delta x)^2$. The right panel of fig. 6.10 confirms this guess as we observe the scaled distributions to almost perfectly overlap. We stress that this scaling is consistent with naïve expectations based on Taylor-expansion type reasoning, although cannot be taken for granted given the way we implement filtering here. It at least represents a strong test of the robustness of the framework that we are exploring here.

6.3.2 Linear regression in log-space

Having discussed the residuals’ scaling with the filter size, let us turn to their modelling. Because the extracted values of η range over 3 orders of magnitude, we will pursue a linear regression model in log-space.

As a simple attempt, we can propose modelling η in terms of only the filtered temperature and baryon density, such that $\eta \equiv \eta(\tilde{T}, \tilde{n})$. This parametrization of a transport property in terms of thermodynamic variables is a common one motivated by viscosity’s microphysical origins. It turns out, however, that this gives very poor results: while we do not show this here, we compared the distributions of model predictions against data—both for the shear viscosity distribution as well as for the corresponding residual distribution—and observed very poor matching. This can be corroborated with a simple correlation plot amongst the variables $\{\eta, \tilde{T}, \tilde{n}\}$, which highlights the weak correlations between these quantities. This may be due to the fact that the filtered temperature \tilde{T} and density \tilde{n} span a range of only one order of magnitude. This is, however, common to most simulations of Kelvin-Helmholtz instabilities we are aware of as the simulation set-up we used is quite standard—e.g. uniform pressure in the initial data. We can try to explain this result by the fact that

⁸This is because $1.2 \approx \log_{10}(16)$.

the Kelvin-Helmholtz instability—used here to sustain the development of turbulence—is inherently a shear-instability whose dynamics are largely controlled by the relative velocity of the two fluids.

Given this, we expand the list of possible explanatory variables used in our regression, including in particular various scalar quantities involving gradients of the Favre-velocity \tilde{u}^μ . This choice is partly motivated by the observation that some of these quantities are much better correlated with the extracted shear viscosity and partly by the fact that we might expect these quantities to actually play a role on theoretical grounds. The full list of possible explanatory variables considered here is

$$\{\tilde{T}, \tilde{n}, \tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}, \det(\tilde{\sigma}), \tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}, \tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu} - \tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}, \tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}/\tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}\} .$$

In addition to the thermodynamic quantities, we include the invariants of the shear tensor, $\det(\tilde{\sigma})$, $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}$, that of the vorticity tensor, $\tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}$, as well as the difference and the quotient between the second invariants of these. We choose to add the last two quantities in the list since the former can be used to identify a vortex in a turbulent flow (Jeong and Hussain, 1995), whilst the latter is a measure of the relative intensity of the shear and vorticity and has recently been used in a discussion of the local-energy cascade in Newtonian hydrodynamic turbulence (Yao, Schnaubelt, et al., 2024).

Moreover, we can make sense of the need to consider such explanatory variables in the following way: the extracted effective transport coefficient capture the energy flowing to or from the scales we have explicitly filtered out, which we expect to happen as non-linearities kick in. Because we extracted these coefficients using formulae inspired from the modelling of viscous fluids and laminar flows, we can expect the intrinsic non-linearity of turbulence to manifest itself by having transport coefficients depending on velocity gradients as well.

Given this list of possible explanatory variables we construct the best model in the following way. We consider all possible sets and subsets of the 7 quantities listed above and for each we run a linear regression routine. Splitting the data into training and validation—using a ratio of 80:20—we identify the best model by contrasting the distributions of extracted data with the model predictions. We use the first Wasserstein distance, also known as Earth mover’s distance, to quantify the quality of a given model⁹: given two distributions X and Y the Wasserstein distance between them can be computed as

$$W_1(X, Y) = \sum_i \|X_{(i)} - Y_{(i)}\|_2 . \quad (6.16)$$

⁹We note that similar results can be obtained using Pearson’s correlation coefficient as a quality factor, although we found it less capable of discriminating between models.

where $\|X\|_2$ denotes the L_2 -norm of a quantity, X . In essence, once we construct the regression models, we look for the one that minimizes the distance between the data and the model predictions in a distributional sense.

6.3.3 Model interpretation and viability

In fig. 6.11 we summarize the results of this analysis. The first thing we note is that the best model we obtain does not involve thermodynamical quantities at all: the model involves only three of the quantities listed above, namely $\det(\tilde{\sigma})$, $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu} - \tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}$, $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}/\tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}$. This is reported in the left panel of fig. 6.11, where we show a scatter plot between the model and the data, while the annotation box reports the regression coefficients of the model. The middle panel of the figure instead shows a comparison between the distributions of the model and the data, where we observe that this model does not fully capture the longer tails of the data's distribution. Finally, we show in the right panel of the figure a comparison between the actual residual under consideration and the model's predictions. A strong matching between the two is seen. For the sake of clarity, let us comment on how we build the model's predictions for the residuals. We consider the validation set and for each sample point we build the prediction for η given the pointwise values of the model's regressors. Then we construct the residual predictions as $\tilde{\pi}_{model}^{\mu\nu} = \eta_{model} \tilde{\sigma}^{\mu\nu}$. Ultimately, the results shown in the right panel are the ones we truly care about. We see that the two distributions largely overlap. However, this simple model cannot quite explain the skew in the distribution data, nor its slightly longer tails. Nonetheless, the improvement with respect to the constant coefficient case above is significant and can be appreciated by comparing this panel with the middle one in fig. 6.9.

The data underlying fig. 6.11 has been filtered using a filter-size of $L = 8\Delta x$. It is also important, however, to check how the best model constructed in this way changes as we vary the filter size—both in the set of explanatory variables used and in the specific values of the coefficients. We have then performed the same process on data filtered with sizes $L = 2, 4\Delta x$ and found that the best set of explanatory variables is unchanged across the different cases and the regression coefficients take almost identical values. The only feature that changes is the offset, which increases by ≈ 0.6 as we double the filter size. This is consistent with the $(L/\Delta x)^2$ scaling shown in fig. 6.10 and provides a useful check on the robustness of the proposed modelling.

We conclude this section by commenting on an additional requirement for a model such as this to be viable. In fact, we need to make sure that the model we construct is regular in the laminar limit. Given the best model above, this would mean that the value for η should not diverge in the limit where $\tilde{\sigma} \rightarrow 0$. At a first glance, this might appear problematic since the regression coefficient for $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}/\tilde{\omega}_{\mu\nu}\tilde{\omega}^{\mu\nu}$ is negative. Nonetheless, we recall that $\det(\tilde{\sigma}) \sim \tilde{\sigma}^3$, while $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu} \sim \tilde{\sigma}^2$. As such, given that the




FIGURE 6.11: Best regression model in log space for η . From left to right we show: i) scatter-plot of the extracted data vs. model predictions ii) probability distributions of extracted η vs its modelling iii) probability distributions of the extracted residual vs its modelling.

coefficient in front of the $\det(\tilde{\sigma})$ regressor is positive and larger than that of $\tilde{\sigma}_{\mu\nu}\tilde{\sigma}^{\mu\nu}$, we assert that the model constructed is actually regular in the laminar limit.

6.4 The Equation of State Residual

So far in this chapter we have repeatedly mentioned, and then by-passed, the fact that filtering may impact on the thermodynamics, as well. In particular, we mentioned the fact that EoS residuals may introduce a disconnect between the microphysics and the EoS input in a numerical simulation. The aim of this section is to fill in this gap by taking a closer look at such residuals. We will also take the first step towards understanding the implications these may have if not accounted for.

To begin with, let us look back at eq. (6.6) and note that there are two residual contributions entering in the total trace of the filtered stress-energy tensor. We see a first one coming from the trace of $\tilde{s}^{\mu\nu}$, which we model in section 6.5, plus a second arising from the equation of state non-linearities. In fact, following from the discussion in Celora, Andersson, et al., 2021, we can write the filtered pressure as

$$\langle p \rangle = -\tilde{\rho} + \tilde{\mu}\tilde{n} + \tilde{T}\tilde{s} + M. \quad (6.17)$$

where M is the EoS residual, $\tilde{\rho}$ and \tilde{n} are the energy and number density (measured by the Favre observer \tilde{u}^μ), while \tilde{T} , $\tilde{\mu}$ are defined with respect to a meso-scale EoS of the form $\tilde{p} = \tilde{p}(\tilde{n}, \tilde{\rho})$. For example, if we choose to work with a ‘Γ-law’ EoS, the meso temperature is obtained as $\tilde{T} = \tilde{p}/\tilde{n}$. As non-linearities in the Gibbs relation can be interpreted as ‘entropy-like’ contributions associated with the fluctuations, there is some freedom in the choice of the filtered thermodynamic potential (e.g. the filtered entropy \tilde{s}). This choice will affect the extracted chemical potential $\tilde{\mu}$ and temperature \tilde{T} and, in turn, the EoS residual M . The most natural choice would be to work with

the same functional form for the equation of state, but this is a-priori a non-unique choice. See Celora, Andersson, et al., 2024 for a more in-depth discussion of this topic.

First of all, we consider the individual residual contributions to the trace of the filtered stress-energy tensor. These are plotted in fig. 6.12 for a snapshot at around $t = 10$ and filtering with size $L = 8 \Delta x$. For this figure we chose to work with the same EoS as used for the micro-model, that is a ‘ Γ -law’ EoS. We then observe that for this choice the EoS residual M is always negative, so we plot in fig. 6.12 the absolute value. In contrast, the $\tilde{\Pi}$ residual is always positive as is their sum, so that we have a net positive increase in the pressure term. This can also be appreciated looking at the relative distributions in the left-most panel of fig. 6.12. Having seen that the EoS




FIGURE 6.12: Visualizing the different residuals contributing to the trace of the stress-energy tensor. In the left-most panel we show distributions corresponding to the EoS residual, the isotropic stresses residual and the sum of the two. In the panels to its right we plot the same quantities individually. Data used in this figure has been filtered with filter-size $L = 8 \Delta x$ using snapshots around $t = 10$.

residuals M are generically non-negligible—their magnitude is of the same order as the others—it is natural to ask the ‘null-hypothesis’ question: what is the impact of neglecting this? This question is relevant in a broader sense as the effect of filtering will inevitably enter any numerical simulation due to the implicit filtering associated with the numerical discretization. As such we want to investigate whether these residuals can have a ‘measurable’ impact on the EoS parameters. To do so we extract locally the effective adiabatic index

$$\Gamma_1 = \left(\frac{\partial \log p}{\partial \log n} \right)_{\mathcal{S}} \quad (6.18)$$

where \mathcal{S} is the specific entropy, but first we need to rewrite this in terms of variables we have access to. Choosing to work with the energy and number densities (ρ, n) we arrive at

$$\Gamma_1 = \frac{n}{p} \left[\left(\frac{\partial p}{\partial n} \right)_{\rho} + \left(\frac{\partial p}{\partial \rho} \right)_n \left(\frac{\partial \rho}{\partial n} \right)_{\mathcal{S}} \right] = \frac{n}{p} \left[\left(\frac{\partial p}{\partial n} \right)_{\rho} + \frac{p + \rho}{n} \left(\frac{\partial p}{\partial \rho} \right)_n \right]. \quad (6.19)$$

As a quick sanity-check one can verify that for a Γ -law equation of state this expression gives $\Gamma_1 = \Gamma$ as expected. Then we consider the (corresponding) values for p, ρ, n

explored in the METHOD simulation and construct a smooth bi-variate cubic spline approximation for $p(n, \varepsilon)$. With this we can then take derivatives and evaluate eq. (6.19) at each grid-point. We do so for the micro-model data—that is the simulation output of METHOD—and compare it to the value used in the numerical code, $\Gamma = 4/3$. In the left panel of fig. 6.13 we plot the relative difference between the two and observe that this is at maximum of the order of 10^{-11} .

As for the meso-model, we do the same but assuming the null hypothesis, that is pretending that M is not there. We then follow the same logic as above but now in terms of the filtered pressure $\langle p \rangle$, energy density and number density $\tilde{\rho}, \tilde{n}$. Assuming that M has no impact at all is equivalent to say that $\langle p \rangle$ is related to $\tilde{\rho}, \tilde{n}$ by the same equation of state as at the micro-level. We test this by extracting Γ_1 locally using eq. (6.19)—although replacing now $p \rightarrow \langle p \rangle$, $\rho \rightarrow \tilde{\rho}$, $n \rightarrow \tilde{n}$ —and comparing it to the null-hypothesis value of $4/3$. The relative difference between the two is plotted in the right panel of fig. 6.13. In sharp contrast with the values obtained for the unfiltered data, we observe differences up to the *percent level*.

Let us try to spell out what possible consequences this may have by considering a binary neutron star merger simulation with finest resolution of $\mathcal{O}(10\text{m})$. The results shown suggest that unaccounted-for turbulent dynamics happening on scales of $\mathcal{O}(1\text{m})$ could induce up to percent level errors in the EoS parameters. While we would still need to quantify how much of a bias this introduces in EoS inferences and parameter estimation, it is clear this is an issue we need to explore in more detail. This is even more true if we consider that we expect turbulent dynamics to be induced at much smaller scales than $\mathcal{O}(1\text{m})$ in an actual merger, see Radice and Hawke, 2024. Having turbulent dynamics at smaller scales than $\mathcal{O}(1\text{m})$ would mean that we can consider the $\mathcal{O}(1\text{m})$ -simulation as effectively introducing a filter with size larger than $\approx 10 \Delta x$. As the residuals—including the EoS one, M —scale with the filter-size, the potential systematic biases this may introduce could be even larger than reported here.

6.5 Modelling the Isotropic Stresses and Momentum Flux Residuals

In this section we report on the modelling of the isotropic stresses and momentum flux residuals, following the same logic we took in section 6.3 for the anisotropic stresses. First, we observe that the distributions of the remaining residuals scale with the filter-size in the same fashion as for the anisotropic stresses: doubling the filter-size corresponds to a rigid shift of the distributions to larger values (the shift being ≈ 1.2 in log-space as above). Second, the distributions corresponding to the relevant gradients—the expansion rate $\tilde{\theta}$ and the temperature gradients $\tilde{\Theta}_\mu \tilde{\Theta}^\mu$ —do not change as we varied the filter-size. As in the case of the effective shear viscosity then, we




FIGURE 6.13: Comparing the 1st adiabatic index Γ_1 , computed locally using eq. (6.19), to the expected value of $4/3$. The left panel corresponds to the fine-scale data, with a largest relative difference on the order of 10^{-11} . The right panel corresponds to the filtered data, explicitly ignoring the EoS residual. For this case we observe differences up to percent level. The data underlying this figure has been filtered with filter size $L = 8 \Delta x$ using snapshots around $t = 10$.

expect the distributions of ζ and κ to scale with $(L\Delta x)^2$. We confirm this by plotting in fig. 6.14 the corresponding re-scaled distributions and find these are practically indistinguishable.

When it comes to the modelling, the only modification with respect to the shear viscosity case is that now we include additional quantities in the list of possible residuals: the new ones considered here are

$$\{\tilde{\theta}, \tilde{a}_\mu \tilde{a}^\mu, \tilde{\Theta}_\mu \tilde{\Theta}^\mu, (D_\mu \tilde{n})(D^\mu \tilde{n}), (D_\mu \tilde{n})\tilde{\Theta}^\mu, \dot{\tilde{n}}, \dot{\tilde{T}}\}.$$

These have been added as we found spatial gradients of the number density and the temperature to be effective at modelling the effective heat conductivity (cf. fig. 6.15). As for the bulk viscosity, we unsurprisingly found that adding the expansion rate to the list of regressors significantly improved the model's performance (cf. fig. 6.16). We also note that the best-fit exponent obtained for the expansion is close to -1 . In practice then, the dependence of the isotropic residuals $\tilde{\Pi}$ on the expansion rate almost cancels out. This is consistent with the poor correlation observed between $\tilde{\Pi}$ and $\tilde{\theta}$ and casts some doubts on the validity of interpreting the isotropic residuals as an effective bulk viscous pressure. Noting that this is not necessarily problematic, we also stress that in order to reach a firm conclusion would require using more data as well as more refined statistical models (and/or closure schemes). Given the scope of this work we leave this for the future.




FIGURE 6.14: Comparing distributions of ζ (left) and κ (right) at different filter sizes. The almost perfect overlap of the re-scaled distributions demonstrate that the coefficients (and the residuals) follow the same scaling as discussed for η .




FIGURE 6.15: Equivalent to fig. 6.11 but focusing on the heat conductivity. From left to right we show: i) scatter-plot of the extracted data vs. model predictions ii) probability distributions of extracted κ vs its modelling iii) probability distributions of the extracted residual vs its modelling.




FIGURE 6.16: Same as fig. 6.11 but focusing on the bulk viscosity. From left to right we show: i) scatter-plot of the extracted data vs. model predictions ii) probability distributions of extracted ζ vs its modelling iii) probability distributions of the extracted residual vs its modelling.

6.6 Extracting the Shear Viscosity, Component-Wise

In the main body of this work, we extracted the magnitude of the shear viscosity coefficient by ‘squaring’ the relation in eq. (6.7). In a sense, it would be natural to ask what would happen had we extracted the coefficient differently. For example, one could think of extracting η using the same relation as above but in a component-wise fashion: this would mean extracting 5 more independent distributions that eventually need to be combined in some way.

However, it is important to stress at this point that (in our view) this appears to be somewhat problematic, particularly because of the effort we made to set up the observers and perform the filtering in a covariant fashion. Phrasing this differently, we prefer the ‘squaring’ option as it allows us to stay ‘as covariant as possible’, while extracting η component-wise is at odds with one of the main arguments for performing the filtering in the way we do it in this work. Nonetheless, and mainly because we find this to be instructive, let us expand here on the results we get if we extract η from each component of eq. (6.7) independently.

In fig. 6.17 we compare the distributions for the absolute values of η obtained via ‘squaring’ vs component-wise. The first thing we note is that the various component-wise distributions are qualitatively similar to each other. We take this as an indication of the fact that while we expect numerical errors to have a larger impact on the component-wise distributions—for example, when a particular component of the shear tensor is small but, say, its second invariant is not—these are somehow distributed isotropically in a statistical/distributional sense. Second, when comparing the component-wise distributions against that obtained via squaring, we note that they all have the same mean but the latter has smaller variance. Stated differently, the component-wise distributions have longer tails with respect to the one obtained via squaring—which may be due to the impact of the above-mentioned numerical errors. While finding significantly different distributions would not be inherently problematic, the fact that we do not is reassuring. In fact, we may even suggest that the values obtained via squaring are reasonably consistent with those obtained in component-wise fashion while being free from the coordinate dependence of the latter.

Having said that, we feel the urge to comment on the sign to be assigned to η . As discussed in Celora, Andersson, et al., 2021, the sign of the effective transport coefficients cannot be constrained or fixed using thermodynamics arguments based on the second law of thermodynamics. In fact, we may even say that a positive shear viscosity corresponds to a positive energy cascade *locally*, where the energy is transferred to smaller scales, with the opposite being the case for negative values. It is also clear that in extracting the coefficient via ‘squaring’ we are only considering its magnitude, while the sign is not a priori fixed. We would then need a procedure/logic to assign the sign that is similarly covariant. One reasonable way of doing this is the

following: we can diagonalize the anisotropic stress residuals $\tilde{\pi}^{\mu\nu}$ and consider the direction associated with the maximum eigenvector (in absolute value). We could then transform the shear tensor to the basis given by the eigenvectors of $\tilde{\pi}^{\mu\nu}$, check whether $\tilde{\sigma}^{\mu\nu}$ is aligned or anti-aligned to $\tilde{\pi}^{\mu\nu}$ in this particular direction and assign a value to η accordingly. Arguably, this would be a geometric way of fixing the sign.

In fig. 6.18 we show the distributions of positive and negative values obtained this way and compare them to what we obtain if we follow the component-wise procedure. While the distributions differ for different components, we observe that those corresponding the $(0, 1)$, $(1, 1)$ and $(0, 2)$, $(1, 2)$ components are reasonably similar. We also observe generically that distributions of positive values have a larger spread. In the figure we chose to plot the histograms and also report in the legend the positive and negative counts: we can observe then a larger number of negative values in all cases with the exception of the $(0, 0)$ distributions. The bottom-right panel of the figure shows instead the distributions obtained via squaring and the logic described above for fixing the sign. Also in this case we observe a distribution of positive values with a larger variance and an overall larger number of negative values. In a loose sense then, this can be considered as a confirmation of the correctness of the logic used for fixing the sign, given that it also gives a larger amount of negative values that are generically smaller in magnitude.

We conclude this section by commenting on why we focused on modelling the magnitude of the shear viscosity only in section 6.3. In fact, it is not quite clear how seriously we should take the negative values extracted: some works in the literature simply discard these arguing they represent conservative spatial fluxes at the filter scale rather than interscale interactions (see e.g. Vela-Martín, 2022), even though the recent analysis in Yao, Schnaubelt, et al., 2024 appears to suggest we may want to take them seriously. We also stress that the simulations used here are 2+1 dimensional, while for real applications we will need to perform the analysis on data coming from 3+1 simulations. Besides, the phenomenology of 2-dimensional turbulence is quite different from the 3-dimensional case. In particular, we expect real 3-dimensional isotropic and homogeneous turbulence to have a net positive energy cascade. Assuming this would correspond to an overall predominance of positive over negative values—as reported in Yao, Schnaubelt, et al., 2024—and given the proof-of-concept nature of this work, we leave the exploration of these aspects for the future.

6.7 Extension to MHD

Above, a framework is presented for filtering an ideal, fine-scale fluid and modelling the result as a coarser, dissipative fluid. Generalising this idea, we are mapping one effective field theory (hydrodynamic model) to another. In principle, if we find that the




FIGURE 6.17: Extracting η component-wise: comparing the magnitude distributions obtained in a component-wise fashion (one per independent component) vs that via squaring. Data underlying this figure has been filtered with filter-size $L = 8\Delta x$ using snapshots around $t = 10$.




FIGURE 6.18: Comparing the histograms for positive and negative values of η , component-wise and via squaring. For each panel we report the total counts of positive vs negative values, so to be able to appreciate whether there are more positive than negative values or vice versa. Data underlying this figure has been filtered with filter size $L = 8\Delta x$ and snapshots around $t = 10$.

additional degrees of freedom (parameters) of the filtered model can be well-fitted to parameters of the starting model, then there exists a ‘strong’ mapping between the two.

Let us consider, for instance, ideal magnetohydrodynamics (iMHD) with an ideal fluid. The fine-scale SET is now

$$T^{\mu\nu} = \rho u^\mu u^\nu + p h^{\mu\nu} + F^{\mu\alpha} F_\alpha^\nu - \frac{1}{4} g^{\mu\nu} F_{\alpha\beta} F^{\alpha\beta}. \quad (6.20)$$

We may propose, as one would probably expect, that at the coarse level we have the filtered SET of a non-ideal fluid with (filtered) EM fields:

$$\bar{T}^{\mu\nu} = \tilde{\rho} \tilde{u}^\mu \tilde{u}^\nu + (\tilde{p} + \Pi) \tilde{h}^{\mu\nu} + q^{(\mu} \tilde{u}^{\nu)} + \pi^{\mu\nu} + \bar{F}^{\mu\alpha} \bar{F}_\alpha^\nu - \frac{1}{4} \tilde{g}^{\mu\nu} \bar{F}_{\alpha\beta} \bar{F}^{\alpha\beta}. \quad (6.21)$$

Or, perhaps more boldly, we may propose that the subgrid behaviour due to the presence of EM fields on the fine-scale may be captured solely by dissipative terms on the coarse scale, such that the filtered SET is simply that of a non-ideal fluid, as before:

$$\bar{T}^{\mu\nu} = \tilde{\rho} \tilde{u}^\mu \tilde{u}^\nu + (\tilde{p} + \Pi) \tilde{h}^{\mu\nu} + q^{(\mu} \tilde{u}^{\nu)} + \pi^{\mu\nu} \quad (6.22)$$

In either case, we will need to calculate the Faraday tensor in order to construct the fine-scale SET. If the vector potential, A^μ , is evolved by a simulation then we can calculate the Faraday tensor $F^{\mu\nu}$ as $F^{\mu\nu} = \nabla^\mu A^\nu - \nabla^\nu A^\mu$. The electromagnetic vector fields $\{E^\mu, B^\mu\}$ are also then given by $\mathbf{B} = \nabla \times \mathbf{A}$ and $\mathbf{E} = \partial_t \mathbf{A}$, at least for ideal electromagnetism. If the magnetic field alone is evolved in a simulation then the electric field may be calculated as $E^\mu = \epsilon^{\mu\alpha\beta} u_\alpha B_\beta$ and the Faraday tensor constructed from the EM vector fields.

6.8 Summary

Binary neutron star mergers are extremely rich and dynamical events whose quantitative modelling requires using numerical relativity simulations, particularly for the latest stages of the inspiral, the actual merger and the post-merger phase. The expected sensitivities of future third generation detectors, as well as the fact that these will also be able to detect the post-merger gravitational wave signal, justifies the ongoing efforts to improve the realism of said simulations. A key issue in this respect is the need to resolve (at least in principle) the full range of scales involved in the turbulent flow that develops in the merger remnant—e.g. due to the Kelvin-Helmholtz instability (Price and Rosswog, 2006).

Because this is (at the very least) not practical, recent years have witnessed various efforts to extend the large-eddy simulations’ strategy from Newtonian physics to relativistic (magneto-)hydrodynamics (Radice, 2020; Carrasco, Viganò, et al., 2020;

Viganò, Aguilera-Miret, et al., 2020a). While these efforts have already delivered impressive results (Aguilera-Miret, Viganò, et al., 2020; Aguilera-Miret, Viganò, et al., 2021; Palenzuela, Aguilera-Miret, et al., 2022; Aguilera-Miret, Palenzuela, et al., 2023; Izquierdo, Bezares, et al., 2024), there are a number of non-trivial aspects involved that are not yet fully understood. Certainly one of these is the fact that all practical implementations so far break covariance, both in the way the filtering operation is performed and in the specific closure scheme used (Radice and Hawke, 2024).

In this respect, in Celora, Andersson, et al., 2021 we have discussed some of these issues and put forward a coherent theoretical framework for performing large-eddy simulations while retaining compatibility with the covariance principle of general relativity. A higher-level discussion of the reasons for tackling such an issue in this way is provided in our recent paper (Celora, Andersson, et al., 2024). In this work however, we have taken a practical step forward by presenting the very first implementation of a fully covariant filtering scheme in relativity.

Being the first of its kind, this work is primarily intended as a demonstration that the logic discussed in Celora, Andersson, et al., 2021 is practically viable. We presented simulations of Kelvin-Helmholtz driven turbulence and used these to discuss the key features of the scheme. In particular, we dynamically identify a suitable observer that moves with the bulk of the flow and perform the filtering in the spatial directions identified by this observer. We demonstrate that, as expected from theoretical grounds, filtering impacts on the stress-energy tensor by introducing effective dissipative terms and that these scale with the size of the filter kernel as naïvely expected. As an illustration of the viability of the scheme, we have then provided a first ‘a-priori’ calibration of the effective transport coefficients that can be extracted directly from simulation data using a simple eddy-viscosity-type model. Finally, we have looked at the impact that filtering may have on the thermodynamics, finding indications that, if not properly modelled, this may introduce up to percent level differences in the EoS parameters. These results demonstrate the kind of questions we may try to answer by working within this framework.

The purpose of this work is also to introduce and present the codebase we have built¹⁰. In this respect, it is important to stress that the results presented here are not expected to be of direct use for real large-scale simulations of compact-object mergers. This is mainly because the simulations used here have two spatial dimensions, and do not contain electromagnetism. It is well known that three-dimensional turbulence is phenomenologically quite different from its two-dimensional counterpart. For example, vortices stretch each other in three-dimensional turbulence whilst in two-dimensional turbulence they simply move past each other. Two-dimensional turbulence may exhibit both a direct and inverse cascade of energy due to the additional conservation of

¹⁰The first release of the code used in this work is publicly available at <https://github.com/Lagrangian-filtering/Lagrangian-filtering>

enstrophy.¹¹ Similarly, magnetic turbulence exhibits an inverse cascade that amplifies large-scale magnetic fields, giving magnetar remnants their enormous magnetic field strength.

Nonetheless, we stress that the key routines have been written in such a way to work in $3 + 1$ dimensions as well. Similarly, the scheme has been structured in a pipeline-like fashion to easily allow for future improvements and extensions. For example, future work will range from extensions to account for electromagnetic effects, to a systematic exploration of various closure schemes and ‘a-posteriori’ tests of these. To this end, we note that our recent work (Hatton and Hawke, 2024) provides a computationally efficient formulation of dissipative hydrodynamics that can be used to implement a subgrid model of the type we derive here.

In summary, this work represents a first step towards a practical implementation of the covariant filtering scheme from Celora, Andersson, et al., 2021 and provides us with a tool that can be used to investigate a number of open issues regarding the modelling of relativistic turbulence.

¹¹Enstrophy is a scalar measure of (square averaged) vorticity, in an analogous way to energy being a scalar measure of (square averaged) velocity (for uniform mass).

Chapter 7

Other Avenues

In this chapter, we present further results from simulations of dissipative hydrodynamics. These pieces of work have not yet come to full fruition but are in a proof-of-concept stage currently. They provide viable avenues for further work in the direction of this thesis.

Firstly, in section 7.1, we introduce the numerics of the BDNK model and compare its simulation results to those of the MIS & MISCE models.

Then, we show preliminary results from performing simulations with an explicit large-eddy scheme in section 7.2. This uses the subgrid closure model derived in chapter 6 that mimics true dissipative hydrodynamics to capture unresolved dynamics.

7.1 Simulations with the BDNK Model

In this section, we present results from simulations using the the Bemfica-Disconzi-Noronha-Kovtun (BDNK) model of dissipative, relativistic hydrodynamics (Kovtun, 2019; Bemfica, Disconzi, et al., 2022). We also make both qualitative and quantitative comparisons with the MIS and MISCE models simulated previous in chapters 4 and 5.

Following the numerical work of Pandya, Pretorius & Most (Pandya, Most, et al., 2022b; Pandya and Pretorius, 2021) in particular, we seek to implement a first-order system of equations that has the advantage of provable stability, causality & hyperbolicity, as well as a smaller system size (compared to MIS) because the dissipative terms have constitutive relations rather than obeying relaxation equations.

We write our conserved vector as

$$\begin{pmatrix} D \\ S_j \\ E \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi + A)W^2 v_j + (q_j + q_k v^k v_j)W + \pi_j^0 \\ (\rho + p + \Pi + A)W^2 - (p + \Pi + A) + 2q_k v^k W + \pi_0^0 \end{pmatrix}, \quad (7.1)$$

whilst our flux vector is

$$\begin{pmatrix} Dv^i \\ S_j v^i + (p + \Pi)\delta_j^i + W(q^i v_j - q_k v^k v_j v^i) + \pi_j^i \\ (E + p)v^i + W(q^i - q_k v^k v^i) + \pi_0^i \end{pmatrix}, \quad (7.2)$$

and the source vector is

$$\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}. \quad (7.3)$$

Note that we have a new variable to define, and that some of our familiar variables now take on a different form:

$$A = \tau_\epsilon [u^\mu \nabla_\mu \rho + (\rho + p) \nabla_\mu u^\mu], \quad (7.4a)$$

$$\Pi = -\zeta \Theta + \tau_\Pi / \tau_\epsilon A, \quad (7.4b)$$

$$q^\mu = \tau_q (\rho + p) u^\nu \nabla_\nu u^\mu + \beta_\epsilon \Delta^{\mu\nu} \nabla_\nu \rho + \beta_n \Delta^{\mu\nu} \nabla_\nu n, \quad (7.4c)$$

$$\pi^{\mu\nu} = -\eta \left[\Delta^{\mu\alpha} \Delta^{\nu\beta} \nabla_\alpha u_\beta + \Delta^{\mu\alpha} \Delta^{\nu\beta} \nabla_\beta u_\alpha - \frac{2}{3} \Delta^{\mu\nu} \Delta^{\alpha\beta} \nabla_\alpha u_\beta \right]. \quad (7.4d)$$

Our C2P procedure is given by the following steps:

1. Compute A , Π , q_μ and π_ν^μ from the values of the primitives at the previous timestep. Note that this will require both temporal and spatial derivatives.
2. Guess $\bar{\chi}$ ($= p + \Pi + A - 2q_k v^k W - \pi_k^k$) and $\bar{\sigma}_j$ ($= (q_j + q_k v^k v_j)W + \pi_j^0$).
3. Compute $v^2 = \gamma^{ij}(S_i - \sigma_i)(S_j - \sigma_j)/(E + \chi)^2$
4. Compute $W = (1 - v^2)^{-1/2}$
5. Compute $n = D/W$
6. Compute $\rho + p = (E + \chi)/W^2 - (\Pi + A)$
7. Compute v_j from $(S_j - \sigma_j)/[(\rho + p + \Pi + A)W^2]$.
8. Invert our EOS to get p from n , $\rho + p$.
9. Compute $\chi = p + \Pi + A - 2q_k v^k W - \pi_k^k$, $\sigma_j = (q_j + q_k v^k v_j)W + \pi_j^0$ from our quantities and compare to our guesses.
10. Recompute A , Π , q_μ and π_ν^μ with the updated primitives set by this C2P scheme.

We have 6 free parameters: $\{\kappa, \zeta, \eta\}$ & $\{\tau_q, \tau_\Pi, \tau_\epsilon\}$. All act to moderate the strength of the additional, non-ideal terms (despite the notation, none are reciprocal timescales as for the MIS case). From some of these parameters we calculate two more coefficients, β_ϵ & β_n :

$$\beta_\epsilon = \tau_q \left(\frac{\partial p}{\partial \rho} \right)_n + \frac{\kappa T(\rho + p)}{n} \left(\frac{\partial(\mu/T)}{\partial \rho} \right)_n, \quad (7.5a)$$

$$\beta_n = \tau_q \left(\frac{\partial p}{\partial n} \right)_\epsilon + \frac{\kappa T(\rho + p)}{n} \left(\frac{\partial(\mu/T)}{\partial n} \right)_\epsilon, \quad (7.5b)$$

Given that our equation of state is $p = (\Gamma - 1)(\rho - n)$, the chemical potential is $\mu = (\Gamma - 1)(\rho - n)n^{-2}$ and temperature $T = p\rho^{-1}$ we have $\mu T^{-1} = \rho n^{-2}$ and hence

$$\beta_\epsilon = \tau_q(\Gamma - 1) + \frac{\kappa T(\rho + p)}{n^3}, \quad (7.6a)$$

$$\beta_n = -\tau_q(\Gamma - 1) - \frac{\kappa p(\rho + p)}{n^4} \equiv -\tau_q(\Gamma - 1) - \frac{\kappa T^4(\rho + p)}{p^3}. \quad (7.6b)$$

7.1.1 An Alternative Evolution Scheme

The above scheme is not that adopted in the latest work by Pandya, Pretorius & Most ("PPM") (Pandya, Most, et al., 2022b). Instead, they undertake a clever approach tied together by the need to *not* finite difference in time when computing the first derivatives needed to close the system.

We usually think of three sets of variables; the conserved variables that are evolved via a balance law form and associated high resolution shock capturing scheme (involving a total derivative of a flux plus a source), the primitive variables (any minimal independent set of variables from which everything else can be computed), and the auxiliary variables (anything else that is useful to store).

The PPM paper uses similar language, but the meaning of the terms is somewhat different. To try and be precise here we will call the *evolved* variables as those within the state vector that are explicitly updated at each (sub-)timestep, and the *basic* variables as the set which need updating at the end of each (sub-)timestep.

The conserved variables are $\mathbf{q} = (D, S, E)$. The primitive variables are $\mathbf{w} = (n, v, \rho)$. In the PPM paper, the evolved variables are $\mathbf{e} = (D, S, E, n, v, \rho) = (\mathbf{q}, \mathbf{w})$. The basic variables are $\mathbf{b} = (n, v, \rho, \partial_t n, \partial_t v, \partial_t \rho) = (\mathbf{w}, \partial_t \mathbf{w})$. The equations of motion can be written as

$$\partial_t \mathbf{e} + \partial_j \mathbf{f}^{(j)}(\mathbf{e}, \mathbf{w}, \partial_t \mathbf{w}) = \mathbf{s}(\mathbf{e}, \mathbf{w}). \quad (7.7)$$

In particular, the equations of motion for the conserved variables (the first half of the state vector) remain the same as before, whilst the equations of motion for the

primitive variables are trivial (as the time derivatives of the primitives are assumed known at this stage, and appear in the basic variables).

The key to the scheme is the conserved-to-primitive transformation, which should really be thought of as an ‘evolved-to-basic’ transformation. After updating one (sub-)timestep, the evolved variables are known, and we have to re-compute the basic and auxiliary variables. First we apply the boundary conditions to the evolved variables so they are known everywhere, including in the ghost zones. Next, compute the spatial derivatives of the primitive variables (as the primitive variables are known everywhere, as they are the second half of the evolved variables). Finally, write

$$\mathbf{q} \equiv \mathbf{q}(\partial_t \mathbf{w}; \mathbf{w}, \partial_j \mathbf{w}). \quad (7.8)$$

The point here is that the conserved variables (which are known – they have been updated in the evolution) are functions of the primitive variables (known similarly), the spatial derivatives of the primitives (computed by finite differencing above, so known), and the time derivatives of the primitives (unknown). As the number of primitive time derivatives ($\partial_t \mathbf{w}$) matches the number of conserved variables (\mathbf{q}), we have a system of nonlinear algebraic equations for which a root-finding method can be employed. Using, as an initial guess, the previous value of the time derivative, its value at the new time can be computed.




FIGURE 7.1: The evolution of discontinuous initial data (a shocktube) for the ideal and bulk-viscous ($\zeta = 1 \times 10^{-2}$) cases. The BDNK formulation is compared to the MIS model, showing excellent visual agreement in both cases.




FIGURE 7.2: The evolution of a spinning ‘rotor’ with initial data similar to that in Pandya, Most, et al., 2022b using the BDNK formulation with a shear viscosity coefficient of $\eta = 1 \times 10^{-3}$. The weak shear viscosity allows small-scale vortex structure to develop but eventually causes breaking of the fluid.

Writing the system explicitly

We will write out the system explicitly here for clarity, also reverting to our nomenclature for the various quantities that aligns with METHOD but still referring to evolved and basic variables when it is useful to. Our conserved vector is now in the same form as before but extended to contain (unusually) primitive fluid quantities

$$\begin{pmatrix} D \\ S_j \\ E \\ v_j \\ p \\ \rho \end{pmatrix} = \begin{pmatrix} nW \\ (\rho + p + \Pi + A)W^2 v_j + (q_j + q_k v^k v_j)W + \pi_j^0 \\ (\rho + p + \Pi + A)W^2 - (p + \Pi + A) + 2q_k v^k W + \pi_0^0 \\ v_j \\ p \\ \rho \end{pmatrix} \quad (7.9)$$




FIGURE 7.3: The long-term evolution of the Kelvin-Helmholtz instability using the BDNK formulation of non-ideal hydrodynamics with no non-ideal effects and 800×800 cells.

whilst our flux vector is

$$\begin{pmatrix} Dv^i \\ S_j v^i + (p + \Pi) \delta_j^i + W(q^i v_j - q_k v^k v_j v^i) + \pi_j^i \\ (E + p) v^i + W(q^i - q_k v^k v^i) + \pi_0^i \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad (7.10)$$

and the source vector is

$$\begin{pmatrix} 0 \\ 0 \\ 0 \\ \partial_t v_j \\ \partial_t p \\ \partial_t \rho \end{pmatrix}. \quad (7.11)$$

The primitive variables appearing in the second-half of the conserved vector are sourced simply from their time derivatives, which are also needed to calculate the non-ideal terms in the first-half of the conserved vector (Π , A , q_j and π_j^i).




FIGURE 7.4: The long-term evolution of the Kelvin Helmholtz instability using the BDNK formulation of non-ideal hydrodynamics with a viscosity parameter of $\eta = 1 \times 10^{-4}$ and 800×800 cells. The small but finite shear viscosity here initially suppresses the breaking of the interface but then causes large-scale features to develop (roll-up, vortices) with fine-scale details within them.

Our conserved-to-primitive scheme now reads significantly differently from previous formulations’.

1. Apply boundary conditions (BCs) so that the conserved variables are correct over our entire domain (interior and exterior).
2. Trivially recover our primitive variables $\{v_j, p, \rho\}$ from the conserved vector.
3. Compute all necessary spatial derivatives of these primitive variables (these calculations are only required for the interior region but will require values of the primitives in the exterior due to the finite differencing).
4. Guess four time derivatives ($\partial_t\{v_j, \rho\}$) required to calculate the non-ideal terms.
5. Calculate $\partial_t n$ using the continuity of mass equation: $u_\mu \nabla^\mu n = -n \nabla^\mu u_\mu$. This leads to $\partial_t n = -v^i \partial_i n - n \Theta / W$ where Θ is the 4-divergence ($= \nabla^\mu u_\mu$).
6. We now have all required quantities to calculate the entire conserved vector! ‘Simply’ root-find on the difference between the value of the conserved vector’s elements and their values computed algebraically (including using the time-derivative guesses).

7. Compute $\partial_t p$ using the chain-rule on our EoS and our calculated values of $\partial_t \rho$ and $\partial_t n$: $\partial_t p = (\Gamma - 1)(\partial_t \rho - \partial_t n)$.
8. We can now update the flux and source vectors with correct values.

We've carried out a number of tests of the BDNK formulation using both of these methods of solving the system. This is ongoing work but the short answer is that both methods are stable only within narrow parameter value ranges (of $\{\kappa, \zeta, \eta\}$ & $\{\tau_q, \tau_\Pi, \tau_\epsilon\}$). In particular, the second approach is highly unstable for weak dissipation. This is unsurprising but problematic. We attribute this instability to a couple of factors:

1. The root-find tries to match a small change in the conserved quantities to a tiny viscosity coefficient multiplying a time derivative, giving a highly sensitive and often artificially large time derivative. This issue is somewhat avoided in their (PPM) work by use of an algebraic C2P, made possible by heavy simplifications such as asserting that the fluid is conformal. We would like to keep our approach more general than this. One alternative is to avoid the root-find entirely by using approximate algebraic expressions for the time derivatives as described in appendix A. This, however, introduces errors into the time derivatives of at least $\mathcal{O}(\delta^2)$ where $\delta = \{\kappa, \zeta, \eta\}, \{\tau_q, \tau_\Pi, \tau_\epsilon\}$. These errors seem to accumulate and eventually destabilize the system anyway.
2. The spatial finite differencing may also be contributing to the problems facing us here. The feedback into the time derivatives could lead to instabilities. Whilst the flux calculations in METHOD are made using a high-order WENO scheme (see section 3.4), the spatial differentials in the non-ideal terms are currently calculated using low-order, central finite differencing. Increasing the order of this spatial differencing doesn't seem to improve the behaviour, but perhaps the more advanced CWENO methods adopted by PPM would.

We show the results of a number of simulations using familiar initial data with the BDNK model in figs. 7.1, 7.3 and 7.4. We also show results from a new test in fig. 7.2, one similar to that used in Pandya, Most, et al., 2022b. It is a simulation of a rotating circle of fluid - a 'rotor' - that mimics, in a highly simplified way, how a spinning compact object such as a neutron star may be affected by the viscosity.

Direct comparison to the other dissipative models in this thesis is difficult because the dynamical hydrodynamic frame of BDNK means an adjustment of certain primitive variables is necessary before their values represent equivalent physical quantities as they do in other formulations with a static frame. However, in fig. 7.1, we do make a favourable comparison to the MIS model for a shocktube test, and in the other simulations we at least see the *expected* behaviour of a dissipative fluid. Further work

more thoroughly analysing the BDNK model and the results it produces in simulations would be an interesting endeavor.

7.2 Subgrid Modelling with Dissipative Hydrodynamics

Building on our filtering work outlined in chapter 6, we will now discuss the steps required for a practical implementation of the subgrid closure obtained therein. This will serve as an *a posteriori* test of our previous work and (hopefully) demonstrate a viable approach to performing covariant large-eddy simulations. To this end, we present preliminary results from performing simulations of this nature.

We will be simulating a dissipative fluid at the meso-scale with stress-energy tensor (SET) given by

$$\tilde{T}^{ab} = (\tilde{\varepsilon} + \tilde{p}) \tilde{u}^a \tilde{u}^b + \tilde{p} g^{ab} + 2\tilde{q}^{(a} \tilde{u}^{b)} + \tilde{s}^{ab} \quad (7.12)$$

where we use the tilde symbol to designate these as meso-scale quantities that *haven't* been explicitly filtered. We chose to model the residuals as

$$\tilde{s}^{ab} = \frac{1}{3} \tilde{\Pi} \tilde{\perp}^{ab} + \tilde{\pi}^{ab}, \quad \tilde{\Pi} = -\zeta \tilde{\theta}, \quad \tilde{\pi}^{ab} = -\eta \tilde{\sigma}^{ab}, \quad \tilde{q}^a = -\kappa \tilde{T} \tilde{\perp}^{ab} \left(\frac{1}{\tilde{T}} \nabla_b \tilde{T} + \tilde{a}_b \right), \quad (7.13)$$

where \tilde{T} is the meso temperature, $\tilde{\perp}^{ab} = g^{ab} + \tilde{u}^a \tilde{u}^b$ is the orthogonal projector relative to the Favre observer, while $\tilde{\theta}$, $\tilde{\sigma}^{ab}$ and \tilde{a}^a are, respectively, the expansion rate scalar, the shear rate tensor and the acceleration relative to the Favre observer:

$$\tilde{a}_a = \tilde{u}^b \nabla_b \tilde{u}_a, \quad (7.14a)$$

$$\tilde{\theta} = \tilde{\perp}^{ab} \nabla_a \tilde{u}_b, \quad (7.14b)$$

$$\tilde{\sigma}_{ab} = \frac{1}{2} \left(\tilde{\perp}_a^c \tilde{\perp}_b^d + \tilde{\perp}_b^c \tilde{\perp}_a^d \right) \nabla_c \tilde{u}_d - \frac{1}{3} \tilde{\theta} \tilde{\perp}_{ab}. \quad (7.14c)$$

From the conservation of energy, momentum and mass we have the usual divergence equations:

$$\nabla_a \tilde{T}^{ab} = 0, \quad \nabla_a \tilde{n}^a = 0, \quad (7.15a)$$

where $\tilde{n}^a = \tilde{n} \tilde{u}^a$. Recasting these, using the ‘Valencia’ formulation, into conservation-law equations of the form

$$\partial_t \mathbf{U} + \partial_i \mathbf{F}^{(i)} = \mathbf{S} = \mathbf{0} \quad (7.16)$$

where

$$\mathbf{U} = \begin{pmatrix} D \\ S_j \\ \tau \end{pmatrix} = \begin{pmatrix} nW \\ (\varepsilon + p + \Pi)W^2 v_j + W(q_0 v_j + q_j) + \pi_{0j} \\ (\varepsilon + p + \Pi)W^2 + 2q_0 W - (p + \Pi - \pi_{00}) - D \end{pmatrix}, \quad \mathbf{F}^{(i)} = \begin{pmatrix} Dv^i \\ S_j^i \\ S^i - Dv^i \end{pmatrix} \quad (7.17a)$$

and

$$S_j^i = (\varepsilon + p + \Pi)W^2 v^i v_j + W(q^i v_j + q_j v^i) + (p + \Pi)\delta_j^i + \pi_j^i. \quad (7.18)$$

The dissipative coefficients $\{\zeta, \eta, \kappa\}$ will now be given by the functional form arrived at through statistical fitting. Specifically, by comparison of the non-ideal terms in this first-order model with the residuals produced by explicitly filtering the micro-scale SET. For example, one set that gave near-maximal correlation score (or rather minimized the Wasserstein distance between the two distributions) is given (for $\tilde{L} = L/\Delta x = 8$) by

$$\eta = 10^{-1.644} \cdot (\det \tilde{\sigma})^{+0.149} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab} - \tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.045} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab}/\tilde{\omega}_{ab}\tilde{\omega}^{ab})^{-0.132}, \quad (7.19a)$$

$$\kappa = 10^{-4.530} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab})^{+0.146} \cdot \tilde{n}^{+0.291} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab} - \tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.227} \cdot (D_a \tilde{n} D^a \tilde{n})^{+0.397} \cdot (D_a \tilde{n} \Theta^a)^{-0.367}, \quad (7.19b)$$

$$\zeta = 10^{-3.922} \cdot (\tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.107} \cdot \tilde{T}^{+0.409} \cdot \tilde{n}^{+0.497} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab} - \tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.459} \cdot \tilde{\theta}^{-0.847}. \quad (7.19c)$$

We observed a scaling of these coefficients $\propto \tilde{L}^2$. Hence, we recalibrate these expressions to be valid for the case when $\tilde{L} = 1$ by adjusting the off-sets down by $\log_{10}(8^2) \approx 1.8$ and then introduce the \tilde{L}^2 scaling explicitly to obtain

$$\eta = 10^{-3.444} \cdot \tilde{L}^2 \cdot (\det \tilde{\sigma})^{+0.149} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab} - \tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.045} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab}/\tilde{\omega}_{ab}\tilde{\omega}^{ab})^{-0.132}, \quad (7.20a)$$

$$\kappa = 10^{-6.330} \cdot \tilde{L}^2 \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab})^{+0.146} \cdot \tilde{n}^{+0.291} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab} - \tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.227} \cdot (D_a \tilde{n} D^a \tilde{n})^{+0.397} \cdot (D_a \tilde{n} \Theta^a)^{-0.367}, \quad (7.20b)$$

$$\zeta = 10^{-5.722} \cdot \tilde{L}^2 \cdot (\tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.107} \cdot \tilde{T}^{+0.409} \cdot \tilde{n}^{+0.497} \cdot (\tilde{\sigma}_{ab}\tilde{\sigma}^{ab} - \tilde{\omega}_{ab}\tilde{\omega}^{ab})^{+0.459} \cdot \tilde{\theta}^{-0.847}. \quad (7.20c)$$

Calculations of all quantities on the list of regression variables have been implemented in METHOD, as well as these functional forms for the coefficients and the first-order gradient model¹ used for the closure scheme. The first thing to check is what values these coefficients take during evolution, without actually performing the subgrid simulation yet. The result of this check is shown through a combination of figures.

¹Essentially, relativistic Navier-Stokes.




FIGURE 7.5: The evolution of the Kelvin-Helmholtz instability (KHI) and onset of turbulence for an ideal fluid. The number density, n , is plotted over a low-resolution domain of 200×200 cells. No subgrid model is active but this simulation is the one used to calculate the model's coefficients in the plots (figs. 7.6 to 7.9) that follow.




FIGURE 7.6: The subgrid model coefficient κ plotted over the domain of the turbulent KHI simulation seen above. The spatial correlation with ‘sharp’ features (strong gradients) in the fluid flow is clear, in particular at the shocks.




FIGURE 7.7: The same plot but for the coefficient ζ . Some spatial structure is seen but much of the domain is ‘dark’ (zero) because of the clipping we perform when the expansion $\tilde{\theta}$ diverges, which has a negative power fitted in our subgrid model.




FIGURE 7.8: For the coefficient η , spikes in its value are seen along the discontinuities, and there is clearly spatial correlation with the vortices.

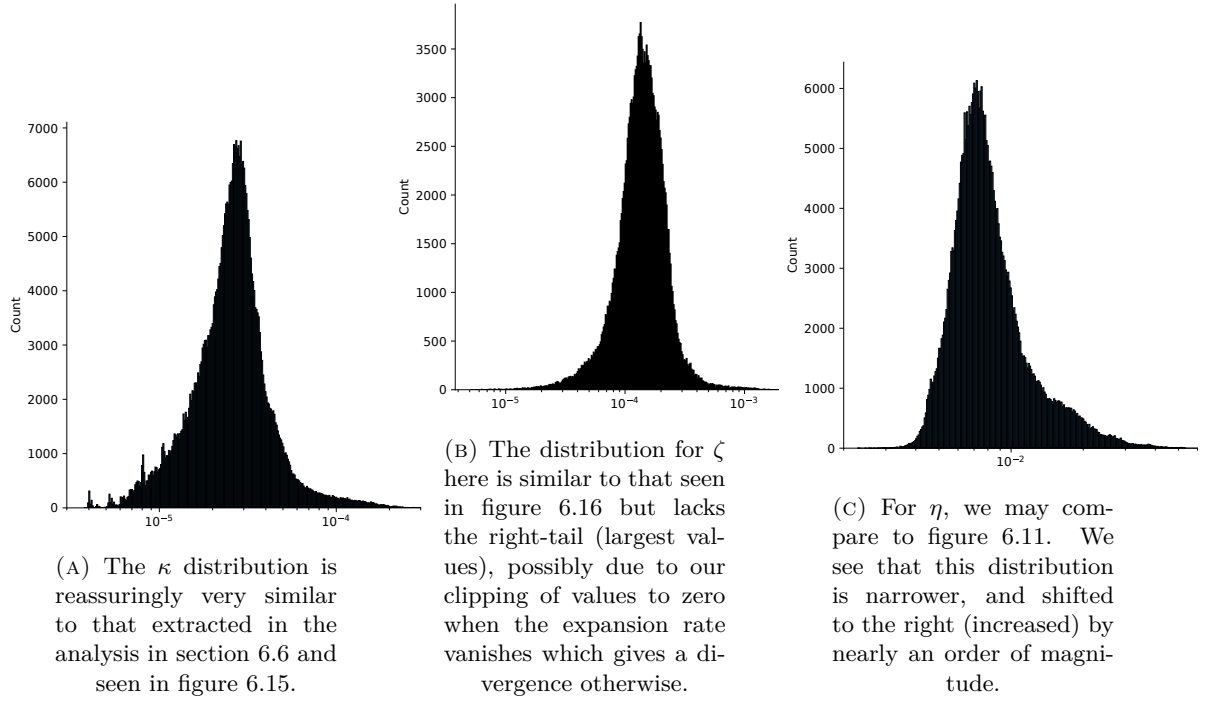


FIGURE 7.9: The distribution of values of κ , ζ and η in the subgrid model for a 200×200 resolution box simulation with initial data creating a Kelvin-Helmholtz instability.

Firstly, in figure 7.5 we show the number density evolution for a low-resolution simulation with $N_x = N_y = 200$. This is simply an ideal fluid and can be directly compared to its high-resolution ($N_x = N_y = 800$) counterpart fig. 6.1 seen in chapter 6. Then, we calculate the coefficient values at every point during this simulation, and plot them in figures 7.6, 7.7 and 7.8. Finally, in figure 7.9 we show histograms of each of the dissipation coefficients' values, conglomerated across all times during the simulation.

We should comment on both the spatial and numerical distribution of the coefficients' values. Visually, the spatial distribution is clearly correlated with the fluid flow. In particular, the coefficient values spike along discontinuities in the number density and/or velocity of the fluid, and the shape of the vortices can be seen in them. One caveat to this is for ζ , where we have clipped its value to zero whenever we find it has diverged. This divergence is due to the expansion scalar $\tilde{\theta}$ being present in its fitted form and having a negative power, leading to divergence when it is zero. Reassuringly, the coefficients' numerical distributions largely match closely to those derived in chapter 6 and seen in figures 6.11, 6.15 and 6.16. Again, ζ is missing its high-tail, likely due to clipping and the η distribution is narrower and right-shifted compared to expectation. Whilst we would like to get to the bottom of these issues, we will not spend more time on them now and will proceed anyway. The next step, therefore, is to switch-on the subgrid model and *see what happens...*

7.2.1 Assessing the SubGrid Model

When trying to perform the same simulation as above, now with the subgrid closure model active, we immediately run into numerical stability issues. These may be due to the fact that we are using a first-order gradient model of dissipative hydrodynamics (essentially, relativistic Navier-Stokes) which is known to be inherently unstable, although it is also (at least) exacerbated by the fact that we now have a non-physical subgrid closure model acting as genuine dissipation. Due to the fitted form of the coefficients, they now take on a (albeit sensible) range of values which can vary significantly on small lengthscales, rather than being constant as in all previous simulations. Hence, we consider alternative models of dissipative hydrodynamics to use as our subgrid closure scheme. However, because we derived this subgrid closure scheme using the first-order gradient model outlined above, we cannot stray far from it or the fitted coefficients will be invalidated. This likely rules out the BDNK model, but makes the MIS and MISCE models potentially viable alternatives. The MIS model removes the pathologies of the Landau & Eckart models by adding additional modes into the system through relaxation equations. If we use small relaxation times, then the MIS model will quickly relax the dissipative variables to their Navier-Stokes forms—the same ones they have directly in the model above. This may allow us to perform a simulation with essentially the same subgrid closure scheme, except for some small deviations.

In fact, we actually find that changing our time-integrator from an explicit Runge-Kutta one to the semi-implicit IMEX scheme SSP2(2,2,2) (see section 3.5.3) stabilizes the numerical evolution. Whilst ostensibly good news, this does go against the ethos of subgrid modelling, which is to cut computational cost by being able to use lower resolutions whilst still capturing the relevant microscale physics. Being forced to use a more costly numerical scheme for the sake of stability is a shame, but one we will have to accept for now until we are able to understand why the explicit integration scheme is unstable and remedy that. We show the first results using the subgrid model in fig. 7.10. Compared to fig. 7.5, the same simulation with no subgrid model active, there are significant differences. We also compare fig. 7.10 to a high-resolution (800×800) simulation with no subgrid model.

It is tricky to assess the success of the subgrid model using a visual analysis, or by inspecting individual features or points within the domain. This is because the subgrid model is designed to capture *on average* the unresolved behaviour. A full assessment of the subgrid model is now outside the scope of this thesis but we present here a few possible measures of its success:

1. A comparison of the kinetic energy spectra (of the type shown here) between a low-resolution simulation with the subgrid model and the high-resolution




FIGURE 7.10: Low resolution (200x200) simulation of the KHI for an ideal fluid with our subgrid model active. Compared to the case with no subgrid model, seen in fig. 7.5, there are small but significant visual differences. In particular, the shocks are noticeably shifted in position due to the inclusion of (effective) dissipative effects that modify their propagation speeds. Also, the maximum number density is slightly suppressed, which is unsurprising given that dissipation smooths gradients in fluid properties.

simulation without it, at least for the wave-modes resolvable in the lower-resolution simulation. Whilst the subgrid model cannot significantly extend the narrower inertial range of the lower-resolution simulation, it may be able to compensate for the overall faster decay of the spectrum for the lower-resolution simulation.

2. An analysis of components of the stress-energy tensor, particular the T^{tt} component often identified with the total energy of the system. Both numerical dissipation and dissipation due to the subgrid model will cause this quantity to decay over time. In fact, this particular consideration and the previous one both further motivate the use of a subgrid model with negative effective dissipation coefficients. In both cases, the decay will be faster for a lower-resolution simulation. This means that a subgrid model acting in a low-resolution simulation needs to generate energy at the smallest lengthscales (or at least move it from there to the longer lengthscales) to match the decay-rate of a higher-resolution simulation.
3. Effective equation of state comparison. We showed in figure 6.13 that filtering introduces deviations from the fine-scale equation of state, specifically in the value of Γ used in its parametrization. Our subgrid model may be tested (or calibrated) by requiring that deviations in the effective equation of state for a simulation with the subgrid model active are smaller than these.

Chapter 8

Conclusions & Outlook

Conclusions

Relativistic, dissipative hydrodynamics has been a field of research for over half a century. It has found applications in the modelling of viscous accretion disks, the quark-gluon plasma created in heavy ion collisions, and, most recently, the hot, dense, out-of-equilibrium matter created in neutron star mergers. With the next generation of gravitational wave detectors on the horizon, it is growing ever more important to include next-to-leading order effects into BNS merger simulations. Heat conduction and viscosity will affect the gravitational and electromagnetic signals we detect, so modelling them accurately is necessary to produce numerical results which may be faithfully compared to future observations.

However, due the onset of turbulence, its interaction with the magnetic field of a remnant, and the out-of-equilibrium processes which mergers enable, small-scale dynamics are still important in determining large-scale behaviour. Unfortunately, direct numerical simulation of these small-scale dynamics will be computationally out of reach for the foreseeable future. Therefore, subgrid modelling and large-eddy schemes, which are well-established in non-relativistic fields of research, must now be applied to relativistic fluids evolving in strong gravity environments. Fortunately, subgrid models for fluids often resemble non-ideal ones, and so through a re-interpretation of the components of non-ideal hydrodynamic models, we may capture this otherwise-unresolved behaviour. These are the motivations behind this research project, and the work contained in this thesis.

In chapter 4 we introduced the Müller-Israel-Stewart formalism and simulated this well-established theory of non-ideal fluids in special relativity. Its numerical limitations for near-equilibrium matter led us in chapter 5 to develop a new formulation for SRDHD. This allowed us to perform significantly more efficient simulations whilst capturing the same dissipative effects. Lightweight additions (such as this one) to

existing models are crucial for including more physics into already very costly simulations of BNS mergers.

In chapter 6 we motivated the need for explicit large-eddy schemes to be used in binary neutron star merger simulations. With the first fully-covariant scheme of its kind, we performed explicit filtering of a microscale fluid and showed that the residuals introduced by this process may be well-captured by a dissipative fluid description at the coarser mesoscale. This does require a statistical fitting of the non-ideal coefficients to thermodynamic and flow properties of the fluid.

In chapter 7 we explored further uses of dissipative hydrodynamic models. In particular, we introduced the recently-popularised BDNK formulation and performed some of the first simulations of viscous fluids using it. We then showed results from an adaptive prototype code that is able to switch between different models on-the-fly, gaining the benefits and avoiding the pitfalls present in each model individually. Finally, we showed preliminary results of applying explicit large-eddy schemes to a turbulent fluid using the sub-grid closure obtained in chapter 6. This shows early promise of being an effective approach to capture otherwise unresolved physics and one that can be applied in BNS mergers.

Outlook

This thesis has explored a number formulations and uses of dissipative hydrodynamics, and provides a strong basis for performing simulations. However, whilst our work has been clearly motivated by neutron star physics and neutron star merger events, it has not been directly applied to an astrophysical scenario. This motivates a fairly clear path towards doing so and providing new insight in this field.

Firstly, we would need to move our models of dissipative hydrodynamics into full general relativity, and couple them to a spacetime evolution code. To perform realistic BNS merger simulations, we would also need to model electromagnetism, radiation and cooling, as well as adopting a realistic equation of state. Whilst this sounds like a considerable task (and it is), thankfully open-source software like the Einstein Toolkit, LORENE and CompOSE provide much of what we'd need. This would then allow us to give quantitative insights into how dissipation affects BNS mergers.

Fortunately, the covariant nature of our LES scheme means that our closure model can be directly lifted into a curved spacetime. However, in order to extend our work on turbulence modelling so that it is more representative of neutron star dynamics, we would need to make the jump to three spatial dimensions and introduce electromagnetism into the stress-energy tensor and hence equations of motion. This introduces additional phenomenology particularly relevant for violent post-merger dynamics where energy and angular momentum cascade between scales and small-scale

dynamos may lead to powerful, large-scale magnetic fields. Subgrid modelling of turbulent dynamics in neutron stars, and the knock-on effect this has on observables, will no doubt be a continued direction of research in the coming years - one that is highly appealing to continue to be a part of.

Appendix A

Approximating Time Derivatives

To approximate time derivatives, in particular for the MISCE model, we make use of the following expansion

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

Note that this is valid in the limit of small dissipation, which is the expected regime when we use MISCE (along with small timescales). Here we will derive this expansion to next order, for instance to apply it to the BDNK formulation where, in general, the dissipative contributions to the state & flux vectors are not small compared to the ideal fluid variables. 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 ($\partial_\mu \dot{\mathbf{w}}$, $\partial_\mu \dot{\mathbf{w}}$).

Consider a system, such as the BDNK model, that can be written in the balance-law form

$$\partial_t [\mathbf{q}_0(\mathbf{w}) + \mathbf{q}_1(\mathbf{w}, \dot{\mathbf{w}}, \dot{\mathbf{w}})] + \partial_i [\mathbf{f}_0^i(\mathbf{w}) + \mathbf{f}_1^i(\mathbf{w}, \dot{\mathbf{w}}, \dot{\mathbf{w}})] = \mathbf{s} = \mathbf{0}. \quad (\text{A.2})$$

In general, we have

$$\partial_t \mathbf{w} = (\partial_{\mathbf{w}} \mathbf{q})^{-1} \partial_t \mathbf{q} = -(\partial_{\mathbf{w}} \mathbf{q})^{-1} \partial_i \mathbf{f}^i = \mathbf{A} \mathbf{B} \quad (\text{A.3})$$

where

$$\mathbf{A} = [\partial_{\mathbf{w}} \mathbf{q}_0 + \delta \partial_{\mathbf{w}} \mathbf{q}_1]^{-1}, \quad \mathbf{B} = [\partial_i \mathbf{f}_0^i + \delta \partial_i \mathbf{f}_1^i] \quad (\text{A.4})$$

and we have explicitly written the (possibly) small prefactor, δ . After some manipulation, and using the assertion that δ is indeed small we arrive at the purely

spatial forms

$$A = \left[\left(\mathcal{I} - \delta \left(\frac{\partial q_0}{\partial w} \right)^{-1} \frac{\partial q_1}{\partial w} \right) \left(\frac{\partial q_0}{\partial w} \right)^{-1} \right], \quad B = \left[\frac{\partial f_0^i}{\partial w} w' + \delta \left(\frac{\partial f_1^i}{\partial w} w' + \frac{\partial f_1^i}{\partial \dot{w}} w' + \frac{\partial f_1^i}{\partial w'} w'' \right) \right]. \quad (\text{A.5})$$

Note that we can choose to use the form of \mathbf{A} in (A.4) and invert the sum of matrices, rather than using the approximate small- δ trick that leads to A in (A.5). Similarly, we can choose the expression for \mathbf{B} from (A.4) which makes use of the fluxes themselves directly, or we may use its form in (A.5) which requires evaluation of second-order spatial differences of the primitive variables.

Appendix B

Extracting Numerical Viscosity

Finite numerical resolution introduces (hopefully small) errors and means that the results of simulations always represent an approximation to the system they aim to solve. One way that these errors appear is in the form of numerical viscosity - a diffusive effect due to the discrete approximation we must use for the momentum advection terms in fluid equations of motion. Understanding this effect is important, especially when ensuring that viscous effects are genuine and not numerical.

In chapter 6, we modelled the effect of using coarse resolutions using viscosity. We showed that the viscous residuals scaled as the square of the coarsening factor. Here, we will use a simple test problem to extract the numerical viscosity and how it scales with resolution.

In section 4.1.3.3 we introduced the reduced equation

$$\partial_t v_y = 2\eta \partial_x^{(2)} v_y - 4\eta^2 \tau_\pi \partial_x^{(4)} v_y \quad (\text{B.1})$$

for the evolution of the y -directed velocity v_y across a one-dimensional domain in x with shear viscosity η . If we take the leading-order piece of our CE-expansion such that $\tau_\pi \rightarrow 0$ then we simply have

$$\partial_t v_y = 2\eta \partial_x^{(2)} v_y \quad (\text{B.2})$$

which has the form of the heat equation but with the heat conductivity κ replaced with 2η . In figure fig. B.1, we simulate this equations and show the smearing of a discontinuity given by the initial data

$$v_y = C_1 \Theta(a) + C_2 \quad (\text{B.3})$$




FIGURE B.1: The smearing of an initial discontinuity due to numerical viscosity for a number of resolutions. The simulation ran until time $t = 20.0$ with an ideal fluid description. This behaviour matches the analytic solution of equation eq. (B.2) which is described by a broadening error function.

where $\Theta(x)$ is the Heaviside step function defined as

$$\Theta(x) = \begin{cases} 0 & x < 0 \\ 1 & x \geq 0 \end{cases} \quad (\text{B.4})$$

and the values used for the constants are $C_1 = 0.2$, $C_2 = -0.1$ and $a = 0$.

We plot the results for a number of resolutions and with an *ideal* fluid description—that is with $\eta = 0$. Despite having set η to 0, the discontinuity does indeed decay as if there were genuine viscosity present, which was simulated and shown in fig. 4.4 in section 4.1.3.3. This is due to numerical viscosity which gives an effective shear viscosity η_{eff} . This decay is more rapid for lower resolutions where η_{eff} is larger.

The particular analytic solution to eq. (B.2) is

$$v_y(x, t) = \frac{C_1}{2} \left[1 - \operatorname{erf} \left(\frac{a - x}{\sqrt{8\eta_{eff}t}} \right) \right] + C_2 \quad (\text{B.5})$$




FIGURE B.2: The power-law scaling of the effective viscosity η_{eff} with grid spacing Δx for the smearing of an initial discontinuity as seen in fig. B.1. As expected, the effective viscosity is proportion to the grid spacing and hence vanishes in the continuum limit of infinite resolution.

and hence by ‘reverse-engineering’ this expression, we can obtain an expression for η_{eff} which is

$$\eta_{eff} = \frac{(a - x)^2}{8t \left[\operatorname{erf}^{-1} \left(\frac{2(C_2 - v_y)}{C_1} + 1 \right) \right]^2}. \quad (\text{B.6})$$

Therefore, we can use the analytic solution (eq. (B.5)) to extract an effective shear viscosity coefficient that captures the effect of numerical dissipation as if it were genuine dissipation. By doing this for a number of resolutions, we are able to see how the value of η_{eff} scales with resolution. This gives us a simple way of obtaining a fit for the dissipative coefficient in terms of resolution, which can be seen in fig. B.2. We find that η_{eff} scales as $\eta_{eff} \propto \Delta x^{0.38}$, which reassuringly means that in the limit of infinite resolution $\Delta x, \eta_{eff} \rightarrow 0$.

Whilst this is a useful technique, it only works when we have an analytic solution to our test problem. This is why comparison between low and high resolution simulations is needed for understanding turbulence modelling and fitting large-eddy schemes. This

is what we did in chapter 6 and has been done by others (Viganò, Aguilera-Miret, et al., [2020b](#); Carrasco, Viganò, et al., [2020](#)).

Bibliography

- [1] B. P. Abbott et al. “A gravitational-wave standard siren measurement of the Hubble constant”. en. In: *Nature* 551.7678 (Nov. 2017). Number: 7678 Publisher: Nature Publishing Group, pp. 85–88. ISSN: 1476-4687. DOI: [10.1038/nature24471](https://doi.org/10.1038/nature24471). URL: <https://www.nature.com/articles/nature24471> (visited on 07/20/2022).
- [2] B. P. Abbott et al. “Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A”. en. In: *The Astrophysical Journal* 848.2 (Oct. 2017). Publisher: American Astronomical Society, p. L13. ISSN: 2041-8205. DOI: [10.3847/2041-8213/aa920c](https://doi.org/10.3847/2041-8213/aa920c). URL: <https://doi.org/10.3847/2041-8213/aa920c> (visited on 07/20/2022).
- [3] B. P. Abbott et al. “GW170817: Measurements of Neutron Star Radii and Equation of State”. en. In: *Physical Review Letters* 121.16 (Oct. 2018), p. 161101. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.121.161101](https://doi.org/10.1103/PhysRevLett.121.161101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.161101> (visited on 02/28/2024).
- [4] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. en. In: *Physical Review Letters* 119.16 (Oct. 2017), p. 161101. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101> (visited on 08/22/2022).
- [5] B. P. Abbott et al. “GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$ ”. en. In: *The Astrophysical Journal Letters* 892.1 (Mar. 2020). Publisher: American Astronomical Society, p. L3. ISSN: 2041-8205. DOI: [10.3847/2041-8213/ab75f5](https://doi.org/10.3847/2041-8213/ab75f5). URL: <https://doi.org/10.3847/2041-8213/ab75f5> (visited on 07/20/2022).

- [6] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. en. In: *Physical Review Letters* 116.6 (Feb. 2016), p. 061102. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102> (visited on 05/05/2021).
- [7] B. P. Abbott et al. “Properties of the Binary Neutron Star Merger GW170817”. en. In: *Physical Review X* 9.1 (Jan. 2019), p. 011001. ISSN: 2160-3308. DOI: [10.1103/PhysRevX.9.011001](https://doi.org/10.1103/PhysRevX.9.011001). URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.011001> (visited on 02/28/2024).
- [8] R. Abbott et al. “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences”. In: *The Astrophysical Journal Letters* 915.1 (July 2021), p. L5. ISSN: 2041-8205, 2041-8213. DOI: [10.3847/2041-8213/ac082e](https://doi.org/10.3847/2041-8213/ac082e). URL: <https://iopscience.iop.org/article/10.3847/2041-8213/ac082e> (visited on 08/22/2022).
- [9] K. Ackley et al. “Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational-wave detector in the global network”. en. In: *Publications of the Astronomical Society of Australia* 37 (2020), e047. ISSN: 1323-3580, 1448-6083. DOI: [10.1017/pasa.2020.39](https://doi.org/10.1017/pasa.2020.39). URL: https://www.cambridge.org/core/product/identifier/S1323358020000399/type/journal_article (visited on 02/28/2024).
- [10] A. Aguayo-Ortiz, S. Mendoza, and D. Olvera. “A direct Primitive Variable Recovery Scheme for hyperbolic conservative equations: The case of relativistic hydrodynamics”. en. In: *PLOS ONE* 13.4 (Apr. 2018). Ed. by Iratxe Puebla, e0195494. ISSN: 1932-6203. DOI: [10.1371/journal.pone.0195494](https://doi.org/10.1371/journal.pone.0195494). URL: <https://dx.plos.org/10.1371/journal.pone.0195494> (visited on 07/14/2022).
- [11] Ricard Aguilera-Miret, Carlos Palenzuela, Federico Carrasco, and Daniele Viganò. “Role of turbulence and winding in the development of large-scale, strong magnetic fields in long-lived remnants of binary neutron star mergers”. en. In: *Physical Review D* 108.10 (Nov. 2023), p. 103001. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.108.103001](https://doi.org/10.1103/PhysRevD.108.103001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.103001> (visited on 04/09/2024).
- [12] Ricard Aguilera-Miret, Daniele Viganò, Federico Carrasco, Borja Miñano, and Carlos Palenzuela. “Turbulent magnetic-field amplification in the first 10 milliseconds after a binary neutron star merger: Comparing high-resolution and large-eddy simulations”. en. In: *Physical Review D* 102.10 (Nov. 2020), p. 103006. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.102.103006](https://doi.org/10.1103/PhysRevD.102.103006).

- URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.103006> (visited on 04/09/2024).
- [13] Ricard Aguilera-Miret, Daniele Viganò, and Carlos Palenzuela. “Universality of the turbulent magnetic field in hypermassive neutron stars produced by binary mergers”. en. In: *arXiv:2112.08406 [astro-ph, physics:gr-qc]* (Dec. 2021). arXiv: 2112.08406. URL: <http://arxiv.org/abs/2112.08406> (visited on 02/01/2022).
- [14] Vasilios Alexiades, Geneviève Amiez, and Pierre-Alain Gremaud. “Super-time-stepping acceleration of explicit schemes for parabolic problems”. en. In: *Communications in Numerical Methods in Engineering* 12.1 (Jan. 1996), pp. 31–42. ISSN: 10698299. DOI: 10.1002/(SICI)1099-0887(199601)12:1<31::AID-CNM950>3.0.CO;2-5. URL: [https://onlinelibrary.wiley.com/doi/10.1002/\(SICI\)1099-0887\(199601\)12:1%3C31::AID-CNM950%3E3.0.CO;2-5](https://onlinelibrary.wiley.com/doi/10.1002/(SICI)1099-0887(199601)12:1%3C31::AID-CNM950%3E3.0.CO;2-5) (visited on 02/17/2025).
- [15] Mark G. Alford, Luke Bovard, Matthias Hanauske, Luciano Rezzolla, and Kai Schwenzer. “Viscous Dissipation and Heat Conduction in Binary Neutron-Star Mergers”. en. In: *Physical Review Letters* 120.4 (Jan. 2018), p. 041101. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.120.041101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.041101> (visited on 06/27/2022).
- [16] Pau Amaro Seoane et al. “The effect of mission duration on LISA science objectives”. en. In: *General Relativity and Gravitation* 54.1 (Jan. 2022), p. 3. ISSN: 0001-7701, 1572-9532. DOI: 10.1007/s10714-021-02889-x. URL: <https://link.springer.com/10.1007/s10714-021-02889-x> (visited on 08/28/2024).
- [17] N Andersson, G L Comer, and I Hawke. “A variational approach to resistive relativistic plasmas”. In: *Classical and Quantum Gravity* 34.12 (June 2017), p. 125001. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/1361-6382/aa6b37. URL: <https://iopscience.iop.org/article/10.1088/1361-6382/aa6b37> (visited on 08/01/2024).
- [18] N Andersson, K Dionysopoulou, I Hawke, and G L Comer. “Beyond ideal magnetohydrodynamics: resistive, reactive and relativistic plasmas”. In: *Classical and Quantum Gravity* 34.12 (June 2017), p. 125002. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/1361-6382/aa6b3a. URL: <https://iopscience.iop.org/article/10.1088/1361-6382/aa6b3a> (visited on 08/01/2024).

- [19] N Andersson, I Hawke, K Dionysopoulou, and G L Comer. “Beyond ideal magnetohydrodynamics: from fibration to $3 + 1$ foliation”. In: *Classical and Quantum Gravity* 34.12 (June 2017), p. 125003. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/1361-6382/aa6b39](https://doi.org/10.1088/1361-6382/aa6b39). URL: <https://iopscience.iop.org/article/10.1088/1361-6382/aa6b39> (visited on 08/01/2024).
- [20] N. Andersson and G. L. Comer. “A covariant action principle for dissipative fluid dynamics: From formalism to fundamental physics”. In: *Classical and Quantum Gravity* 32.7 (Apr. 2015). arXiv:1306.3345 [gr-qc, physics:math-ph], p. 075008. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/32/7/075008](https://doi.org/10.1088/0264-9381/32/7/075008). URL: <http://arxiv.org/abs/1306.3345> (visited on 08/01/2022).
- [21] Luis Antón et al. “Numerical 3+1 General Relativistic Magnetohydrodynamics: A Local Characteristic Approach”. In: *The Astrophysical Journal* 637.1 (Jan. 2006), p. 296. ISSN: 0004-637X. DOI: [10.1086/498238](https://doi.org/10.1086/498238). URL: <https://dx.doi.org/10.1086/498238>.
- [22] R. Arnowitt, S. Deser, and C. W. Misner. “Dynamical Structure and Definition of Energy in General Relativity”. en. In: *Physical Review* 116.5 (Dec. 1959), pp. 1322–1330. ISSN: 0031-899X. DOI: [10.1103/PhysRev.116.1322](https://doi.org/10.1103/PhysRev.116.1322). URL: <https://link.aps.org/doi/10.1103/PhysRev.116.1322> (visited on 07/31/2024).
- [23] Andrew Aspden, Nikos Nikiforakis, Stuart Dalziel, and John Bell. “Analysis of implicit LES methods”. en. In: *Communications in Applied Mathematics and Computational Science* 3.1 (Dec. 2008), pp. 103–126. ISSN: 2157-5452, 1559-3940. DOI: [10.2140/camcos.2008.3.103](https://doi.org/10.2140/camcos.2008.3.103). URL: <http://msp.org/camcos/2008/3-1/p06.xhtml> (visited on 08/16/2022).
- [24] W. Baade and F. Zwicky. “Cosmic Rays from Super-Novae”. en. In: *Proceedings of the National Academy of Sciences* 20.5 (May 1934), pp. 259–263. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.20.5.259](https://doi.org/10.1073/pnas.20.5.259). URL: <http://www.pnas.org/lookup/doi/10.1073/pnas.20.5.259> (visited on 05/05/2021).
- [25] W. Baade and F. Zwicky. “On Super-Novae”. en. In: *Proceedings of the National Academy of Sciences* 20.5 (May 1934), pp. 254–259. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.20.5.254](https://doi.org/10.1073/pnas.20.5.254). URL: <http://www.pnas.org/lookup/doi/10.1073/pnas.20.5.254> (visited on 05/05/2021).
- [26] Rudolf Baier, Paul Romatschke, Dam Thanh Son, Andrei O Starinets, and Mikhail A Stephanov. “Relativistic viscous hydrodynamics, conformal invariance, and holography”. en. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 100–100. ISSN: 1029-8479. DOI: [10.1088/1126-6708/2008/04/100](https://doi.org/10.1088/1126-6708/2008/04/100). URL: <http://stacks.iop.org/1126->

- 6708/2008/i=04/a=100?key=crossref.bf083c128d7b8752ac086978b91a77a2 (visited on 04/23/2021).
- [27] Luca Baiotti. “Gravitational waves from neutron star mergers and their relation to the nuclear equation of state”. en. In: *Progress in Particle and Nuclear Physics* 109 (Nov. 2019), p. 103714. ISSN: 01466410. DOI: 10.1016/j.pnpnp.2019.103714. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0146641019300493> (visited on 08/27/2024).
- [28] Luca Baiotti and Luciano Rezzolla. “Binary neutron star mergers: a review of Einstein’s richest laboratory”. en. In: *Reports on Progress in Physics* 80.9 (Sept. 2017), p. 096901. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/1361-6633/aa67bb. URL: <https://iopscience.iop.org/article/10.1088/1361-6633/aa67bb> (visited on 06/15/2022).
- [29] M. Baldo and G. F. Burgio. “The Nuclear Symmetry Energy”. In: *Progress in Particle and Nuclear Physics* 91 (Nov. 2016). arXiv:1606.08838 [astro-ph, physics:nucl-th], pp. 203–258. ISSN: 01466410. DOI: 10.1016/j.pnpnp.2016.06.006. URL: <http://arxiv.org/abs/1606.08838> (visited on 08/27/2024).
- [30] Dinshaw S. Balsara. “Riemann Solver for Relativistic Hydrodynamics”. en. In: *Journal of Computational Physics* 114.2 (Oct. 1994), pp. 284–297. ISSN: 00219991. DOI: 10.1006/jcph.1994.1167. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0021999184711673> (visited on 08/20/2024).
- [31] Francesc Banyuls, José A. Font, José Ma. Ibáñez, José Ma. Martí, and Juan A. Miralles. “Numerical $\{3 + 1\}$ General Relativistic Hydrodynamics: A Local Characteristic Approach”. In: *The Astrophysical Journal* 476.1 (Feb. 1997), p. 221. ISSN: 0004-637X. DOI: 10.1086/303604. URL: <https://dx.doi.org/10.1086/303604>.
- [32] P. Barrère, J. Guilet, A. Reboul-Salze, R. Raynaud, and H.-T. Janka. *A new scenario for magnetar formation: Tayler-Spruit dynamo in a proto-neutron star spun up by fallback*. Number: arXiv:2206.01269 arXiv:2206.01269 [astro-ph]. June 2022. DOI: 10.48550/arXiv.2206.01269. URL: <http://arxiv.org/abs/2206.01269> (visited on 06/17/2022).
- [33] Thomas W. Baumgarte and Stuart L. Shapiro. “Numerical integration of Einstein’s field equations”. en. In: *Physical Review D* 59.2 (Dec. 1998), p. 024007. ISSN: 0556-2821, 1089-4918. DOI: 10.1103/PhysRevD.59.024007. URL: <https://link.aps.org/doi/10.1103/PhysRevD.59.024007> (visited on 07/31/2024).

- [34] Kris Beckwith and James M. Stone. “A Second Order Godunov Method for Multidimensional Relativistic Magnetohydrodynamics”. In: *The Astrophysical Journal Supplement Series* 193.1 (Mar. 2011). arXiv:1101.3573 [astro-ph, physics:gr-qc, physics:physics], p. 6. ISSN: 0067-0049, 1538-4365. DOI: 10.1088/0067-0049/193/1/6. URL: <http://arxiv.org/abs/1101.3573> (visited on 08/18/2022).
- [35] Fábio S. Bemfica, Marcelo M. Disconzi, Vu Hoang, Jorge Noronha, and Maria Radosz. “Nonlinear Constraints on Relativistic Fluids Far from Equilibrium”. en. In: *Physical Review Letters* 126.22 (June 2021), p. 222301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.126.222301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.222301> (visited on 05/14/2024).
- [36] Fábio S. Bemfica, Marcelo M. Disconzi, and Jorge Noronha. “Causality and existence of solutions of relativistic viscous fluid dynamics with gravity”. en. In: *Physical Review D* 98.10 (Nov. 2018), p. 104064. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.104064. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.104064> (visited on 02/11/2025).
- [37] Fábio S. Bemfica, Marcelo M. Disconzi, and Jorge Noronha. “Causality of the Einstein-Israel-Stewart Theory with Bulk Viscosity”. en. In: *Physical Review Letters* 122.22 (June 2019), p. 221602. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.122.221602. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.221602> (visited on 06/01/2024).
- [38] Fábio S. Bemfica, Marcelo M. Disconzi, and Jorge Noronha. “First-Order General-Relativistic Viscous Fluid Dynamics”. en. In: *Physical Review X* 12.2 (May 2022), p. 021044. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.12.021044. URL: <https://link.aps.org/doi/10.1103/PhysRevX.12.021044> (visited on 05/14/2024).
- [39] Fábio S. Bemfica, Marcelo M. Disconzi, and Jorge Noronha. “Nonlinear causality of general first-order relativistic viscous hydrodynamics”. en. In: *Physical Review D* 100.10 (Nov. 2019), p. 104020. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.104020. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.104020> (visited on 08/23/2022).
- [40] Peter L Bender, Mitchell C Begelman, and Jonathan R Gair. “Possible LISA follow-on mission scientific objectives”. In: *Classical and Quantum Gravity* 30.16 (Aug. 2013), p. 165017. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/30/16/165017. URL:

- <https://iopscience.iop.org/article/10.1088/0264-9381/30/16/165017> (visited on 08/28/2024).
- [41] L. C. Berselli, Traian Iliescu, and William J. Layton. *Mathematics of large eddy simulation of turbulent flows*. Scientific computation. Berlin ; New York: Springer, 2006. ISBN: 978-3-540-26316-6.
- [42] Emanuele Berti et al. *Snowmass2021 Cosmic Frontier White Paper: Fundamental Physics and Beyond the Standard Model*. arXiv:2203.06240 [gr-qc, physics:hep-ph, physics:hep-th]. Mar. 2022. URL: <http://arxiv.org/abs/2203.06240> (visited on 02/28/2024).
- [43] Luiz Bevilacqua, Augusto C. N. R. Galeão, and Flavio P. Costa. “On the significance of higher order differential terms in diffusion processes”. en. In: *Journal of the Brazilian Society of Mechanical Sciences and Engineering* 33.2 (June 2011), pp. 166–175. ISSN: 1678-5878. DOI: 10.1590/S1678-58782011000200007. URL: http://www.scielo.br/scielo.php?script=sci_arttext&pid=S1678-58782011000200007&lng=en&nrm=iso&tlng=en (visited on 05/11/2021).
- [44] Nigel T. Bishop and Luciano Rezzolla. “Extraction of gravitational waves in numerical relativity”. en. In: *Living Reviews in Relativity* 19.1 (Dec. 2016), p. 2. ISSN: 2367-3613, 1433-8351. DOI: 10.1007/s41114-016-0001-9. URL: <http://link.springer.com/10.1007/s41114-016-0001-9> (visited on 06/27/2024).
- [45] Rajesh Biswas, Ashutosh Dash, Najmul Haque, Shi Pu, and Victor Roy. “Causality and stability in relativistic viscous non-resistive magneto-fluid dynamics”. en. In: *Journal of High Energy Physics* 2020.10 (Oct. 2020), p. 171. ISSN: 1029-8479. DOI: 10.1007/JHEP10(2020)171. URL: [https://doi.org/10.1007/JHEP10\(2020\)171](https://doi.org/10.1007/JHEP10(2020)171) (visited on 03/02/2021).
- [46] N. Bucciantini, B. D. Metzger, T. A. Thompson, and E. Quataert. “Short gamma-ray bursts with extended emission from magnetar birth: jet formation and collimation: SGRBEEs from protomagnetar spin-down”. en. In: *Monthly Notices of the Royal Astronomical Society* 419.2 (Jan. 2012), pp. 1537–1545. ISSN: 00358711. DOI: 10.1111/j.1365-2966.2011.19810.x. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2011.19810.x> (visited on 02/27/2025).
- [47] G.F. Burgio, H.-J. Schulze, I. Vidaña, and J.-B. Wei. “Neutron stars and the nuclear equation of state”. en. In: *Progress in Particle and Nuclear Physics* 120 (Sept. 2021), p. 103879. ISSN: 01466410. DOI: 10.1016/j.ppnp.2021.103879. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0146641021000338> (visited on 08/27/2024).

- [48] Claudio Canuto, Mohammed Yousuff Hussaini, and Alfio Quarteroni. *Spectral Methods: Evolution to Complex Geometries and Applications to Fluid Dynamics*. eng. Scientific Computation. Berlin Heidelberg: Springer e-books, 2007. ISBN: 978-3-540-30728-0.
- [49] Federico Carrasco, Daniele Viganò, and Carlos Palenzuela. “Gradient subgrid-scale model for relativistic MHD large-eddy simulations”. In: *Physical Review D* 101.6 (Mar. 2020). Publisher: American Physical Society, p. 063003. DOI: [10.1103/PhysRevD.101.063003](https://link.aps.org/doi/10.1103/PhysRevD.101.063003). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.063003> (visited on 07/12/2022).
- [50] B. Carter. *Conductivity with causality in relativistic hydrodynamics: the regular solution to Eckart’s problem*. Conference Name: Gravitation and Cosmology Pages: 58-65 ADS Bibcode: 1988grco.conf...58C. Jan. 1988. URL: <https://ui.adsabs.harvard.edu/abs/1988grco.conf...58C> (visited on 08/23/2022).
- [51] B. Carter. *The canonical treatment of heat conduction and superfluidity in relativistic hydrodynamics*. Conference Name: A Random Walk in Relativity and Cosmology Pages: 48-62 ADS Bibcode: 1985rwrc.conf...48C. Jan. 1985. URL: <https://ui.adsabs.harvard.edu/abs/1985rwrc.conf...48C> (visited on 08/23/2022).
- [52] Carlo Cattaneo. “Sulla conduzione del calore”. In: *Atti Sem. Mat. Fis. Univ. Modena* 3 (1948), pp. 83–101.
- [53] T. Celora, N. Andersson, I. Hawke, and G. L. Comer. “Covariant approach to relativistic large-eddy simulations: The fibration picture”. en. In: *Physical Review D* 104.8 (Oct. 2021), p. 084090. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.104.084090](https://link.aps.org/doi/10.1103/PhysRevD.104.084090). URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.084090> (visited on 05/14/2024).
- [54] T. Celora, M. J. Hatton, I. Hawke, and N. Andersson. *Covariant approach to relativistic large-eddy simulations: Lagrangian filtering*. Version Number: 1. 2024. DOI: [10.48550/ARXIV.2405.13593](https://arxiv.org/abs/2405.13593). URL: <https://arxiv.org/abs/2405.13593> (visited on 09/01/2024).
- [55] T. Celora, I. Hawke, P. C. Hammond, N. Andersson, and G. L. Comer. “Formulating bulk viscosity for neutron star simulations”. en. In: *Physical Review D* 105.10 (May 2022), p. 103016. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.105.103016](https://link.aps.org/doi/10.1103/PhysRevD.105.103016). URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.103016> (visited on 06/16/2022).

- [56] Thomas Celora, Nils Andersson, Ian Hawke, Greg L. Comer, and Marcus J. Hatton. *A higher-level large-eddy filtering strategy for general relativistic fluid simulations*. arXiv:2407.18012 [astro-ph, physics:gr-qc, physics:physics]. July 2024. URL: <http://arxiv.org/abs/2407.18012> (visited on 08/23/2024).
- [57] Carlo Cercignani and Gilberto Medeiros Kremer. *The Relativistic Boltzmann Equation: Theory and Applications*. en. Basel: Birkhäuser Basel, 2002. ISBN: 978-3-0348-9463-0 978-3-0348-8165-4. DOI: [10.1007/978-3-0348-8165-4](https://doi.org/10.1007/978-3-0348-8165-4). URL: <http://link.springer.com/10.1007/978-3-0348-8165-4> (visited on 05/11/2021).
- [58] Michail Chabanov and Luciano Rezzolla. *Impact of bulk viscosity on the post-merger gravitational-wave signal from merging neutron stars*. Version Number: 1. 2023. DOI: [10.48550/ARXIV.2307.10464](https://doi.org/10.48550/ARXIV.2307.10464). URL: <https://arxiv.org/abs/2307.10464> (visited on 07/30/2024).
- [59] Michail Chabanov, Luciano Rezzolla, and Dirk H Rischke. “General-relativistic hydrodynamics of non-perfect fluids: 3+1 conservative formulation and application to viscous black hole accretion”. en. In: *Monthly Notices of the Royal Astronomical Society* 505.4 (July 2021), pp. 5910–5940. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stab1384](https://doi.org/10.1093/mnras/stab1384). URL: <https://academic.oup.com/mnras/article/505/4/5910/6276727> (visited on 05/11/2023).
- [60] Sydney Chapman and T. G. Cowling. *The mathematical theory of non-uniform gases: an account of the kinetic theory of viscosity, thermal conduction, and diffusion in gases*. 3rd ed. Cambridge mathematical library. Cambridge ; New York: Cambridge University Press, 1990. ISBN: 978-0-521-40844-8.
- [61] A. I. Chugunov and D. G. Yakovlev. “Shear viscosity and oscillations of neutron star crust”. en. In: *Astronomy Reports* 49.9 (Sept. 2005), pp. 724–738. ISSN: 1063-7729, 1562-6881. DOI: [10.1134/1.2045323](https://doi.org/10.1134/1.2045323). URL: <http://link.springer.com/10.1134/1.2045323> (visited on 04/09/2024).
- [62] Riccardo Ciolfi. “The key role of magnetic fields in binary neutron star mergers”. en. In: *General Relativity and Gravitation* 52.6 (June 2020), p. 59. ISSN: 0001-7701, 1572-9532. DOI: [10.1007/s10714-020-02714-x](https://doi.org/10.1007/s10714-020-02714-x). URL: <https://link.springer.com/10.1007/s10714-020-02714-x> (visited on 07/29/2024).
- [63] Riccardo Ciolfi and Jay Vijay Kalinani. “Magnetically Driven Baryon Winds from Binary Neutron Star Merger Remnants and the Blue Kilonova of 2017 August”. en. In: *The Astrophysical Journal* 900.2 (Sept. 2020). Publisher: American Astronomical Society, p. L35. ISSN: 2041-8205. DOI: [10.3847/2041-8213/abb240](https://doi.org/10.3847/2041-8213/abb240). URL: <https://doi.org/10.3847/2041-8213/abb240> (visited on 07/20/2022).

- [64] Riccardo Ciolfi et al. “General relativistic magnetohydrodynamic simulations of binary neutron star mergers forming a long-lived neutron star”. en. In: *Physical Review D* 95.6 (Mar. 2017), p. 063016. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.95.063016](https://doi.org/10.1103/PhysRevD.95.063016). URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.063016> (visited on 08/23/2022).
- [65] J W Cornelisse. “LISA mission and system design”. In: *Classical and Quantum Gravity* 13.11A (Nov. 1996), A251–A258. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/13/11A/034](https://doi.org/10.1088/0264-9381/13/11A/034). URL: <https://iopscience.iop.org/article/10.1088/0264-9381/13/11A/034> (visited on 08/28/2024).
- [66] K. Danzmann. “LISA mission overview”. en. In: *Advances in Space Research* 25.6 (Jan. 2000), pp. 1129–1136. ISSN: 02731177. DOI: [10.1016/S0273-1177\(99\)00973-4](https://doi.org/10.1016/S0273-1177(99)00973-4). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0273117799009734> (visited on 08/28/2024).
- [67] Karsten Danzmann and Albrecht R Diger. “LISA technology concept, status, prospects”. In: *Classical and Quantum Gravity* 20.10 (May 2003), S1–S9. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/20/10/301](https://doi.org/10.1088/0264-9381/20/10/301). URL: <https://iopscience.iop.org/article/10.1088/0264-9381/20/10/301> (visited on 10/25/2023).
- [68] Ashutosh Dash, Masoud Shokri, Luciano Rezzolla, and Dirk H. Rischke. “Charge diffusion in relativistic resistive second-order dissipative magnetohydrodynamics”. en. In: *Physical Review D* 107.5 (Mar. 2023), p. 056003. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.107.056003](https://doi.org/10.1103/PhysRevD.107.056003). URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.056003> (visited on 05/11/2023).
- [69] Soumi De et al. “Tidal Deformabilities and Radii of Neutron Stars from the Observation of GW170817”. en. In: *Physical Review Letters* 121.9 (Aug. 2018), p. 091102. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.121.091102](https://doi.org/10.1103/PhysRevLett.121.091102). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.091102> (visited on 02/28/2024).
- [70] L. Del Zanna et al. “Relativistic viscous hydrodynamics for heavy-ion collisions with ECHO-QGP”. en. In: *The European Physical Journal C* 73.8 (Aug. 2013). arXiv: 1305.7052, p. 2524. ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-013-2524-5](https://doi.org/10.1140/epjc/s10052-013-2524-5). URL: <http://arxiv.org/abs/1305.7052> (visited on 01/28/2021).

- [71] G. S. Denicol, H. Niemi, E. Molnár, and D. H. Rischke. “Derivation of transient relativistic fluid dynamics from the Boltzmann equation”. In: *Physical Review D* 85.11 (June 2012). Publisher: American Physical Society, p. 114047. DOI: [10.1103/PhysRevD.85.114047](https://link.aps.org/doi/10.1103/PhysRevD.85.114047). URL: <https://link.aps.org/doi/10.1103/PhysRevD.85.114047> (visited on 09/27/2023).
- [72] Eleonora Di Valentino et al. “In the Realm of the Hubble tension — a Review of Solutions”. In: *Classical and Quantum Gravity* 38.15 (July 2021). arXiv:2103.01183 [astro-ph, physics:gr-qc, physics:hep-ph], p. 153001. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/1361-6382/ac086d](https://arxiv.org/abs/2103.01183). URL: <http://arxiv.org/abs/2103.01183> (visited on 07/20/2022).
- [73] Tobias Dieselhorst, William Cook, Sebastiano Bernuzzi, and David Radice. “Machine Learning for Conservative-to-Primitive in Relativistic Hydrodynamics”. en. In: *Symmetry* 13.11 (Nov. 2021). arXiv:2109.02679 [astro-ph, physics:physics], p. 2157. ISSN: 2073-8994. DOI: [10.3390/sym13112157](http://arxiv.org/abs/2109.02679). URL: <http://arxiv.org/abs/2109.02679> (visited on 03/02/2023).
- [74] Kyriaki Dionysopoulou, Daniela Alic, Carlos Palenzuela, Luciano Rezzolla, and Bruno Giacomazzo. “General-relativistic resistive magnetohydrodynamics in three dimensions: Formulation and tests”. en. In: *Physical Review D* 88.4 (Aug. 2013), p. 044020. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.88.044020](https://link.aps.org/doi/10.1103/PhysRevD.88.044020). URL: <https://link.aps.org/doi/10.1103/PhysRevD.88.044020> (visited on 05/16/2023).
- [75] Lipei Du and Ulrich Heinz. “(3+1)-dimensional dissipative relativistic fluid dynamics at non-zero net baryon density”. en. In: *Computer Physics Communications* 251 (June 2020), p. 107090. ISSN: 0010-4655. DOI: [10.1016/j.cpc.2019.107090](https://www.sciencedirect.com/science/article/pii/S0010465519303996). URL: <https://www.sciencedirect.com/science/article/pii/S0010465519303996> (visited on 02/05/2021).
- [76] Matthew D. Duez, Yuk Tung Liu, Stuart L. Shapiro, Masaru Shibata, and Branson C. Stephens. “Evolution of magnetized, differentially rotating neutron stars: Simulations in full general relativity”. en. In: *Physical Review D* 73.10 (May 2006), p. 104015. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.73.104015](https://link.aps.org/doi/10.1103/PhysRevD.73.104015). URL: <https://link.aps.org/doi/10.1103/PhysRevD.73.104015> (visited on 07/15/2022).
- [77] Matthew D. Duez, Yuk Tung Liu, Stuart L. Shapiro, and Branson C. Stephens. “General relativistic hydrodynamics with viscosity: Contraction, catastrophic collapse, and disk formation in hypermassive neutron stars”. en. In: *Physical*

- Review D* 69.10 (May 2004), p. 104030. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.69.104030. URL: <https://link.aps.org/doi/10.1103/PhysRevD.69.104030> (visited on 04/09/2024).
- [78] Matthew D. Duez et al. “Comparison of momentum transport models for numerical relativity”. en. In: *Physical Review D* 102.10 (Nov. 2020), p. 104050. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.104050. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.104050> (visited on 08/24/2022).
- [79] Carl Eckart. “The Thermodynamics of Irreversible Processes. III. Relativistic Theory of the Simple Fluid”. en. In: *Physical Review* 58.10 (Nov. 1940), pp. 919–924. ISSN: 0031-899X. DOI: 10.1103/PhysRev.58.919. URL: <https://link.aps.org/doi/10.1103/PhysRev.58.919> (visited on 05/07/2021).
- [80] Gregory L. Eyink and Theodore D. Drivas. “Cascades and Dissipative Anomalies in Relativistic Fluid Turbulence”. en. In: *Physical Review X* 8.1 (Feb. 2018), p. 011023. ISSN: 2160-3308. DOI: 10.1103/PhysRevX.8.011023. URL: <https://link.aps.org/doi/10.1103/PhysRevX.8.011023> (visited on 05/22/2024).
- [81] Joshua A. Faber and Frederic A. Rasio. “Binary Neutron Star Mergers”. en. In: *Living Reviews in Relativity* 15.1 (July 2012), p. 8. ISSN: 1433-8351. DOI: 10.12942/lrr-2012-8. URL: <https://doi.org/10.12942/lrr-2012-8> (visited on 07/20/2022).
- [82] C. Freiburghaus, S. Rosswog, and F.-K. Thielemann. “r-Process in Neutron Star Mergers”. en. In: *The Astrophysical Journal* 525.2 (Oct. 1999). Publisher: IOP Publishing, p. L121. ISSN: 0004-637X. DOI: 10.1086/312343. URL: <https://iopscience.iop.org/article/10.1086/312343/meta> (visited on 07/20/2022).
- [83] S. Gandolfi, J. Carlson, S. Reddy, A. W. Steiner, and R. B. Wiringa. “The equation of state of neutron matter, symmetry energy and neutron star structure”. en. In: *The European Physical Journal A* 50.2 (Feb. 2014), p. 10. ISSN: 1434-6001, 1434-601X. DOI: 10.1140/epja/i2014-14010-5. URL: <http://link.springer.com/10.1140/epja/i2014-14010-5> (visited on 08/27/2024).
- [84] L. Gavassino, M. Antonelli, and B. Haskell. “When the entropy has no maximum: A new perspective on the instability of the first-order theories of dissipation”. en. In: *Physical Review D* 102.4 (Aug. 2020), p. 043018. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.043018. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.043018> (visited on 05/12/2025).

- [85] L. Gavassino, M. Disconzi, and J. Noronha. “Universality classes of relativistic fluid dynamics: Applications”. en. In: *Physical Review D* 109.9 (May 2024), p. 096041. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.109.096041](https://doi.org/10.1103/PhysRevD.109.096041). URL: <https://link.aps.org/doi/10.1103/PhysRevD.109.096041> (visited on 02/12/2025).
- [86] R. A. Gingold and J. J. Monaghan. “Smoothed particle hydrodynamics: theory and application to non-spherical stars”. In: *Monthly Notices of the Royal Astronomical Society* 181.3 (Dec. 1977), pp. 375–389. ISSN: 0035-8711. DOI: [10.1093/mnras/181.3.375](https://doi.org/10.1093/mnras/181.3.375). URL: <https://doi.org/10.1093/mnras/181.3.375> (visited on 07/19/2022).
- [87] K. Glampedakis, N. Andersson, and S. K. Lander. “Hydromagnetic equilibrium in non-barotropic multifluid neutron stars”. In: *Monthly Notices of the Royal Astronomical Society* 420.2 (Feb. 2012), pp. 1263–1272. ISSN: 0035-8711. DOI: [10.1111/j.1365-2966.2011.20112.x](https://doi.org/10.1111/j.1365-2966.2011.20112.x). URL: <https://doi.org/10.1111/j.1365-2966.2011.20112.x> (visited on 08/01/2022).
- [88] A. Goldstein et al. “An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A”. In: *The Astrophysical Journal Letters* 848.2 (Oct. 2017), p. L14. ISSN: 2041-8205, 2041-8213. DOI: [10.3847/2041-8213/aa8f41](https://iopscience.iop.org/article/10.3847/2041-8213/aa8f41). URL: <https://iopscience.iop.org/article/10.3847/2041-8213/aa8f41> (visited on 02/13/2025).
- [89] Sigal Gottlieb, David I. Ketcheson, and Chi-Wang Shu. “High Order Strong Stability Preserving Time Discretizations”. en. In: *Journal of Scientific Computing* 38.3 (Mar. 2009), pp. 251–289. ISSN: 0885-7474, 1573-7691. DOI: [10.1007/s10915-008-9239-z](http://link.springer.com/10.1007/s10915-008-9239-z). URL: <http://link.springer.com/10.1007/s10915-008-9239-z> (visited on 08/11/2022).
- [90] Sigal Gottlieb and Chi-Wang Shu. “Total variation diminishing Runge-Kutta schemes”. en. In: *Mathematics of Computation* 67.221 (1998), pp. 73–85. ISSN: 0025-5718, 1088-6842. DOI: [10.1090/S0025-5718-98-00913-2](https://www.ams.org/mcom/1998-67-221/S0025-5718-98-00913-2/). URL: <https://www.ams.org/mcom/1998-67-221/S0025-5718-98-00913-2/> (visited on 08/11/2022).
- [91] P. Hammond, I. Hawke, and N. Andersson. “Impact of nuclear reactions on gravitational waves from neutron star mergers”. en. In: *Physical Review D* 107.4 (Feb. 2023), p. 043023. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.107.043023](https://link.aps.org/doi/10.1103/PhysRevD.107.043023). URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.043023> (visited on 05/14/2024).

- [92] P. Hammond, I. Hawke, and N. Andersson. “Thermal aspects of neutron star mergers”. en. In: *Physical Review D* 104.10 (Nov. 2021), p. 103006. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.104.103006](https://doi.org/10.1103/PhysRevD.104.103006). URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.103006> (visited on 06/16/2022).
- [93] Matthias Hanauske et al. “Neutron Star Mergers: Probing the EoS of Hot, Dense Matter by Gravitational Waves”. en. In: *Particles* 2.1 (Mar. 2019). Number: 1 Publisher: Multidisciplinary Digital Publishing Institute, pp. 44–56. ISSN: 2571-712X. DOI: [10.3390/particles2010004](https://doi.org/10.3390/particles2010004). URL: <https://www.mdpi.com/2571-712X/2/1/4> (visited on 07/20/2022).
- [94] Alice Harpole. “Multiscale modelling of neutron star oceans”. PhD thesis. University of Southampton, June 2018. URL: <https://eprints.soton.ac.uk/422175/>.
- [95] Ami Harten. “High resolution schemes for hyperbolic conservation laws”. en. In: *Journal of Computational Physics* 49.3 (Mar. 1983), pp. 357–393. ISSN: 0021-9991. DOI: [10.1016/0021-9991\(83\)90136-5](https://doi.org/10.1016/0021-9991(83)90136-5). URL: <https://www.sciencedirect.com/science/article/pii/0021999183901365> (visited on 07/05/2022).
- [96] B. Haskell, N. Andersson, and G. L. Comer. “Dynamics of dissipative multifluid neutron star cores”. In: *Physical Review D* 86.6 (Sept. 2012). Publisher: American Physical Society, p. 063002. DOI: [10.1103/PhysRevD.86.063002](https://doi.org/10.1103/PhysRevD.86.063002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.063002> (visited on 08/01/2022).
- [97] Marcus J. Hatton and Ian Hawke. *A dissipative extension to ideal hydrodynamics*. arXiv:2405.09495 [astro-ph, physics:gr-qc, physics:physics]. May 2024. URL: <http://arxiv.org/abs/2405.09495> (visited on 05/22/2024).
- [98] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. “Observation of a Rapidly Pulsating Radio Source”. en. In: *Nature* 217.5130 (Feb. 1968), pp. 709–713. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/217709a0](https://doi.org/10.1038/217709a0). URL: <http://www.nature.com/articles/217709a0> (visited on 05/05/2021).
- [99] William A Hiscock and Lee Lindblom. “Stability and causality in dissipative relativistic fluids”. en. In: *Annals of Physics* 151.2 (Dec. 1983), pp. 466–496. ISSN: 0003-4916. DOI: [10.1016/0003-4916\(83\)90288-9](https://doi.org/10.1016/0003-4916(83)90288-9). URL: <https://www.sciencedirect.com/science/article/pii/0003491683902889> (visited on 04/23/2021).
- [100] William A. Hiscock and Lee Lindblom. “Generic instabilities in first-order dissipative relativistic fluid theories”. en. In: *Physical Review D* 31.4 (Feb. 1985), pp. 725–733. ISSN: 0556-2821. DOI: [10.1103/PhysRevD.31.725](https://doi.org/10.1103/PhysRevD.31.725). URL: <https://link.aps.org/doi/10.1103/PhysRevD.31.725> (visited on 05/05/2021).

- [101] J. Hoyos, A. Reisenegger, and J. A. Valdivia. “Magnetic field evolution in neutron stars: one-dimensional multi-fluid model”. en. In: *Astronomy & Astrophysics* 487.3 (Sept. 2008). Number: 3 Publisher: EDP Sciences, pp. 789–803. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361:200809466](https://doi.org/10.1051/0004-6361/200809466). URL: <https://www.aanda.org/articles/aa/abs/2008/33/aa09466-08/aa09466-08.html> (visited on 08/01/2022).
- [102] Werner Israel. “Nonstationary irreversible thermodynamics: A causal relativistic theory”. en. In: *Annals of Physics* 100.1 (Sept. 1976), pp. 310–331. ISSN: 0003-4916. DOI: [10.1016/0003-4916\(76\)90064-6](https://doi.org/10.1016/0003-4916(76)90064-6). URL: <https://www.sciencedirect.com/science/article/pii/0003491676900646> (visited on 04/23/2021).
- [103] Werner Israel and J.M. Stewart. “On transient relativistic thermodynamics and kinetic theory. II”. en. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 365.1720 (Feb. 1979), pp. 43–52. ISSN: 0080-4630. DOI: [10.1098/rspa.1979.0005](https://doi.org/10.1098/rspa.1979.0005). URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1979.0005> (visited on 08/01/2022).
- [104] Manuel R. Izquierdo, Miguel Bezares, Steven Liebling, and Carlos Palenzuela. *Large Eddy Simulations of Magnetized Mergers of Black Holes and Neutron Stars*. arXiv:2403.09770 [astro-ph, physics:gr-qc]. Mar. 2024. URL: <http://arxiv.org/abs/2403.09770> (visited on 04/09/2024).
- [105] Amaresh Jaiswal and Victor Roy. “Relativistic Hydrodynamics in Heavy-Ion Collisions: General Aspects and Recent Developments”. en. In: *Advances in High Energy Physics* 2016 (2016), pp. 1–39. ISSN: 1687-7357, 1687-7365. DOI: [10.1155/2016/9623034](https://doi.org/10.1155/2016/9623034). URL: <https://www.hindawi.com/journals/ahp/2016/9623034/> (visited on 02/20/2025).
- [106] A. Janiuk and Y.-F. Yuan. “The role of black hole spin and magnetic field threading the unstable neutrino disk in gamma ray bursts”. In: *Astronomy and Astrophysics* 509 (Jan. 2010), A55. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/200912725](https://doi.org/10.1051/0004-6361/200912725). URL: <http://www.aanda.org/10.1051/0004-6361/200912725> (visited on 02/27/2025).
- [107] Jinhee Jeong and Fazle Hussain. “On the identification of a vortex”. en. In: *Journal of Fluid Mechanics* 285 (Feb. 1995), pp. 69–94. ISSN: 0022-1120, 1469-7645. DOI: [10.1017/S0022112095000462](https://doi.org/10.1017/S0022112095000462). URL: https://www.cambridge.org/core/product/identifier/S0022112095000462/type/journal_article (visited on 05/23/2024).

- [108] Jay V. Kalinani et al. “Implementing a new recovery scheme for primitive variables in the general relativistic magnetohydrodynamic code Spritz”. en. In: *Physical Review D* 105.10 (May 2022), p. 103031. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.105.103031](https://doi.org/10.1103/PhysRevD.105.103031). URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.103031> (visited on 07/04/2022).
- [109] Victoria M. Kaspi and Andrei Beloborodov. “Magnetars”. en. In: *Annual Review of Astronomy and Astrophysics* 55.1 (Aug. 2017). arXiv:1703.00068 [astro-ph], pp. 261–301. ISSN: 0066-4146, 1545-4282. DOI: [10.1146/annurev-astro-081915-023329](https://doi.org/10.1146/annurev-astro-081915-023329). URL: <http://arxiv.org/abs/1703.00068> (visited on 06/01/2022).
- [110] Wolfgang Kastaun, Jay Vijay Kalinani, and Riccardo Ciolfi. “Robust recovery of primitive variables in relativistic ideal magnetohydrodynamics”. en. In: *Physical Review D* 103.2 (Jan. 2021), p. 023018. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.103.023018](https://doi.org/10.1103/PhysRevD.103.023018). URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.023018> (visited on 07/05/2022).
- [111] Kenta Kiuchi, Pablo Cerdá-Durán, Koutarou Kyutoku, Yuichiro Sekiguchi, and Masaru Shibata. “Efficient magnetic-field amplification due to the Kelvin-Helmholtz instability in binary neutron star mergers”. In: *Physical Review D* 92.12 (Dec. 2015). arXiv:1509.09205 [astro-ph, physics:gr-qc, physics:nucl-th], p. 124034. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.92.124034](https://doi.org/10.1103/PhysRevD.92.124034). URL: <http://arxiv.org/abs/1509.09205> (visited on 07/11/2022).
- [112] Kenta Kiuchi, Koutarou Kyutoku, Yuichiro Sekiguchi, and Masaru Shibata. “Global simulations of strongly magnetized remnant massive neutron stars formed in binary neutron star mergers”. In: *Physical Review D* 97.12 (June 2018). Publisher: American Physical Society, p. 124039. DOI: [10.1103/PhysRevD.97.124039](https://doi.org/10.1103/PhysRevD.97.124039). URL: <https://link.aps.org/doi/10.1103/PhysRevD.97.124039> (visited on 07/11/2022).
- [113] Kenta Kiuchi, Alexis Reboul-Salze, Masaru Shibata, and Yuichiro Sekiguchi. “A large-scale magnetic field produced by a solar-like dynamo in binary neutron star mergers”. en. In: *Nature Astronomy* 8.3 (Feb. 2024), pp. 298–307. ISSN: 2397-3366. DOI: [10.1038/s41550-024-02194-y](https://doi.org/10.1038/s41550-024-02194-y). URL: <https://www.nature.com/articles/s41550-024-02194-y> (visited on 08/27/2024).
- [114] Nicolas Kovensky, Aaron Poole, and Andreas Schmitt. “Building a realistic neutron star from holography”. en. In: *Physical Review D* 105.3 (Feb. 2022), p. 034022. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.105.034022](https://doi.org/10.1103/PhysRevD.105.034022).

- URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.034022> (visited on 06/16/2022).
- [115] Nicolas Kovensky and Andreas Schmitt. “Holographic quarkyonic matter”. en. In: *Journal of High Energy Physics* 2020.9 (Sept. 2020), p. 112. ISSN: 1029-8479. DOI: 10.1007/JHEP09(2020)112. URL: [https://link.springer.com/10.1007/JHEP09\(2020\)112](https://link.springer.com/10.1007/JHEP09(2020)112) (visited on 08/11/2022).
- [116] Pavel Kovtun. “First-order relativistic hydrodynamics is stable”. en. In: *Journal of High Energy Physics* 2019.10 (Oct. 2019), p. 34. ISSN: 1029-8479. DOI: 10.1007/JHEP10(2019)034. URL: [https://link.springer.com/10.1007/JHEP10\(2019\)034](https://link.springer.com/10.1007/JHEP10(2019)034) (visited on 04/25/2022).
- [117] M. Kramer et al. “Strong-field tests of gravity using pulsars and black holes”. en. In: *New Astronomy Reviews* 48.11-12 (Dec. 2004), pp. 993–1002. ISSN: 13876473. DOI: 10.1016/j.newar.2004.09.020. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1387647304000909> (visited on 08/10/2022).
- [118] M. Kramer et al. “Tests of General Relativity from Timing the Double Pulsar”. en. In: *Science* 314.5796 (Oct. 2006), pp. 97–102. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1132305. URL: <https://www.science.org/doi/10.1126/science.1132305> (visited on 08/10/2022).
- [119] Kallash Kumar. “The Chapman-Enskog Solution of the Boltzmann Equation: A Reformulation in Terms of Irreducible Tensors and Matrices”. en. In: *Australian Journal of Physics* 20.3 (1967), p. 205. ISSN: 0004-9506. DOI: 10.1071/PH670205. URL: <http://www.publish.csiro.au/?paper=PH670205> (visited on 05/11/2021).
- [120] L. D. Landau and E. M. Lifshitz. *Fluid mechanics*. engus. 2nd ed., 2nd English ed., rev. Course of theoretical physics v. 6. Oxford, England ; New York: Pergamon Press, 1987. ISBN: 978-0-08-033933-7 978-0-08-033932-0.
- [121] J. M. Lattimer. “The Physics of Neutron Stars”. en. In: *Science* 304.5670 (Apr. 2004), pp. 536–542. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1090720. URL: <https://www.sciencemag.org/lookup/doi/10.1126/science.1090720> (visited on 05/18/2021).
- [122] J.M. Lattimer. “Neutron Stars and the Nuclear Matter Equation of State”. en. In: *Annual Review of Nuclear and Particle Science* 71.1 (Sept. 2021), pp. 433–464. ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev-nucl-102419-124827. URL: <https://>

- [//www.annualreviews.org/doi/10.1146/annurev-nucl-102419-124827](http://www.annualreviews.org/doi/10.1146/annurev-nucl-102419-124827)
(visited on 06/17/2022).
- [123] James M. Lattimer. “The Nuclear Equation of State and Neutron Star Masses”. In: *Annual Review of Nuclear and Particle Science* 62.1 (Nov. 2012). arXiv:1305.3510 [astro-ph, physics:nucl-th], pp. 485–515. ISSN: 0163-8998, 1545-4134. DOI: [10.1146/annurev-nucl-102711-095018](https://doi.org/10.1146/annurev-nucl-102711-095018). URL: <http://arxiv.org/abs/1305.3510> (visited on 07/15/2022).
- [124] William H. Lee, Enrico Ramirez-Ruiz, and Glenn van de Ven. “SHORT GAMMA-RAY BURSTS FROM DYNAMICALLY ASSEMBLED COMPACT BINARIES IN GLOBULAR CLUSTERS: PATHWAYS, RATES, HYDRODYNAMICS, AND COSMOLOGICAL SETTING”. en. In: *The Astrophysical Journal* 720.1 (Aug. 2010). Publisher: American Astronomical Society, pp. 953–975. ISSN: 0004-637X. DOI: [10.1088/0004-637x/720/1/953](https://doi.org/10.1088/0004-637x/720/1/953). URL: <https://doi.org/10.1088/0004-637x/720/1/953> (visited on 06/16/2022).
- [125] Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge: Cambridge University Press, 2002. ISBN: 978-0-521-00924-9. DOI: [10.1017/CB09780511791253](https://doi.org/10.1017/CB09780511791253). URL: <https://www.cambridge.org/core/books/finite-volume-methods-for-hyperbolic-problems/97D5D1ACB1926DA1D4D52EAD6909E2B9> (visited on 04/27/2021).
- [126] Xu-Dong Liu, Stanley Osher, and Tony Chan. “Weighted Essentially Non-oscillatory Schemes”. en. In: *Journal of Computational Physics* 115.1 (Nov. 1994), pp. 200–212. ISSN: 00219991. DOI: [10.1006/jcph.1994.1187](https://doi.org/10.1006/jcph.1994.1187). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0021999184711879> (visited on 08/07/2024).
- [127] Xiao-Hui Liu, Zhen-Hua Li, Jing-Zhao Qi, and Xin Zhang. “Galaxy-scale Test of General Relativity with Strong Gravitational Lensing”. In: *The Astrophysical Journal* 927.1 (Mar. 2022), p. 28. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/ac4c3b](https://doi.org/10.3847/1538-4357/ac4c3b). URL: <https://iopscience.iop.org/article/10.3847/1538-4357/ac4c3b> (visited on 08/10/2022).
- [128] Hao Lu and Christopher J. Rutland. “Structural subgrid-scale modeling for large-eddy simulation: A review”. en. In: *Acta Mechanica Sinica* 32.4 (Aug. 2016), pp. 567–578. ISSN: 1614-3116. DOI: [10.1007/s10409-016-0556-4](https://doi.org/10.1007/s10409-016-0556-4). URL: <https://doi.org/10.1007/s10409-016-0556-4> (visited on 07/11/2022).
- [129] Hou-Jun Lü, Bing Zhang, Wei-Hua Lei, Ye Li, and Paul D Lasky. “THE MILLISECOND MAGNETAR CENTRAL ENGINE IN SHORT GRBs”. In: *The Astrophysical Journal* 805.2 (May 2015), p. 89. ISSN: 1538-4357. DOI: [10.1088/0004-637X/805/2/89](https://doi.org/10.1088/0004-637X/805/2/89). URL:

- <https://iopscience.iop.org/article/10.1088/0004-637X/805/2/89>
(visited on 02/27/2025).
- [130] L. B. Lucy. “A numerical approach to the testing of the fission hypothesis.” en. In: *The Astronomical Journal* 82 (Dec. 1977), pp. 1013–1024. ISSN: 0004-6256. DOI: 10.1086/112164. URL: <https://ui.adsabs.harvard.edu/abs/1977AJ.....82.1013L/abstract> (visited on 07/19/2022).
- [131] Ibanez Jose Ma. “Riemann Solvers in Relativistic Hydrodynamics: Basics and Astrophysical Applications”. eng. In: *Journal of The Korean Astronomical Society* 34.4 (2001). Publisher: The Korean Astronomical Society, pp. 191–201. ISSN: 1225-4614. URL: <https://www.koreascience.or.kr/article/JAK0200102637740684.page> (visited on 06/06/2022).
- [132] Michele Maggiore et al. “Science Case for the Einstein Telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2020.03 (Mar. 2020). arXiv:1912.02622 [astro-ph, physics:gr-qc], pp. 050–050. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2020/03/050. URL: <http://arxiv.org/abs/1912.02622> (visited on 07/30/2024).
- [133] Juliette M. Mammei, Charles J. Horowitz, Jorge Piekarewicz, Brendan T. Reed, and Concettina Sfienti. “Neutron Skins: Weak Elastic Scattering and Neutron Stars”. en. In: *Annual Review of Nuclear and Particle Science* (July 2024). ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev-nucl-102122-024207. URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-nucl-102122-024207> (visited on 08/27/2024).
- [134] Cristina Manuel and Laura Tolos. “Shear viscosity due to phonons in superfluid neutron stars”. en. In: *Physical Review D* 84.12 (Dec. 2011), p. 123007. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.84.123007. URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.123007> (visited on 04/09/2024).
- [135] L. G. Margolin and W. J. Rider. “The design and construction of implicit LES models”. en. In: *International Journal for Numerical Methods in Fluids* 47.10-11 (Apr. 2005), pp. 1173–1179. ISSN: 0271-2091, 1097-0363. DOI: 10.1002/fld.862. URL: <https://onlinelibrary.wiley.com/doi/10.1002/fld.862> (visited on 08/16/2022).
- [136] L. G. Margolin, W. J. Rider, and F. F. Grinstein. “Modeling turbulent flow with implicit LES”. en. In: *Journal of Turbulence* 7 (Jan. 2006), N15. ISSN: 1468-5248. DOI: 10.1080/14685240500331595. URL: <http://www.tandfonline.com/doi/abs/10.1080/14685240500331595> (visited on 08/16/2022).

- [137] Jose M. Marti and Ewald Mueller. “Numerical hydrodynamics in special relativity”. en. In: *Living Reviews in Relativity* 2.1 (Dec. 1999). arXiv:astro-ph/9906333, p. 3. ISSN: 2367-3613, 1433-8351. DOI: [10.12942/lrr-1999-3](https://doi.org/10.12942/lrr-1999-3). URL: <http://arxiv.org/abs/astro-ph/9906333> (visited on 08/05/2022).
- [138] B. D. Metzger, D. Giannios, T. A. Thompson, N. Bucciantini, and E. Quataert. “The protomagnetar model for gamma-ray bursts: Protomagnetar model for GRBs”. en. In: *Monthly Notices of the Royal Astronomical Society* 413.3 (May 2011), pp. 2031–2056. ISSN: 00358711. DOI: [10.1111/j.1365-2966.2011.18280.x](https://doi.org/10.1111/j.1365-2966.2011.18280.x). URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2011.18280.x> (visited on 02/27/2025).
- [139] S Miranda-Aranguren, M A Aloy, and T Rembiasz. “An HLLC Riemann solver for resistive relativistic magnetohydrodynamics”. en. In: *Monthly Notices of the Royal Astronomical Society* 476.3 (May 2018), pp. 3837–3860. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/sty419](https://doi.org/10.1093/mnras/sty419). URL: <https://academic.oup.com/mnras/article/476/3/3837/4867992> (visited on 05/16/2023).
- [140] Miquel Miravet-Tenés, Pablo Cerdá-Durán, Martin Obergaulinger, and José A Font. “Assessment of a new sub-grid model for magnetohydrodynamical turbulence – II. Kelvin–Helmholtz instability”. en. In: *Monthly Notices of the Royal Astronomical Society* 527.1 (Oct. 2023), pp. 1081–1092. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stad3237](https://doi.org/10.1093/mnras/stad3237). URL: <https://academic.oup.com/mnras/article/527/1/1081/7326766> (visited on 07/30/2024).
- [141] Miquel Miravet-Tenés, Pablo Cerdá-Durán, Martin Obergaulinger, and José A Font. “Assessment of a new sub-grid model for magnetohydrodynamical turbulence. I. Magnetorotational instability”. In: *Monthly Notices of the Royal Astronomical Society* 517.3 (Dec. 2022), pp. 3505–3524. ISSN: 0035-8711. DOI: [10.1093/mnras/stac2888](https://doi.org/10.1093/mnras/stac2888). URL: <https://doi.org/10.1093/mnras/stac2888> (visited on 01/16/2023).
- [142] Miquel Miravet-Tenés and Martin E. Pessah. *An Effective Model for Magnetic Field Amplification by the Magnetorotational and Parasitic Instabilities*. arXiv:2411.05064 [astro-ph]. Nov. 2024. DOI: [10.48550/arXiv.2411.05064](https://doi.org/10.48550/arXiv.2411.05064). URL: <http://arxiv.org/abs/2411.05064> (visited on 02/27/2025).
- [143] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. San Francisco: W. H. Freeman, 1973. ISBN: 978-0-7167-0334-1 978-0-7167-0344-0.

- [144] E. Molnár, H. Niemi, and D. H. Rischke. “Numerical tests of causal relativistic dissipative fluid dynamics”. en. In: *The European Physical Journal C* 65.3-4 (Feb. 2010). arXiv: 0907.2583, pp. 615–635. ISSN: 1434-6044, 1434-6052. DOI: [10.1140/epjc/s10052-009-1194-9](https://doi.org/10.1140/epjc/s10052-009-1194-9). URL: <http://arxiv.org/abs/0907.2583> (visited on 03/30/2021).
- [145] Elias R Most et al. “Projecting the likely importance of weak-interaction-driven bulk viscosity in neutron star mergers”. en. In: *Monthly Notices of the Royal Astronomical Society* 509.1 (Nov. 2021), pp. 1096–1108. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stab2793](https://doi.org/10.1093/mnras/stab2793). URL: <https://academic.oup.com/mnras/article/509/1/1096/6378903> (visited on 05/14/2024).
- [146] Elias R. Most. “Impact of a mean field dynamo on neutron star mergers leading to magnetar remnants”. en. In: *Physical Review D* 108.12 (Dec. 2023), p. 123012. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.108.123012](https://doi.org/10.1103/PhysRevD.108.123012). URL: <https://link.aps.org/doi/10.1103/PhysRevD.108.123012> (visited on 08/27/2024).
- [147] Elias R. Most and Jorge Noronha. “Dissipative Magnetohydrodynamics for Non-Resistive Relativistic Plasmas”. en. In: *Physical Review D* 104.10 (Nov. 2021). arXiv:2109.02796 [astro-ph, physics:gr-qc, physics:nucl-th, physics:physics], p. 103028. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.104.103028](https://doi.org/10.1103/PhysRevD.104.103028). URL: <http://arxiv.org/abs/2109.02796> (visited on 07/08/2022).
- [148] Elias R. Most et al. *Emergence of microphysical viscosity in binary neutron star post-merger dynamics*. en. arXiv:2207.00442 [astro-ph, physics:gr-qc, physics:nucl-th]. July 2022. URL: <http://arxiv.org/abs/2207.00442> (visited on 07/08/2022).
- [149] Ariadna Murguía-Berthier et al. “The Fate of the Merger Remnant in GW170817 and Its Imprint on the Jet Structure”. In: *The Astrophysical Journal* 908.2 (Feb. 2021), p. 152. ISSN: 0004-637X, 1538-4357. DOI: [10.3847/1538-4357/abd08e](https://doi.org/10.3847/1538-4357/abd08e). URL: <https://iopscience.iop.org/article/10.3847/1538-4357/abd08e> (visited on 02/13/2025).
- [150] J. R. Oppenheimer and G. M. Volkoff. “On Massive Neutron Cores”. en. In: *Physical Review* 55.4 (Feb. 1939), pp. 374–381. ISSN: 0031-899X. DOI: [10.1103/PhysRev.55.374](https://doi.org/10.1103/PhysRev.55.374). URL: <https://link.aps.org/doi/10.1103/PhysRev.55.374> (visited on 08/10/2022).

- [151] Feryal Özel and Paulo Freire. “Masses, Radii, and the Equation of State of Neutron Stars”. en. In: *Annual Review of Astronomy and Astrophysics* 54.1 (Sept. 2016), pp. 401–440. ISSN: 0066-4146, 1545-4282. DOI: [10.1146/annurev-astro-081915-023322](https://doi.org/10.1146/annurev-astro-081915-023322). URL: <https://www.annualreviews.org/doi/10.1146/annurev-astro-081915-023322> (visited on 08/27/2024).
- [152] Dany Page, James M. Lattimer, Madappa Prakash, and Andrew W. Steiner. “Minimal Cooling of Neutron Stars: A New Paradigm”. en. In: *The Astrophysical Journal Supplement Series* 155.2 (Dec. 2004). Publisher: IOP Publishing, p. 623. ISSN: 0067-0049. DOI: [10.1086/424844](https://doi.org/10.1086/424844). URL: <https://iopscience.iop.org/article/10.1086/424844/meta> (visited on 06/17/2022).
- [153] Carlos Palenzuela, Luis Lehner, Oscar Reula, and Luciano Rezzolla. “Beyond ideal MHD: towards a more realistic modelling of relativistic astrophysical plasmas”. en. In: *Monthly Notices of the Royal Astronomical Society* 394.4 (Apr. 2009), pp. 1727–1740. ISSN: 00358711, 13652966. DOI: [10.1111/j.1365-2966.2009.14454.x](https://doi.org/10.1111/j.1365-2966.2009.14454.x). URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2009.14454.x> (visited on 05/16/2023).
- [154] Carlos Palenzuela, Steven Liebling, and Borja Miñano. “Large eddy simulations of magnetized mergers of neutron stars with neutrinos”. en. In: *Physical Review D* 105.10 (May 2022), p. 103020. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.105.103020](https://doi.org/10.1103/PhysRevD.105.103020). URL: <https://link.aps.org/doi/10.1103/PhysRevD.105.103020> (visited on 04/09/2024).
- [155] Carlos Palenzuela et al. “Turbulent magnetic field amplification in binary neutron star mergers”. en. In: *Physical Review D* 106.2 (July 2022), p. 023013. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.106.023013](https://doi.org/10.1103/PhysRevD.106.023013). URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.023013> (visited on 04/09/2024).
- [156] Alex Pandya, Elias R. Most, and Frans Pretorius. “Causal, stable first-order viscous relativistic hydrodynamics with ideal gas microphysics”. en. In: *Physical Review D* 106.12 (Dec. 2022), p. 123036. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.106.123036](https://doi.org/10.1103/PhysRevD.106.123036). URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.123036> (visited on 05/14/2024).
- [157] Alex Pandya, Elias R. Most, and Frans Pretorius. “Conservative finite volume scheme for first-order viscous relativistic hydrodynamics”. en. In: *Physical Review D* 105.12 (June 2022), p. 123001. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.105.123001](https://doi.org/10.1103/PhysRevD.105.123001). URL:

- <https://link.aps.org/doi/10.1103/PhysRevD.105.123001> (visited on 05/14/2024).
- [158] Alex Pandya and Frans Pretorius. “Numerical exploration of first-order relativistic hydrodynamics”. en. In: *Physical Review D* 104.2 (July 2021), p. 023015. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.104.023015](https://link.aps.org/doi/10.1103/PhysRevD.104.023015). URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.023015> (visited on 05/14/2024).
- [159] Lorenzo Pareschi and Giovanni Russo. “Implicit–Explicit Runge–Kutta Schemes and Applications to Hyperbolic Systems with Relaxation”. en. In: *Journal of Scientific Computing* 25.1 (Oct. 2005), pp. 129–155. ISSN: 0885-7474, 1573-7691. DOI: [10.1007/s10915-004-4636-4](https://link.springer.com/10.1007/s10915-004-4636-4). URL: <https://link.springer.com/10.1007/s10915-004-4636-4> (visited on 08/14/2022).
- [160] J. Peitz and S. Appl. “3+1 formulation of non-ideal hydrodynamics”. en. In: *Monthly Notices of the Royal Astronomical Society* 296.2 (May 1998), pp. 231–244. ISSN: 0035-8711, 1365-2966. DOI: [10.1046/j.1365-8711.1998.01259.x](https://academic.oup.com/mnras/article-lookup/doi/10.1046/j.1365-8711.1998.01259.x). URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1046/j.1365-8711.1998.01259.x> (visited on 01/07/2021).
- [161] Jochen Peitz and Stefan Appl. “Dissipative fluid dynamics in the 3 + 1 formalism”. en. In: *Classical and Quantum Gravity* 16.3 (Mar. 1999), pp. 979–989. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/16/3/025](https://iopscience.iop.org/article/10.1088/0264-9381/16/3/025). URL: <https://iopscience.iop.org/article/10.1088/0264-9381/16/3/025> (visited on 04/23/2021).
- [162] A. Y. Potekhin, A. F. Fantina, N. Chamel, J. M. Pearson, and S. Goriely. “Analytical representations of unified equations of state for neutron-star matter”. en. In: *Astronomy & Astrophysics* 560 (Dec. 2013). Publisher: EDP Sciences, A48. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/201321697](https://www.aanda.org/articles/aa/abs/2013/12/aa21697-13/aa21697-13.html). URL: <https://www.aanda.org/articles/aa/abs/2013/12/aa21697-13/aa21697-13.html> (visited on 07/15/2022).
- [163] Alexander Y. Potekhin, José A. Pons, and Dany Page. “Neutron Stars—Cooling and Transport”. en. In: *Space Science Reviews* 191.1 (Oct. 2015), pp. 239–291. ISSN: 1572-9672. DOI: [10.1007/s11214-015-0180-9](https://doi.org/10.1007/s11214-015-0180-9). URL: <https://doi.org/10.1007/s11214-015-0180-9> (visited on 07/15/2022).
- [164] Aviral Prakash et al. “Signatures of deconfined quark phases in binary neutron star mergers”. en. In: *arXiv:2106.07885 [astro-ph, physics:gr-qc]* (June 2021). arXiv: 2106.07885. URL: <http://arxiv.org/abs/2106.07885> (visited on 06/22/2021).

- [165] D. J. Price and S. Rosswog. “Producing Ultrastrong Magnetic Fields in Neutron Star Mergers”. en. In: *Science* 312.5774 (May 2006), pp. 719–722. ISSN: 0036-8075, 1095-9203. DOI: [10.1126/science.1125201](https://doi.org/10.1126/science.1125201). URL: <https://www.science.org/doi/10.1126/science.1125201> (visited on 08/23/2024).
- [166] M Punturo et al. “The Einstein Telescope: a third-generation gravitational wave observatory”. In: *Classical and Quantum Gravity* 27.19 (Oct. 2010), p. 194002. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002). URL: <https://iopscience.iop.org/article/10.1088/0264-9381/27/19/194002> (visited on 10/25/2023).
- [167] J. Qian. “Generalization of the Kolmogorov $-5/3$ law of turbulence”. en. In: *Physical Review E* 50.1 (July 1994), pp. 611–613. ISSN: 1063-651X, 1095-3787. DOI: [10.1103/PhysRevE.50.611](https://doi.org/10.1103/PhysRevE.50.611). URL: <https://link.aps.org/doi/10.1103/PhysRevE.50.611> (visited on 08/03/2022).
- [168] David Radice. “Binary Neutron Star Merger Simulations with a Calibrated Turbulence Model”. en. In: *Symmetry* 12.8 (July 2020), p. 1249. ISSN: 2073-8994. DOI: [10.3390/sym12081249](https://doi.org/10.3390/sym12081249). URL: <https://www.mdpi.com/2073-8994/12/8/1249> (visited on 03/23/2022).
- [169] David Radice. “General-Relativistic Large-Eddy Simulations of Binary Neutron Star Mergers”. en. In: *The Astrophysical Journal* 838.1 (Mar. 2017). arXiv:1703.02046 [astro-ph, physics:gr-qc], p. L2. ISSN: 2041-8213. DOI: [10.3847/2041-8213/aa6483](https://doi.org/10.3847/2041-8213/aa6483). URL: <http://arxiv.org/abs/1703.02046> (visited on 06/30/2022).
- [170] David Radice and Ian Hawke. “Turbulence modelling in neutron star merger simulations”. en. In: *Living Reviews in Computational Astrophysics* 10.1 (Feb. 2024), p. 1. ISSN: 2365-0524. DOI: [10.1007/s41115-023-00019-9](https://doi.org/10.1007/s41115-023-00019-9). URL: <https://link.springer.com/10.1007/s41115-023-00019-9> (visited on 05/14/2024).
- [171] Adriana R. Raduta, Flavia Nacu, and Micaela Oertel. “Equations of state for hot neutron stars”. en. In: *The European Physical Journal A* 57.12 (Dec. 2021), p. 329. ISSN: 1434-601X. DOI: [10.1140/epja/s10050-021-00628-z](https://doi.org/10.1140/epja/s10050-021-00628-z). URL: <https://doi.org/10.1140/epja/s10050-021-00628-z> (visited on 07/15/2022).
- [172] Peter B. Rau and Ira Wasserman. “Relativistic finite temperature multifluid hydrodynamics in a neutron star from a variational principle”. In: *Physical Review D* 102.6 (Sept. 2020). Publisher: American Physical Society, p. 063011. DOI: [10.1103/PhysRevD.102.063011](https://doi.org/10.1103/PhysRevD.102.063011). URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.063011> (visited on 08/01/2022).

- [173] D. G. Ravenhall, C. J. Pethick, and J. R. Wilson. “Structure of Matter below Nuclear Saturation Density”. en. In: *Physical Review Letters* 50.26 (June 1983), pp. 2066–2069. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.50.2066](https://doi.org/10.1103/PhysRevLett.50.2066). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.50.2066> (visited on 06/17/2022).
- [174] Raphaël Raynaud, Jérôme Guilet, Hans-Thomas Janka, and Thomas Gastine. “Magnetar formation through a convective dynamo in protoneutron stars”. In: *Science Advances* 6.11 (Mar. 2020). Publisher: American Association for the Advancement of Science, eaay2732. DOI: [10.1126/sciadv.aay2732](https://doi.org/10.1126/sciadv.aay2732). URL: <https://www.science.org/doi/full/10.1126/sciadv.aay2732> (visited on 06/17/2022).
- [175] Alexis Reboul-Salze, Jérôme Guilet, Raphaël Raynaud, and Matteo Bugli. *MRI-driven $\alpha - \Omega$ dynamos in protoneutron stars*. Number: arXiv:2111.02148 arXiv:2111.02148 [astro-ph, physics:physics]. Nov. 2021. DOI: [10.48550/arXiv.2111.02148](https://doi.org/10.48550/arXiv.2111.02148). URL: <http://arxiv.org/abs/2111.02148> (visited on 06/17/2022).
- [176] David Reitze et al. *Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO*. arXiv:1907.04833 [astro-ph, physics:gr-qc]. July 2019. URL: <http://arxiv.org/abs/1907.04833> (visited on 02/28/2024).
- [177] L. Rezzolla and Olindo Zanotti. *Relativistic hydrodynamics*. First edition. Oxford: Oxford University Press, 2013. ISBN: 978-0-19-852890-6.
- [178] Luciano Rezzolla et al. “THE MISSING LINK: MERGING NEUTRON STARS NATURALLY PRODUCE JET-LIKE STRUCTURES AND CAN POWER SHORT GAMMA-RAY BURSTS”. en. In: *The Astrophysical Journal* 732.1 (Apr. 2011). Publisher: American Astronomical Society, p. L6. ISSN: 2041-8205. DOI: [10.1088/2041-8205/732/1/L6](https://doi.org/10.1088/2041-8205/732/1/L6). URL: <https://doi.org/10.1088/2041-8205/732/1/L6> (visited on 07/20/2022).
- [179] B. Ripperda et al. “General-relativistic Resistive Magnetohydrodynamics with Robust Primitive-variable Recovery for Accretion Disk Simulations”. In: *The Astrophysical Journal Supplement Series* 244.1 (Sept. 2019), p. 10. ISSN: 1538-4365. DOI: [10.3847/1538-4365/ab3922](https://doi.org/10.3847/1538-4365/ab3922). URL: <https://iopscience.iop.org/article/10.3847/1538-4365/ab3922> (visited on 05/16/2023).
- [180] Gabriel S. Rocha, David Wagner, Gabriel S. Denicol, Jorge Noronha, and Dirk H. Rischke. “Theories of Relativistic Dissipative Fluid Dynamics”. en. In: *Entropy* 26.3 (Feb. 2024), p. 189. ISSN: 1099-4300. DOI: [10.3390/e26030189](https://doi.org/10.3390/e26030189). URL: <https://www.mdpi.com/1099-4300/26/3/189> (visited on 06/01/2024).

- [181] P.L. Roe. “Approximate Riemann solvers, parameter vectors, and difference schemes”. en. In: *Journal of Computational Physics* 43.2 (Oct. 1981), pp. 357–372. ISSN: 00219991. DOI: [10.1016/0021-9991\(81\)90128-5](https://doi.org/10.1016/0021-9991(81)90128-5). URL: <https://linkinghub.elsevier.com/retrieve/pii/0021999181901285> (visited on 08/08/2024).
- [182] Paul Romatschke. “New Developments in Relativistic Viscous Hydrodynamics”. en. In: *International Journal of Modern Physics E* 19.01 (Jan. 2010). arXiv: 0902.3663, pp. 1–53. ISSN: 0218-3013, 1793-6608. DOI: [10.1142/S0218301310014613](https://doi.org/10.1142/S0218301310014613). URL: <http://arxiv.org/abs/0902.3663> (visited on 06/22/2021).
- [183] S Rosswog and P Diener. “SPHINCS_BSSN: a general relativistic smooth particle hydrodynamics code for dynamical spacetimes”. In: *Classical and Quantum Gravity* 38.11 (June 2021), p. 115002. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/1361-6382/abee65](https://doi.org/10.1088/1361-6382/abee65). URL: <https://iopscience.iop.org/article/10.1088/1361-6382/abee65> (visited on 08/02/2024).
- [184] S. Rosswog and P. Diener. *SPHINCS_BSSN: Numerical Relativity with Particles*. Version Number: 2. 2024. DOI: [10.48550/ARXIV.2404.15952](https://doi.org/10.48550/ARXIV.2404.15952). URL: <https://arxiv.org/abs/2404.15952> (visited on 08/02/2024).
- [185] Stephan Rosswog. “SPH Methods in the Modelling of Compact Objects”. en. In: *Living Reviews in Computational Astrophysics* 1.1 (Oct. 2015), p. 1. ISSN: 2365-0524. DOI: [10.1007/lrca-2015-1](https://doi.org/10.1007/lrca-2015-1). URL: <https://doi.org/10.1007/lrca-2015-1> (visited on 07/20/2022).
- [186] Stephan Rosswog, Peter Diener, and Francesco Torsello. “Thinking Outside the Box: Numerical Relativity with Particles”. en. In: *Symmetry* 14.6 (June 2022), p. 1280. ISSN: 2073-8994. DOI: [10.3390/sym14061280](https://doi.org/10.3390/sym14061280). URL: <https://www.mdpi.com/2073-8994/14/6/1280> (visited on 08/02/2024).
- [187] Stephan Rosswog, Francesco Torsello, and Peter Diener. “The Lagrangian numerical relativity code SPHINCS_BSSN_v1.0”. In: *Frontiers in Applied Mathematics and Statistics* 9 (Oct. 2023), p. 1236586. ISSN: 2297-4687. DOI: [10.3389/fams.2023.1236586](https://doi.org/10.3389/fams.2023.1236586). URL: <https://www.frontiersin.org/articles/10.3389/fams.2023.1236586/full> (visited on 08/02/2024).
- [188] Wolfram Schmidt. “Large Eddy Simulations in Astrophysics”. en. In: *Living Reviews in Computational Astrophysics* 1.1 (Dec. 2015), p. 2. ISSN: 2367-3621, 2365-0524. DOI: [10.1007/lrca-2015-2](https://doi.org/10.1007/lrca-2015-2). URL: <http://link.springer.com/10.1007/lrca-2015-2> (visited on 05/23/2024).

- [189] Andreas Schmitt and Peter Shternin. “Reaction Rates and Transport in Neutron Stars”. en. In: *The Physics and Astrophysics of Neutron Stars*. Ed. by Luciano Rezzolla, Pierre Pizzochero, David Ian Jones, Nanda Rea, and Isaac Vidaña. Astrophysics and Space Science Library. Cham: Springer International Publishing, 2018, pp. 455–574. ISBN: 978-3-319-97616-7. DOI: [10.1007/978-3-319-97616-7_9](https://doi.org/10.1007/978-3-319-97616-7_9). URL: https://doi.org/10.1007/978-3-319-97616-7_9 (visited on 06/16/2022).
- [190] Stuart L. Shapiro and Saul A. Teukolsky. *Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects*. en. 1st ed. Wiley, July 1983. ISBN: 978-0-471-87316-7 978-3-527-61766-1. DOI: [10.1002/9783527617661](https://onlinelibrary.wiley.com/doi/book/10.1002/9783527617661). URL: <https://onlinelibrary.wiley.com/doi/book/10.1002/9783527617661> (visited on 05/05/2021).
- [191] Masaru Shibata, Sho Fujibayashi, and Yuichiro Sekiguchi. “Long-term evolution of neutron-star merger remnants in general relativistic resistive magnetohydrodynamics with a mean-field dynamo term”. en. In: *Physical Review D* 104.6 (Sept. 2021), p. 063026. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.104.063026](https://link.aps.org/doi/10.1103/PhysRevD.104.063026). URL: <https://link.aps.org/doi/10.1103/PhysRevD.104.063026> (visited on 07/11/2022).
- [192] Masaru Shibata and Takashi Nakamura. “Evolution of three-dimensional gravitational waves: Harmonic slicing case”. en. In: *Physical Review D* 52.10 (Nov. 1995), pp. 5428–5444. ISSN: 0556-2821. DOI: [10.1103/PhysRevD.52.5428](https://link.aps.org/doi/10.1103/PhysRevD.52.5428). URL: <https://link.aps.org/doi/10.1103/PhysRevD.52.5428> (visited on 07/31/2024).
- [193] Chi-Wang Shu. “Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws”. en. In: *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations: Lectures given at the 2nd Session of the Centro Internazionale Matematico Estivo (C.I.M.E.) held in Cetraro, Italy, June 23–28, 1997*. Ed. by Bernardo Cockburn, Chi-Wang Shu, Claes Johnson, Eitan Tadmor, and Alfio Quarteroni. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, 1998, pp. 325–432. ISBN: 978-3-540-49804-9. DOI: [10.1007/BFb0096355](https://doi.org/10.1007/BFb0096355). URL: <https://doi.org/10.1007/BFb0096355> (visited on 08/05/2022).
- [194] J. Smagorinsky. “GENERAL CIRCULATION EXPERIMENTS WITH THE PRIMITIVE EQUATIONS: I. THE BASIC EXPERIMENT”. EN. In: *Monthly Weather Review* 91.3 (Mar. 1963). Publisher: American Meteorological Society Section: Monthly Weather Review, pp. 99–164. ISSN: 1520-0493, 0027-0644. DOI: [10.1175/1520-0493\(1963\)091<0099:GCEWTP>2.3.CO;2](https://journals.ametsoc.org/view/journals/mwre/91/3/1520-0493_1963_091_0099_gcewtp_2_3_co_2.xml). URL: https://journals.ametsoc.org/view/journals/mwre/91/3/1520-0493_1963_091_0099_gcewtp_2_3_co_2.xml (visited on 07/15/2022).

- [195] William M. Spinella. “A systematic investigation of exotic matter in neutron stars.” PhD thesis. San Diego State University, 2017. URL: <https://www.proquest.com/docview/1979346397?pq-origsite=gscholar&fromopenview=true>.
- [196] J.M. Stewart. “On transient relativistic thermodynamics and kinetic theory”. en. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 357.1688 (Oct. 1977), pp. 59–75. ISSN: 0080-4630. DOI: 10.1098/rspa.1977.0155. URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1977.0155> (visited on 06/25/2024).
- [197] Gilbert Strang. “On the Construction and Comparison of Difference Schemes”. In: *SIAM Journal on Numerical Analysis* 5.3 (1968). _eprint: <https://doi.org/10.1137/0705041>, pp. 506–517. DOI: 10.1137/0705041. URL: <https://doi.org/10.1137/0705041>.
- [198] Makoto Takamoto and Shu-ichiro Inutsuka. “A Fast Numerical Scheme for Causal Relativistic Hydrodynamics with Dissipation”. en. In: *Journal of Computational Physics* 230.18 (Aug. 2011). arXiv: 1106.1732, pp. 7002–7017. ISSN: 00219991. DOI: 10.1016/j.jcp.2011.05.030. URL: <http://arxiv.org/abs/1106.1732> (visited on 03/10/2021).
- [199] T. M. Tauris and E. P. J. van den Heuvel. *Formation and evolution of compact stellar X-ray sources*. Vol. 39. Pages: 623-665 Publication Title: Compact stellar X-ray sources ADS Bibcode: 2006csxs.book..623T. Apr. 2006. DOI: 10.48550/arXiv.astro-ph/0303456. URL: <https://ui.adsabs.harvard.edu/abs/2006csxs.book..623T> (visited on 08/27/2024).
- [200] Christopher Thompson and Robert C. Duncan. “Neutron Star Dynamos and the Origins of Pulsar Magnetism”. In: *The Astrophysical Journal* 408 (May 1993). ADS Bibcode: 1993ApJ...408..194T, p. 194. ISSN: 0004-637X. DOI: 10.1086/172580. URL: <https://ui.adsabs.harvard.edu/abs/1993ApJ...408..194T> (visited on 04/11/2024).
- [201] Slava G. Turyshev. “Experimental Tests of General Relativity”. en. In: *Annual Review of Nuclear and Particle Science* 58.1 (Nov. 2008), pp. 207–248. ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev.nucl.58.020807.111839. URL: <https://www.annualreviews.org/doi/10.1146/annurev.nucl.58.020807.111839> (visited on 08/10/2022).

- [202] Stefan Typel. “CompOSE”. en. In: *The European Physical Journal A* 57.10 (Oct. 2021), p. 284. ISSN: 1434-601X. DOI: [10.1140/epja/s10050-021-00593-7](https://doi.org/10.1140/epja/s10050-021-00593-7). URL: <https://doi.org/10.1140/epja/s10050-021-00593-7> (visited on 07/15/2022).
- [203] Alberto Vela-Martín. “Subgrid-scale models of isotropic turbulence need not produce energy backscatter”. In: *Journal of Fluid Mechanics* 937 (2022). Edition: 2022/02/22 Publisher: Cambridge University Press, A14. ISSN: 0022-1120. DOI: [10.1017/jfm.2022.123](https://doi.org/10.1017/jfm.2022.123). URL: <https://www.cambridge.org/core/product/7CBFF497BCB29C17F243C7FA2600D741>.
- [204] Daniele Viganò, Ricard Aguilera-Miret, Federico Carrasco, Borja Miñano, and Carlos Palenzuela. “General relativistic MHD large eddy simulations with gradient subgrid-scale model”. en. In: *Physical Review D* 101.12 (June 2020), p. 123019. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.101.123019](https://doi.org/10.1103/PhysRevD.101.123019). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.123019> (visited on 08/16/2023).
- [205] Daniele Viganò, Ricard Aguilera-Miret, Federico Carrasco, Borja Miñano, and Carlos Palenzuela. “GRMHD large eddy simulations with gradient subgrid-scale model”. en. In: *Physical Review D* 101.12 (June 2020). arXiv:2004.00870 [astro-ph, physics:gr-qc], p. 123019. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.101.123019](https://doi.org/10.1103/PhysRevD.101.123019). URL: <http://arxiv.org/abs/2004.00870> (visited on 05/25/2022).
- [206] D. Wagner and L. Gavassino. “Regime of applicability of Israel-Stewart hydrodynamics”. en. In: *Physical Review D* 109.1 (Jan. 2024), p. 016019. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.109.016019](https://doi.org/10.1103/PhysRevD.109.016019). URL: <https://link.aps.org/doi/10.1103/PhysRevD.109.016019> (visited on 05/14/2024).
- [207] Yunpeng Wang, Zelong Yuan, Xiaoning Wang, and Jianchun Wang. “Constant-coefficient spatial gradient models for the sub-grid scale closure in large-eddy simulation of turbulence”. In: *Physics of Fluids* 34.9 (Sept. 2022), p. 095108. ISSN: 1070-6631. DOI: [10.1063/5.0101356](https://doi.org/10.1063/5.0101356). URL: <https://doi.org/10.1063/5.0101356> (visited on 07/28/2023).
- [208] A. J. Wright and I. Hawke. “Resistive and Multi-fluid RMHD on Graphics Processing Units”. In: *The Astrophysical Journal Supplement Series* 240.1 (Jan. 2019), p. 8. ISSN: 0067-0049, 1538-4365. DOI: [10.3847/1538-4365/aaf1b0](https://doi.org/10.3847/1538-4365/aaf1b0). URL: <https://iopscience.iop.org/article/10.3847/1538-4365/aaf1b0> (visited on 05/14/2024).
- [209] Alex James Wright. “Non-ideal simulations of neutron star mergers”. PhD thesis. University of Southampton, Sept. 2020. URL: <https://eprints.soton.ac.uk/448034/>.

- [210] Alex James Wright and Ian Hawke. “A resistive extension for ideal MHD”. en. In: *Monthly Notices of the Royal Astronomical Society* 491.4 (Feb. 2020). arXiv: 1906.03150, pp. 5510–5523. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/stz2779](https://doi.org/10.1093/mnras/stz2779). URL: <http://arxiv.org/abs/1906.03150> (visited on 04/27/2021).
- [211] Hanxun Yao, Michael Schnaubelt, Alexander S. Szalay, Tamer A. Zaki, and Charles Meneveau. “Comparing local energy cascade rates in isotropic turbulence using structure-function and filtering formulations”. In: *Journal of Fluid Mechanics* 980 (2024). Edition: 2024/02/06 Publisher: Cambridge University Press, A42. ISSN: 0022-1120. DOI: [10.1017/jfm.2023.1066](https://doi.org/10.1017/jfm.2023.1066). URL: <https://www.cambridge.org/core/product/4E439515B02957031B3182A93F7554E2>.