

UNIVERSITY OF SOUTHAMPTON

Non-ideal simulations of neutron star mergers

by

Alex James Wright

 [0000-0002-5953-4221](#)

 [AlexJamesWright](#)

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the

[Faculty of Engineering and Physical Sciences](#)
[Next Generation Computational Modelling Group](#)

November 2020

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND PHYSICAL SCIENCES
NEXT GENERATION COMPUTATIONAL MODELLING GROUP

Doctor of Philosophy

by Alex James Wright

As gravitational wave astronomy begins to mature, the need for precise and efficient numerical modelling of compact object mergers grows. Simulations of mergers require relativistic models capable of capturing complex dynamics that act over disparate length and time scales, where improving the quality of the solution always comes at the cost of additional computation.

In the vast majority of merger simulations, the equations defining the interaction of the matter and electromagnetic fields make a number of specific assumptions. In this work, we provide several means of efficiently evolving systems of equations that move away from this *ideal* framework. So-called *non-ideal* models of magnetohydrodynamics (MHD) may be dissipative or describe the evolution of multiple fluid species. Here, we present both kinds, describe why they are more physically accurate than the models that are generally used for merger simulations, and demonstrate the drawbacks inherent in their use.

To alleviate the computational cost incurred by some non-ideal models, we develop a number of methods. First, we demonstrate how resistive and multiple fluid models of MHD can be accelerated by $21\times$ by executing the most computationally expensive tasks on graphics processing units.

We then focus our attention on resistive MHD, developing an approximation that can extend the ideal models that are currently used, emulating a full resistive model within a fraction of the time. After modifying this approximation to work in general relativity, we then use it for resistive simulations of neutron star mergers. These simulations demonstrate conclusively that a physically realistic magnitude of resistivity for neutron star matter can have dramatic effects upon the evolution of merger simulations, with large impacts on key observables.

Finally, we present two novel, hybrid models of MHD that allow for significantly improved accuracy and computational efficiency when evolving resistive MHD. With all these tools, the computational costs preventing the widespread adoption of non-ideal models for merger simulations have been dramatically reduced, allowing for higher resolution and more realistic, resistive descriptions of neutron star matter for future merger simulations.

Publications

Parts of this thesis have been published in collaboration with Ian Hawke. A description of the numerics and convergence properties of the code `METHOD`, featured in Chapters 2 and 3, along with its execution on graphics processing units, Chapter 4, has been published under the title of “Resistive and multi-fluid RMHD on graphics processing units” in the Astrophysical Journal Supplement Series (Wright and Hawke, 2019a).

As the various models of MHD and the numerical methods are referred to and utilised repeatedly throughout, we have contained them within a single chapter. Furthermore, to motivate Chapters 4 and 5, the testing and validation of `METHOD` also lies within a single chapter. We believe this allows for easier reference, and leads to a more sensible layout. As a result, some sections of the publication Wright and Hawke (2019a) exist in separate chapters.

In addition, the work of Chapter 5 has been published as “A resistive extension for ideal RMHD” in the Monthly Notices of the Royal Astronomical Society (Wright and Hawke, 2019b).

Contents

List of Figures	xi
List of Tables	xiii
Declaration of Authorship	xv
Personal acknowledgements	xvii
Further acknowledgements	xix
Acronyms	xxi
1 Motivation	1
1.1 Introduction	1
1.2 Mergers in nature	2
1.3 Mergers in silicon	3
1.3.1 A history in theory	3
1.3.2 Current successes	7
1.3.3 Modern limitations	11
1.3.3.1 Resolution	11
1.3.3.2 MHD model	12
1.4 Problems to address	14
2 Mathematics and numerics	17
2.1 Introduction	17
2.2 Neutron stars are fluids	18
2.2.1 Euler equations	18
2.2.2 Ideal magnetohydrodynamics	19
2.2.3 Resistive magnetohydrodynamics	22
2.2.4 Two-fluid electro-magnetohydrodynamics	23
2.3 Numerical relativity	26
2.3.1 Arnowitt-Deser-Misner decomposition	27
2.3.2 Gauge freedom	28
2.3.3 Time derivative of spatial metric determinant	29
2.3.4 Weak hydrodynamic coupling	30
2.3.5 The Einstein Toolkit and <code>Whisky</code>	30
2.4 Hyperbolic conservation laws and finite volume methods	31
2.4.1 Conservation laws	31
2.4.2 The Riemann problem	32
2.4.3 Finite volume approach	32
2.5 Flux methods	33
2.6 Time-integration schemes	35

2.6.1	Explicit schemes	36
2.6.2	Semi-implicit schemes	36
2.7	Primitive Recovery	39
2.8	Divergence Cleaning	41
3	Validating our METHOD	43
3.1	Trust building exercises	43
3.1.1	Orszag-Tang vortex	44
3.1.2	Brio-Wu shock tube	44
3.1.3	Self-similar current sheet	49
3.1.4	Field loop advection	51
3.1.5	Kelvin-Helmholtz instability	51
3.1.5.1	Effects of resolution	53
3.1.5.2	Model dependence	53
3.2	Computational demands	58
3.3	Conclusions	60
4	Massively-parallel processing	63
4.1	Parallel programming	63
4.1.1	Parallel implementation	65
4.1.2	CUDA overview	65
4.2	Host and device memory	66
4.2.1	Memory coalescence	66
4.2.2	Memory optimisation	67
4.3	Performance	68
4.4	Discussion	72
5	Chapman-Enskog analysis	75
5.1	Scalar equation expansion	76
5.1.1	Stability	77
5.2	System expansion	78
5.3	Numerical implementation	81
5.4	Stability criterion	84
5.5	Convergence with σ	86
5.6	Stability	97
5.7	Performance	99
5.8	Discussion	101
6	General relativistic REGIME	103
6.1	Derivation	104
6.2	Implementation	107
6.2.1	Matrix derivatives	107
6.2.2	Inverting numerically	107
6.2.3	Matrix terms	109
6.2.4	A small hack	110
6.2.5	Stability	111
6.3	Testing	112
6.3.1	Unit testing and test driven development	112
6.3.2	Tests problems	113
6.4	Binary neutron star mergers	114
6.4.1	Simulation set-up	114
6.4.2	Results	118
6.4.3	Low resolution simulation	118

6.4.4	Medium resolution simulations	121
6.4.4.1	Overview of evolution	122
6.4.4.2	Delayed rBNS collapse	126
6.4.4.3	Post-collapse dynamics	133
6.4.5	Highly magnetised case	136
6.5	Discussion	139
7	A hybrid approach to resistive MHD	141
7.1	Introduction	141
7.2	A special relativistic example	142
7.2.1	Hybrid++	143
7.2.2	Results	144
7.3	A general relativistic example	148
7.3.1	Ideal atmosphere	149
7.3.2	Results	149
7.4	Discussion	151
8	Final thoughts	153
8.1	Overview	153
8.2	Looking forward	155
A		157
A.1	Magnetic freezing	157
A.2	Overview of METHOD	157
A.2.1	General Design Features	157
A.2.2	Available Physics Models	160
A.2.3	Testing Paradigms	160
A.2.3.1	Test-Driven Development	161
A.2.3.2	Regression Testing	161
A.2.3.3	Code Validation and Unit Testing	161

List of Figures

3.1	Results for the relativistic Orszag-Tang vortex.	45
3.2	Divergence of magnetic fields in Orszag-Tang vortex.	45
3.3	Rotational symmetry of METHOD 's solutions.	46
3.4	Brio-Wu shock tube for resistive MHD.	47
3.5	Brio-Wu convergence with resolution.	48
3.6	Brio-Wu shock tube for two fluid model.	49
3.7	Brio-Wu shock tube for two-fluid, showing two-fluid effect.	50
3.8	Single and two-fluid self-similar current sheet solution.	50
3.9	Final state of the magnetic field advection test.	51
3.10	Convergence of the magnetic field advection test.	52
3.11	Dependence of Kelvin-Helmholtz vortex formation on resolution.	54
3.12	Kelvin-Helmholtz vorticity with resolution.	55
3.13	Dependence on resolution of magnetic evolution in the Kelvin-Helmholtz instability.	56
3.14	Dependence on model of magnetic evolution in the Kelvin-Helmholtz instability.	56
3.15	Power spectrum dependence on the model choice of the Kelvin-Helmholtz instability.	57
3.16	Performance comparison of single fluid ideal and resistive MHD, and the two-fluid model	59
4.1	Cost of evolution per cell on a graphics processing unit.	69
4.2	Parallel speed-up of graphics processing units.	70
4.3	Relative cost of contiguous memory rearrangement on a graphics processing unit.	71
5.1	Convergence of Brio-Wu with conductivity for REGIME.	86
5.2	Error growth of Brio-Wu with conductivity for REGIME.	87
5.3	Convergence of current sheet with conductivity for REGIME.	88
5.4	Error growth of current sheet with conductivity for REGIME.	89
5.5	Solution of resistive reconnection for rMHD and REGIME.	91
5.6	Final state of magnetic fields strength for resistive reconnection for rMHD and REGIME.	93
5.7	Gaussian fits of current density in resistive reconnection.	94
5.8	Reconnection rate for rMHD and REGIME.	95
5.9	REGIME comparison of final state of Kelvin-Helmholtz.	95
5.10	REGIME comparison of magnetic evolution for Kelvin-Helmholtz instability.	96
5.11	REGIME comparison of power spectrum for Kelvin-Helmholtz instability.	97
5.12	Final state of the density for the Brio-Wu shock-tube problem with increasing resolution. For simulations with $N_x > 1498$ the stability criterion is violated and the evolution becomes unstable.	98
5.13	Stability criterion for 1D, 2D and 3D test problems. Points represent the resolution of the last stable simulation for the indicated initial data, with the gradient highlighted in the legend, and the corresponding value for the softening parameter, δ . This suggests that the criterion, Equation 5.42, implies stability when using a softening factor of $\delta = 1$	99
5.14	Performance comparison for REGIME with rMHD.	100

6.1	Confirmation of tests passing in the output file. Due to the non-physical initial data, we have to kill the simulation after the tests are completed.	112
6.2	Convergence of the current sheet with the conductivity, σ , at a final time of $T = 8$. The convergence of GRR and REGIME with the conductivity is very similar. . .	114
6.3	Convergence of the Brio-Wu shock tube with the conductivity, σ , at a final time of $T = 0.8$. The oscillations left of the shock are caused by the non-staggered vector potential formulation in Whisky . Clearly, GRR successfully diffusing the magnetic field, with similar results to the METHOD code.	115
6.4	Low resolution BNS density slices	119
6.5	Evolution of low resolution merger	120
6.6	Evolution of poloidal and toroidal magnetic fields	121
6.7	Density slices for medium resolution merger	122
6.8	Matter evolution and GW signals for the medium resolution run	123
6.9	Magnetic evolution of medium resolution merger	125
6.10	Symmetry of rest-mass density	127
6.11	Symmetry of magnetic norm	128
6.12	Angle-averaged quantities in x - z slice	129
6.13	Angle-averaged pressure contributions in the x - z slice	130
6.14	Effective centrifugal pressure gradient evolution in the x - z slice	131
6.15	Pressure gradients in the x - z slice	132
6.16	Toroidal and poloidal magnetic field in the x - z slice	132
6.17	Evolution of mean toroidal and poloidal magnetic field	133
6.18	Post-collapse magnetic field lines.	134
6.19	Angle-averaged quantities post-collapse	135
6.20	Post-collapse magnetic field slices.	135
6.21	Matter evolution of highly magnetised BNS	137
6.22	Magnetic configuration of highly magnetised BNS	138
7.1	Hybrid penalty method	144
7.2	Hybrid current sheet error	145
7.3	Final state with functional conductivity	146
7.4	Final conductivity profile	147
7.5	Performance comparison of Hybrid++	148
7.6	Hybrid merger rest mass density slice	150
7.7	Hybrid merger maximum rest mass density	150
7.8	Hybrid merger magnetic evolution	151
A.1	Class structure of METHOD	158
A.2	Inheritance structure of METHOD	159

List of Tables

2.1	The Butcher notation for second-order IMEX coefficients.	38
3.1	Initial data for the single fluid Orszag-Tang vortex.	44
4.1	Profile of routines on a graphics processing unit.	71

Declaration of Authorship

I, Alex James Wright, declare that this thesis entitled ‘Non-ideal simulations of neutron star mergers’ and the work presented in it are 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 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: “Resistive and multi-fluid RMHD on graphics processing units” ([Wright and Hawke, 2019a](#)), and “A resistive extension for ideal MHD” ([Wright and Hawke, 2019b](#)).

Signed:

Date:

Personal acknowledgements

Unfortunately, there is not enough space in a single page to accurately portray my gratitude to all those who have helped me throughout my PhD. Help can come in many forms, from those who have directly taught or inspired me, to those who have provided a welcome distraction from anything science related. I am sure that not everybody who has supported me through this journey will fully understand the impact they have had on my time in Southampton, but I will try to highlight those who require special attention here.

As is customary in thesis acknowledgements, the first to be thanked must be my family—they are the ones who have seen, first hand, the highs and lows that come with research. Always patient, each member of our ever-growing family has seen all four of my seasons in a single day, a direct result of success and elation, often followed quickly by doubt and despair. But the foundation provided by a supportive and loving family has made the whole process a pleasure to share with them. A small part of this work is theirs.

I have also been fortunate enough to meet a number of people in Southampton from which I have gained a huge amount. Andrei and Tom have provided snooker on demand, and the realisation that we have had our last, post-work Friday session is a source of sadness. Office 2019 offered perhaps all too frequent distractions, but I am sure that I will never find a better group of office mates. Equal, maybe, but not better. All members of the third cohort of the NGCM have helped to create the most welcoming introduction to Southampton, and random, high-pitched calls of “pub” will go sorely missed. I have also been lucky enough to live with a number of people who have made exploring a new place exciting and generally unpredictable: Ryan, Laura, Harry, Florin, Cathy, Lili, Timo and Pete, you have all played your part. A special thanks must also go to Marta. Having been the closest to me throughout this process, you have learned when I needed support, and importantly when to ignore me. No puedo esperar para nuestro futuro juntos, gordi. Tecueme.

Finally, I owe my eternal gratitude to my supervisor, Ian Hawke. Ian is a fountain of knowledge and educated guesses, whose ability to argue both sides of a point is second only to his ability to make his students feel comfortable and inspired. I am certain that without The Hawk’s endless support and guidance, this project would not have gone much further than the second chapter. Although it is my name on the front page, this is perhaps unjust, and the work presented in this thesis is hugely indebted to the time and wisdom that Ian has provided. I am sorry that I never paid you the owed pint for the stability bet, Ian, but hope we can celebrate the end of this project properly once the world opens up again.

Further acknowledgements

This research was performed using financial support from the EPSRC Centre for Doctoral Training in Next Generation Computational Modelling grant EP/ L015382/1. A huge thanks must also go to the maintainers of the IRIDIS High Performance Computing Facility at the University of Southampton, and its associated support services at the University of Southampton—particularly to David Hempston for handling my endless requests for additional resources.

In addition, I would like to acknowledge the use of the following software, distributions, and platforms that have made my life, and this PhD, far easier than it could have been:

- Anaconda, and specifically NumPy (Oliphant, 2006), SciPy (Virtanen et al., 2020), Spyder, Matplotlib (Hunter, 2007) and Seaborn (Waskom et al., 2014), for all things python: www.anaconda.com
- Atom for code development: www.atom.io
- Remote Atom for remote code development: www.atom.io/packages/remote-atom
- Cactus and the Einstein Toolkit for general relativistic simulations: www.einsteintoolkit.org
- Whisky for GRMHD (Giacomazzo and Rezzolla, 2007)
- PyCactus for Cactus data analysis: www.github.com/wokast/PyCactus (public version)
- Git and GitHub for countless disaster recoveries
- Travis CI for continuous integration and testing: www.travis-ci.org
- Doxygen for documentation www.doxygen.nl
- Cminpack for rootfinding in C and C++: www.devernay.free.fr/hacks/cminpack
- Googletest for unit testing in C++: www.github.com/google/googletest

Acronyms

ADM	Arnowitt-Deser-Misner
AMR	Adaptive mesh refinement
BBH	Binary black-hole
BH	Black-hole
BHNS	Black-hole neutron star [binary]
BMP	Binary merger product
CFD	Computational fluid dynamics
CFL	Courant-Friedrichs-Lewy [condition]
CPU	Central processing unit
EM	Electromagnetic
EoS	Equation of state
FFE	Force-free electrodynamics
FLOPS	Floating point operations per second
FVS	Flux vector splitting
GPU	Graphics processing unit
GR	General relativity
GRMHD	General relativistic magnetohydrodynamics
GRR	General relativistic REGIME
GW	Gravitational waves
HD	Hydrodynamics
HLL	Harten-Lax-van Leer
HMNS	Hypermassive neutron star
HPC	High performance computing
HRSC	High-resolution shock capturing
iBMP	Ideal (MHD) binary merger product
iBNS	Ideal MHD binary neutron star [merger]
IMEX	Implicit-explicit
iMHD	Ideal magnetohydrodynamics
KHI	Kelvin-Helmholtz instability
LES	Large eddy simulations
METHOD	Multi-fluid electro-magnetohydrodynamics [code]
MHD	Magnetohydrodynamics
NS	Neutron star
OO	Object-oriented
OTV	Orszag-Tang vortex

PDE	Partial differential equations
PIC	Particle-in-cell
QCD	Quantum chromodynamics
rBMP	Resistive MHD binary merger product
rBNS	Resistive MHD binary neutron star [merger]
REGIME	Resistive extension upgrade for ideal MHD
RK	Runge-Kutta
R/LHS	Right/Left hand side
RMHD	Relativistic magnetohydrodynamics
rMHD	Resistive magnetohydrodynamics
sGRB	Short-duration γ -ray burst
SSP	Strong-stability preserving
TVD	Total variation diminishing
WENO	Weighted essentially non-oscillatory

1

Motivation

Come with me and you'll be in a world of pure imagination.

– Willy Wonka, *Charlie and the Chocolate Factory*

This thesis begins with an introduction to neutron star (NS) mergers, and an overview of the available literature, historical and current. Section 1.2 presents the phenomena that NS mergers are believed to be sources of, followed by a review of the progress and limitations of modern merger simulations, Section 1.3. With this in mind, in Section 1.4 we present the problems with current simulations that this work addresses, and end this chapter with an outline for the remainder of the thesis.

1.1 Introduction

Gravitational wave (GW) astronomy is beginning to enter maturity. One hundred years after Einstein predicted the existence of GWs (Einstein, 1916), the first direct observation of them was made by the LIGO collaboration (Abbott et al., 2016), who measured the distortions in spacetime produced by two merging stellar-mass black-holes (BH). Decades of progress in developing the hypersensitive equipment necessary for this detection culminated in the Nobel prize for physics being awarded to the founders of the collaboration.

Since their initial discovery, GW detections have become semi-regular, with signals typically being the result of equal-mass binary BHs (BBHs) (Abbott et al., 2019b), two cases of unequal-mass BBHs (Abbott et al., 2020a,b) and one confirmed case of a binary NS (BNS) merger (Abbott et al., 2017a). With these detections, the collaboration have ushered in the era of GW astronomy, and with plans of future detectors under way (Punturo et al., 2010, 2011) this already rich area of science will only continue to develop.

Part of the excitement for GW research is the ability to detect a single source via multiple means. For some astrophysical events, a detection of one kind may be accompanied by another—for example, GWs followed by an electromagnetic (EM) counterpart (Abbott et al., 2017a, 2019a), or a flux of neutrinos coincident with an X-ray detection (Cowen, 2020). Multi-messenger

astronomy is touted to be the future of observational astrophysics, and one promising (and validated) source is that of BNS mergers.

Mergers of NSs produce GW signals with multiple characteristic, ranging from inspiral through to merger and ring-down, and have been shown to emit electromagnetically in the optical and ultra-violet band (Nicholl et al., 2017; Valenti et al., 2017) as well as via X-ray, γ -ray and radio emission (Abbott et al., 2017b,c). Due to the extreme conditions present during a merger event, be it NS-NS or NS-BH, observing these events gives us a unique opportunity to infer properties of systems that are well beyond the reach of terrestrial laboratories, e.g. Margalit and Metzger (2017). In the remainder of this chapter, we will provide an overview of mergers, the physics we hope to probe in observing them, and the simulations and limitations thereof that we can and hope to perform. Whilst many of the properties of BNS and BH-NS merger are similar, this thesis will mainly focus on the former.

1.2 Mergers in nature

The mergers we study in this work are most often the products of a binary star system that was ejected from a star forming region, after which both stars exploded in a supernova leaving two compact objects, either a BH or a NS (Burrows, 2000; Faber and Rasio, 2012). If the two remnant stars are in a tight enough orbit to merge within the Hubble time, i.e. $P_{orbit} \sim 1$ day (Rosswog, 2014), gravitational radiation will further reduce the size of the orbit at an increasing rate until the two objects merge. At the point the objects merge, the stars typically have an orbital period of milliseconds, radiating $\sim 10^{53}$ erg/s in the form of GWs (Faber and Rasio, 2012).

In addition to gravitational radiation, NS mergers are the sources of a range of electromagnetically dominated behaviour—including, for example, short-duration ($\lesssim 2$ s) γ -ray bursts (sGRB) (Eichler et al., 1989; Nakar, 2007; Tanvir et al., 2013; Abbott et al., 2017b). The strong, poloidal magnetic field of the merger remnant seems to dominate the energy density in the polar regions, making it a likely site for the central engine of an ultra-relativistic outflow that produces an sGRB (Rezzolla et al., 2011).

Neutron stars are known to be endowed with incredibly strong magnetic fields, up to 10^{15} G. Merging neutron stars, however, should be fairly old, and thus there should have been a long period of time for the magnetic fields to decay, resulting in fields of around 10^{8-10} G (Baiotti and Rezzolla, 2017). Neutron stars with the strongest magnetic fields, nicknamed ‘magnetars’, have strengths estimated at around $\sim 10^{14-15}$ G, with estimates utilising the relation between the stars spin-down rate and its magnetic field strength (Kaspi and Beloborodov, 2017). Such immensely strong magnetic fields cannot be accounted for through typical means such as magnetic compression (collapse of a perfectly conducting, magnetised main sequence star) (Spruit, 2008). It is thought that some initial field is amplified exponentially through magnetohydrodynamic (MHD) instabilities, including the Kelvin-Helmholtz instability (KHI) and magneto-rotational instability (MRI) (Baiotti and Rezzolla, 2017), giving rise to the typical magnetic fields strengths seen in magnetars.

Another source of interesting behaviour is the neutron star matter itself. Second only to black-holes, neutron stars are incredibly dense objects, with a central rest-mass density greater than that of a nucleus, $\rho_{max}^{NS} \gtrsim 10^{15} \text{g/cm}^3$. Such high densities cannot be probed in earth-bound laboratories, and hence interactions between neutron stars during mergers can give information about the equation of state (EoS) of matter at nuclear densities, providing guidance for quantum theories of matter (QCD) (Lattimer and Prakash, 2004; Hanauske et al., 2019).

Furthermore, ejecta from a merger is the site for interesting nuclear physics phenomena. Rapid neutron-capture, the so-called *r*-process, should occur in abundance, and in doing so be the source for the majority of the heavy elements. By accurately modelling the ejecta and its composition in merger simulations, we are able to infer the masses of the *r*-process products and the relative abundances of elements throughout the universe (Freiburghaus et al., 1999). In addition, the decay of the unstable, neutron heavy isotopes produced during rapid neutron capture heat up the ejecta, and is one possible cause for the observed delay in prompt X-ray emission for BNS mergers (Desai et al., 2019). It is thought that the energy deposited via radioactive heating may delay the accretion fallback onto the merger remnant, although a self-consistent re-heating model has not been implemented in a merger simulation. Such a delay in the prompt X-ray emission is one mechanism for distinguishing between BHNS and BNS mergers, with extended emission favouring the BHNS system.

The radioactive decay also produces electromagnetic radiation across the optical/IR band (Li and Paczynski, 1998). This counterpart to the gravitational wave signal peaks on the order of days, powering a radioactive transient with a peak luminosity a thousand times greater than a classical nova, a *kilonova* (Metzger et al., 2010b). The combination of a strong gravitational wave signal in tandem with a prompt X-ray or γ -ray burst, and a following optical/IR transient are the indicators for a BNS merger such as GW170817 (Abbott et al., 2017a,b).

1.3 Mergers in silicon

A powerful method to gain information from mergers, particularly in the merger and post-merger phase, is via numerical simulation. Indeed, past the point where tidal deformability directly impacts the orbit of the inspiral, numerical simulations are the only available tool for a detailed analysis of the dynamics (Baiotti and Rezzolla, 2017). In this section, we outline the journey the field has taken to arrive at the accurate, high-resolution merger simulations performed today. We begin with a short history of the development of the relevant theories, before describing the current successes and limitations of the state-of-the-art simulations teams are capable of now.

1.3.1 A history in theory

Simulations of inspiralling and merging compact objects have benefited from a great deal of development in numerical relativity, computational fluid dynamics and MHD theory, and numerical methods over the last half-century. For accurate modelling of the dynamics of the spacetime one generally needs either modified Newtonian gravitation (a.k.a. post-Newtonian), or a full general relativistic description. As we will be interested in later sections in the evolution of a BNS system at early times post-merger, we will need the latter description, and so will focus on full

general relativity. The historical account of what follows is by no means exhaustive, but aims to highlight some of the key insights that have made modern day merger simulations so successful.

Numerical relativity

The laws of Einstein’s gravitation were laid out in 1916, and unsurprisingly these beautifully simple field equations are the basis of numerical relativity:

$$G_{\mu\nu} = 8\pi T_{\mu\nu}. \quad (1.1)$$

Here, $G_{\mu\nu}$ is the Einstein tensor that describes the geometry of the spacetime, and $T_{\mu\nu}$ the stress-energy tensor describing the matter within it.

Equation 1.1 lacks any notion of a coordinate frame, rather it holds for all choices of coordinates we wish to impose, defining instead the spacetime itself. In order to reduce this to a set of partial differential equations (PDEs) that can be solved in time, we must introduce the notion of time and corresponding *space-like* surfaces for an observer. This is the so-called 3 + 1 formalism, pioneered by [Misner, Thorne, and Wheeler \(1974\)](#).

The initial formulation of the equations of motion that received great attention was the Arnowitt-Deser-Misner (ADM) formulation, see Section 2.3.1. This constitutes a set of PDEs that may be solved to give the evolution of the spacetime geometry, subject to a number of constraints—the Hamiltonian and momentum constraints. Simulations using the ADM equations proved to be useful for numerical studies, but ultimately failed to give stable, long term evolutions. Constraint violations from finite precision numerics eventually add up, producing unphysical solutions ([Shinkai and Yoneda, 2000](#)). Furthermore, the evolution equations are not hyperbolic, and thus have an ill-posed initial value problem ([Hilditch et al., 2019](#)). Despite these flaws, two-dimensional axisymmetric simulations of head on collisions of BHs were able to extract information on the gravitational wave signals by 1975 ([Eppley, 1975](#)).

The non-hyperbolicity of Einstein’s equations was recognised and rectified in some cases as early as 1981 ([Friedrich, 1981](#)), motivated by the knowledge that hyperbolic formulations can be shown to demonstrate the existence and uniqueness of solutions, in addition to well-posedness ([Kelly et al., 2001](#)). The initial re-formulation of ADM by [Shibata and Nakamura \(1995\)](#), and later by [Baumgarte and Shapiro \(1998\)](#), demonstrated a much improved stability of the standard system. Using what later became known as the BSSNOK formulation ([Shibata and Nakamura, 1995](#); [Baumgarte and Shapiro, 1998](#); [Nakamura et al., 1987](#)), successful long term evolutions of low amplitude GWs ([Baumgarte and Shapiro, 1998](#)), the first BNS merger ([Shibata, 2000](#)), and a flurry of BBH firsts from 2005-2007 ([Pretorius, 2005](#); [Campanelli et al., 2006](#); [González et al., 2007](#)) soon followed. This was helped by the development of numerical GR gauge choices that can handle regions of strong spacetime curvature ([Thierfelder et al., 2011](#)).

Alternatives to BSSNOK, for example the CCZ4 ([Alic et al., 2012](#)) formulation, exist that may better satisfy the Hamiltonian and momentum constraints for some systems ([Alic et al., 2013](#); [De Pietri et al., 2016](#)), but excluding simulations of extreme mass-ratio binaries, one can consider the evolution of the spacetime more or less a solved problem. At the very least, attention

has moved away from improving numerical GR models and on to improving the realism of the HD, MHD and microphysics.

Hydrodynamics and magnetohydrodynamics

The first simulations of NS mergers were performed using Newtonian hydrodynamics (Zhuge et al., 1994; Davies and Benz, 1994). Nowadays, it is common for the HD equations to be written using the *Valencia* formulation (Marti et al., 1991). In this system, the equations of motion are written in the flux-conservative form in order to take advantage of the fact that the relativistic Euler equations are hyperbolic. This allows for the use of high-resolution shock capturing (HRSC) schemes, originally developed for classical fluid dynamics, to be applied in relativistic HD (Font, 2000), resulting in reliably stable simulations of shocks and preserving the conservative nature of the equations of motion. In Chapter 2 we discuss some of these schemes. For more on relativistic HD see Martí and Müller (1999); Font (2000) and for relativistic MHD see Font (2008); Martí and Müller (2015).

For simulations of mergers where the degree of tidal disruption is limited—i.e. for quasi-circular inspirals—finite volume-based codes are most commonly used. The ability to utilise efficient and reliable numerical methods (see next section) borrowed from classical fluid dynamics simulations allows for high resolution solutions. However, simulations of tidal disruption events, for example close encounters of white dwarfs and BHs (Rosswog et al., 2008), can be severely restricted by imposing a fixed physical domain, as in finite volume methods. Matter from extreme disruption events can be ejected to huge distances very quickly, and grid-based codes would require prohibitively large domains in order to track it. For these kinds of problems, it is more efficient to apply a smoothed particle hydrodynamics (SPH) approach, (see Rosswog (2014) for a review of SPH for compact objects).

SPH treats the evolution of the fluid from a Lagrangian perspective, with calculations of local quantities performed by an averaging over particles within a given radius, or smoothing length (see Rosswog (2010) for a far more extensive review of SPH on given spacetime backgrounds). The benefit of such a description lies in the fact that there is no computation wasted on regions of the simulation that are *uninteresting*—that is, there is no need to evolve an atmosphere—and that SPH particles are not confined to a finite grid, but rather are free to move in an infinite domain. Furthermore, SPH methods have had success with regards to the nuclear EoS. The first SPH study into the effects of a nuclear EoS, the Lattimer-Swesty EoS (Lattimer and Swesty, 1991), was performed in 1999 (Rosswog et al., 1999) in the context of mass ejection in BNS mergers. Since then, microphysical realness has continued to improve, with benefits including low density, temperature dependant EoSs (Rosswog and Davies, 2002) and neutrino leakage schemes (Rosswog and Liebendörfer, 2003). Unfortunately, no SPH code currently exists that can also evolve the spacetime as well as the matter, and so for accurate studies of late-stage BNS mergers, grid-based methods are the only means.

The first simulations of a magnetised BNS merger were performed in Price and Rosswog (2006), with continued success in the later years using full GR with dynamical spacetimes (Liu et al., 2008; Anderson et al., 2008; Chawla et al., 2010). To achieve an accurate simulation of the late-inspiral, merger and ringdown, a code must be able to solve the coupled Einstein-Maxwell-MHD system. The challenge in terms of adding magnetic fields to this comes in maintaining the

divergence-free Maxwell equation,

$$\nabla \cdot \mathbf{B} = 0. \quad (1.2)$$

A number of methods exist for this, including constrained transport (Evans and Hawley, 1988; Balsara and Spicer, 1999; Tóth, 2000) (which reconstructs the magnetic fields in such a way as to preserve this condition to machine precision) and divergence cleaning (Dedner et al., 2002; Komissarov, 2007) (which suppresses any violations, see Section 2.8). An alternative to constrained transport which also maintains $\nabla \cdot \mathbf{B} = 0$ by construction is the vector potential formulation. Introduced in Etienne et al. (2010), the vector potential evolution equation behaves very much like the magnetic field evolution and exists in conservation form. The magnetic field is then computed from the vector potential via $\mathbf{B} = \nabla \times \mathbf{A}$.

A specific choice of EM gauge must be chosen for this formulation. Initially, Etienne et al. (2010) adopted the algebraic gauge to simplify the evolution equation of the vector potential, but later simulations demonstrated that zero-speed, divergence-free violations emerged at mesh refinement boundaries, eventually polluting the entire domain (Etienne et al., 2012). The proposed solution requires evolving the EM scalar potential—i.e. solving $\nabla_\mu \mathcal{A}^\mu = 0$ where \mathcal{A} is the EM four-potential—in what is known as the Lorenz gauge, leading to significantly improved violations. Modern simulations can further damp any violations due to numerical errors by using the generalised Lorenz gauge, $\nabla_\mu \mathcal{A}^\mu = \xi n_\mu \mathcal{A}^\mu$, where ξ is a constant and n_μ is a timelike unit normal (Farris et al., 2012).

When teams model NSs in merger simulations, it is most commonly assumed that NS matter will be well approximated as a perfectly conducting, ideal fluid. That is, the models neglect any effects due to physical viscosity, thermal diffusion or electrical resistivity. The ideal fluid core of a neutron star has been known about for a long time (Baym et al., 1969), and these characteristics are built in to the relativistic Euler and ideal MHD (iMHD) equations that govern merger simulations. Nowadays, in the majority of cases where teams model magnetised neutron star binaries, excluding the addition of different microphysics, the system of equations that are evolved are those of iMHD—although, as we shall see, there is growing interest in alternative formulations. The successes that have been achieved with iMHD are discussed in Section 1.3.2, and the limitations and work on extensions to these models are presented in Section 1.3.3, and is the focus of the later chapters of this thesis.

Numerical methods

The system of 3 + 1 numerical relativity and GRMHD comprises a set of PDEs describing the transport of fluid quantities. The process of solving these sets of equations equates to a Cauchy initial value problem, i.e. once given a set of initial data, produce the state of the system at some time later subject to some boundary conditions. A powerful tool for solving systems of PDEs in which fields vary in both time and space is the so-called *method of lines* (MoL).

MoL reduces the set of PDEs to a set of ordinary differential equations (ODEs) by introducing a spatial discretisation to approximate spatial derivatives. Once the system is in ODE form, it can benefit from any of the well-known and tested ODE integrators—generally the explicit Runge-Kutta (RK) method, first introduced in 1895 (Runge, 1895), see Butcher (1996) for a detailed history and review. RK integrators are multi-stage methods, where the right hand side (RHS)

of the ODE is approximated at a number of points within one (time) step to improve accuracy. When the RHS terms of an RK step can all be computed, the method is explicit. Explicit RK methods have a bounded region of stability (Süli and Mayers, 2003), and so are not suitable for systems with large source terms. So called *stiff* equations require the update to be calculated implicitly to improve stability at the cost of additional computation (see Section 2.6). These kinds of integrators have been used in merger simulations, and reducing the cost imposed by these methods is one of the focuses of this thesis.

With regards to the evolution of the space time, smooth initial data will generally result in a smooth solution at all times (excluding at singularities and matter contributions), and so suitable methods for the spatial discretisation include high-order finite differencing, or spectral methods (Lehner and Pretorius, 2014). Spatially discretising the GRMHD equations, however, requires more care. Smooth initial data can lead to the formation of shocks which become unstable to numerical evolution when using vanilla finite difference methods at discontinuities, leading to the well known Gibbs oscillations at discontinuities (Gottlieb and Gottlieb, 2003). To handle shocks, a number of methods can be used that result in a high degree of accuracy for smooth flows, and a reduced order of accuracy but crucially stable for discontinuous flows. Simple flux limiting at cell boundaries can provide an easy solution and is generally second or third order accurate in space (LeVeque, 2002). The first application of classical fluid dynamic HRSC methods to relativistic dynamics was in Marti et al. (1991), and all modern finite volume based GRHD/MHD codes now utilise similar methods. Methods with these properties are known as modern shock capturing, or high resolution shock capturing (HRSC) schemes, which we discuss in more detail in Section 2.5.

The majority of codes used for merger simulations are structured on a Cartesian grid, often with an extent of many hundreds, if not thousands, of kilometres. At the centre, two neutron stars with a radius of $\sim 10\text{km}$ resolved to $\sim 100\text{m}$ are modelled—evolving the the entire domain at this resolution is a big ask, and not just in terms of the memory footprint. To make simulations more tractable, teams adapt the resolution of the mesh for different regions, improving the accuracy for *interesting* parts and reducing cost in *uninteresting* parts. Modern day mesh refinement dates back to Berger and Oliger (1984) where refined meshes could be dynamically added or removed according to some local estimate of the error. For any fluid evolution simulation, adaptive mesh refinement (AMR) provides a crucial means of capturing behaviour at a range of scales efficiently, and many open source packages now exist allowing relatively simple inclusion into simulations (MacNeice et al., 2000; Schnetter et al., 2004; Liebling, 2002; Zhang et al., 2019)¹.

1.3.2 Current successes

The culmination of the work of the last section have led to the highly accurate and efficient simulations of mergers that we can perform today. Success in these areas has allowed us to investigate a plethora of phenomena, including late inspiral, merger, post-merger, collapse to BH, and long term evolution of the ejecta. Here, we highlight some of the key breakthroughs achieved in numerical simulations of mergers and the physics that we can gain from them.

¹These codes are available at the time of writing at www.cactuscode.org/, www.had.liu.edu/, [www.github.com/AMReX-Codes/amrex](https://github.com/AMReX-Codes/amrex)

The first successful simulations of neutron star mergers using full general relativity were performed at the turn of the new millennium (Shibata, 1999, 2000). Adopting an ideal fluid equation of state for the hydrodynamics, these simulations demonstrated the sensitivity of the binary merger product (BMP) to the initial conditions, particularly the compactness of the progenitor stars. They also gave estimates for the total mass of the disc surrounding the BMP, and identified a relationship between the BMP and the possible GW signals that are emitted.

These initial simulations paved the way for numerous parameter studies, culminating in the publicly available GW template library, CoRe (Dietrich et al., 2018), among many others. Parameter studies are a means of assessing how the initial data used in a merger simulation, such as the total binary mass or spin, affects the resulting observables (Ajith et al., 2014). For example, Dietrich et al. (2017) shows how the NSNS mass-ratio directly alters the imprint on the GW multipoles in the late-inspiral-merger phase—larger mass-ratios lead to a decrease in GW frequencies.

Equation of state

Similarly to the relationship between the mass-ratio and the emitted GWs of a binary, numerous studies have been performed to understand how the EoS of NS matter impacts the output of merger simulations. EoSs relate the thermodynamic properties of the fluid, things like temperature and pressure, to evolved quantities such as the density and internal energy. The EoS of matter at extremely high (i.e. nuclear) densities is one of the biggest unknowns regarding the physics of NSs—the different types of possible EoSs for merger simulations is large, and currently there is no real consensus on which may characterise NS matter best.

EoSs are often characterised simply as stiff, i.e. less compressible matter, or soft, and the stiffness of the matter assumed in a simulation can have a big effect on, for example, the total mass or composition of the ejecta (Sekiguchi et al., 2016), the time of collapse (Kawamura et al., 2016), or even the time of merger (De Pietri et al., 2016). The BMP is also affected, with EoSs that permit larger maximum masses being more likely to produce a hyper-massive NS, rather than a black hole, and stiffer EoSs delaying the time of collapse (Kawamura et al., 2016). Comparisons of GW170817 with numerical results can help constrain EoS models and parameters (Radice et al., 2018b; De et al., 2018), and the maximum mass of NSs (Margalit and Metzger, 2017; Godzieba et al., 2020).

Magnetic fields

The first simulations of magnetised NS mergers were performed using SPH methods and modified Newtonian gravity (Price and Rosswog, 2006). Later, full GRMHD simulations from an Eulerian viewpoint demonstrated where magnetic fields are important. Magnetic fields were shown to have no effect on the inspiral phase (Giacomazzo et al., 2009), but have a dramatic effect on the dynamics post-merger (Giacomazzo et al., 2011), even delaying the collapse to BH (Anderson et al., 2008).

The dynamics of the magnetic fields are complex in a merger simulation, and subject to numerous instabilities. The KHI was identified in the first magnetised BNS merger simulation

as a mechanism for amplifying fields by many orders of magnitude within milliseconds (Price and Rosswog, 2006). The growth of the magnetic field is governed by a number of factors, but should cease when the magnetic energy reaches equipartition with the energy contained within the turbulent eddies (Zrake et al., 2013).

The growth of the magnetic field is heavily dependant upon the resolution of the simulation (Kiuchi et al., 2015)—numerical viscosity acts to suppress the growth of the KH vortices, damping the possible amplification. Even for the highest resolution BNS merger simulations to date (Kiuchi et al., 2017), with a central resolution of $\Delta x = 12.5\text{m}$, there are few signs of convergence with regards to the magnetic amplification, or the ability to sustain a magnetoturbulent state post-merger. As a result, magnetic energy equipartition has not been achieved from realistic initial magnetic field strengths.

The second dominant mechanism for magnetic amplification appears to be MRI. Post-merger, the remnant is surrounded by an accretion disk situated in a large-scale magnetic field. This is the site for the MRI which acts to further amplify the magnetic field whose saturation is independent of the initial strength (Guilet et al., 2020). The longer the delay of the collapse to BH, the longer this process can continue, and as a result simulations that employ stiffer EoSs will tend towards stronger remnant magnetic fields (Kawamura et al., 2016). The turbulence as a result of the MRI also plays an important role in the later stages of evolution related to the r -process (Siegel and Metzger, 2017).

Jets and sGRBs

Strong magnetic fields are an essential ingredient for the formation of jets, posited to be the progenitors of sGRBs (Paczynski, 1986; Eichler et al., 1989; Narayan et al., 1992). The process by which mergers result in sGRBs is still not fully understood, partly because simulations from merger to the jet formation that produces the sGRB require a very long time ($10^2\text{--}3\text{ms}$) (Ciolfi, 2018; Ruiz et al., 2019). The first solid indication from merger simulations that seemed to support the merger-jet-sGRB hypothesis came almost a decade ago (Rezzolla et al., 2011), where simulations indicated the formation of a magnetically dominated polar region with a jet-like structure. Later, it was demonstrated that this magnetic structure that is so important for an incipient jet is independent of the EoS, and appears to be a general feature of BNS merger simulations (Kawamura et al., 2016).

Fast-forwarding to the present day, parameter surveys have been possible, highlighting the role of magnetic field configuration and orientation (Ruiz et al., 2016, 2020), NS spin (Ruiz et al., 2019) and electrical resistivity (Qian et al., 2018) on the BMP and fallback accretion. The general picture (Ciolfi, 2018) amounts to a magnetised BNS (or a fraction of BHNS) merger(s) resulting in magnetar strength magnetic fields, which may be initially extending/confined or aligned/anti-aligned, finally settling to a large-scale dipolar structure post-merger. The most likely BMP is that of a stellar-mass BH surrounded by an accretion torus with a mass of $\sim 0.1M_\odot$ and an accretion rate of $\sim 0.1 - 1M_\odot s^{-1}$. Whilst the formation of a jet from a long lived neutron star has not been entirely ruled out, the higher ejecta mass in the atmosphere makes a jet breakout seem less likely than for a scenario in which the remnant collapses at earlier times (Ciolfi et al., 2019). This is consistent with estimates on the delay time of collapse of the GW170817 remnant,

with bounds of $t_{BH} \lesssim 0.9 - 1.1\text{s}$ (Murguia-Berthier et al., 2020; Murguia-berthier et al., 2017; Gill et al., 2019; Lazzati et al., 2020).

Two mechanisms are believed to be at play that may power the jet: $\nu\bar{\nu}$ annihilation in the polar regions, and the Blandford-Znajek mechanism (Blandford and Znajek, 1977). It is not known which of these mechanisms is dominant, or if they contribute equally, although it seems that a purely neutrino driven jet may be unlikely (Murguia-Berthier et al., 2014; Just et al., 2016; Perego et al., 2017). Simulations indicate that the resulting jet has an opening angle of $\gtrsim 10^\circ$, and have been able to constrain the viewing angle of GRB170817A, the counterpart to GW170817, to $\approx 21^\circ$ (Nathanail et al., 2020).

Ejecta

The relativistic outflow from the jet accounts for only part of the total matter ejected from the merger, with the rest being supplied by tidal or shock driven ejecta from the merger process, or baryon-rich winds associated with the BMP (Ciolfi, 2020). The total mass ejected has been shown to have a strong dependence upon, for example, the mass ratio of the progenitors. There appears to be a somewhat linear relation between the mass ratio, the total mass ejected, and its corresponding kinetic energy (Dietrich et al., 2017). Furthermore, larger mass ratios tend to lead to greater accretion rates post-merger (Kawamura et al., 2016). Similar features can be found for different EoSs, with more massive accretion disks resulting from softer EoSs (Palenzuela et al., 2015; Lehner et al., 2016; Sekiguchi et al., 2016).

Typically one expects the mass ejection over a period of $\gtrsim 0.1 - 1\text{s}$ (Ciolfi and Kalinani, 2020), making simulations following the final stages of merger through to total mass ejection and jet formation virtually intractable. Despite this, some very-long-term evolutions—from 100ms (Ciolfi et al., 2019) up to 250ms (Ciolfi and Kalinani, 2020) post-merger—have been achieved, and given insights into the late time behaviour of magnetically driven ejecta. Furthermore, it appears that MHD simulations demonstrate a significant increase in total ejected matter over hydro simulations— $2\times$ (Fernández et al., 2019)—signalling the importance of including magnetic fields in models to be able to derive accurate results. The matter ejected on timescales $\gtrsim 1\text{ms}$, however, appear to be dominated by neutrino cooling or viscous angular momentum redistribution (Fernández et al., 2019).

The r -process and kilonovae

The reason there is so much emphasis on the behaviour of ejecta at late times is for its impact on the formation of heavy elements and the corresponding EM counterpart. There is mounting evidence for mergers as the production site of the r -process elements (Kasen et al., 2017; Rosswog et al., 2018). The composition of the ejecta, and how bound it is, has a large impact on the colour of the kilonova and the time of its peak. Even the rate of heating due to the radioactive decay of the unstable r -process elements impacts upon the orbits of marginally bound matter (Metzger et al., 2010a), potentially delaying fallback on to the BMP and leading to a delay in the prompt X-ray emission (Desai et al., 2019).

Because the composition of the matter is vital to determining the expected behaviour, it is common practice to include weak interactions via neutrino treatments, especially for simulations of the later stages (Palenzuela et al., 2015; Siegel and Metzger, 2017; Foucart et al., 2014; Siegel and Metzger, 2018; Foucart et al., 2019; Decoene et al., 2020). Neutrinos provide a mechanism for the efficient transport of energy, cooling of the accretion disk and generally lead to lower electron fractions, Y_e (Siegel and Metzger, 2017), whilst also heating the polar regions via annihilation to the tune of $\sim 10^{49}$ ergs (Perego et al., 2017). There is also evidence to suggest that hotter BMPs will drive densities down, allowing greater rates of positron capture that in turn increase Y_e (Lehner et al., 2016). Higher electron fractions will result in the neutron capture choking, therefore producing a smaller fraction of third peak r -process nuclei.

The kilonova resulting from the radioactive decay of the neutron heavy, r -process isotopes has also been studied—the longest merger simulations to date indicate that baryon driven winds are the site for the blue, UV kilonova (Ciolfi and Kalinani, 2020). The mass ratio also impacts the peak of the kilonova signal, with larger mass ratios ($q \sim 2$) tending to result in peaks in the near-IR some weeks after merger (Dietrich et al., 2017).

1.3.3 Modern limitations

Despite the successes that simulations of mergers have achieved in addressing different areas of astrophysics, there are a number of factors limiting further progress. In one way or another, these limitations are all predicated on the finite computational resources available to researchers.

1.3.3.1 Resolution

As an example of computational limitations, we mentioned previously that simulations have shed light on how NS magnetic fields are amplified during and after mergers via MHD instabilities. However, in the highest resolution simulations to date (Kiuchi et al., 2014, 2015, 2017) the magnetic field amplifications did not reach saturation even with a resolution of $\Delta x = 12.5$ m. Considering that computational cost scales as $\propto N^4$, where N is number of grid points in one direction, simulations of this scale are infeasible for many teams, with this run totalling 40 million CPU hours on the Japanese K Supercomputer. For the time being, it appears that we should not expect simulations that are able to fully resolve amplification mechanisms, including the KHI and MRI. Furthermore, lower resolution simulations possess a greater amount of numerical diffusion that is not present in better resolved simulations. This additional diffusion can act to suppress the growth of the instabilities (Zhang et al., 2009), and if energy is transferred through an inverse cascade—from small scales to larger scales—can be prevented entirely, purely by numerical effects.

To address this limitation there are a number of avenues besides additional computer time. A simple method that can be used if one is interested primarily in the post-merger dynamics is to simply increase the strength of the magnetic fields in the initial data. Typically, NS magnetic fields are set to $\sim 10^{12} G$ in simulations of mergers, as these are the fields possessed in pulsar binaries (Lorimer, 2008). In Ciolfi et al. (2019), the authors impose a seed magnetic field of $B_{\text{max}} = 10^{16} G$, and whilst old NSs are not thought to be endowed with such strong fields,

doing this should make up for the reduced amplification on account of the insufficient resolution post-merger. This will be the topic of a later section, 6.4.5.

Alternatively, one can formalise the effect that the unresolved motion would have on the dynamics of the matter and electromagnetic fields, where the unresolved motion accounts for anything that occurs on scales less than the finest grid-spacing, $l \lesssim \Delta x$. Commonly used in simulations of classical turbulence, subgrid models aim at modelling the unresolved motion, or the effects thereof, in order to emulate the results of a higher resolution simulation. For large eddy simulations (LES), the equations of motion of the fluid are explicitly defined in terms of resolved and unresolved quantities, and an assumption about the connection between the two scales is made to close the system. Initially proposed in Smagorinsky (1963) for simulations of the atmosphere, Smagorinsky’s idea has been widely applied to systems in which turbulence or vortical behaviour is important.

The first LES merger simulations were performed using HD by Radice (2017), and investigated the impact on the structure of the BMP remnant, emitted GWs and neutrino emissions of the unresolved eddies. The extension of LES to MHD was done a few years later by Viganò et al. (2019), and further developed for relativistic fluids (Viganò et al., 2020), an important step towards full MHD LES merger simulations.

An alternative scheme for achieving magnetar strength magnetic fields in low resolution merger simulations is by artificially amplifying magnetic fields via a carefully constructed source term in the equations of motion. This was the work of Giacomazzo et al. (2015), whereby the RHS of the vector potential evolution equation picks up an additional term acting to increase its amplitude. Introducing a subgrid model in this way, whilst not constrained in a rigorous mathematical way, does provide a means of limiting the effect that an insufficient resolution has on the post-merger dynamics.

1.3.3.2 MHD model

The particular choice of model will drastically alter the output of a given simulation. We have discussed how the post-merger dynamics can be largely affected by the inclusion, or not, of magnetic fields, and whilst the complexity of the equations governing iMHD increase somewhat over pure hydro models, there is little difference in the numerical methods either use. For example, identical time integrators can be used for hydro and iMHD simulations, and other than the inclusion of a handful of terms in the stress-energy tensor, the evolution of the space-time remains unaltered also. As a result, there will be only a relatively small additional computational cost when adding the evolution of the magnetic (or vector potential) field over the cost of a hydro simulation, partly because the evolution of the spacetime takes up a significant fraction of the total runtime.

The equations of ideal MHD assume a number of approximations from a full kinetic theory. As with all fluid models, the typical length scales of the system under investigation are assumed to be much larger than the mean free path between the plasma’s constituents. In addition, there is assumed to be zero (physical) viscosity, no heat conduction and zero resistivity (Font, 2008). For NS matter, which is generally considered to be perfectly conducting, this is typically regarded as a good approximation. With respect to the atmosphere, however, there is no real

consensus on the correct formulation of MHD. Far from the merger, one expects the EM waves to be modelled via Maxwell’s vacuum equations (Dionysopoulou et al., 2013, 2015), whereas in the low density, magnetically dominated magnetosphere of the BMP, the inertia of the EM fields will dominate and the equations of force-free electrodynamics (FFE) may be more applicable (Blandford and Znajek, 1977; Contopoulos et al., 1999; Gruzinov, 2007; Dionysopoulou, 2016). The picture is further complicated by the addition of the artificial atmosphere in finite volume simulations, generally used to allow for stable evolutions in low density regions (Etienne et al., 2010).

To address the problems in the atmosphere, some teams have utilised resistive descriptions of MHD. Studies of pulsar phenomenology and magnetar outbursts have utilised FFE—the limit of resistive MHD (rMHD) whereby the electromagnetic inertia dominates over the fluid inertia (Spitkovsky, 2006; McKinney, 2006; Pétri, 2016; Carrasco et al., 2019)—in addition to the magnetospheres of isolated binary black holes (Palenzuela et al., 2010).

Notably, Lehner et al. (2012) developed a hybrid approach, in which iMHD is used to model the interior of isolated NSs, whilst coupled to the FFE limit of resistive MHD in the exterior. Using this approach, they are able to more accurately model the electromagnetic outbursts of a collapsing hypermassive NS (HMNS). The iMHD solution at the edge of the HMNS is used as boundary conditions for the FFE equations, but no back-reaction of the FFE solution on the star is present. An extension of their work corrects for this (Ponce et al., 2014), by instead evolving the equations of resistive MHD throughout the domain and switching between different limits via the definition of Ohm’s law. By defining the phenomenological Ohm’s law to depend upon the local rest-mass density, they are able to model the ideal NS matter and the magnetically dominated magnetosphere using one self-consistent description in two different limits. With this set-up, they were able to perform resistive merger simulations, demonstrating that orbital energy from the binary is efficiently transferred into an pre-merger, EM counterpart.

Resistive descriptions may also apply to the interior of NSs. Further investigations into the electrical properties of NS matter have been carried out, and a study by Harutyunyan and Sedrakian (2016) shows a dependence of the finite electrical conductivity on local parameters: density; temperature; and magnetic field strength/orientation. Later calculations (Harutyunyan et al., 2018) suggest the time and length scales that these features should be important on are not conducive to being seen in merger simulations. They have, however, assumed smooth profiles in both the matter and EM fields, something which certainly is not always true in mergers where shocks inevitably develop.

To understand the importance of resistivity on the matter in BNS simulations, Dionysopoulou et al. (2015) performed fully resistive, GRMHD simulations of BNS mergers with a finite conductivity profile for the matter, evolving the system for ~ 20 ms post-merger. Their analysis shows that alterations in the transfer of angular momentum result in a delayed collapse of the remnant in a resistive simulation. Similar to the aforementioned force-free studies of Lehner et al. (2012) and Ponce et al. (2014), they were also able to analyse the electromagnetic counterpart that follows the merger as their model reduced to the vacuum Maxwell equations far from the merger.

Despite this, problems related to the stiffness of the electric field evolution in full rMHD means the value of the resistivity is set to be unphysically large in order to ease numerical treatment². In the resistive mergers of Ponce et al. (2014) and Dionysopoulou et al. (2015), conductivities of $\sigma_{\max} \sim 10^{11} \text{s}^{-1}$ are referred to as the *ideal* limit of rMHD, despite these values being at least six orders of magnitude too resistive for NS interiors. Furthermore, in the simulations of Dionysopoulou et al. (2015), the initial NS orbits are kicked towards merger in order to keep simulations runtimes reasonable, making uncertain their conclusions about the realistic effect of resistivity. Due to the comparatively long timescale of resistive diffusion to the matter dynamics (Harutyunyan et al., 2018), longer term simulations of mergers are needed to assess the impact of a physically realistic resistivity.

Looking further into the future of merger simulations, one can argue that evolving separately the various charged species will provide a more accurate description of the underlying physics. Indeed, this is the aim of the work of Andersson et al. (2017b,a). The full set of their evolution equations comprise of four separate fluids—positive, negative, and neutral charge carriers and entropy—allowing for the evolution of heat, charge separation and the effects of a scalar resistivity. Unfortunately, the full equations are not in a numerically suitable form for long term evolution, as in their most general form that permits entrainment they cannot be written in flux-conservative form.

Simplifications of the multi-fluid equations can reduce the number of fluids. Amano (2016) developed a resistive two-fluid model—only positive and negative carriers—for potential use in astrophysical settings. Similarly to the electric fields in rMHD, it appears that in multi-fluid models the charge-mass ratio of the constituent species can result in potential stiffness in the equations of motion, leading to the same requirement of implicit time integration schemes such as those presented in Pareschi and Russo (2005) (see Section 2.6.2). Once again, evolving multi-fluid models such as these using the same numerical methods as for rMHD merger simulations will result in excessive simulation runtimes, and severely limit the physical accuracy one can achieve with regards to initial data or resolution.

1.4 Problems to address

To briefly summarise the previous section, physically motivated and precise simulations of mergers are computationally expensive. Including EoSs motivated by nuclear theory combined with weak interactions and neutrino emission, resistive and viscous diffusion and dissipation, interactions between constituent charge species and sufficient resolution to capture dynamics down to the dissipative scales is insurmountable for at least the short to mid-term future. A theorist’s dream is a numericist’s nightmare.

In Section 3.2, we will demonstrate some of this, focussing on different models of MHD. We will also demonstrate the limitations brought on by finite resolution, specifically by describing its effect on the KHI, known to be an important feature in simulations of mergers.

In a few words, this thesis addresses the restrictions imposed on astrophysical simulations by limited computational resources and combats these restrictions in three ways. The first regards

²Similar stiffness issues can also arise in neutrino treatments in GRMHD, however we will not discuss neutrino treatments or their numerical difficulties in this thesis.

utilising massively parallel processors to speed up simulations of resistive and two-fluid relativistic MHD (RMHD). At the time of writing, seven of the top 10 most powerful supercomputers around the world are powered by graphics processing units (GPUs), up from five only two years ago. Progress in deep learning, arguably one of the greatest scientific achievements of this millennium, can be partially attributed to the computational power afforded by large systems of GPUs (and tensor processing units), and machine learning models would be severely limited if this resource was not available.

In keeping with this trend, a number of HD/MHD codes have been ported to run on GPUs. All of these examples of GPU capable HD/MHD codes evolve systems which can be suitably integrated using explicit schemes, but there is uncertainty on how the additional data required for an implicit integration will affect parallel speed-ups. This can have a big impact on the future of merger simulations, as if resistive and multi-fluid models are to be widely adopted they must be efficient to run on the next generation of supercomputers. We investigate this in Chapter 4, utilising the special relativistic MHD code that has been developed in tandem with this thesis, porting a number of the routines to the GPU.

The second approach takes inspiration from the techniques of sub-grid scale modelling, and applies it to the case of rMHD. The limiting factor in rMHD is the potentially stiff RHS of the electric field evolution equation, see Section 2.2.3, which becomes very large when the fluid is highly conducting, as is the case for most of the matter in a BNS merger simulation. In Chapter 5, we develop a first order correction in the resistivity to the equations of special relativistic iMHD, and perform a number of test problems to identify the appropriate convergence and stability properties of the resulting model.

Later, in Chapter 6, we extend our model for applications in general relativity, and perform full merger simulations with it. The ability to evolve the model explicitly for large conductivities means there is very little constraining the initial data or possible values of resistivity, and so we perform these simulations with parameters taken from studies of NS conductivities (Harutyunyan and Sedrakian, 2016). This represents the first attempt at modelling a BNS merger using both a realistic magnitude for the resistivity within the NS and typical merger initial data. From these simulations, we infer a number of differences regarding our results with realistic resistivities and those generated with iMHD. We provide direct comparisons of our resistive merger result with an iMHD model using identical initial data, and discuss the potential impact of the results on observables, and the future of merger simulations.

Finally, the work in Chapter 7 takes inspiration from the hybrid models of MHD mentioned in the previous sections. We present two hybrid approaches as a proof-of-concept for the potential of resistive hybrid models for astrophysical simulations. The first hybrid approach combines the benefits of iMHD and rMHD, switching between the models dynamically based on local properties of a simulation, and enhanced with the resistive extension developed in Chapter 5. The second aims at improving the efficiency of the merger simulation of Chapter 6. Noting that there is very little variation in the conserved quantities in the atmosphere, we treat this region without computing the resistive correction, thereby saving on potentially unnecessary computation.

2

Mathematics and numerics

Computers can make very fast, very accurate mistakes.

– Anon.

Numerical methods can be dated back centuries—for example, to techniques used to solve single equations in which no closed form solution can be found ([Wallis, 1685](#)). Initially the interest of mathematicians, the field has developed over the 20th and 21st centuries at an extraordinary rate—much to the benefit of the physical sciences.

In what follows, we give a brief overview of how one models neutron stars numerically and the mathematical models that exist to this end. Following this, we tackle general conservation laws and the infamous Riemann problem, and then set out the methods used today in solving this problem. Later sections relate to specific methodologies that we implement in our numerical, multi-fluid code `METHOD`, that has been developed in tandem with this thesis. We discuss this code in detail in [Chapter 3](#).

2.1 Introduction

Modern day numerical methods are indebted to [Von Neumann and Goldstine \(1947\)](#), whose attempts to understand the accuracy of applied linear algebra operations spawned the type of numerical analysis that is essentially a requirement when developing new methods today.

The persistent development of an expansive range of techniques for any number of physical problems has resulted in computational methods becoming an integral part of the scientific method—originally, such methods were of the domain of mathematics and physics, but more recently fields such as biology and chemistry have developed their own *in silico* techniques to generate new data for problems from protein folding and modelling the circulatory system, to tumour recognition in X-rays. The application, and furthermore the necessity, of robust, reliable, and computationally efficient algorithms in many fields of modern day research is altogether inescapable, and to such an extent that there is significant pressure on research institutions and

universities alike to invest in and maintain exceptionally powerful high performance computing (HPC) suites.

We have touched upon a key notion that reappears frequently when considering numerical methods for different systems: computational expense. By expense, we refer to the number of arithmetic calculations required to execute a simulation from start to end, whereby more expensive simulations require more computations to be performed.

2.2 Neutron stars are fluids

The end goal of this thesis is to understand and to ease the computational demands set by non-ideal models of MHD for use in BNS merger simulations—particularly that of rMHD—and to then use these models to investigate physically realistic resistivities. As such, in this section we present the models of MHD that are the basis of the following work.

Each of the models presented in the following sections represents a different set of approximations and thus different physics. Each model will (approximately) agree with reality to varying degrees and only be suitable for modelling certain systems. We make clear the benefits and limitations of each model. Henceforth, we shall assume that the fluids we discuss are compressible (i.e. a gas to a physicist), and shall introduce and relax various approximations such as electrical resistance and relativistic effects as we progress. We also assume charge neutrality throughout.

It is also worth noting that all the models of MHD considered in this chapter are those for dynamic processes. Whilst (magneto)hydrostatics are important in understanding the structure of isolated compact objects, in order to understand the evolution of chaotic, cataclysmic events one must consider full time-dependence. Furthermore, in the present section we have limited our scope to that of the special relativistic limit of GR, however the same techniques for deriving iMHD apply in a general relativistic framework, for example. We discuss extensions of these models to GR, specifically a resistive extension to iMHD, in Chapter 6.

Finally, we use a mostly-positive metric and adopt the standard Einstein summation notation in which we sum over repeated indices (unless otherwise stated)—Latin indices indicate the three spatial dimensions, $(a, b, i, j) \in \{1, 2, 3\}$, and Greek indices indicate the four space-time dimensions, $(\mu, \nu) \in \{0, 1, 2, 3\}$.

2.2.1 Euler equations

The simplest model we present is that of the Euler equations. These represent an uncharged, non-relativistic fluid flowing in 3D defined by the Cartesian axes x_i for $i \in \{1, 2, 3\}$.

$$\partial_t \begin{pmatrix} \rho \\ \rho u_i \\ \mathcal{E} \end{pmatrix} + \partial_j \begin{pmatrix} \rho u_j \\ \rho u_j u_i + \delta_{ij} p \\ (\mathcal{E} + p) u_j \end{pmatrix} = 0, \quad (2.1)$$

We have used the following relations: ρ , \mathcal{E} , p and u_i represent the mass density, energy density, fluid pressure and fluid 3-velocity in the i^{th} direction respectively, and δ_{ij} is the Kronecker delta.

The total energy density is defined as $\mathcal{E} = \rho e + \rho u^2/2$, with e the specific internal energy. We also use the shorthand $\partial_i \equiv \partial/\partial x_i$ to represent the partial derivative of some quantity along the x_i direction, and $\partial_t \equiv \partial/\partial t$ the partial derivative with respect to time.

It is useful to consider the first term on the LHS as the vector of conserved quantities. Indeed, this is one way to derive the Euler equations, by considering the conservation of mass, momentum, and energy of the fluid through a fluid element. Any changes in the values of the conserved quantities in a fluid element are due to a flux of that quantity through the cell as a result of advection and internal (and possibly external) forces..

To close system (2.1) we need an EoS relating the thermodynamic quantities such as pressure, density and internal energy. Whilst this may be exceedingly complex in reality, it does not add much to the way in which we must solve these systems (see [Mignone and McKinney \(2007\)](#), for example), and so for our purposes, as is commonly done in numerical relativity, we adopt the ideal Γ -law EoS, $p = \rho e(\Gamma - 1)$. This is the EoS we shall use in all subsequent models.

The equations given in (2.1) are very simple, so much so that they are only (if ever) used to model the most elementary of astrophysical systems—for much of astrophysics the limitations of the Euler equations render them wildly inappropriate for realistic modelling, even after including gravity. We present them here as an intuitive reference of later, more realistic models. The general structure of the more realistic alternatives are essentially unchanged from those that are presented in this section.

The equations generally have the following form: the first equation is related to the change in the total mass (or particle number density) in some volume; the next component(s) describes how the momentum of that volume evolves due to an energy flux and some internal pressure gradient in the fluid; the next describes how the total/kinetic energy in the system changes due to work done on the fluid element and the motion of the fluid.

Additional equations may be needed to describe the evolution of electric and magnetic fields, or the evolution of separate fluid species for example, but the general layout and interpretation of the equations in the system remains the same. Thus, we present the simpler Euler equations to give a guideline on the common themes of the following systems and their respective interpretations—despite the systems of equations appearing more complex, they represent many of the same ideas.

2.2.2 Ideal magnetohydrodynamics

Many astrophysical systems require a description of electric and magnetic fields, in addition to the hydrodynamics. This is the domain of MHD, and what follows is an outline of some MHD models in special relativity that we will use repeatedly throughout this thesis.

A common model of MHD used in much of astrophysical modelling is that of an ideal, or perfect, magnetofluid. Perfect fluids exhibit no dissipative behaviours—that is, they are perfect conductors of electricity, and do not conduct heat or possess viscosity. As in the Newtonian case presented in Section 2.2.1, the iMHD equations are derived by considering the conservation of various quantities.

Much of what follows in this and subsequent sections has been adapted from [Palenzuela et al. \(2009\)](#); [Antón et al. \(2010\)](#); [Dionysopoulou \(2016\)](#), with varying notation. Let N^μ and $T^{\mu\nu}$ be the rest-mass current density and stress-energy tensor for a perfect magnetofluid, where

$$N^\mu = \rho u^\mu, \quad (2.2)$$

$$T^{\mu\nu} = \rho h^* u^\mu u^\nu + \eta^{\mu\nu} p^* - b^\mu b^\nu, \quad (2.3)$$

$$(2.4)$$

and ρ is the rest mass-energy density, the 4-velocity is given by u^μ , $\eta^{\mu\nu}$ is the flat space-time metric, b^μ is the magnetic field 4-vector and h^* and p^* is the specific enthalpy and fluid pressure, both modified to include contributions from the magnetic field (defined later). The magnetic 4-vector, b^μ , is related to the magnetic field in the lab frame, B^i , via the following relations:

$$b^0 = W B_i v^i, \quad (2.5)$$

$$b^i = B^i/W + b^0 v^i, \quad (2.6)$$

$$b^2 = B_i B^i/W^2 + (B_i v^i)^2, \quad (2.7)$$

for which the 4-velocity $u^\mu = W(1, v^i)$, with $W = 1/\sqrt{1 - v_i v^i}$ as the Lorentz factor for a 3-velocity given by v^i . Conservation of Baryon number (or rest-mass) and of total energy and momentum are imposed via the following relations:

$$\partial_\mu N^\mu = 0, \quad (2.8)$$

$$\partial_\mu T^{\mu\nu} = 0. \quad (2.9)$$

If we consider first the conservation of mass given by Equation (2.8), we get the following:

$$\partial_t(\rho W) + \partial_i(\rho W v^i) = 0, \quad (2.10)$$

in which we have used $\partial_0 \equiv \partial_t$. Now consider energy-momentum conservation, Equation (2.9). The time component, $\nu = 0$, gives

$$\partial_\mu [\rho h^* W u^\mu + \eta^{\mu 0} - b^\mu b^0] = 0, \quad (2.11)$$

and performing the sum over time and space, defining $U = \rho h^* W^2 - p^* - (b^0)^2$ and rearranging, gives

$$\partial_t U + \partial_i [(U + p^*) v^i + b^0 B^i/W] = 0. \quad (2.12)$$

This is the conservation-law form of the evolution of the energy density. The same analysis may be applied to the spatial components, $\nu = i$, of Equation (2.9), giving the conservation of momentum density:

$$\partial_t(\rho h^* W^2 v^i) + \partial_j (\rho h^* W^2 v^i v^j + p \delta^{ij} - b^i b^j) = 0. \quad (2.13)$$

To describe the evolution of the magnetic fields, we use the Maxwell equations of electromagnetism,

$$\partial_\nu {}^*F^{\mu\nu} = 0, \quad (2.14)$$

where the Maxwell dual tensor has the form ${}^*F^{\mu\nu} = u^\mu b^\nu - u^\nu b^\mu$. In the same way as for the energy-momentum equations, we may separate the spatial components to arrive at the form for the evolution of the magnetic fields. The spatial components of Equation (2.14) give the time evolution of the magnetic fields,

$$\partial_t B^i + \partial_j (v^j B^i - v^i B^j) = 0. \quad (2.15)$$

Attention must also be given to the time component of Equation (2.14), which reduces to the null divergence of the magnetic fields, i.e. no magnetic monopoles:

$$\partial_i B^i = 0. \quad (2.16)$$

Numerical errors cannot be avoided in computational modelling on account of the finite accuracy of the approximations of mathematical functions, and furthermore the finite precision of machines. Consequently, there is no guarantee that condition (2.16) will be satisfied—far from it, deviation from this ideal will produce unphysical magnetic fields and these effects will propagate into the fluid's behaviour. It is necessary, therefore, to constrain the evolution of the system in such a manner as to enforce the condition imposed by Equation (2.16). Multiple methods exist to accomplish this. A common choice is that of constrained transport or the vector potential formulation, which preserve the divergence-free constraint on the magnetic field up to machine precision (Londrillo and Del Zanna, 2004; Etienne et al., 2010).

Alternative methods introduce an additional scalar field which captures these numerical errors and drives them exponentially to zero (Dedner et al., 2002). These scalar fields (one for the magnetic field and one for the electric field, if required) are introduced into the equations of motion in the form of source terms, which we discuss later in this chapter in Section 2.8. This scheme, known as *divergence cleaning*, has been implemented in METHOD due to its reliability and simplicity.

The full equations of special relativistic iMHD in conservative form are,

$$\partial_t \begin{pmatrix} D \\ S^j \\ \tau \\ B^k \\ \phi \end{pmatrix} + \partial_i \begin{pmatrix} Dv^i \\ S^j v^i + p^* \delta^{ij} - b^j B^i / W \\ \tau v^i + p^* v^i - b^0 B^i / W \\ v^i B^k - v^k B^i + \delta^{ik} \phi \\ B^i \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ -\kappa \phi \end{pmatrix}, \quad (2.17)$$

where we relate the conserved and primitive quantities via

$$D = \rho W, \quad (2.18)$$

$$S^j = \rho h^* W^2 v^j - b^0 b^j, \quad (2.19)$$

$$\tau = \rho h^* W^2 - p^* - (b^0)^2 - D, \quad (2.20)$$

and note the relations,

$$W = 1/\sqrt{1 - v_i v^i} \quad (2.21)$$

$$p^* = p + b^2/2 \quad (2.22)$$

$$h^* = h + b^2/\rho. \quad (2.23)$$

As in system (2.1), the first conserved quantity, D , corresponds to the rest-mass-energy density of some moving fluid in some volume with velocity v^i and the second, S^j , is the relativistic momentum density in the j^{th} direction. τ is a kinetic energy density-like term, constructed from the total energy density and rest mass density, $\tau = U - D$, which is a more convenient form to evolve than the energy density for systems in which the internal energy may be small with respect to the rest-mass density (Dionysopoulou, 2016). Next, we have the equations for the evolution of the magnetic field 3-vector and for the unphysical field, ϕ , which maintains the divergence-free constraint $\nabla \cdot \mathbf{B} = 0$ (see Section 2.8).

The final difference to note in the progression to the relativistic MHD from those of a Newtonian fluid is the necessity of a scheme to transform the conserved quantities into the primitive quantities, namely (ρ, v^i, p) . We discuss this further in Section 2.7.

2.2.3 Resistive magnetohydrodynamics

The previous section outlined the form of the equations that describe a perfectly conducting magnetofluid, i.e. zero electrical resistivity. Whilst this formulation has had great success and been used extensively in astrophysical simulations, there are situations in which it may not accurately describe all behaviour. To this end, there has been work, albeit somewhat limited, into the extension of MHD to non-ideal models—some of which have been mentioned in Section 1.3.3.2. In the following section, we present the equations of special relativistic rMHD, in which we introduce a phenomenological Ohm's law that includes a finite electrical resistivity.

For more detailed discussions about arriving at the resistive MHD equations see Dionysopoulou et al. (2013); Palenzuela et al. (2009). We now present the special relativistic formulation of the resistive MHD equations for three dimensions:

$$\partial_t \begin{pmatrix} D \\ S_i \\ \tau \\ B^i \\ E^i \\ \psi \\ \phi \\ \varrho \end{pmatrix} + \partial_k \begin{pmatrix} Dv^k \\ S_i^k \\ S^k - Dv^k \\ -\epsilon^{ijk} E_j + \delta_i^k \phi \\ \epsilon^{ijk} B_j + \delta_i^k \psi \\ E^k \\ B^k \\ J^k \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ -J^i \\ \varrho - \kappa \psi \\ -\kappa \phi \\ 0 \end{pmatrix}. \quad (2.24)$$

The conserved vector $\{D, S_i, \tau, B_i, E_i, \psi, \phi, \varrho\}$ conveys the same information as in the ideal case, Equation (2.17), with the addition of the electric fields E_i , and the additional scalar field

ψ . ψ is analogous to ϕ in that it maintains the divergence constraint on the electric fields via Gauss' law, $\nabla \cdot \mathbf{E} = \varrho$, where ϱ is the charge density, see Section 2.8. One can determine the charge density, ϱ , either by its explicit evolution according to equations (2.24), or by computing $\nabla \cdot \mathbf{E} = \varrho$ —as we have discontinuous electric fields in the type of test problems we present here, we decide to evolve ϱ explicitly to avoid any issues with large gradients at shock fronts.

The conserved variables are related to the primitive quantities via

$$\begin{pmatrix} D \\ S_i \\ \tau \\ J_i \\ S_{ij} \end{pmatrix} = \begin{pmatrix} \rho W \\ \rho h W^2 v_i + \epsilon_{ijk} E^j B^k \\ \rho h W^2 - p + \frac{1}{2}(E^2 + B^2) - \rho W \\ \varrho v_i + W \sigma [E_i + \epsilon_{ijk} v^j B^k - (v_k E^k) v_i] \\ \rho h W^2 v_i v_j + [p + \frac{1}{2}(E^2 + B^2)] \delta_{ij} - E_i E_j - B_i B_j \end{pmatrix}, \quad (2.25)$$

where we note the inclusion of a term representing the conductivity, σ —which may depend upon local fluid quantities such as density, e.g. the functional conductivity in Section 7.2.2—in the definition of the 3-charge current, J_i .

Whilst the form of the equations is close to those of ideal MHD, the resistive equations differ numerically. The evolution of the electric fields is dependant upon the generalized form of Ohm's law—that is, the relationship between electrical current and the electric and magnetic fields. This relation is given in Equation (2.25) for the form of the charge current density, J_i . The generalised form for this has been taken from [Dionysopoulou et al. \(2013\)](#), and we note that a more realistic prescription would be derived from the micro-physics of the system we are investigating ([Andersson, 2012](#)). However, for our purposes, this simple form will reproduce much of the behaviour that we are interested in, and is common for resistive, astrophysical models [Palenzuela et al. \(2009\)](#); [Amano \(2016\)](#); [Qian et al. \(2016\)](#).

The form we have chosen for J_i is proportional to the conductivity, σ . As a result, for regions of high conductivity the RHS term of the electric fields will be large and its evolution equation will become stiff, potentially posing numerical difficulties. For this reason in the ideal limit, as $\sigma \rightarrow \infty$, it becomes necessary to use integrators which maintain improved stability in stiff regions. We discuss these methods in Section 2.6.2, but it suffices to say these methods come with a large computational cost.

2.2.4 Two-fluid electro-magnetohydrodynamics

In Section 1.3.3.2 we described potential extensions to the single-fluid MHD framework and how these systems may one day be utilised in merger simulations. Whilst there has been work in describing a system of four interdependent fluids for the purposes of NS simulations ([Andersson et al., 2017a,b,c](#))—namely electrons, positrons/protons, neutrons and entropy—with our current understanding, the equations are not in a form that is suitable for numerics. As a compromise, and to understand the difficulties that may arise in the more complex four-fluid model, we present here a two-fluid description of an electron-positron/proton plasma, adapted from [Amano \(2016\)](#) with different notation.

As we are considering a system of two fluids, the conserved variables must be some kind of sum over the individual species, therefore, we label the species with a subscript s , where $s \in \{e, p\}$ labels negative and positive charge carriers. We shall also assume, unless otherwise stated, that the sum over these species is implied such that, for example, the relativistic mass-energy density

$$D = \rho_s W_s = \sum_s \rho_s W_s = \rho_e W_e + \rho_p W_p, \quad (2.26)$$

where ρ_s is the rest-mass density of species s , and W_s is the Lorentz factor of species s .

With this in mind, we write the general form of the hyperbolic conservation law plus a source term as

$$\partial_t \mathbf{U} + \partial_j \mathbf{F}^j(\mathbf{U}) = \mathbf{\Psi}(\mathbf{U}). \quad (2.27)$$

The conserved vector for the two fluid model is then

$$\mathbf{U} = \begin{pmatrix} D \\ S_i \\ \tau \\ \overline{D} \\ \overline{S}_i \\ \overline{\tau} \\ B^i \\ E^i \\ \psi \\ \phi \end{pmatrix} = \begin{pmatrix} \rho_s W_s \\ \rho_s h_s W_s^2 v_{s,i} + \epsilon_{ijk} E^j B^k \\ (\rho_s h_s W_s^2 - p_s) + (E^2 + B^2)/2 - \rho_s W_s \\ \mu_s \rho_s W_s \\ \mu_s \rho_s h_s W_s^2 v_{s,i} \\ \mu_s (\rho_s h_s W_s^2 - p_s - \rho_s W_s) \\ - \\ - \\ - \\ - \end{pmatrix}, \quad (2.28)$$

where the following definitions apply: the charge-mass ratio is $\mu = q/m$, where the charge and mass of a constituent particle species are q and m respectively, p is the hydrodynamic pressure, $h = 1 + e + p/\rho$ is the specific enthalpy, S_i is the total relativistic momentum in the i^{th} direction, τ is a kinetic energy-like term, and ψ and ϕ are scalar fields for Maxwell's divergence constraints. Unbarred forms of the standard relativistic hydrodynamic variables $\{D, S_i, \tau\}$ correspond to the sums of these quantities over the two fluid species, whereas the barred counterparts $\{\overline{D}, \overline{S}_i, \overline{\tau}\}$ are summed according to a weighting factor given by the charge-mass ratio, i.e. $\overline{D} = \mu_e \rho_e W_e + \mu_p \rho_p W_p$, minus any contribution from the electromagnetic fields. Due to the opposing charges of the two species, we can consider this the difference in their respective rest mass densities, for example.

The flux vector is defined by

$$\mathbf{F}^j(\mathbf{U}) = \begin{pmatrix} \rho_s W_s v_s^j \\ S_{s,i}^j \\ S^j - D_s v_s^j \\ \bar{D}_s v_s^j \\ \bar{S}_s^j v_{s,i} + \delta_i^j p_s \mu_s \\ (\bar{\tau}_s + \mu_s p_s) v_s^j \\ \epsilon^{ijk} E_k + \delta_i^j \phi \\ -\epsilon^{ijk} B_k + \delta_i^j \psi \\ E^j \\ B^j \end{pmatrix}, \quad (2.29)$$

and the source vector is

$$\mathbf{\Psi}(\mathbf{U}) = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \omega_p^2 [W E_i + \epsilon_{ijk} u^j B^k - \eta (J_i - \varrho u_i)] \\ \omega_p^2 [u_i E^i - \eta (\varrho - \varrho_0 W)] \\ 0 \\ -J^i \\ \varrho - \kappa \psi \\ -\kappa \phi \end{pmatrix}, \quad (2.30)$$

where we have defined the following variables:

$$\begin{pmatrix} S_{s,i}^j \\ \varrho \\ J_i \\ W \\ u^i \\ \varrho_0 \\ \omega_p^2 \end{pmatrix} = \begin{pmatrix} \rho_s h_s W_s^2 v_{s,i} v^j + \delta_i^j p_s - [E_i E^j + B_i B^j] + \delta_i^j [E^2 + B^2]/2 \\ \mu_s \rho_s W_s \\ \mu_s \rho_s W_s v_{s,i} \\ \mu_s^2 \rho_s W_s / \omega_p^2 \\ \mu_s^2 \rho_s W_s v_s^i / \omega_p^2 \\ W \varrho - J_i u^i \\ \mu_s^2 \rho_s \end{pmatrix}, \quad (2.31)$$

in which $\eta = 1/\sigma$ is the resistivity, κ is a constant term due to the divergence cleaning (Section 2.8), ϱ is the fluid charge density with ϱ_0 the rest charge density, and ω_p is the plasma frequency. Note that in all the equations above, we are assuming the sum over species.

As with the previous MHD models, we have chosen to use the Lagrangian multiplier method (Dedner et al., 2002; Palenzuela et al., 2009) to constrain the divergence conditions set by Maxwell's equations. Amano posits that the divergence cleaning method may be unstable for the two fluid model (Amano, 2016), and instead promotes the use of constrained transport as the method of choice. As we shall see in Chapter 3, we found no issues using the divergence cleaning method.

In a single fluid description of a plasma, one implicitly confines the constituent charges to be perfectly coupled—there is no relative velocity between protons and electrons, for example. With a multi-fluid description, however, this condition is relaxed. In deriving the equations of motions of the constituent species, one explicitly factors in a difference in their observed velocities. This process allows us to capture the effects of charge separation as a result of the decoupling of the species’ motion. This two-fluid effect occurs on scales that are smaller than the plasma skin depth (Saupe and Sovinec, 2016), $\lambda_p = c/\omega_p$, where c and $\omega_p = \sqrt{-\mu_e\mu_p\rho}$ are respectively the speed of light and the plasma frequency. We should therefore expect to see this effect when the resolution of the simulation $\Delta x \lesssim (-\mu_e\mu_p\rho)^{-1/2}$.

In addition to the equations of motion, care must be taken when constructing initial data for the two-fluid model—there are constraints that must be accounted for when generalising resistive MHD problems to incorporate into multi-fluid set ups. First, a basic assumption of MHD is that systems contain no net charge, $\int_{\Omega} \varrho(\mathbf{x}) d^3x = 0$. This relation must hence be satisfied by the two species we are considering. As it is not essential that both species have the same charge-mass ratio, this results in the following condition for the rest-mass densities:

$$\rho_e = \frac{m_e}{m_p + m_e} \rho, \quad (2.32)$$

$$\rho_p = \frac{m_p}{m_p + m_e} \rho. \quad (2.33)$$

We will also assume that the species are initially in thermal equilibrium, and so the hydrodynamic pressures must be equivalent— $p_e/2 = p = p_p/2$. As is also the case for single fluid MHD, we must ensure that the magnetic fields are divergence free, and the electric fields satisfy Gauss’ law. If the system has a non-uniform magnetic field, it is also necessary to ensure that the current density satisfies Ampere’s law. As a result, the velocities of the fluid species may not necessarily be identical to each other on account of differing charge-mass ratios.

Finally, we note how the necessary tolerance of the primitive recovery procedure for the resistive, two-fluid model restricts us to certain design choices for the variables types. Due to the high precision required to minimise error growth, all variables when coded must be double precision floating point numbers. We discuss why this is the case and its impact in Section 2.7.

2.3 Numerical relativity

A large number of astrophysical systems experience strong gravitational fields and contain relativistic matter. For some of these systems, a Newtonian or post-Newtonian approach to gravitational interactions may be sufficiently accurate. For other systems, however, including neutron star or black hole mergers and the collapse of objects to BHs (Baumgarte and Shapiro, 2003), gravitational interactions are strong enough or dynamic enough to warrant a full description of general relativity.

Much of the work on numerical solutions to Einstein’s field equations dates back to black-hole merger simulations, and the field is now sufficiently mature that open source projects such as the Einstein Toolkit (Löffler et al., 2012) (see Section 2.3.5) are capable of reliably evolving complex

spacetimes. Of course, there are still systems which pose a challenge for current numerical approaches to GR, such as extreme mass-ratio inspirals, but we do not touch upon them here.

As we will be performing simulations of neutron star mergers, and much of our discussion is based on these kinds of simulations, we will now outline some of the basic principles behind numerical relativity. For more comprehensive reviews, the reader is directed towards the origin of these ideas in [Arnowitt et al. \(2008\)](#), to [Baumgarte and Shapiro \(2002\)](#) from which much of this section is based, and to [Alcubierre \(2008\)](#) and [Baumgarte and Shapiro \(2010\)](#). Throughout this section we will be using geometric units where $G = c = 1$.

2.3.1 Arnowitt-Deser-Misner decomposition

Einstein's theory of general relativity describes the mechanism of gravitation using the notion of spacetime curvature. Spacetime is the unification of space and time that is central to GR, however when dealing with GR numerically it is far more convenient to consider a "3+1" formulation, in which one explicitly defines a time-coordinate and three spatial coordinates.

The decision we must make is the way we slice the spacetime. The four-dimensional spacetime is foliated into a family of non-intersecting, spacelike hypersurfaces, Σ , with some constant coordinate time, $t = \text{const}$. We define the time-like unit normal to the hypersurface as $n_\mu = -\alpha \nabla_\mu t$, which points in the direction of increasing t and where the lapse is strictly positive, $\alpha > 0$. The induced spatial metric on the hypersurface is then defined as

$$\gamma_{\mu\nu} = g_{\mu\nu} + n_\mu n_\nu, \quad (2.34)$$

where $g_{\mu\nu}$ is the spacetime metric. The time vector then describes how the coordinates change from one hypersurface to the next,

$$t^\mu = \alpha n^\mu + \beta^\mu, \quad (2.35)$$

where the shift vector, β , is purely spatial, so $n_\mu \beta^\mu = 0$. The shift vector corresponds to the spatial change in the coordinates with respect to n^μ , and the lapse determines the change in proper time between the slices. Solving this for the components of the normal vector gives,

$$n^\mu = \alpha^{-1}(1, -\beta^i), \quad (2.36)$$

and as the normal is time-like, $n_\mu n^\mu = -1$, one can easily show

$$n_\mu = (-\alpha, 0, 0, 0). \quad (2.37)$$

With these definitions, the line element is as follows:

$$ds^2 = -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt). \quad (2.38)$$

So far we have described a specific choice of foliation, and one can derive how the resultant slices are embedded in the four dimensional spacetime. This includes terms such as the three dimensional connection, ${}^{(3)}\Gamma_{jk}^i$, and the extrinsic curvature, K_{ij} . At this point, one must use

Einstein's field equations to determine how the spacetime evolves,

$$G_{\mu\nu} = 8\pi T_{\mu\nu}. \quad (2.39)$$

A series of projections of these equations onto Σ and n^μ result in the ADM equations. These are the basis for a number of formalisms, and relate the dynamics of the foliated curvature quantities to the stress energy tensor, $T^{\mu\nu}$.

The projections result in the following set of constraint equations:

$$R + K^2 - K_{ij}K^{ij} = 16\pi\rho_E, \quad (2.40a)$$

$$D_j K_i^j - D_i K = 8\pi S_i, \quad (2.40b)$$

and the following evolution equations:

$$\partial_t \gamma_{ij} = -2\alpha K_{ij} + D_i \beta_j + D_j \beta_i, \quad (2.41a)$$

$$\begin{aligned} \partial_t K_{ij} = & -D_i D_j \alpha + \alpha(R_{ij} - 2K_{ik}K_j^k + KK_{ij}) \\ & - \alpha 8\pi \left[S_{ij} - \frac{\gamma_{ij}}{2}(S - \rho_E) \right] \\ & + \beta^k D_k K_{ij} + K_{ik} D_j \beta^k + K_{kj} D_i \beta^k, \end{aligned} \quad (2.41b)$$

where K is the trace of the extrinsic curvature, R is the trace of the three dimensional Ricci tensor, $\rho_E \equiv n^\mu n^\nu T_{\mu\nu} = \tau + \rho W$ is the Eulerian energy density, D_i is the three dimensional covariant derivative, S_i and S_{ij} are as before and $S \equiv S_i^i$.

These represent the full set of ADM evolution equations, and have been used successfully for dynamic simulations of black holes, for example. Nowadays, however, alternative formulations are used for evolutions that have improved stability properties and are more suitable, particularly, for long term simulations ([Shibata and Nakamura, 1995](#); [Baumgarte and Shapiro, 1998](#); [Alic et al., 2012](#)).

2.3.2 Gauge freedom

Note how there are no expressions in Equations (2.40) or (2.41) for the lapse or shift functions—in fact, this is the case for all 3+1 formulations. This reflects the coordinate freedom of GR, and manifests itself in a choice for α and β^i . There are a number of options one can choose for these, but for the sake of brevity, we will discuss a common choice that we use in later simulations of NS mergers.

Gauge choices can affect the numerical stability of a simulation, and a good choice should typically allow for longer evolutions. [Rezzolla and Zanotti \(2013\)](#) state three key things that should result from a good choice of gauge: singularities should be avoided; coordinate distortions as a result of spatial curvatures should be counteracted; and the choice should lead to a numerically inexpensive prescription of the lapse and shift functions.

The most popular slicing condition—i.e. choice for α —is the 1 + log slicing condition. In this, the choice for α is reduced to solving the following hyperbolic equation (Bona et al., 1995),

$$\partial_t \alpha - \beta^k \partial_k \alpha = -f(\alpha) \alpha^2 (K - K_0), \quad (2.42)$$

where $f(\alpha) = 2/\alpha^{-1}$, and $K_0 = K(t = 0)$. With this method, the growth in the lapse near a singularity is suppressed, thus achieving the singularity-avoiding condition.

Similarly for the shift condition—choice of β^i —each specific choice can have varying numerical properties. Typically, what is tried is to keep the $\tilde{\Gamma}^i$ s approximately constant (Alcubierre and Bru, 2001). Here, $\tilde{\Gamma}^i = \tilde{\gamma}^{ij} \tilde{\gamma}^{kl} \partial_t \gamma_{jk}$, where the conformal three-metric is related to the physical three-metric via $\tilde{\gamma}_{ij} = \gamma^{-1/3} \gamma_{ij}$. This behaviour is true of the *Gamma-driver shift condition* (Alcubierre and Bru, 2001), which in its more modern form is

$$\partial_t \beta^i - \beta^j \partial_j \beta^i = \frac{3}{4} \Phi^i, \quad (2.43a)$$

$$\partial_t \Phi^i - \beta^j \partial_j \Phi^i = \partial_t \tilde{\Gamma}^i - \beta^j \partial_j \tilde{\Gamma}^i - \eta \Phi^i, \quad (2.43b)$$

where Φ is an auxiliary variable and η is some damping term.

These two conditions are generally sufficient to evolve spacetimes numerically. It is for this reason that in Chapter 6, all general relativistic simulations are evolved using these methods.

2.3.3 Time derivative of spatial metric determinant

As we will need $\partial_t \sqrt{\gamma}$ later on, it is useful to address how to compute it now. One way would be to compute it via finite differencing, but this will require additional computation and additional memory. Another (and less error prone) way is to use the relation, $\nabla_a \sqrt{-g} = 0$:

$$0 = \nabla_a \sqrt{-g} = \frac{1}{\sqrt{-g}} \partial_a (\sqrt{-g} \sqrt{-g}), \quad (2.44)$$

and so,

$$0 = \partial_a (\sqrt{-g} \sqrt{-g}) \quad (2.45a)$$

$$= 2\sqrt{-g} \partial_a \sqrt{-g}. \quad (2.45b)$$

Using the relation $\sqrt{-g} = \alpha \sqrt{\gamma}$, we get

$$0 = \partial_a \sqrt{-g} \quad (2.46a)$$

$$= \partial_a (\alpha \sqrt{\gamma}) \quad (2.46b)$$

$$= \alpha \partial_a \sqrt{\gamma} + \sqrt{\gamma} \partial_a \alpha \quad (2.46c)$$

which, when evaluated for $a = t$, gives the relation

$$\partial_t \sqrt{\gamma} = -\frac{\sqrt{\gamma}}{\alpha} \partial_t \alpha. \quad (2.47)$$

For many gauges in 3+1 numerical relativity, the time derivative of the lapse is computed directly via, for example, Equation (2.42) of the previous section.

2.3.4 Weak hydrodynamic coupling

We know that the motion of matter will directly affect the motion of the spacetime in a simulation, however for the purpose of a later analysis in Section 6.1 we will make an assumption that the matter couples only weakly. The effect of this is as follows: if we take some curvature parameter that is used for the spacetime evolution, say the trace of the extrinsic curvature K , and some vector of matter variables, $\mathbf{w} = (\rho, v_i, p)$, then the former is constant with respect to changes in the latter. That is,

$$\frac{\partial K}{\partial \mathbf{w}} = 0, \quad (2.48)$$

but we will not assume that they are constant in space or time, so $\partial_t \sqrt{\gamma} \neq 0$, for example. This approximation extends to all the parameters in the spacetime evolution.

2.3.5 The Einstein Toolkit and Whisky

The previous sections touch upon what is required for the numerical evolution of spacetime, and yet the current state-of-the-art goes quite a lot further. Fortunately, however, it is not necessary to develop an implementation of these ideas from scratch, and there are a number of ready-made and well-tested packages that provide what we need. One example of this, and what we will be using for later simulations of binary mergers, is the Einstein Toolkit (ET) (Löffler et al., 2012).

Based upon the Cactus computational framework¹, ET provides a set of modules, or thorns, that implement various features from physics models (magnetohydrodynamics, equation of state, etcetera) to boundary conditions, data input/output, and spacetime evolution. These tools are open source, and allow research groups to conveniently study astrophysical systems such as binary black hole or neutron star mergers, isolated neutron stars, and even solve elliptic equations such as Poisson's equation with general RHSs.

This community driven project removes much of the hard work regarding large scale astrophysical simulations. Many thorns may be used as a black box, abstracting away the difficulties of solving complex sets of equations, and allowing the user to focus on a single aspect of a model, or analysis. In addition, through the use of the **Carpenter** driver for adaptive mesh refinement (Schnetter et al., 2004), it is effortless to run very high resolution simulations across a large number of HPC nodes efficiently, something that is essential when modelling behaviour at disparate scales.

¹<http://cactuscode.org/>

2.4 Hyperbolic conservation laws and finite volume methods

In Section 1.3.1 we mentioned the notion of finite volume methods for grid-based, computational fluid dynamics. As we have seen in earlier sections, the system of fluid equations that are relevant to merger simulations can be cast into conservation form. Here, we highlight this class of equations and the foundational work that goes into solving them. These ideas are covered in great depth in [Laney \(1998\)](#) and [LeVeque \(2002\)](#).

2.4.1 Conservation laws

The form of the equations that are used in merger simulations are given in their balance form, namely

$$\partial_t \mathbf{U} + \partial_i \mathbf{F}^i(\mathbf{U}) = \mathbf{\Psi}(\mathbf{U}), \quad (2.49)$$

in which the conserved vector is denoted by \mathbf{U} and is of degree N , and \mathbf{F} and $\mathbf{\Psi}$ are functions of the conserved vector for the flux and source terms, respectively. Presented in Equation (2.49) is the so-called strong form of a hyperbolic conservation law inclusive of a source term, a form which, if not already, will soon become familiar. Without loss of generality, we may exclude the source term and consider only the contribution of the fluid flux to the change in the conserved quantities. Then, limiting attention to only the x -direction, we may rewrite Equation (2.49) in quasi-linear form:

$$\partial_t \mathbf{U} + \underline{\underline{A}}(\mathbf{U}) \partial_x \mathbf{U} = 0. \quad (2.50)$$

The Jacobian of the system is given by $\underline{\underline{A}}(\mathbf{U}) = \partial \mathbf{F}(\mathbf{U}) / \partial \mathbf{U}$, and has corresponding eigenvalues denoted by λ_i , $i \in \{1, 2, \dots, N\}$. The system (2.50) is defined to be strictly hyperbolic if all N eigenvalues are real and distinct. In the context of fluid mechanics, the values prescribed by λ_i correspond to the velocities of perturbations to the i^{th} component of \mathbf{U} , thus the condition that the system is strictly hyperbolic ensures that perturbations travel at physical velocities. Furthermore, if we reorder the eigenvalues in ascending magnitude such that, now,

$$|\lambda_1| < |\lambda_2| < \dots < |\lambda_N|, \quad (2.51)$$

we can characterize the stiffness of the system given by (2.49). If the system is such that $|\lambda_N| \gg |\lambda_1|$, the system is said to be stiff, and this may pose some numerical difficulties. A large range in speeds of the perturbations corresponds to dynamics occurring on largely different length- and timescales, which makes computationally efficient means of capturing the behaviour on these different scales difficult. We discuss this in more depth in Section 2.6.

2.4.2 The Riemann problem

The Riemann problem is a class of problems that involves two piecewise constant states separated by a sharp interface. Physically, this may represent a shock tube in which a gas is separated via some partition into two regions of differing pressures, densities, etcetera. The problem poses the question, ‘How does this system evolve?’

Riemann problems occur naturally in computational fluid dynamics, as the domain of computation is subdivided into numerous computational cells. These cells represent some average quantity (see 2.4.3) across the cell which will generally differ from the cell’s neighbours. Therefore, at the interface of each cell there is a sharp change in some quantity, known as a jump discontinuity.

The differing pressures, for example, on either side of a jump discontinuity produce a shock that propagates through the system, the speed of which is given by the Rankine-Hugoniot condition. Consider such a discontinuity in a one-dimensional system located at some position $X(t)$ which propagates with time. Due to the discontinuous nature of our set up, a strong form conservation equation of the type (2.49) does not make sense. Instead, we consider the weak form which we determine via integration over the domain $[x_L, x_R]$,

$$\frac{d}{dt} \int_{x_L}^{x_R} \mathbf{U}(t, x) dx = \mathbf{F}(\mathbf{U}(t, x_L)) - \mathbf{F}(\mathbf{U}(t, x_R)), \quad (2.52)$$

where, again, we limit to one dimension and assume zero source. We can now inspect what happens either side of a shock, located such that $(x_L, x_R) = (X(t) - \epsilon, X(t) + \epsilon)$. This gives the relation

$$\frac{d}{dt} \int_{X(t)-\epsilon}^{X(t)} \mathbf{U}(t, x) dx + \frac{d}{dt} \int_{X(t)}^{X(t)+\epsilon} \mathbf{U}(t, x) dx = \mathbf{F}(\mathbf{U}(t, X(t) - \epsilon)) - \mathbf{F}(\mathbf{U}(t, X(t) + \epsilon)), \quad (2.53)$$

which, after applying the derivative and taking the limit $\epsilon \rightarrow 0$ reduces to

$$\frac{dX(t)}{dt} [\mathbf{U}] = [\mathbf{F}], \quad (2.54)$$

where we have use the notation $[P]$ to represent the difference in P across the shock. Here, $S = dX(t)/dt$ is the velocity of the propagating shock. One can use this result when considering the properties on either side of a moving shock in a coordinate frame that is co-moving with it. Taking the Euler equations (2.1) in the x -direction, for example, we can categorise the states on either side of the shock solely in terms of the ratio of the pressures, p_R/p_L , a result known as the Rankine-Hugoniot relations (Anderson, 1990).

2.4.3 Finite volume approach

We now address a method for the solution of the initial value problem for conservation laws. More specifically, we are concerned with systems in which we have some initial state $\mathbf{U}_0(\mathbf{x}) = \mathbf{U}(t = 0, \mathbf{x})$ which is governed by some set of conservation laws, such as those given by (2.49), and we are interested in some solution at a later time, t .

Initially we must choose the form of our domain, Ω , which we define here to be one-dimensional. Suppose we separate Ω in to discrete cells $I_i = [x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}]$, where $\Delta x_i = x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}$. Moreover, suppose we discretize time such that $t_n = n\Delta t$. We can now integrate some arbitrary cell, I_i , over space and time, that is, over the volume $\mathcal{V} = [t^n, t^{n+1}] \times [x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}]$. Doing this, and defining the spatial and temporal averages of the conserved and flux vector respectively,

$$\mathbf{Q}_i^n = \frac{1}{\Delta x_i} \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \mathbf{U}(t_n, \mathbf{x}) \, dx, \quad (2.55)$$

$$\mathbf{f}_{i+\frac{1}{2}}^n = \frac{1}{\Delta t_n} \int_{t_n}^{t_{n+1}} \mathbf{F}(t, \mathbf{x}_{i+\frac{1}{2}}) \, dt, \quad (2.56)$$

$$(2.57)$$

we arrive at the *numerical conservation law* (Laney, 1998),

$$\mathbf{Q}_i^{n+1} = \mathbf{Q}_i^n - \frac{\Delta t}{\Delta x} (\mathbf{f}_{i+\frac{1}{2}}^n - \mathbf{f}_{i-\frac{1}{2}}^n). \quad (2.58)$$

We may interpret this in terms of changes in cell averages—the total change in the quantities \mathbf{U} in some cell volume in one timestep is a result of the net flux of \mathbf{U} through the faces of that cell. We can see that by construction this scheme will preserve the conservative nature of the equations of motion, as the flux exiting one cell boundary is identical to the flux entering the neighbouring cell. This, of course, may be extended to include a source term, such that the changes in the conserved quantities are due to the flux of the quantities and some source/sink—for example due to dissipative effects, combustion, etcetera.

From the form of Equation (2.58), we can easily see how to solve any initial value problem posed. Provided we have some initial state that spans the entire domain, i.e. we have $\mathbf{Q}_i^0 \forall i$, and any relevant boundary conditions, we take the flux difference at the cell faces to generate \mathbf{Q}_i^1 and repeat.

The problem now is in the construction of the fluxes. From the current state of the system we have \mathbf{f}_i^n , the temporally averaged flux at the cell centre, $x = x_i$, but what we require is the difference in the face-centred values of the flux, $\mathbf{f}_{i+\frac{1}{2}}^n$, due to the current state of neighbouring cells. We can, therefore, define some numerical flux function, \mathcal{F} , such that, for example, $\mathbf{f}_{i+\frac{1}{2}}^n = \mathcal{F}(\mathbf{Q}_{i+1}^n, \mathbf{Q}_i^n)$. Now we only need to find the form of \mathcal{F} which is a function of the known values of \mathbf{Q}_i . It is worth noting that in general there is no constraint on the number of arguments from neighbouring cells, and as a rule of thumb, the larger the stencil (i.e. the more neighbouring cells considered in the construction of the numerical flux function) the higher the possible order of accuracy, and the larger the potential of stability issues. In the next section we outline some methods generally adopted in fluid codes, and detail the scheme used in **METHOD**.

2.5 Flux methods

As alluded to previously, there are a host of different approximate Riemann solvers, each with unique benefits and drawbacks, that aim to approximately solve for the flux at cell faces. Approximate solvers are ubiquitous when dealing with finite volume methods, as they offer comparable

results to exact solvers with significantly reduced computation. There can, however, be large differences in performance and accuracy between the different solvers.

One of the earliest types of approximate solvers was developed by [Roe \(1981\)](#), and its early success deemed it a pillar of fluid codes for some time. The analysis of the travelling waves is performed via an eigen-decomposition of the Jacobian. For large systems, such as those of multi-fluid MHD, performing the eigenanalysis is a daunting task—this limits the application of Roe-type methods to simple systems such as Euler’s equations (2.1).

Here, we outline the simple Lax-Friedrichs method, discuss alternative methods that are commonly used in relativistic codes, and describe the scheme applied in **METHOD**. For the purpose of this discussion we will limit ourselves to a system where $N = 1$, but the methods conferred generalise to systems of any size. Consider a one-dimensional system in which we approximate a cell averaged conserved quantity, Q_i , with corresponding flux vector $F_i = F(Q_i)$. One possible numerical flux at the boundary may be the arithmetic average of the neighbours’ cell-centred flux vectors,

$$f_{i+1/2} = \frac{F(Q_i) + F(Q_{i+1})}{2}, \quad (2.59)$$

giving

$$Q_i^{n+1} = Q_i^n - \frac{\Delta t}{2\Delta x} [F(Q_i) + F(Q_{i+1})], \quad (2.60)$$

but it is known that for most hyperbolic systems this is numerically unstable. A better approximation that is conditionally stable for hyperbolic problems can be made, which results in

$$Q_i^{n+1} = \frac{Q_{i+1}^n + Q_{i-1}^n}{2} - \frac{\Delta t}{2\Delta x} [F(Q_i) + F(Q_{i+1})]. \quad (2.61)$$

It is easy to show that the corresponding form of the numerical flux is given by

$$\mathcal{F}_i = \frac{1}{2} \left[F(Q_{i+1}) + F(Q_i) - \frac{\Delta x}{\Delta t} (Q_{i+1}^n - Q_i^n) \right]. \quad (2.62)$$

Equation (2.62) is the so-called Lax-Friedrichs flux approximation, and its stability originates from the inclusion of an artificial viscosity that scales with the size of the jump in the conserved variables. As a result, simulations which use this flux approximation can appear to smooth the shock front, diffusively damping discontinuities.

Additional forms for the flux approximation may come in the form of the HLL+ methods ([Harten et al., 1983](#)). This class of schemes is named after Harten, Lax and van Leer of whom were the proponents of the first instance of this type of method, then improved upon by [Toro et al. \(1994\)](#). The original HLL scheme assumes the flux to be the result of a set of travelling waves separating discontinuous states. Where as the HLL method failed to accurately capture contact discontinuities, the HLLC scheme (C for ‘contact’) is more accurate at the expense of considering additional waves in the system. Methods for MHD-HLLC can be found in [Honkila and Janhunen \(2007\)](#). The code used later in this thesis for merger simulations, **Whisky**, uses a method based on the HLLE scheme ([Einfeldt, 1988](#)).

[Sangeeth and Mandal \(2018\)](#) provides a clear and concise overview of the various HLL+ schemes and provides a method for maintaining stability around shocks—a common problem for these schemes. However, due to HLL+ schemes becoming potentially unstable around shocks in multi-dimensional simulations, and the comparatively expensive Roe methods requiring complex eigenanalysis for large systems, we implement in `METHOD` an alternative scheme developed by [Shu \(1997\)](#) of the name *flux vector splitting* (FVS). FVS is relatively inexpensive and simple to implement, robust, captures shocks well, is easily expandable to higher orders, and is constructed to minimise numerical oscillations at shocks when applied with weighted, essentially non-oscillatory (WENO) reconstructions—for more details on WENO reconstructions see [Shu \(1997\)](#).

FVS is conceptually different to determining the flux at a cell face when compared to HLL or Roe-type methods—a summary of which can be found in [Kong \(2011\)](#). Instead of attempting to (approximately) solve the Riemann problem, FVS reduces the total flux to left- and right-going components. The separated components are then reconstructed across a section of the domain, and the difference between upwind and downwind contributions found to approximate the net flux for a given cell.

It follows from Equation (2.58) that we can decompose the flux at a cell boundary to up- and downwind components, namely $f_i = f_i^+ + f_i^-$. To do this we use Lax-Friedrichs splitting,

$$f_i^\pm = \frac{1}{2}(f_i \pm \alpha Q_i), \quad (2.63)$$

where $\alpha = \max_p |\lambda_p|$ and λ_p is the p^{th} eigenvalue of the Jacobian, A , of the system. Note that $\partial_{Q_i} f_i^+ \geq 0$ and $\partial_{Q_i} f_i^- \leq 0$. In what follows, if the order of the WENO scheme used requires a stencil spanning s compute cells (compute cells comprise all cells in the domain, including the unphysical ghost cells), α is the largest magnitude of velocity of the waves in these s cells. Alternatively, for simplicity, we may define $\alpha = 1$, effectively increasing the dissipation of the method. The effect of this is limited in MHD simulations, however, as electromagnetic waves propagate at $c = 1$.

To determine the net flux of the conserved quantities at each cell face, we take the difference of the upwind and downwind fluxes at the cell boundary. In other words, first determine the upwind reconstruction $f_{i+1/2}^+$ and downwind reconstruction $f_{i+1/2}^-$, and approximate the flux through the face of the i^{th} cell as the difference of these values, recalling that the upwind and downwind contributions will have opposing signs due to opposing directions. We therefore need to calculate $f_{i+1/2} = f_{i+1/2}^+ + f_{i+1/2}^-$. The total net flux through some cell i is then computed as the difference between the flux at the cell faces.

2.6 Time-integration schemes

Modern techniques of solving systems of conservation laws are focussed on maintaining high order accuracy (typically $2^{\text{nd}} - 4^{\text{th}}$) for smooth solutions whilst maintaining stability for discontinuous data. Such schemes are known as high resolution shock-capturing, and to minimise the onset of spurious numerical oscillations at discontinuous shocks, schemes are chosen that are total variation diminishing (TVD). A method is said to be TVD ([Gottlieb and Shu, 1998](#)) if the

following inequality holds:

$$TV(Q^{n+1}) \leq TV(Q^n) \quad (2.64)$$

where the total variation of the solution is defined as

$$TV(Q^n) = \sum_j \left| Q_{j+1}(t^n) - Q_j(t^n) \right|. \quad (2.65)$$

Methods in which the norms satisfy $\|Q^{n+1}\| \leq \|Q^n\|$ are also known as strong stability preserving (SSP).

2.6.1 Explicit schemes

There are a large number of proven SSP-RK integrators that are generally used for fluid evolution codes (Gottlieb and Shu, 1998), with some recursively calculated up to arbitrary order (Gottlieb et al., 2009). Given the ODE

$$\frac{dQ}{dt} = \mathcal{L}(Q), \quad (2.66)$$

where $\mathcal{L}(Q)$ is the numerical approximation of the flux and source terms, at least for the systems we are interested in, the general form of an explicit RK scheme is given by

$$Q^{(i)} = \sum_{k=0}^{i-1} \left[\alpha_{ik} Q^{(k)} + \Delta t \beta_{ik} \mathcal{L}(Q^{(k)}) \right], \quad i = 1, \dots, m \quad (2.67)$$

$$Q^{(0)} = Q^n, \quad Q^{(m)} = Q^{n+1}, \quad (2.68)$$

where m is the number of stages in the RK scheme and $\alpha_{ik}, \beta_{ik} \geq 0$ are positive constants. Schemes of this type reduce to complex combinations of different sized Euler steps, and we present here the second order scheme used in `METHOD`:

$$Q^{(1)} = Q^n + \Delta t \mathcal{L}(Q^n) \quad (2.69)$$

$$Q^{(m)} = \frac{1}{2} [Q^n + Q^{(1)} + \Delta t \mathcal{L}(Q^{(1)})] \quad (2.70)$$

$$Q^{n+1} = Q^{(m)}. \quad (2.71)$$

For dealing with (non-stiff) sources, we combine the numerical flux approximation with the source term evaluated at the cell centre, $\mathcal{L}(Q_i) = \mathcal{F}_i + \Psi_i$, and proceed as normal. This allows for simple, high order time integration of systems with source terms, without reverting to operator splitting techniques, although operator splitting up to second order has also been implemented.

2.6.2 Semi-implicit schemes

We have previously made reference to stiff and non-stiff source terms, now we will attempt to quantify what we mean by this. There is no strict, widely accepted definition stiffness, and so we

shall make an argument that is relevant to the systems we are studying. Take a simple system of the form

$$\frac{d\mathbf{Q}}{dt} = \mathcal{L}(\mathbf{Q}), \quad (2.72)$$

which we can write as,

$$\frac{d\mathbf{Q}}{dt} = -K\mathbf{Q}. \quad (2.73)$$

We can think of the RHS of this equation as acting on some timescale governed by the factor K . In other words, we have

$$\frac{d\mathbf{Q}}{dt} = \frac{-\mathbf{Q}}{\tau}, \quad (2.74)$$

with τ as the timescale. Performing a simple, explicit Euler step on this gives

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n - \frac{\Delta t}{\tau} \mathbf{Q}^n, \quad (2.75)$$

where Δt is the size of the timestep. Rearranging this gives the final form

$$\mathbf{Q}^{n+1} = \left(1 - \frac{\Delta t}{\tau}\right)^{n+1} \mathbf{Q}^0. \quad (2.76)$$

We can see from this that if $\tau \ll \Delta t$, the solution diverges. We can infer then that the timescale the RHS acts upon must be comparable to the timestep of the simulation (or larger), $\tau \gtrsim \Delta t$. As a result, if the timescale τ becomes small, we must limit the size of the timestep. For systems in which the timescale becomes very small, for example in highly conducting simulations of resistive MHD, the timestep would have to be so small that simulations would take far too long to execute.

Conversely, if we instead use the implicit Euler step, i.e. we use

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n - \frac{\Delta t}{\tau} \mathbf{Q}^{n+1}, \quad (2.77)$$

then the final form of the solution becomes

$$\mathbf{Q}^{n+1} = \frac{\mathbf{Q}^0}{\left(1 + \frac{\Delta t}{\tau}\right)^{n+1}}. \quad (2.78)$$

Clearly, we can see that the size of the solution is bounded when $\tau \ll \Delta t$. In fact, the implicit Euler method is unconditionally stable, meaning that there is no restriction on the size of the timestep or the timescale of the source in order to maintain stability. Of course, for very large timesteps, implicit methods will still have poor accuracy, even if they are stable. One other thing worth noting for implicit schemes is that in the update, Equation (2.77), the solution appears on both sides of the equation. This generally requires a root finding procedure to determine, deeming a single implicit step more computationally expensive than an explicit step.

Considering again a balance law of the form,

$$\frac{d\mathbf{Q}}{dt} = \mathcal{F}(\mathbf{Q}) + \Psi(\mathbf{Q}), \quad (2.79)$$

where the source contribution may be large, we would like an integration scheme that is both efficient—calculates as many terms explicitly as possible—but also benefits from the stability properties of implicit schemes.

For this, a series of semi-implicit time integrators (Pareschi and Russo, 2005), known as IMEX (Implicit-Explicit) schemes, have been constructed. The following method aims to solve the source implicitly to aid stability whilst treating the flux terms on the LHS explicitly, keeping computational costs at a minimum.

Rewriting Equation (2.79), we have

$$\frac{d\mathbf{Q}}{dt} = \mathcal{F}(\mathbf{Q}) + \frac{\mathbf{R}(\mathbf{Q})}{\tau_s}, \quad (2.80)$$

where we have defined the source contribution as $\Psi(\mathbf{Q}) = \mathbf{R}(\mathbf{Q})/\tau_s$ that acts on a timescale τ_s , and $\mathbf{R}(\mathbf{Q})$ as some N -vector of conserved and primitive quantities. With this, an IMEX scheme takes the general form

$$\mathbf{Q}^{(i)} = \mathbf{Q}^n + \Delta t \sum_{j=1}^{i-1} \tilde{a}_{ij} \mathcal{F}(\mathbf{Q}^{(j)}) + \Delta t \sum_{j=1}^s a_{ij} \Psi(\mathbf{Q}^{(j)}), \quad (2.81)$$

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n + \Delta t \sum_{j=1}^s \tilde{w}_j \mathcal{F}(\mathbf{Q}^{(j)}) + \Delta t \sum_{j=1}^s w_j \Psi(\mathbf{Q}^{(j)}). \quad (2.82)$$

In the above formulation, \tilde{w}_j and w_j are the j^{th} elements of the two s -vectors of coefficients, \tilde{w} and w , and \tilde{a}_{ij} and a_{ij} are the elements of two $s \times s$ matrices \tilde{A} and A respectively. IMEX schemes are often classified using the notation $\text{SSP}k(s, \sigma, p)$, where k is the order of the SSP scheme, and (s, σ, p) are the number of implicit and explicit stages, and the order of accuracy of the IMEX scheme.

The matrices \tilde{A} and A are constructed in such a way that the resulting equations for the flux contribution are explicit, and conversely are implicit for the source terms. Implicit and explicit terms are enforced by ensuring non-zero and zero diagonal components in the respective matrices, A and \tilde{A} —therefore, $a_{ij} = 0$ for $i > j$ and $\tilde{a}_{ij} = 0$ for $i \geq j$.

The coefficients are derived by comparing the numerical solution with the Taylor expansion of the exact solution, as in other RK methods. Using the standard Butcher notation of

$$\begin{array}{c|c} \tilde{c} & \tilde{A} \\ \hline & \tilde{w}^T \end{array} \quad \begin{array}{c|c} c & A \\ \hline & w^T \end{array}$$

the second order accurate $\text{SSP}2(2,2,2)$ coefficients are

$$\begin{array}{c|cc} 0 & 0 & 0 \\ 1 & 1 & 0 \\ \hline & 1/2 & 1/2 \end{array} \quad \begin{array}{c|cc} \gamma & \gamma & 0 \\ 1-\gamma & 1-2\gamma & \gamma \\ \hline & 1/2 & 1/2 \end{array}$$

TABLE 2.1: The Butcher notation for second-order IMEX coefficients.

where $\gamma = 1 - 1/\sqrt{2}$.

For the form given in Equation (2.80), the resulting scheme is given by

$$\mathbf{Q}^{n+1} = \mathbf{Q}^n + \Delta t [\mathcal{F}(\mathbf{Q}^{(1)}) + \mathcal{F}(\mathbf{Q}^{(2)})] + \Delta t [\Psi(\mathbf{Q}^{(1)}) + \Psi(\mathbf{Q}^{(2)})], \quad (2.83)$$

in which

$$\mathbf{Q}^{(1)} = \mathbf{Q}^n + \Delta t \gamma \Psi(\mathbf{Q}^{(1)}), \quad (2.84)$$

$$\mathbf{Q}^{(2)} = \mathbf{Q}^n + \Delta t \mathcal{F}(\mathbf{Q}^{(1)}) + \Delta t [(1 - 2\gamma)\Psi(\mathbf{Q}^{(1)}) + \gamma\Psi(\mathbf{Q}^{(2)})]. \quad (2.85)$$

Determining the solutions for equations (2.84) and (2.85) clearly require a non-linear root-find—this is guaranteed due to the construction of the implicit scheme. The most common tool for this purpose is an N -dimensional secant method (Palenzuela et al., 2009; Dionysopoulou, 2016), 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 (2.83), but also for every new estimate of the solution used in the root-find presented in equations (2.84) and (2.85) until a root has been found. The recovery of the primitive variables is outlined in Section 2.7 and model specific, but for resistive, relativistic fluids will involve an additional one-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. We expect systems described by the single-fluid ideal MHD equations to have size $N = 9$. This is increased to $N = 14$ for resistive MHD, and higher still for full GR multi-fluid simulations. Consequently 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 less than for explicit schemes.

2.7 Primitive Recovery

Once we have a set of evolution equations such as those given in Section 2.2, we need to establish a method for converting the conserved state vector \mathbf{Q} to the primitive variables $\mathbf{W} = \{\rho, v_x, v_y, v_z, p, B_x, B_y, B_z, E_x, E_y, E_z\}^T$, such that we can calculate the fluxes. As the primitive variables can not be expressed in closed form in terms of the conserved quantities, they require a root-finding procedure. As we are trying to motivate resistive and multi-fluid MHD, we direct the reader to Antón et al. (2010) for the ideal case.

The two-fluid model presented in 2.2.4 is composed of two resistive fluids coupled electromagnetically, thus we can recover the primitive variables by separating the two fluids and applying the standard resistive primitive recovery to each. First, the conserved fields must be decomposed into contributions from each species. We note that for some conserved variable V the charge average V is $\bar{V} = \mu_s V_s$, and so for the rest-mass density, D , we see that

$$D_1 = \frac{\bar{D} - \mu_2 D}{\mu_1 - \mu_2}, \quad (2.86)$$

$$D_2 = \frac{\bar{D} - \mu_1 D}{\mu_2 - \mu_1}. \quad (2.87)$$

For the momentum and kinetic energy terms, we must first subtract the EM field contribution as normal in resistive MHD. For example, we can define $\tilde{S}_i = \bar{S}_i - \epsilon_{ijk} E^j B^k$ and then proceed as before with

$$\tilde{S}_{i,1} = \frac{\bar{S}_i - \mu_2 \tilde{S}_i}{\mu_1 - \mu_2}, \quad (2.88)$$

$$\tilde{S}_{i,2} = \frac{\bar{S}_i - \mu_1 \tilde{S}_i}{\mu_2 - \mu_1}. \quad (2.89)$$

Now, $\tilde{S}_{1,i}$ represents the purely hydrodynamical momentum of species 1 in the i^{th} direction, i.e. $\tilde{S}_{1,i} = \rho_1 h_1 W_1^2 v_{1,i}$.

With these, we can continue in the normal resistive MHD primitive recovery following, for example, the method outlined in Palenzuela et al. (2009) or in Dionysopoulou et al. (2013)—both methods have been implemented in `METHOD` as in stiff regions one method may converge whereas the other may not given the same data, but we show only the latter for brevity.

Once the two fluids have been decoupled, we apply the following procedure to each species:

1. Remove the electromagnetic contributions from momentum and kinetic energy terms, giving $\tilde{\tau}$ and \tilde{S}_i for this species.
2. Using the definition $x \equiv \rho h W^2$, compute

$$\begin{aligned} \bar{v}^2 &= \frac{\tilde{S}_i \tilde{S}^i}{x^2}, \\ \bar{W} &= \frac{1}{\sqrt{1 - \bar{v}^2}}, \\ \bar{\rho} &= D / \bar{W}, \\ \bar{h} &= x / (\bar{\rho} \bar{W}^2), \\ \bar{v}_i &= \frac{\tilde{S}_i}{\bar{\rho} \bar{h} \bar{W}^2}, \\ \bar{p} &= \frac{(\Gamma - 1)(\bar{\rho} - \bar{h})}{\Gamma}, \end{aligned} \quad (2.90)$$

and impose the physical constraints $(p, \rho) > 0$, $(h, W) \geq 1$ and $0 \leq v^2 < 1$. A guess is required for the value of x —a sufficiently accurate guess may use the values of (ρ, h, W) at the previous timestep.

3. Calculate the value of the residual function

$$f(x) = \left(1 - \frac{\Gamma - 1}{\bar{W}^2 \Gamma}\right)x + \left(\frac{\Gamma - 1}{\bar{W} \Gamma} - 1\right)D - \tilde{\tau}. \quad (2.91)$$

4. Using a Newton-secant method, solve $f(x) = 0$, and use the relations given in step 2 to determine the rest of the primitive variables

It should be noted that this method of primitive recovery is less expensive than is required for the ideal MHD case, which involves finding the root of a two-dimensional problem.

Tolerance issues

Whichever primitive recovery method is implemented, the absolute tolerance to which the residual is to be minimised is important. For example, when evolving the two-fluid model, numerical errors in the velocity generated in the primitive recovery result in further errors in the charge current density. These errors are then multiplied by the square of the plasma frequency, ω_p^2 , in the source term, which can be as large as 10^6 in the problems we are interested in. As a result, small errors accumulated during the conservative-to-primitive variable transformation are amplified in a manner that does not occur in single fluid descriptions. It is for this reason that the tolerance is set to 10^{-13} for the two-fluid primitive recovery procedure in **METHOD** (see Section 3). This limit on the tolerance has implications on the possible accelerations that can be achieved using graphics processing units, see Section 4.1, as we must ensure that all variables are evolved using double precision floating point numbers.

2.8 Divergence Cleaning

As mentioned in Sections 2.2.2 and 2.2.3, the time evolution of the electromagnetic fields are given by Maxwell's equations which describe the temporal change of one field due to a change in flux of the other. In order to constrain the evolution of these such that the fields are physically valid, however, we must also ensure that the two divergence constraints are satisfied, namely

$$\partial_i B^i = 0, \quad \text{and} \quad (2.92a)$$

$$\partial_i E^i = \varrho, \quad (2.92b)$$

where ϱ is the charge density. Deviations from Equation (2.92a) correspond to solutions in which magnetic monopoles exist, and from Equation (2.92b) reflect a net creation (or destruction) of charge. Unchecked, these deviations will begin to significantly impact upon the solution of the simulation as a whole, and over time will result in unphysical behaviour elsewhere.

There are numerical methods that are formulated in such a way as to permanently ensure that, at every timestep in the simulation, the divergence constraints are satisfied. For example, the constrained transport technique developed initially by [Evans and Hawley \(1988\)](#) and later built upon by [Balsara and Spicer \(1999\)](#) and [Tóth \(2000\)](#) (a comparison of which can be found in

Balsara and Kim (2004)) ensures these constraints are satisfied. More often than not, however, it is sufficient to simply ensure that deviations from equations (2.92) are kept small.

As a result, the method of hyperbolic divergence cleaning (Dedner et al., 2002; Komissarov, 2007; Palenzuela et al., 2009) is frequently employed as a computationally efficient method of achieving physically realistic solutions for the electromagnetic fields, and is implemented in **METHOD**. This technique does not forbid violations of the constraints, but acts to suppress these violations and drive them exponentially to zero over some short timescale. Providing the magnitude of the timescale is sufficient, they can be kept small enough so as to only impact minimally on the simulation. Furthermore, it has been shown that this formulation is strongly hyperbolic (Hilditch et al., 2019).

We introduce two additional (unphysical) scalar fields to the conserved vector of variables, namely ψ and ϕ . These supplement the Maxwell equations, (2.14), which are now written as

$$\partial_\nu(F^{\mu\nu} + g^{\mu\nu}\psi) = I^\mu - \kappa\psi n^\mu, \quad (2.93a)$$

$$\partial_\nu(*F^{\mu\nu} + g^{\mu\nu}\phi) = -\kappa\phi n^\mu, \quad (2.93b)$$

in which we have defined $F^{\mu\nu} = n^\mu E^\nu - n^\nu E^\mu + \epsilon^{\mu\nu\rho} B_\rho$ as the Maxwell tensor, with $n_\mu = (-1, 0, 0, 0)$ and $E^0 = B^0 = 0$, and the electric four-current, $I^\nu = (-\varrho, J^i)$. Of course, in Minkowski space, $g_{\mu\nu} = \eta_{\mu\nu} = (-1, 1, 1, 1)$.

In a fashion similar to Section 2.2.2, we can investigate the special case $\mu = 0$, which gives us the following pair of equations describing the evolution of the new scalar fields:

$$\partial_t\psi + \partial_i E^i = q - \kappa\psi, \quad (2.94a)$$

$$\partial_t\phi + \partial_i B^i = -\kappa\phi. \quad (2.94b)$$

We can see that the evolution of the scalar fields is described by a wave equation, and that the solution will decay if $\kappa > 0$. Any deviations of the electromagnetic fields from Maxwell's constraints are ultimately driven away at the speed of light ($v = 1$), and their magnitude diminished exponentially according to the timescale κ^{-1} . The simplicity of this technique results in very few changes to existing models or code bases being needed—the flux vector of the electromagnetic fields gain one additional term, and the conserved vector gains two additional fields. As a result, this is the scheme implemented in **METHOD** where, unless otherwise stated, we use the timescale $\kappa^{-1} = 1$.

3

Validating our METHOD

If you don't understand it, invoke magnetic fields. If you still don't understand it, invoke turbulence.

– Astrophysics proverb

Chapters 4, 5 and 7 build upon the current state-of-the-art models that are used in current merger simulations. To test our progress, we have developed a numerical fluid code, capable of evolving special relativistic, ideal and resistive MHD for both single and two-fluid models, named **METHOD**—a multi-fluid electromagneto-hydrodynamics solver available open source through GitHub¹. In addition to what is presented in this thesis, **METHOD** is being extended to include vector potential formulations of iMHD, and is now MPI-capable, allowing for use on distributed memory clusters.

METHOD has been designed from the outset to be easily extensible to new numerical methods and physical models, allowing for a fast investigation of new ideas that would be much more cumbersome for much larger code bases. In the subsequent chapters, it will serve as a basis for execution on GPUs, a resistive extension to iMHD, and a hybrid ideal/resistive MHD model. In order to be satisfied in the solutions that this code produces, we will first present a number of tests to validate it, assess its performance for different models, and touch upon the some of the known features of the KHI and its role in merger simulations. A more detailed description of **METHOD**'s functionality and structure can be found in the Appendix, Section A.2.

3.1 Trust building exercises

In building **METHOD** we have implemented a number of testing designs (described in Section A.2.3) aimed at ensuring the various routines give the expected output. Once we are confident that the elements are passing their respective unit tests, we can perform some standardised simulations for comparison with the literature, allowing us to test the code as a whole. There are a host of routine MHD simulations that gain a considerable amount of attention when validating numerical fluid

¹www.github.com/AlexJamesWright/METHOD

codes. Typically, we need to ensure that the code has the corresponding properties: that solutions converge at the expected rates with resolution; they can handle shocks and discontinuities well; they have good energy preserving properties; they predict the correct phenomenology, e.g. the effects of resistivity or multi-fluid effects; any constraints on the solutions are satisfied well.

3.1.1 Orszag-Tang vortex

To investigate the divergence cleaning properties of MHD codes, many turn to the test case initially posed in [Orszag and Tang \(1979\)](#), now known as the Orszag-Tang vortex (OTV). Very quickly in the simulation of these vortices discontinuities begin to appear in all variables owing to the nonlinear nature of the MHD equations, thus testing the method’s shock handling capabilities also.

There are a number of different initial data that may be used for the OTV, but for our purposes we take the relativistic MHD set up from [Beckwith and Stone \(2011\)](#) and [Muddle \(2015\)](#). The simulation consists of a two-dimensional plane with the initial data given by table 3.1—all other data is zero. For this set up, we use iMHD with $\Gamma = 5/3$.

ρ	v_x	v_y	p	B_x	B_y
$\frac{25}{36\pi}$	$\frac{-\sin 2\pi y}{2}$	$\frac{\sin 2\pi x}{2}$	$\frac{5}{12\pi}$	$\frac{-\sin 2\pi y}{\sqrt{4\pi}}$	$\frac{\sin 4\pi x}{\sqrt{4\pi}}$

TABLE 3.1: Initial data for the single fluid Orszag-Tang vortex.

Figures 3.1 show that the results for the pressure and rest-mass density are indistinguishable from [Muddle \(2015\)](#) despite the significant reduction in resolution. We can see the formation of multiple interacting shocks that propagate throughout the system, a feature that makes this problem challenging for MHD codes.

We can assess the affect the divergence cleaning method has by showing the distribution of deviations from $\nabla \cdot \mathbf{B} = 0$. From Figure 3.2 we see the deviations from a divergence free magnetic field occur around shocks. When these deviations occur, they are advected away at the speed of light and suppressed, thus tending ϕ , and therefore $\nabla \cdot \mathbf{B}$, to zero. The differences between plots 3.2(a) and 3.2(b) is the timescale, $\kappa^{-1} = 1$ and 10^6 respectively. We see there is very little difference between the results for the two timescales, indicating that even a slow suppression timescale is very effective at reducing these deviations.

Note also that the initial conditions for the OTV are symmetric upon a π -rotation. We can see how well our model maintains this symmetry by taking the final density field, for example, and subtracting it from its π -rotated counterpart considering, Figure 3.3. The mean of the deviation from a perfectly symmetric solution is $|\overline{\Delta\rho_{\text{rot}}}| \approx 8 \times 10^{-6}$. The rotational symmetry is therefore accurate up to the fifth decimal place after almost 600 iterations—this is impressive considering the tolerance of the primitive recovery is 10^{-8} .

3.1.2 Brio-Wu shock tube

Brio-Wu-type shock tubes ([Brio and Wu, 1988](#)) are a common type of validation test for high resolution shock capturing schemes. The set-up consists of a one dimensional domain of length




FIGURE 3.1: Results of the relativistic Orszag-Tang vortex for ideal-MHD. Simulation end time is $T = 1.0$ with a domain of 200×200 cells.




FIGURE 3.2: In the above figures we plot the log of the resulting divergence of the magnetic field for the same simulation as Figures 3.1 but with two divergence timescales. Plots 3.2(a) and 3.2(b) correspond to values of $\kappa^{-1} = 1$ and $\kappa^{-1} = 10^6$ respectively. We note that the average value across the whole domain was $\langle \nabla \cdot \mathbf{B} \rangle \sim 10^{-3.8450}$ and $10^{-3.7956}$ respectively.

L , partitioned into two separate states and separated by a discontinuity. Upon removal of the partition, the two states interact and a series of waves propagate through the system.

Due to the simple nature of the set-up, the problem essentially consists of a single Riemann




FIGURE 3.3: Presented here are the deviations from a perfect rotational symmetry for the rest-mass density. Even after almost 600 time steps the errors are 1 part in 62500.

problem. As a result, this is one of the few instances in fluid dynamic simulations in which one can compare their outputs against an exact solution. For ideal MHD, we retrieve the exact solution from a code built by [Giacomazzo](#) as part of his PhD thesis. Of course, this will not give us the exact solution in the resistive case, but we can assert that resistive implementation tends to the exact solution in the ideal limit, $\sigma \rightarrow \infty$. For resistive set-ups we compare solutions to those in the literature.

Single fluid

To examine the behaviour in the single fluid, resistive regime we use the set-up from [Palenzuela et al. \(2009\)](#), and vary the magnitude of the resistivity. The initial state is given by equations (3.1) and we use $\Gamma = 2$ for consistency with [Dionysopoulou \(2016\)](#); [Palenzuela et al. \(2009\)](#); [Dionysopoulou et al. \(2013\)](#).

$$(\rho, p, B_y) = \begin{cases} (1, 1, 0.5) & \text{for } 0 \leq x < L/2, \\ (0.125, 0.1, -0.5) & \text{for } L/2 \leq x \leq L. \end{cases} \quad (3.1)$$

The behaviour in the ideal limit is assessed by observing how the solution varies with the system's conductivity. This corresponds to a larger source term in the resistive description, as

the source vector $\mathbf{S}(\mathbf{U}) \propto \sigma$, therefore giving us the opportunity to assess the robustness of the numerical methods we are using.




FIGURE 3.4: Dependence of y -direction magnetic field on the conductivity for the single-fluid Brio-Wu shock tube test. The exact solution corresponds to $\sigma \rightarrow \infty$, the ideal limit of MHD. Conductivities greater than $\sigma \sim 10^4$ produce no distinguishable differences in the solution—at this point, the diffusion of the magnetic fields is predominantly numerical, rather than resistive, and thus dominated by resolution.

Using a fully explicit time integrator, that of second order operator splitting, it is possible to run simulations with conductivities no greater than $\sigma \sim 10^6$, albeit with ludicrous Courant factors. Beyond this, it is necessary to use the semi-implicit IMEX integrator of Section 2.6.2—equipped with this integrator, **METHOD** can execute rMHD simulations up to $\sigma \sim 10^{15}$. Note that this assumes a constant conductivity throughout the domain, whereas in reality one expects there to be fluctuations in the electrical properties of the plasma dependant upon the local state of that region (temperature, density, magnetic field orientation/strength, etcetera). **METHOD** does have the capability of functional conductivities, as demonstrated in Chapter 7.

From Figure 3.4 we can see that in the limit of high conductivity resistive MHD tends to the exact solution given by ideal MHD. This is to be expected, as this is the definition of iMHD (zero resistivity). Furthermore, the ability of the shock capturing scheme used to accurately describe discontinuities can be seen to improve with an increase in spatial (and temporal, due to the CFL condition) resolution. The convergence to the exact solution with resolution for a number of primitive variables is depicted in Figure 3.5.

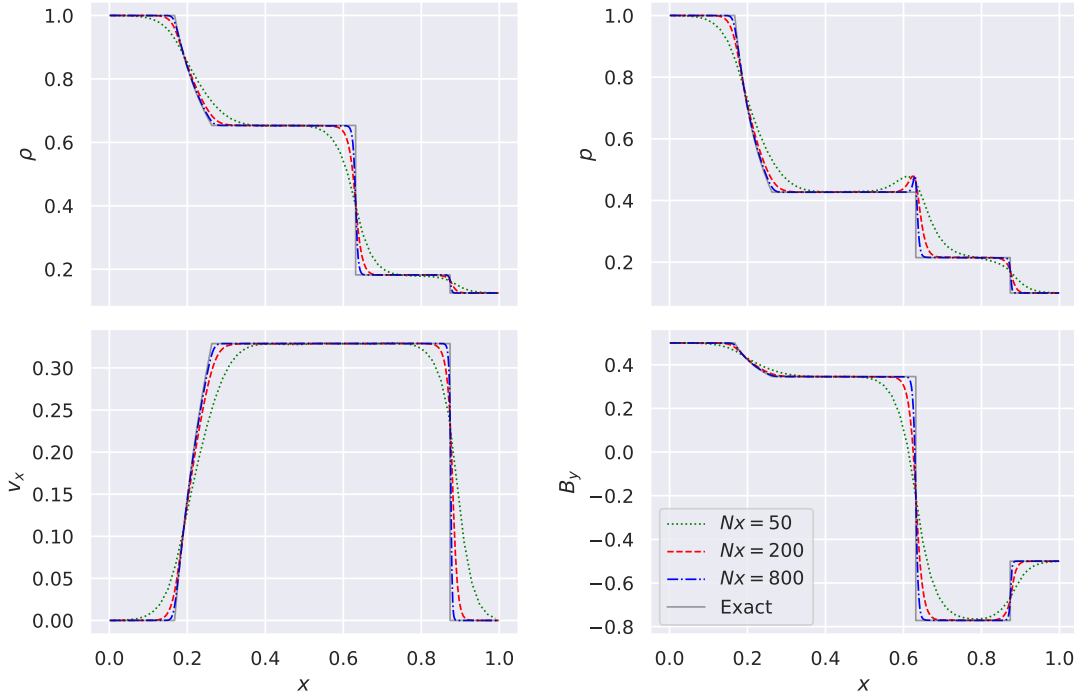


FIGURE 3.5: Convergence of the resistive solution in the ideal limit, $\sigma = 10^6$. Nx is the number of cells in the x -direction.

Two-fluid

[Amano \(2016\)](#) gives the initial data for the Brio-Wu shock tube generalised for the case of a two-fluid plasma of the type described in Section 2.2.4. This is an interesting example as it highlights some behaviour that is not present in single fluid descriptions of plasmas. With the considerations at the end of Section 2.2.4 in mind, the initial data is taken from [Amano \(2016\)](#):

$$(\rho, p, B_x, B_y) = \begin{cases} (1, 1, 0.5, 1) & \text{for } 0 \leq x < L/2, \\ (0.125, 0.1, 0.5, -1) & \text{for } L/2 \leq x \leq L. \end{cases} \quad (3.2)$$

We want the two-fluid model to recover the behaviour predicted by iMHD for large σ , but this means considering the scales at which the two-fluid effect can be resolved. By setting the charge-mass ratios to higher values we are effectively decreasing the length scale at which the behaviour is present. Once this length scale is small with respect to the computational grid size, Δx , we can no longer resolve behaviour of the separate fluids and the results match those of single fluid MHD. We can see this in Figure 3.6. Simulations have been run using $N_x = 1600$, and a moderate conductivity of $\sigma = 10^4$.

Instead, setting the magnitude of the charge-mass ratio to a smaller value, the finer grid sizes are now sufficient to resolve the weakly coupled fluid motions. This is manifest in Figure 3.7(a) as a series of dispersion waves visible via the oscillatory deviations from the reference solution. The most coarse resolution does not resolve the two fluid effect well, but as resolution increases to $N = 1600$ we can see the effects become more marked. This solution matches that given by [Amano \(2016\)](#).




FIGURE 3.6: Numerical solutions to the shock-tube test found in [Amano \(2016\)](#). The domain contained 1600 cells and was evolved until $T = 0.4$ with $\mu_p = 10^4 = -\mu_e$ and $\sigma = 10^4$.

The description of two fluid plasmas given in Section 2.2.4 makes no assumption that the charged species possess identical masses, only that we provide the relevant charge-mass ratios. We can therefore also simulate the behaviour of an electron-proton-like plasma, in which we define $m_p/m_e = 100$. Figure 3.7(b) shows the expected solution which matches Amano's. The plasma skin depth is now $10\times$ less, and thus for the same resolution we should not be able to resolve the two fluid effect as easily.

3.1.3 Self-similar current sheet

First presented in [Komissarov \(2007\)](#), another set up that permits comparison with an exact solution is that of a current sheet. For a 1D fluid in equilibrium, $p = \text{const}$, with B_y the only component of the magnetic field and the magnetic pressure much less than that of the hydrodynamic (HD) pressure, the evolution of the magnetic field is described by the diffusion equation,

$$\frac{\partial B_y}{\partial t} - \eta \frac{\partial^2 B_y}{\partial x^2} = 0, \quad (3.3)$$

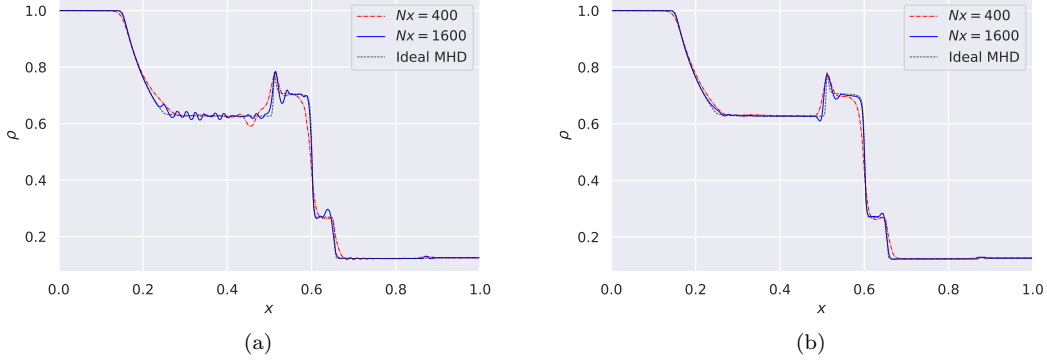


FIGURE 3.7: Presented here is the rest-mass density for the two-fluid Brion-Wu shock tube test. Left, Figure 3.7(a), shows the results for the case of an electron-positron plasma, i.e. $\mu_p = 10^3 = -\mu_e$, and right, Figure 3.7(b), for an electron-proton plasma of $\mu_p = 10^3 = \mu_e/100$. For both figures, the solution generated by the single fluid, ideal MHD model is shown for reference. We have used 1600 grid points and evolved the system up to $T = 0.4$.

where $\eta = 1/\sigma$ is the resistivity, a solution of which is

$$B_y(x, t) = B_0 \operatorname{erf}\left(\frac{\sqrt{\sigma x^2}}{2t}\right), \quad (3.4)$$

in which erf is the error function. This gives us an opportunity to compare the numerical solution in the resistive regime with an exact solution.

To generate our results, we have used the initial data corresponding to $T = 1.0$ in Equation (3.4), and set $(\rho, p, B_0) = (1.0, 50.0, 1.0)$. We use a moderate resistivity $\eta = 1/\sigma = 1/100$, and compute the solution on a domain spanning $x \in [-1.5, 1.5]$. Using a total of just $N = 50$ grid points, we can already see that the solution for both the single and two-fluid models has converged well (Figure 3.8). Due to the coarse numerical grid, two-fluid effects can not be resolved and we recover resistive MHD.

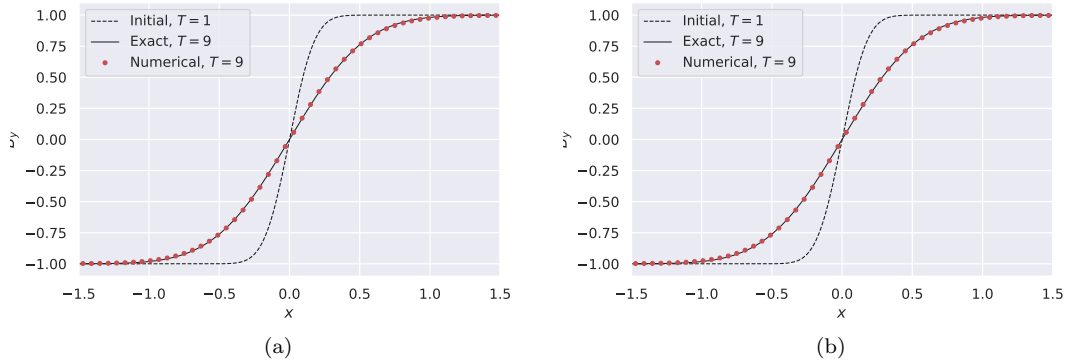


FIGURE 3.8: Comparison of the exact and numerical solutions for the self-similar current sheet for single (left) and two fluid (right) models using $N = 50$ cells and $\sigma = 10^2$. The two fluid model simulates a pair plasma, $\mu_p = 10^3 = -\mu_e$.

3.1.4 Field loop advection

Convergence tests are useful tools when assessing code behaviour. It is expected that as the resolution of the simulation is increased, the numerical solution should converge to the true solution as a function of the resolved scale. To show the numerical convergence of METHOD, we present the field loop advection test, adapted from [Gardiner and Stone \(2008\)](#). In this set-up, the fluid is set in hydrodynamic equilibrium and a weak magnetic field loop is advected through the system. Comparing the solution after one crossing time to the initial data for a range of spatial resolutions gives an indicator of the order of convergence of the algorithm.




FIGURE 3.9: The initial state (left) and the final state (right) after one crossing time for the z -direction magnetic field, B_z , using 128×128 grid points for the field loop advection test.

The hydrodynamic system is initialised with $(\rho, v_x, v_y, v_z, p) = (3, 0.1, 0.1, 0, 1)$ in a two-dimensional domain with boundaries at $(-0.5, 0.5) \times (-0.5, 0.5)$, and composed of $N \times N$ grid points. The z -direction magnetic field is then set as

$$B_z = \begin{cases} B_0(R - r) & \text{for } r \leq R \\ 0 & \text{for } r > R \end{cases} \quad (3.5)$$

with $B_0 = 10^{-3}$.

For this test, we compared initial and final states after one crossing time, $T = 10$, with periodic boundaries. The results are shown in Figure 3.9 for the single-fluid resistive model using the IMEX integrators with $\sigma = 10^5$ and with $\Gamma = 2.0$. The error is calculated with the L_1 -norm for all cells in the system and plotted as a function of the spatial resolution, where we note that $\Delta x = \Delta y$. Figure 3.10 shows the expected second order convergence, which is limited by the order of the time integrator.

3.1.5 Kelvin-Helmholtz instability

As we are motivating using more physically accurate models to simulate neutron star mergers, it makes sense to investigate the differences between the models in contexts that may arise during these types of events. Cataclysmic events, such as mergers, are exceedingly dynamic, and involve




FIGURE 3.10: The convergence for a range of resolutions for the field loop advection test. The convergence is as expected at second order.

capturing events over a range of length and timescales. At the point of merger, fluid instabilities are likely present and will grow exponentially, resulting in turbulent flows. It may be the case that these turbulent dynamos are the factories for the intense magnetic fields we see in magnetars (see Section 1.3.2).

To replicate this behaviour we simulate a magnetised KHI (Mignone et al., 2009; Beckwith and Stone, 2011; Radice, 2013). We set up a two-dimensional domain where $x \in [-0.5, 0.5]$ and $y \in [-1.0, 1.0]$ with a fluid of constant pressure, $p = 1.0$, flowing in the x -direction, with the velocity v_x given as

$$v_x = \begin{cases} v_{shear} \tanh\left(\frac{y-0.5}{a}\right) & \text{if } y > 0.0 \\ -v_{shear} \tanh\left(\frac{y+0.5}{a}\right) & \text{if } y \leq 0.0. \end{cases} \quad (3.6)$$

We choose $v_{shear} = 0.5$ as the velocity of the shear layer with characteristic thickness $a = 0.01$. The instability is seeded by a velocity perturbation in the y -direction such that

$$v_y = \begin{cases} A_0 v_{shear} \sin(2\pi x) e^{\frac{-(y-0.5)^2}{l^2}} & \text{if } y > 0.0 \\ -A_0 v_{shear} \sin(2\pi x) e^{\frac{-(y+0.5)^2}{l^2}} & \text{if } y \leq 0.0, \end{cases} \quad (3.7)$$

where the amplitude of the perturbation, $A_0 = 0.1$, decreases exponentially over a length scale $l = 0.1$. Finally, the density is set as

$$\rho = \begin{cases} \rho_0 + \rho_1 \tanh\left(\frac{y-0.5}{a}\right) & \text{if } y > 0.0 \\ \rho_0 - \rho_1 \tanh\left(\frac{y+0.5}{a}\right) & \text{if } y \leq 0.0, \end{cases} \quad (3.8)$$

so that $\rho = 1.0$ in regions of $v_x = 0.5$, and $\rho = 0.1$ in regions of $v_x = -0.5$, thus $(\rho_0, \rho_1) = (0.55, 0.45)$. When we investigate the evolution of the magnetic fields, we simply add a seed field perpendicular to the fluid, $B_z = 0.1$. This is significantly less than the field used in Mignone

et al. (2009) as we are interested in how the fluid motion affects the magnetic field, not vice versa, and so choose a weak field that will have a limited effect on the fluid dynamics.

For the initial data for the two fluid simulation we assume the charge species are co-moving, and thus have the same velocity as given for the single fluid set-up. Furthermore, we use a charge mass ratio of $\mu_p = 10^2 = -\mu_e$, and set $\Gamma = 4/3$. When using resistive models, the conductivity is $\sigma = 10$ as we are not interested in the limiting behaviour, but rather the differences between the models in their respective regimes.

3.1.5.1 Effects of resolution

A range of resolutions is presented for ideal MHD in Figure 3.11 at snapshots of $T = 1, 3$ and 6. It is clear that the KHI vortices develop sooner for simulations with a finer grid spacing—a coarse resolution results in increased numerical diffusion, reducing the growth rate of the instability. This is down to diffusion having the effect of suppressing vortex formation.

We can study the vorticity directly by looking at Figure 3.12. In the first three plots we have the magnitude of the vorticity, $|\nabla \times \mathbf{v}|$, at the simulation end time, $T = 6$, for each resolution, and the mean vorticity, $\langle |\nabla \times \mathbf{v}| \rangle$, throughout the domain as a function of time in the final plot. Clearly higher resolution simulations resolve a greater number of vortices, seen as a larger number of filaments. In addition, the vortices that are resolved are stronger in magnitude, visible also in the final plot. The fact that the vorticity increases for each resolution at $T = 0$ demonstrates that we have not fully resolved the shear layer—this will also apply in merger simulations.

The effect of the resolution on the evolution of the magnetic fields is shown in Figure 3.13. The greater vorticity of the high resolution simulation allows for a more efficient folding of the magnetic field lines, leading to a larger amplification when compared to the coarser domains. The result at the end of the linear growth stage is a magnetic field of greater intensity, and higher energy density.

The features that we have just described are common between all MHD models used. The dependence of the growth of vortices on the resolution is a general feature of all models. This is because, regardless of the model, the growth of the KHI is fastest for the highest wavenumber perturbations, and high resolutions will capture the evolution of higher wavenumber modes.

It should be noted that it is common for the linear growth of the KHI to show dependence on the numerical methods that are being used. Radice (2013) provides an analysis of different flux methods based on the WENO reconstruction (of which we have implemented the third order WENO3 here), and shows that at high enough resolutions additional vortices are formed and grow, provided the numerical scheme can resolve the initial contact discontinuity.

3.1.5.2 Model dependence

Following on from the discussion about the effects of resolution, we now provide a direct comparison of the three MHD models we have implemented. We show the resulting temporal evolution of the peak magnetic field strength and associated magnetic energy density Figure 3.14. Comparing the results of the same KHI set ups for the resistive and ideal models in the high resolution




FIGURE 3.11: Comparison of the results of the ideal MHD KHI for (top to bottom) 128×256 , 256×512 and 512×1024 grid points, at (left to right) $T = 1$, $T = 3$ and $T = 6$.

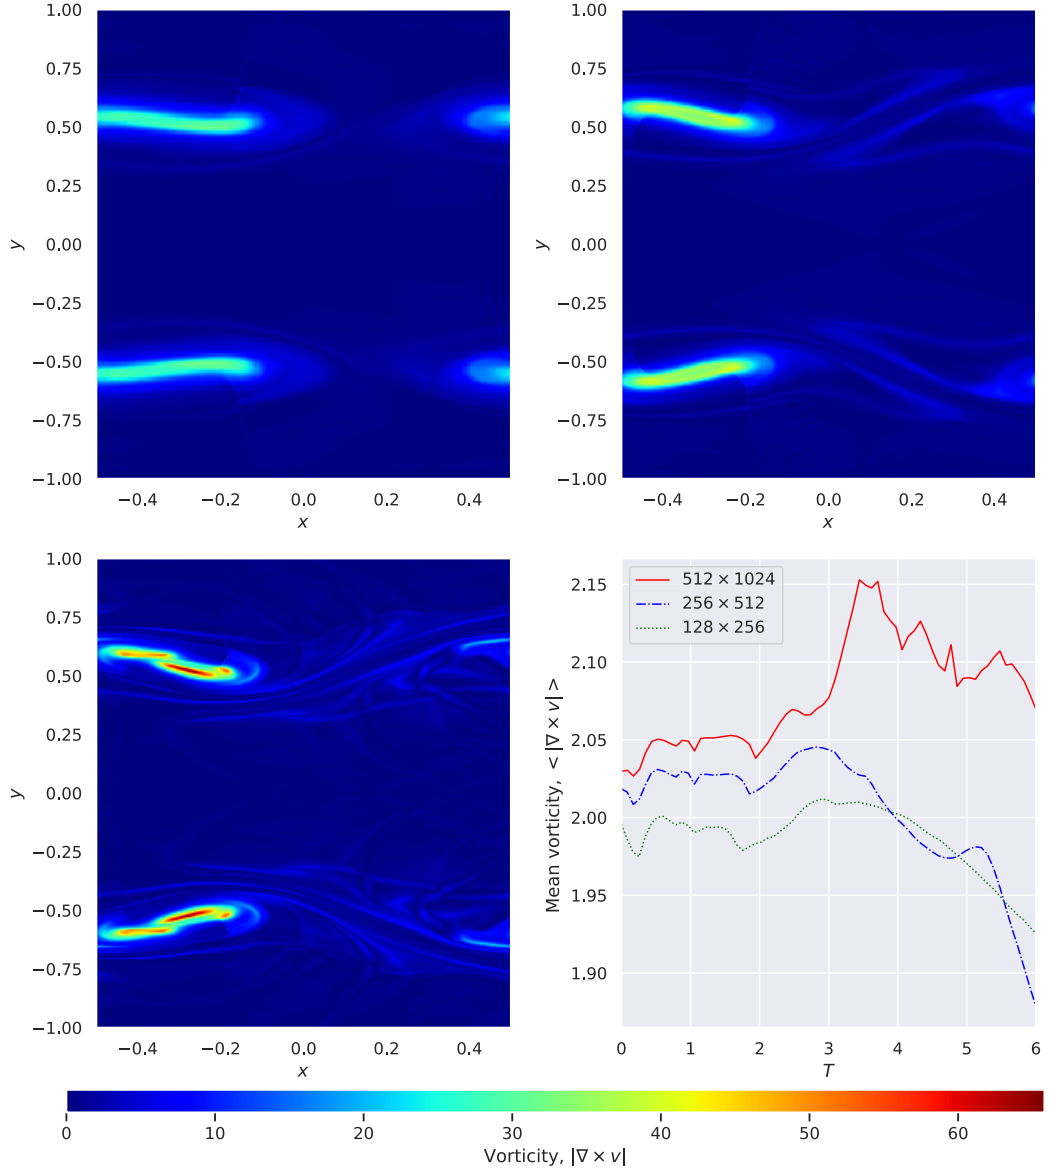


FIGURE 3.12: Final vorticity at $T = 6$ for (upper left) 128×256 , (upper right) 256×512 and (lower left) 512×1024 grid points, and (lower right) the temporal evolution of the vorticity for each resolution.

runs (512×1024), we can see that the magnetic field strength is significantly more amplified in the ideal case than for either of the two resistive models. This result agrees with those presented in [Dionysopoulou et al. \(2015\)](#), as the amplification during the merger phase as a result of the KHI was larger in the ideal simulation.

An alternative way to compare the models is through the use of their integrated power spectrum, which in discrete form is given as

$$|P_u(k)|^2 = \sum_{j=0}^{N_y-1} |\hat{u}(k, y_j)|^2 \Delta y. \quad (3.9)$$




FIGURE 3.13: Evolution of the magnetic energy density (top) and the peak magnetic field strength (bottom) for the ideal MHD KHI for different resolutions.

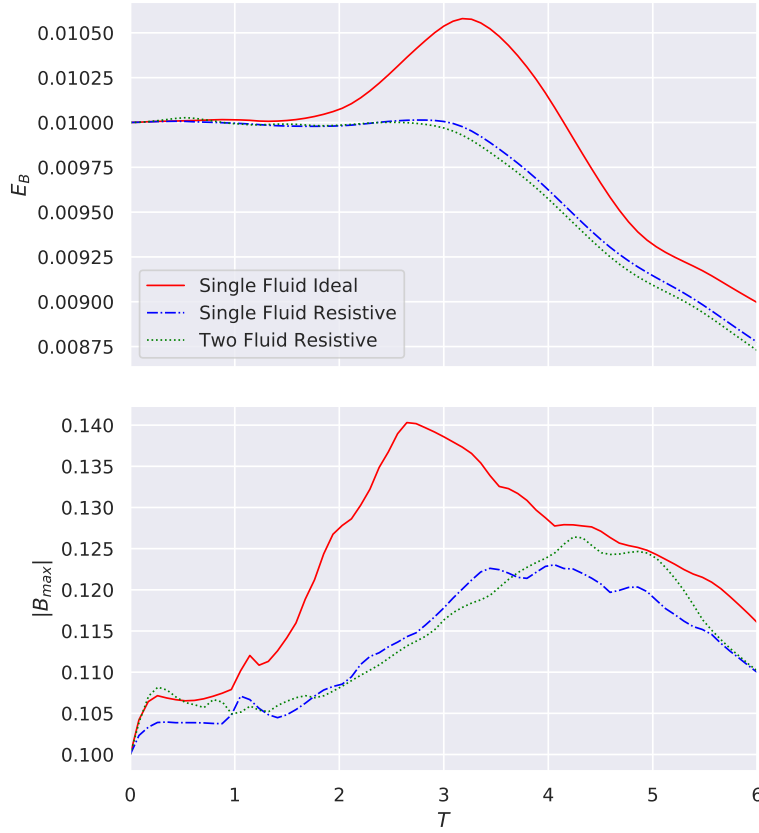


FIGURE 3.14: Evolution of the magnetic energy density (top) and the peak magnetic field strength (bottom) for each MHD model. The initial data is the relativistic KHI.

The one-dimensional Fourier transform of the variable $u(x, y)$ is given by

$$\hat{u}(k, y) = \frac{1}{N_x} \sum_{n=0}^{N_x-1} u(x_n, y) \exp \frac{-2\pi i k n}{N_x}, \quad (3.10)$$

and N_x, N_y is number of grid points in the x - and y -direction respectively. The power spectrum is then normalised such that

$$\sum_{k=1}^{k_s} |P_u(k)|^2 dk = 1, \quad (3.11)$$

where k_s is the Nyquist frequency, given by $k_s = N_x/2$ for the resolutions we have used for the KHI.

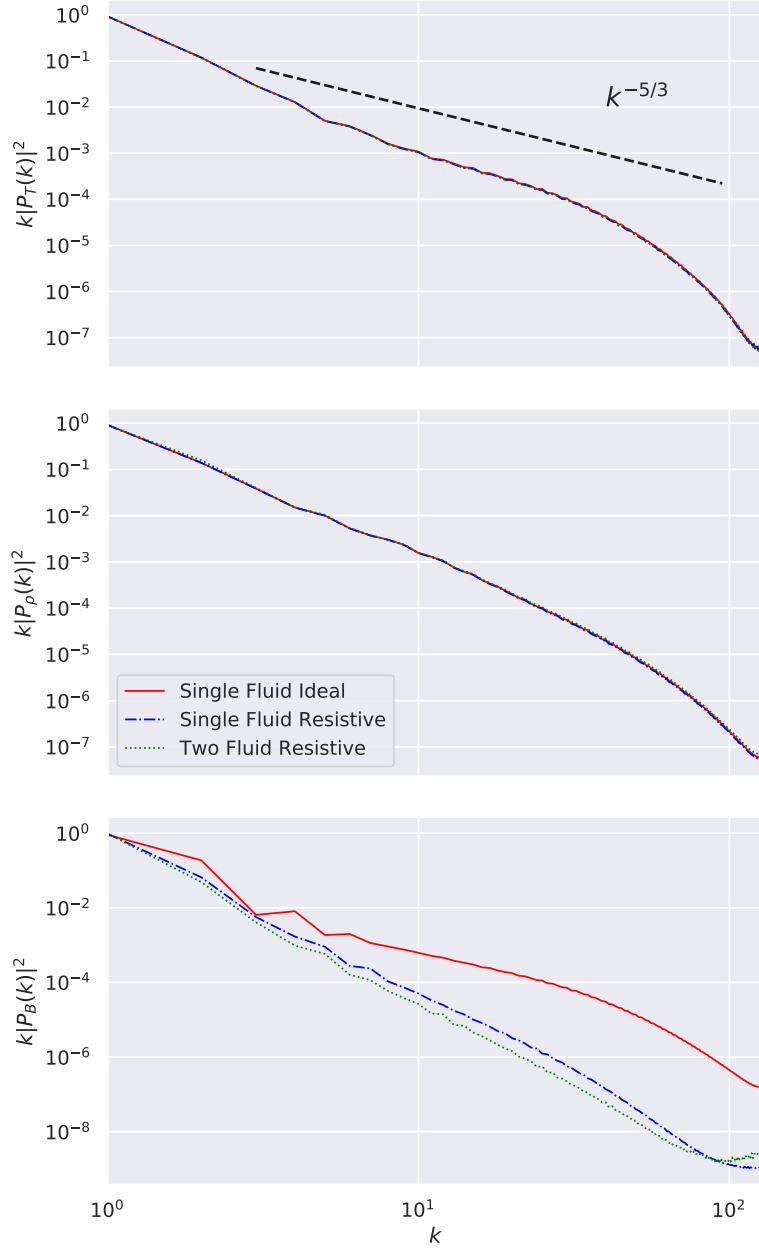


FIGURE 3.15: The power spectrum of the kinetic energy density (top), rest-mass density (middle) and magnetic energy density (bottom) for each model at $T = 3.0$. Initial data is that of the relativistic KHI.

Figure 3.15 shows various power spectra for the three models. First, we show the spectrum for the kinetic energy density. As expected, all three models follow Kolmogorov's 5/3 law ([Qian](#),

1994), and we can interpret this as a continuous flux of energy from larger scale vortices to smaller ones.

The combination of the spectra for the kinetic energy density and the rest-mass density shows that the three models give virtually indistinguishable results in terms of the structure of the fluid, differing only on the smallest scales, and by a marginal amount. This is to be expected, as although there are differences in their electromagnetic fields, the field strength is small enough so as not to significantly disturb the fluid’s motion.

The differences in terms of magnetic fields, however, are drastic. By $T = 3$, far more energy has been transferred to the magnetic fields for the ideal MHD simulation than for the two resistive models. As the magnetic field is locked to the fluid motion in the ideal limit (see Section A.1), the vortices are far more efficient at folding the field lines, leading to a greater amplification. This explains why we see such a large increase in the magnetic energy density throughout the ideal simulation, as seen in Figure 3.14. And whilst the differences between the resistive single and two-fluid resistive models are minor on large scales, we see the behaviour of the two models is different on the smaller scales (as we saw for high resolution simulations of the Brio-Wu shock tube, see Section 3.1.2).

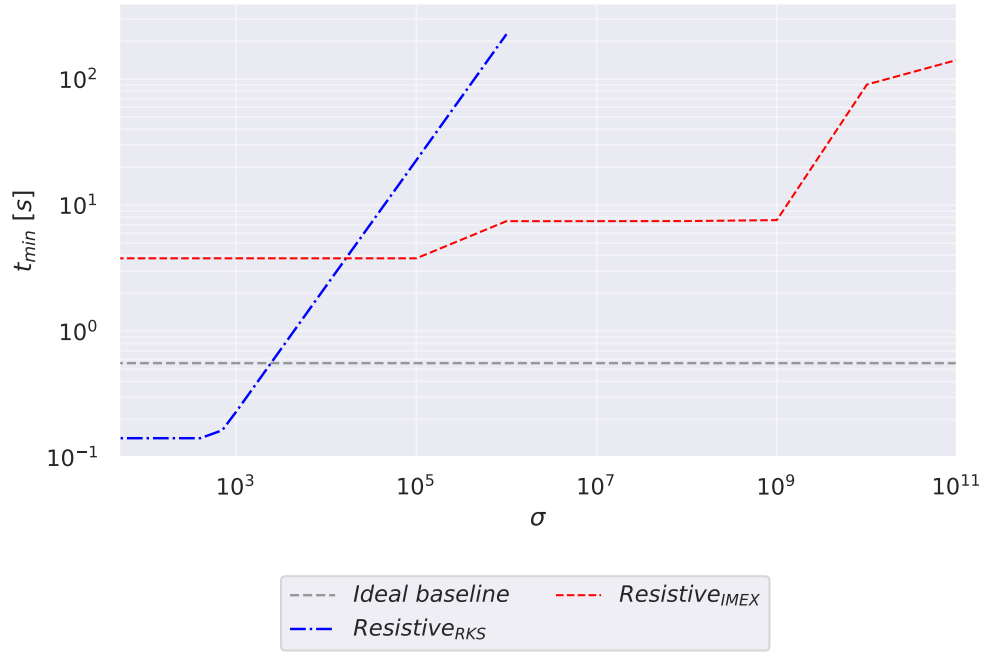
3.2 Computational demands

In Section 1.3.3.2 we motivated the study of non-ideal models of MHD for simulating mergers. We are interested now in the different behaviours the various models exhibit, with a view of alleviating the cost of running such simulations in later chapters.

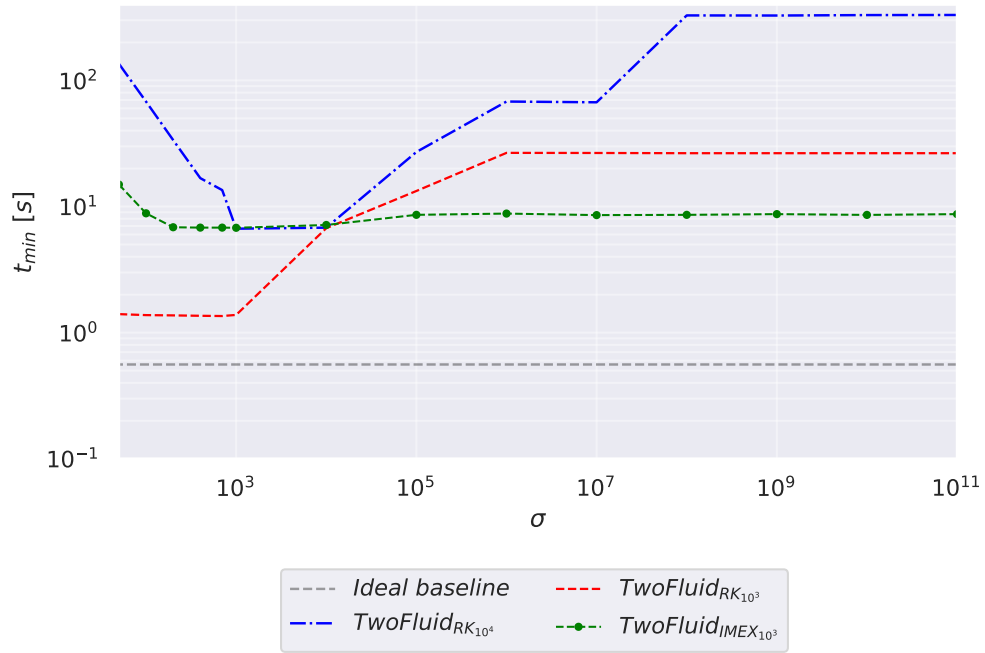
To this end, we therefore give an account of the work required to execute each model. Figures 3.16 show the results for the Brio-Wu initial data (Section 3.1.2) using both the explicit, operator split RK integrator and the IMEX schemes. All times shown are the total time for execution using the largest timestep possible whilst producing indistinguishable results, and are therefore the fastest execution times possible. We present the execution time of the simulation as a function of the conductivity. We should note that the ideal baseline is given as an indication of the time for an ideal MHD simulation to terminate, which of course does not depend on the conductivity. In Figure 3.16 the time integrator that is used is indicated in the legend, RK corresponds to the second order, operator split Runge-Kutta scheme (Section 2.6.1), whilst IMEX is for the second order, SSP2-IMEX integrator (Section 2.6.2).

First looking at the results for rMHD, Figure 3.16(a), it is clear that for the majority of the parameter space simulations will take longer than for iMHD when using explicit RK methods. We can attribute this to the source term present in the equations of motion. When the source terms become large, smaller timesteps are necessary to maintain a stable evolution, increasing the total time taken. The region for small conductivities, $\sigma \lesssim 10^3$, corresponds to source terms that are not yet too restrictive on the timestep. As a result, because of the reduced amount of computation in the resistive primitive recovery scheme—a one-dimensional root-find compared to two-dimensional for ideal MHD—the execution is faster than the ideal simulations.

It is also apparent that the IMEX schemes are important in maintaining stable evolution for highly conductive simulations. Explicit evolution very quickly becomes unreasonable, with



(a)



(b)

FIGURE 3.16: Performance comparison using the Brio-Wu initial data, Section 3.1.2, for the resistive 3.16(a) and two-fluid 3.16(b) models and a range of conductivities. Models are compared against the ideal MHD execution time, and all times correspond to the fastest possible execution whilst generating indistinguishable results. For the two-fluid model, the charge mass ratio is indicated in the legend.

simulation times tending towards days as conductivities grow. If we recall that the majority of NS matter in a simulation will be approximated well with ideal MHD, we can clearly see the need for IMEX schemes for these kinds of systems—although for very high conductivities ($\sigma \gtrsim 10^6$)

even IMEX schemes fail to be very efficient.

Turning our attention to the two-fluid model, Figure 3.16(b), we can once again see that the execution time is significantly greater than for an ideal simulation. In addition to the conductivity, we must also consider the effect of the charge mass ratio, μ , indicated in the legend by the subscripts. As the source for the two-fluid model is proportional to the square of the charge-mass ratio, this is also a leading contributor to stiffness. As a result, simulations with a larger charge-mass ratio tend to take longer to execute.

A key difference between the results for the two-fluid and resistive models is the effectiveness of the IMEX schemes in improving stability. In the two-fluid case, the IMEX schemes were only successful in reducing execution times for systems with small charge mass-ratios. In the case when $\mu = 10^4$, the source was too large even for the IMEX scheme to properly converge for small Courant factors, and as a result the evolution is only possible with the explicit integrator and a very small timestep.

It is also interesting to note that performing a resistive MHD simulation using the IMEX scheme is around 30% slower than a comparable ideal MHD simulation using the same integrator, despite the faster primitive recovery for the resistive model. The reported slow down must therefore be due to the larger conserved vector, and therefore a higher dimensional root-find in the implicit step. Roughly speaking, increasing the size of the conserved vector by 50% results in at least a 30% slower execution. Extrapolating this to a four fluid model in which the conserved vector might have size $N = 28$, we would expect the model to run around three times as slow as the ideal MHD if we assume linear scaling—although in reality it is likely to scale worse than this. A conservative estimate for a full, four fluid GRMHD model may be an order of magnitude slower to compute than single fluid ideal MHD when both are using second order IMEX integrators (note: this is ignoring the evolution of the spacetime metric).

In summary, we may expect the fastest execution time possible for a four-fluid resistive GRMHD model to be around $100\times$ slower than current ideal MHD models—this is due to a combination of far more conserved variables to evolve and the potential stiffness requiring semi-implicit integrators. It is clear that even if the models were in a computationally viable form, realistic simulations of mergers with multi-fluid GRMHD models are currently too expensive.

3.3 Conclusions

Throughout this chapter we have provided a series of numerical tests that verify the output of our numerical fluid code, *METHOD*. As is common in the literature, we have ensured that the results produced by *METHOD* converge at the expected rates, can handle discontinuous data whilst avoiding spurious numerical oscillations, and shown that the various models produce the expected results in the various physical limits.

We have also seen how sensitive the output of simulations can be to the resolution and the model of MHD being used. When performing NS simulations, researchers are interested in how observables are affected by the inclusion of new physics, for example, modified descriptions of the NS atmosphere (Lehner et al., 2012), EoS (Kawamura et al., 2016) or neutrino treatments (Foucart et al., 2018). However, we have seen in Section 3.1.5 that features such as the magnetic

fields can be hypersensitive to the simulation parameters, and that sufficient resolution and accurate models must be used.

In Section 3.2, we see more clearly the cost of using more complex models of MHD. If it is believed that non-ideal models of MHD are required to further the accuracy of merger simulations, new techniques must be developed that ease their computational demands, allowing still for a sufficient resolution to capture behaviour at many scales.

The aim of the rest of this thesis is to demonstrate a number of methods that reduce the execution time of non-ideal models of MHD. In Chapter 4, we discuss the benefits of utilising massively parallel processors for resistive and multi-fluid MHD simulations. Then, in Chapter 5, we develop a resistive extension to special relativistic, iMHD that approximates the full system, and alleviates some of the stiffness issues that arise for conductive simulations whilst benefiting from drastically improved performance. The resistive extension is then applied in a GR framework in Chapter 6, and used to model NSNS merger simulations with a realistic-scale conductivity. Finally, in Chapter 7, we demonstrate how hybrid models allow for more accurate simulations of rMHD at all conductivities, whilst requiring comparable computational time to iMHD models.

4

Massively-parallel processing

The way the processor industry is going, is to add more and more cores, but nobody knows how to program those things. I mean, two, yeah; four, not really; eight, forget it.

– Steve Jobs

In Chapter 3 we validated the `METHOD` code against a host of common MHD test problems. We have also discussed, in Section 3.2, the computational costs of the different models and the effect these demands have regarding future merger simulations. If we wish to accurately capture behaviour that occurs on a wider range of scales using more complex physical models, we will need to pursue more efficient means of executing them than are currently employed. This chapter constitutes our first attempt at alleviating the computational demand by utilising massively-parallel processors.

The results in the following chapter have been published in the *Astrophysical Journal Supplement Series* ([Wright and Hawke, 2019a](#)).

4.1 Parallel programming

Almost all multi-purpose numerical codes that are used for large scale simulations will use some parallel programming paradigm. Generally, the message passing interface (MPI) is used to deconstruct the domain into multiple smaller sub-domains and send their corresponding data to separate processors for computation ([Gabriel et al., 2004](#)). Within a single processor there may be multiple cores available for use, and computation may be further reduced to run in parallel threads using the OpenMP multi-processing architecture ([OpenMP Architecture Review Board, 2015](#)).

OpenMP and MPI are interfaces to allow parallel computation across multiple processors that operate on the central processing unit (CPU). These are multi-purpose processors and as such provide a wide range of functionality. For numerical simulations, however, computation is typically simple and substantial, and the more complex functionality offered by CPUs is

often unnecessary. Graphics processing units (GPUs), on the other hand, have a much reduced functionality.

CPU cores are generally far better optimised than cores you may find on a GPU, to the degree that a single core on a GPU will take longer to perform the same task. What graphics cards sacrifice in single thread execution, however, they make up for in sheer numbers of compute cores, allowing massively parallel tasks to be completed rapidly. Whilst there are typically up to four cores on a given CPU, massively-parallel hardware can encapsulate thousands, and the peak theoretical floating-point-operations-per-second (FLOPS) performance can be an order of magnitude greater than standard processors.

The success of these ‘accelerator architectures’ has been cultivated in industry. Demand for improved graphics in gaming resulted in the manufacture of a separate co-processor designed to render images—processors which are now used to perform image recognition (Krizhevsky et al., 2012) and to steer self-driving cars (Liu et al., 2017). As a result, the benefits of accelerator technologies have been steadily gaining attention in the scientific community for a number of years. The availability of large numbers of GPUs in high-performance computing clusters is growing, and the peak theoretical performance of graphics cards is increasing exponentially. Now, GPUs are being utilised to solve compute-intensive tasks for a range of systems.

Massively-parallel processors have been tested and verified in a number of fields, most notably in artificial intelligence (AI) as a means for accelerating the training and application of neural networks (Strigl et al., 2010). Areas of research from neuroimaging (Lee et al., 2012) to drug discovery (Gawehn et al., 2018) have seen their computing power substantially improved by porting computationally demanding tasks to GPUs. In this chapter, we investigate the potential benefits of GPU execution in evolving different models of relativistic MHD.

Recently, there has been interest in astrophysics in utilising GPUs. Zink (2011) showed how simulations of ideal GRMHD in static spacetimes may be accelerated by over $200\times$ over a single CPU core, and Lewis demonstrates significant improvement for isolated black hole simulations (Lewis and Pfeiffer, 2018). Now, there are many examples of GPU capable (magneto)hydrodynamic codes for which the execution time of simulations is dropping (Zhang et al., 2018; Bazow et al., 2016; Wong et al., 2011).

Whilst this progress is encouraging for the case of merger simulations, all these implementations use ideal descriptions. There has been no investigation into whether resistive or multi-fluid models may benefit, or how the implementations differ. In this chapter, we will see how the increased size of the conservation system and required computation necessary for resistive, multi-fluid models results in memory management issues that are more troublesome than for ideal model. Nonetheless, we demonstrate how parallel speed-ups against a single core CPU of at least a factor of $21\times$ are possible, and describe design features we believe may accelerate simulations further. To communicate with the graphics card we use the compute unified device architecture (CUDA) API developed by Nvidia (Nickolls et al., 2008).

In Section 2.7, we discussed how the plasma frequency that appears in the source term of the two-fluid model acts to amplify errors that occur during the primitive recovery. Due to the nonlinear nature of the equations of motion, these errors then propagate back into the primitive recovery during the next timestep until large enough errors have accumulated that produce

unphysical values for the primitive quantities (e.g. negative pressures). To minimise this, we have had to constrain the tolerance of the rootfinder to 10^{-13} , and we are therefore limited to using double precision floating point numbers for these models. This limits the potential benefits available by porting to GPUs, as double precision arithmetic can be many factors slower than single precision (Zink, 2011).

4.1.1 Parallel implementation

To investigate the possibility of performance improvements by utilising graphics cards, the most computationally intensive tasks have been ported to the GPU. This form of design implementation is sub-optimal in the long term, as simulation data generally resides on the CPU and must be copied as needed to the GPU. This presents a limitation in the possible speed-up as host-to-device and device-to-host memory transfers are notoriously slow (Nickolls et al., 2008; Bauer et al., 2011). It does, however, ease the process of adapting existing code to become GPU-capable. A more computationally efficient design would be to keep all data on the GPU throughout the simulation and only copy data to the host for outputting. Memory limitations can be handled, along with additional performance improvements, by utilising multiple graphics cards. Data transfer between devices can then be done using CUDA-aware MPI calls. Future iterations of METHOD will be designed in such a way.

In order to achieve good parallel speed-ups by porting functions to the GPU, one needs to identify the bottlenecks in the simulation. For the two resistive models presented here, this is the root-finding procedure required by the IMEX schemes, and specifically the conservative-to-primitive variable transformation required in each iteration. We find that, depending upon the stiffness of the problem, an increase in execution time in excess of $10\times$ is not uncommon for resistive models using IMEX schemes when compared to ideal MHD with explicit time-integrators.

To maximise the chances of performance improvements, we therefore port the time integrator, and any functions it implements, to the GPU. This includes the implicit root-find for the IMEX integrator, flux vector calculations, source vector calculations, and, crucially, the primitive recovery.

4.1.2 CUDA overview

The power of general purpose GPU (GPGPU) programming lies in the number of floating point operations per second (FLOPS) permitted by modern day GPUs (Nickolls et al., 2008). Graphics cards typically have a few thousand cores per device—where a core is a single processor, capable of executing one instruction at a time. The immense number of cores available on a GPU means that they lend themselves very well to scientific programming, in which there are often a large number operations to be performed independently of each other.

Cores are batched together on the GPU (or device) into groups of 32, called streaming multiprocessors (SMs). Batches of 32 threads, called warps, are executed concurrently in which the same instruction is applied by the threads in the warp on different data. Numbers of warps are further abstracted into blocks, and a number of blocks are dispatched on the device at a time.

The distinction of threads into blocks, and blocks into grids eases thread recruitment and the management of the data the threads are operating on. The optimum configuration of threads and blocks is discussed in Section 4.2.2.

While it is easy to run a program on a graphics card, running the program efficiently and achieving good parallel speed-ups can be difficult. Problems that are suited to GPU-execution tend to require a large number of compute operations on a small data set, in addition to each operation being largely independent of each other. Whilst there may be some interdependence between threads (e.g. during the flux reconstruction), fluid and MHD simulations are well suited to parallel programming architectures on account of the large amount of work required to evolve a single cell in the domain.

4.2 Host and device memory

When programming for the CPU, developers rarely need to have an in-depth understanding of the underlying memory management. Modern day computers are extremely efficient at optimising access to data in host memory (RAM), and it is often sufficient to allocate memory and allow the CPU to automatically use caching to improve performance. The same cannot be said when considering GPGPU programming, however.

Nvidia graphics cards have multiple different kinds of memory spaces that the programmer must be familiar with in order to achieve efficient execution. Good performance improvements are achieved only when the correct balance of the various memory resources are used well, and getting this step wrong can significantly hamper possible improvements. We will now explain some of the strategies applied in *METHOD* to optimise memory management.

4.2.1 Memory coalescence

When copying data between the CPU (the host) and device, one must ensure that the data is aligned in a suitable way. Transferring data to device memory essentially boils down to copying a single, one-dimensional array, and in order to optimise this process the developer must be familiar with how data is stored in memory and how the data is accessed for a computation.

In *METHOD*, data is aligned in the following way: the conserved vector is an array with dimensions $N_{\text{cons}} \times N_x \times N_y \times N_z$, which we can access using `cons[var, i, j, k]` where `var`, `i`, `j`, `k` access the `var`th conserved variable in the `i`th, `j`th and `k`th cell in the x , y , z directions, respectively. Due to the way in which memory is allocated in C++, all data is aligned first in the z -direction, followed by y , x and then the variable, `var`—this means that all data in the z -direction is contiguous in memory.

As the flux reconstruction requires interpolating data between neighbouring cells, this is easily performed in the z -direction as data is already contiguous in memory. Reconstructing in the x - and y -directions, however, requires us to rearrange the original array such that data is contiguous in the x - and y -directions, respectively. This rearranging of data is also necessary when copying data to the device for the IMEX integrator, as in this case, data must be contiguous in the conserved variables, `var`.

The data is reorganised on the CPU so that it can be copied in a contiguous array to the device for operation. To reduce the performance hit this reorganising has, we use OpenMP (OpenMP Architecture Review Board, 2015) to parallelise this process. Another possible option to reduce this overhead is to copy all data to the device in its original order and allow threads to copy the relevant data from global memory to shared memory into the correct order. Initial efforts to achieve this slowed simulations by many factors, and instead we opted to stick with the CUDA/OpenMP hybrid option. In Section 4.3 we discuss the resulting performance impact.

4.2.2 Memory optimisation

There are three kinds of device memory that are utilised in **METHOD**—global, shared and register—each with different sizes and latencies. Memory latency is the amount of time it takes from a thread requesting some data, to the data being available for computation. The different memory spaces have different latencies due to the distance from where the threads operate, the SM, to the location of the data on the device.

Global memory is the largest block of memory, and likewise the slowest to access. When data is copied to the device from the host, it is copied into global memory where every thread can access it. Shared memory is smaller in size, typically $\approx 96\text{KB}$ per SM, but around $100\times$ faster to access. Threads in the same block have access to the same shared memory, but cannot see data lying in another block’s shared memory. This restriction means that when threads require data from their neighbours, one must overlap data between blocks. Finally, register memory is thread specific, and around $100\times$ quicker than shared memory.

Handling memory management well in GPGPU programming is the difference between order of magnitude speed-ups or slow-downs, and often a large amount of trial and error testing is required until an optimal balance between shared and global memory is found. The same can be said in terms of the configuration of threads and blocks—a large number of threads in a block will limit the amount of shared memory per thread that is available, but a small number of threads may reduce the opportunity to hide memory latency behind useful computation. The optimal configuration will always be problem specific.

For the resistive MHD models presented here, the best strategy we found is to maximise usage of shared memory for only 32 threads per block. This shared memory is specifically allocated for use in the primitive recovery, as these variables are accessed a very large number of times throughout the course of the IMEX scheme, and the benefits of having quicker access to these variables outweighs the benefits of having them lie in global memory with the possibility of hiding memory latency within blocks. In order to hide latency between blocks, we therefore choose the maximum number of blocks possible with the available memory.

The most notable difference between implementing ideal and resistive MHD models on the GPU is the necessity of an implicit integrator to maintain stability in stiff regions. The process of implementing an implicit step requires some N -dimensional root-finding procedure, in which N is the size of the conserved vector. The package used for this process in **METHOD** is adapted from CMINPACK¹, a C version of the widely used Fortran package, MINPACK.

¹Original source: <https://github.com/devernay/cminpack>

CMINPACK requires the use of a work-array of size no less than $N(3N+13)$ double precision, floating point numbers. The work-array provides the CMINPACK routines with resources to determine the numerical Jacobian of the system, error norms, etcetera, and so requires data that will be accessed multiple times throughout a single step. Ideally, this would lie in shared memory to minimise latency, but due to the size of both resistive systems and the limitations imposed by the size of the shared memory, the work-array must be allocated in global memory. This issue is specific to systems that require implicit integrators, and does not occur for ideal MHD models. As some of the data that is frequently required is located in global memory, one would expect greater latency when using implicit over explicit integrators, and thus for resistive over ideal MHD models.

4.3 Performance

Our attention now turns to the performance of the CPU and GPU implementations of **METHOD**, and the potential benefits of GPU-capable, resistive MHD codes. For the comparison, we make use of the IRIDIS 5 HPC cluster at the University of Southampton—the GPU nodes possess a number of Nvidia GTX1080 (Pascal architecture) (Nvidia, 2016) and V100 (Volta architecture) (Nvidia, 2017) graphics cards. We perform our analysis on a single GTX1080, and a single V100 card to assess performance differences between the two architectures, and compare these against a benchmark, CPU implementation on an INTEL XEON 6138 GOLD using a single core.

To compare the execution times of the serial and parallel implementations, we measure the wall clock time of the main evolution loop for ten iterations. This execution does not involve any time spent in setting up the domain or initial data, or writing data to the disk, and so is representative of the differences one would expect in the use of these methods for large-scale merger simulations where the majority of the execution time is spent on the evolution of the system (as opposed to initialisation or output of data).

We use two different initial data for the comparison—the Brio-Wu shock tube test and the more complex KHI instability, both described in Section 3.1. Furthermore, as we are only concerned with the GPU implementations of the two resistive models, we use the IMEX schemes mentioned in Section 2.6.2—results from both the SSP2(222) and SSP3(322) integrators (Pareschi and Russo, 2005) are presented. All simulations are run using double-precision floating-point accuracy as per the discussion in Section 2.7.

The serial execution time does not vary much between runs as differences depend only upon the current state of the CPU. The execution time for the parallel implementation, however, depends on a number of factors such as usage of shared memory and registers, number of host-device memory copies, and the configuration of threads and blocks. As such, presented here are the optimal configurations that minimise the execution time on the graphics card. We find that the optimal configurations is to minimise the number of threads in a block to 32, as discussed in Section 4.2.2, to allow sufficient usage of shared memory for the primitive recovery. We then maximise the number of blocks per grid to hide latency, relax the condition that serial and parallel output must exactly match (whilst keeping physically valid results), compile device code with the `-O3` flag and switch off CUDA’s fused-multiply-add functionality.





FIGURE 4.1: A measure of the cost-per-cell (execution time in μs) on a V100 graphics card and the CPU for the Brio-Wu test using the SSP3 integrator. Execution on the CPU is independent of the domain size, whilst execution time on the device decreases until the GPU is saturated.

Figure 4.1 shows the average cost-per-cell for the Brio-Wu simulation when using the SSP3 integrator on the CPU and a V100. As expected, the cost-per-cell for the CPU implementation is independent of the size of the domain, all data lies in the same memory (RAM) and so accessing and computation of data is common among all cells. Instead, looking at the GPU implementation, we can see that as the number of cells in the simulation increases, the cost-per-cell reduces drastically.


Smaller domains contain fewer cells, and thus less opportunity to hide memory latency behind useful computation. As a result, for small domains the problem is memory-bound—a significant amount of time is spent waiting for access to various locations in memory and as a result the threads lay idle. As we increase the size of the domain, the bottleneck then becomes the amount of computation required for each time step, and the problem is then compute-bound. We can see the transition of the problem from memory- to compute-bound in Figure 4.1 as the cost-per-cell begins to plateau for the largest systems. At this point, the GPU is saturated, and improvements may only be made via smarter memory management, or additional GPUs.

To see how this impacts the absolute speed-up due to the GPU implementation, we plot the ratio of the serial execution time and parallel execution time (parallel speed-up) for ten time-steps as a function of the total number of grid points in the simulation, Figure 4.2. The results for two generations of cards are shown, the GTX1080 and V100. It is clear by comparing the two cards that the newer generation graphics card, the V100 with the Volta architecture, provides improved speed-ups. It can also be seen that there is little difference between the results of the two initial data and the two IMEX schemes for the older (but still widely used) Pascal architecture. On the other hand, the new Volta architecture offers greater memory bandwidth (nearly $3\times$), and therefore a reduced memory latency—the result is to allow execution time to be dominated by computation. We see this as a much improved parallel speed-up, over $21\times$, specifically for the compute intensive algorithms of the resistive, two-fluid model with the SSP3-IMEX scheme.

The previous result suggests that higher order methods which require additional computation will achieve greater parallel speed-ups than lower order alternatives. The additional computation



(a) Speed up for the GTX1080 graphics card.



(b) Speed up for the V100 graphics card.

FIGURE 4.2: The measure of parallel speed-up, $T_{\text{serial}}/T_{\text{parallel}}$, for the GTX1080 (top) and V100 (bottom) graphics card. Squares and diamonds correspond to the resistive, single fluid Brio-Wu, while pluses and crosses to the two-fluid Kelvin-Helmholtz instability. The speed-up depends strongly upon the total number of cells, reaching a maximum of $21\times$ before memory resources on the device are exhausted. The maximum parallel speed-up for each card is indicated by the dashed line.

required by an IMEX-SSP3(433) scheme, for example, which requires a total of four implicit stages and three flux reconstructions, would result in a greater proportion of the execution time being spent on the graphics card. As a result, the benefits of a GPU-implementation for simulations in which the higher order schemes are used would be more apparent—merger simulations typically employ fourth order methods, for example.

To gain insight into where the performance bottlenecks are in **METHOD**, we present a table with various profiling data, Table 4.1. Data was generated using the Brio-Wu initial data and a relatively large domain on $N_x = 131072$ grid points, the SSP3 integrator using third order FVS reconstruction, and a conductivity of $\sigma = 10^4$. Even with execution on the device, the IMEX scheme makes up the largest proportion of the execution time. The next largest contributor is the rearrangement of data into contiguous arrays for the IMEX scheme, as discussed in Section

4.2.1. The fraction of time taken to rearrange the data is heavily dependant upon the domain size. The expected scaling of the rearrangement on the CPU goes as the number of cells, N_x , but we have seen that the cost-per-cell of the main loop reduces with the number of cells. As a result, as the domain increases in size, more time is spent proportionally rearranging the data when compared to the total execution time (which is dominated by the execution on the device). We can see this visually in Figure 4.3 as a growing fraction of the execution time being spent reordering data. Large scale simulations of neutron star mergers may evolve tens or hundreds of times as many grid points as we show here, and so a significant fraction of the execution time would be performing this rearrangement—in this case, a more efficient scheme for performing this would be highly beneficial.

Performance Profile					
Functionality	IMEX rearrange	IMEX root-find	FVS rearrange	FVS reconstruction	boundary conditions
Proportion of execution [%]	21.7	55.3	1.0	2.2	0.2

TABLE 4.1: Performance profile for the Brio-Wu initial data using the SSP3 integrator and $N_x = 2^{17}$ cells. Results represent the percentage of the total execution time of the main loop. The rest of the execution time is composed of the correction stage, Equation (2.83), primitive recovery that is not within the implicit stages of the IMEX scheme, and data and memory management.




FIGURE 4.3: Cost of rearranging data for IMEX scheme as a percentage of the total execution time. There is a strong dependence on the number of cells in the domain.

Time spent rearranging data for the flux reconstruction, however, is significantly less than that for the IMEX scheme. As only the conserved and flux vectors need rearranging, this results in a total of 3 arrays being reordered in the FVS method. This is compared to a total of 28 arrays

required for the SSP3 scheme, including the conserved, primitive, source and flux vectors for each stage. As a result, there are fewer benefits to be had from optimising the flux reconstruction.

4.4 Discussion

When evolving resistive single or multi-fluid MHD models, one must use a time integrator that solves, at the least, the source terms implicitly—as a result, the majority of the execution time for these kinds of simulations is spent in these schemes. We have shown that the most commonly used class of integrators for this procedure, IMEX-RK schemes, can be accelerated by at least a factor of $21\times$ through execution on the device. Furthermore, we argue that there are potentially greater improvements to be made for higher order methods, as proportionally more time will be spent in the accelerated regions of the simulation.

The greatest improvements in performance are possible when managing the resources of shared memory correctly. The recovery of the primitive variables, necessary for each iteration of the implicit stage in the IMEX scheme, itself requires a root-finding procedure. To minimise memory latency, the state arrays required for the recovery are kept in shared memory—this severely limits the number of threads per block that we can recruit but saves multiple factors in terms of execution time.

The available speed-ups of *METHOD* are limited by its design—*METHOD* has been extended to execute the most computationally expensive functions (time integration) on the device, and inherent in this is a number of host-to-device and device-to-host memory transfers. In order to copy data from contiguous arrays we are required to perform a rearrangement of the data on the host which takes up a significant fraction of the overall execution. To reduce this overhead, and to improve the potential parallel speed-ups, we note that a more efficient design would be to develop the GPU-capabilities from the outset. In this way, all data will be in device memory and only transferred to the host for output. Organising the simulation data in this way may allow threads to load data into contiguous arrays in parallel on the device, which for large domains may reduce the cost of the rearrangement significantly.

Furthermore, even with the acceleration of the time-integrator by execution on the device, the majority of time is still spent in the IMEX root-find. This means that optimisations in this region will likely produce further improvements. This could include reducing the size of the work-array required by the root-finding procedure such that it can fit in shared memory, although greater improvements seem likely by keeping the primitive recovery variables in shared memory at the expense of the IMEX work-array.

The size of the conserved vector has little effect on the possible speed up. The single-fluid, resistive MHD model, which has the smallest conserved vector presented, requires a work array for the time integration that is too large to fit in shared memory. As a result, the speed-up comes primarily from the primitive recovery variables lying in shared memory. The same is true for models with larger conserved vectors, and so more complex, four-fluid models should expect comparable accelerations—the required memory for the primitive recovery grows linearly with the conserved vector size, in contrast to the quadratic dependence for the time integrator. As a

result, the methods employed here should transfer over to the more computationally demanding, general-relativistic multi-fluid models proposed in [Andersson et al. \(2017a\)](#).

The first release of METHOD used for generating these results is available through Zenodo ([Wright, 2018](#)), and the latest iteration is publicly available on GitHub².

²www.github.com/AlexJamesWright/METHOD

5

Chapman-Enskog analysis

Theory is the first term in the Taylor series expansion of practice.

– Thomas Cover

In Section 1.3.3.1 we introduced the idea of sub-grid models as a means of emulating effects due to unresolved fluid motion. In this chapter, we take this idea and extend it to include effects due to alternative models of MHD.

We have seen in Chapter 3 how the various models of MHD that we present in this thesis behave in a range of scenarios. Specifically relating to the KHI, we have seen that the behaviour of the magnetic fields depends to a large extent upon the inclusion of resistivity, or of multiple fluids on the smaller scales of the simulation. Whilst we believe resistive and multi-fluid MHD to be a more accurate description of neutron star systems, using them encounters the numerical difficulties that come with potentially stiff sources.

Many systems of interest that exhibit resistive behaviour do so near the ideal limit. For example, we have presented simulations (Section 3.1) in which the resistivity (in non-dimensionalised units) is as large as $\eta \sim 10^{-1}$, but often $\eta \sim 10^{-2}$ or less shows behaviour that differs significantly from ideal simulations. In astrophysical simulations of mergers, resistivities in matter have been used on the order of $\eta \sim 10^{-6}$ (Ponce et al., 2014; Dionysopoulou et al., 2015), and simulations of jet launching from black-hole accretion have used $\eta \sim 10^{-3}$ (Qian et al., 2018). These values for the resistivity result in the equations of motion of the electric fields being stiff, requiring the use of implicit integrators for stable evolution.

In this chapter, we will see how one can extend the equations of special relativistic iMHD to include an additional source that captures the behaviour of rMHD—dubbed a resistive extension upgrade for ideal magnetohydrodynamics (REGIME). The contribution of the new source term to the change in the conserved variables goes as the inverse of the conductivity, and thus is stiff in the opposing limit of rMHD. As a result, for systems in which we have only weak resistivity, $\eta \lesssim 10^{-2}$, the extension can be evolved efficiently using explicit integrators.

To derive REGIME, we use a Chapman-Enskog-type analysis. This type of analysis has been used successfully in fluid dynamics in determining moments of the Boltzmann equation,

yielding both the Euler and Navier-Stokes equations (Cercignani and Kremer, 2002). In later sections, we apply the same reasoning to the system of equations defining rMHD. What emerges is an alternative formulation that is accurate up to first order in the resistivity, and that is more suitable numerically for highly conducting plasmas.

5.1 Scalar equation expansion

As an example of the Chapman-Enskog type expansion we will use in the next section, we present an analysis of a toy problem taken from LeVeque's book (LeVeque, 2002), and considered as early as 1987 (Liu, 1987). We start with a one dimensional, coupled, scalar system of two variables, $u(x, t)$ and $v(x, t)$, defined by

$$\frac{\partial u}{\partial t} + \frac{\partial v}{\partial x} = 0, \quad (5.1a)$$

$$\frac{\partial v}{\partial t} + a \frac{\partial u}{\partial x} = \frac{f(u) - v}{\epsilon}, \quad (5.1b)$$

where $0 < \epsilon \ll 1$, $a > 0$ and $f(u)$ is some given function. The equilibrium state of this system is defined when the RHS is zero, which occurs when $v = f(u)$ —this is the solution in the limit $\epsilon \rightarrow 0$. At this point, system (5.1) reduces to the scalar conservation equation,

$$\frac{\partial u}{\partial t} + \frac{\partial f(u)}{\partial x} = 0, \quad (5.2)$$

with v computed from the condition $v = f(u)$. Following the reasoning of LeVeque, we assume that in general the solution for the second variable, v , in the limit $\epsilon \rightarrow 0$ can be expanded around this solution,

$$v(x, t) = f(u(x, t)) + \epsilon v_1(x, t) + \epsilon^2 v_2(x, t) + \dots, \quad (5.3)$$

where the equilibrium solution for the second variable is $v_0(x, t) = f(u(x, t))$, and the perturbations on top of this are of the order of increasing powers of ϵ . Substituting this back into the evolution equation of the first variable gives

$$\frac{\partial u}{\partial t} + \frac{\partial f(u)}{\partial x} = -\epsilon \frac{\partial v_1(x, t)}{\partial x} + \mathcal{O}(\epsilon^2). \quad (5.4)$$

This expression looks very much like a conventional conservation system, with some small (order ϵ) added source term governed by $v_1(x, t)$, which is to be determined. Substituting Equation (5.3) again into the evolution of the second variable, Equation (5.1b), and only keeping terms of order $\mathcal{O}(\epsilon^0)$, gives

$$\frac{\partial f(u)}{\partial t} + a \frac{\partial u}{\partial x} = \frac{f(u) - (f(u) + \epsilon v_1(x, t) + \dots)}{\epsilon}, \quad (5.5)$$

$$\frac{\partial u}{\partial t} \frac{\partial f(u)}{\partial u} + a \frac{\partial u}{\partial x} = -v_1(x, t). \quad (5.6)$$

Removing the time dependence using Equation (5.2), we get the final form of the perturbation as

$$v_1(x, t) = \frac{\partial u}{\partial x} \left[\left(\frac{\partial f}{\partial u} \right)^2 - a \right]. \quad (5.7)$$

This gives us a new evolution equation for the first variable in the limit of small ϵ , up to first order in ϵ :

$$\frac{\partial u}{\partial t} + \frac{\partial f(u)}{\partial x} = \epsilon \frac{\partial}{\partial x} \left(\left[a - \left(\frac{\partial f}{\partial u} \right)^2 \right] \frac{\partial u}{\partial x} \right). \quad (5.8)$$

The fact that this balance law contains a second derivative term means that the source plays the role of diffusion. Clearly for $\epsilon = 0$ we recover the standard conservation law, or equilibrium solution, for u . For small but non-zero values of ϵ , the RHS acts to modify this solution, generally smoothing it out. As ϵ becomes larger, higher order moments of Equation (5.3) may begin to become important.

5.1.1 Stability

If we briefly focus on the RHS of Equation (5.8), reducing the system to

$$\frac{\partial u}{\partial t} = \epsilon \frac{\partial}{\partial x} \left(\left[a - \left(\frac{\partial f}{\partial u} \right)^2 \right] \frac{\partial u}{\partial x} \right), \quad (5.9)$$

we can determine the stability of the system as a result of the diffusive correction. Making the assumption that the Jacobian, $\partial_u f$, is not varying in space, we can rearrange this to

$$\frac{\partial u}{\partial t} = \epsilon \left[a - \left(\frac{\partial f}{\partial u} \right)^2 \right] \frac{\partial^2 u}{\partial x^2}, \quad (5.10)$$

comparing with the archetypal diffusion equation,

$$\frac{\partial u}{\partial t} = k \frac{\partial^2 u}{\partial x^2}. \quad (5.11)$$

To ensure $k \geq 0$, we require that

$$-\sqrt{a} \geq \frac{\partial f}{\partial u} \geq \sqrt{a}. \quad (5.12)$$

This relation is known as the subcharacteristic condition (Liu, 1987), and for systems in which this constraint is not satisfied solutions tend to blow up along characteristics (LeVeque, 2002). A more strict statement of this condition also requires that the eigenvalues of the Jacobian of the relaxation system, Equations (5.1), interlace those of the relaxed system, Equation (5.2) (Flåtten and Lund, 2011).

5.2 System expansion

The remainder of this chapter has been published in Monthly Notices of the Royal Astronomical Society ([Wright and Hawke, 2019b](#)).

In the last section, we saw how the Chapman-Enskog expansion can be used to approximate the solution near some equilibrium state for a specific scalar problem. In this section, we perform the same analysis for first order balance laws. First, we introduce the notation that will be used.

In general, we could write system (5.1) in the following way:

$$\partial_t \mathbf{q} + \partial_{(a)} \mathbf{f}^{(a)}(\mathbf{q}, \bar{\mathbf{q}}) = \mathbf{s}(\mathbf{q}, \bar{\mathbf{q}}), \quad (5.13a)$$

$$\partial_t \bar{\mathbf{q}} + \partial_{(a)} \bar{\mathbf{f}}^{(a)}(\mathbf{q}, \bar{\mathbf{q}}) = \frac{\bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}})}{\epsilon}, \quad (5.13b)$$

where we indicate the variables of stiff equations with an over-bar. We denote ϵ as the timescale that the stiff source acts upon, and assume the sum of the fluxes over the three spatial coordinates, $a \in [x, y, z]$.

We are going to apply this expansion to the equations of rMHD. Recall from Section 2.2.3 that the source term for the electric fields in rMHD is proportional to the conductivity, σ . Ideal MHD corresponds to the assumption of perfect conductivity—i.e. that $\sigma \rightarrow \infty$, or equivalently, $\epsilon \rightarrow 0$ —thus, in this limit the equations of rMHD behave in the same way as system (5.1). From this, we can take iMHD as the equilibrium solution of rMHD.

In order to maintain finite solutions in the ideal limit, we require that $\lim_{\epsilon \rightarrow 0} \bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}) = \mathbf{0}$. We will refer to the solution in the limit $\epsilon \rightarrow 0$ as the equilibrium system, as for the models we are interested in here this corresponds to the RHS = 0, and Equation (5.13a) reduces to iMHD. As this limit corresponds to iMHD where $\sigma = \infty$, for large but finite conductivities we can think of the solution of rMHD as some perturbation on top of the solution of iMHD.

To derive the form of this perturbation, we need to define the equilibrium solution. At this point, we assume that the stiff variables have some values given by $\bar{\mathbf{q}}_0$, and so the stiff source in the equilibrium limit must be

$$\bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}_0) = \mathbf{0}. \quad (5.14)$$

In analogy to what we did in the previous section, Equation (5.3), the Chapman-Enskog expansion now requires we express the vector of stiff variables in terms of a series of increasing powers of the small quantity, ϵ . If we expand the solution for the stiff variables about the equilibrium solution, this gives us

$$\bar{\mathbf{q}} = \bar{\mathbf{q}}_0 + \epsilon \bar{\mathbf{q}}_1 + \epsilon^2 \bar{\mathbf{q}}_2 + \mathcal{O}(\epsilon^3), \quad (5.15)$$

in which each new power of ϵ represents some higher order perturbation on top of the previous. Using the Taylor expansion of some vector function as $\mathbf{A}(\bar{\mathbf{q}}_0 + \epsilon \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2)) = \mathbf{A}(\bar{\mathbf{q}}_0) +$

$\epsilon \partial_{\bar{\mathbf{q}}} \mathbf{A}(\bar{\mathbf{q}})|_{\bar{\mathbf{q}}=\bar{\mathbf{q}}_0} \cdot \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2)$, we can expand the stiff variables, Equation (5.13b), about the equilibrium solution, giving

$$\partial_t \bar{\mathbf{q}}_0 + \epsilon \partial_t \bar{\mathbf{q}}_1 + \partial_{(a)} \left[\bar{\mathbf{f}}^{(a)}(\mathbf{q}, \bar{\mathbf{q}}_0) + \epsilon \frac{\partial \bar{\mathbf{f}}^{(a)}(\mathbf{q}, \bar{\mathbf{q}})}{\partial \bar{\mathbf{q}}} \Big|_{\bar{\mathbf{q}}=\bar{\mathbf{q}}_0} \cdot \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2) \right] = \frac{\bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}_0) + \epsilon \frac{\partial \bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}})}{\partial \bar{\mathbf{q}}} \Big|_{\bar{\mathbf{q}}=\bar{\mathbf{q}}_0} \cdot \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2)}{\epsilon}. \quad (5.16)$$

Expanding the non-stiff variables in the same way gives

$$\partial_t \mathbf{q} + \partial_{(a)} \left[\mathbf{f}^{(a)}(\mathbf{q}, \bar{\mathbf{q}}_0) + \epsilon \frac{\partial \mathbf{f}^{(a)}(\mathbf{q}, \bar{\mathbf{q}})}{\partial \bar{\mathbf{q}}} \Big|_{\bar{\mathbf{q}}=\bar{\mathbf{q}}_0} \cdot \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2) \right] = \mathbf{s}(\mathbf{q}, \bar{\mathbf{q}}_0) + \epsilon \frac{\partial \mathbf{s}(\mathbf{q}, \bar{\mathbf{q}})}{\partial \bar{\mathbf{q}}} \Big|_{\bar{\mathbf{q}}=\bar{\mathbf{q}}_0} \cdot \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2). \quad (5.17)$$

We will use the notation \mathbf{A}_0 to represent the vector $\mathbf{A}(\mathbf{q}, \bar{\mathbf{q}}_0)$, evaluated at the equilibrium point—therefore, $\partial \mathbf{s}(\mathbf{q}, \bar{\mathbf{q}})/\partial \bar{\mathbf{q}}|_{\bar{\mathbf{q}}=\bar{\mathbf{q}}_0} \equiv \partial \bar{\mathbf{s}}_0/\partial \bar{\mathbf{q}}$. In the limit as $\epsilon \rightarrow 0$, we see again the need for $\bar{\mathbf{s}}(\mathbf{q}, \bar{\mathbf{q}}_0) \equiv \bar{\mathbf{s}}_0 \rightarrow \mathbf{0}$ in Equation (5.16). Keeping only leading order terms, we get an evolution equation for the equilibrium state of the stiff system up to order $\mathcal{O}(\epsilon^0)$:

$$\partial_t \bar{\mathbf{q}}_0 + \partial_{(a)} \bar{\mathbf{f}}_0^{(a)} = \frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \bar{\mathbf{q}}_1. \quad (5.18)$$

We now wish to determine the form for the first order perturbation of the stiff variables, $\bar{\mathbf{q}}_1$. First, we assume that the equilibrium state of the stiff variables can be fully characterised by the non-stiff variables—that is, $\bar{\mathbf{q}}_0 = \bar{\mathbf{q}}_0(\mathbf{q})$. For the case of resistive and ideal MHD, we know this is true as the electric fields can be expressed as $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$. Next, use the fact that

$$\frac{\partial \bar{\mathbf{q}}_0}{\partial t} = \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \frac{\partial \mathbf{q}}{\partial t}. \quad (5.19)$$

Up to zeroth order in ϵ , Equation (5.17) yields

$$\partial_t \mathbf{q} + \partial_{(a)} \mathbf{f}_0^{(a)} = \mathbf{s}_0. \quad (5.20)$$

Now, we can remove the time dependence of the first order perturbation and re-arrange for $\bar{\mathbf{q}}_1$:

$$\bar{\mathbf{q}}_1 = \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left[\frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \left(\mathbf{s}_0 - \frac{\partial \mathbf{f}_0^{(a)}}{\partial x^{(a)}} \right) + \frac{\partial \bar{\mathbf{f}}_0^{(a)}}{\partial x^{(a)}} \right] \quad (5.21)$$

where we note that the non-stiff flux evaluated at the equilibrium point is only a function of the non-stiff variables, $\mathbf{f}(\mathbf{q}, \bar{\mathbf{q}}_0(\mathbf{q})) \equiv \mathbf{f}(\mathbf{q}) \equiv \mathbf{f}$.

With this, we can see the perturbations on top of the equilibrium system. First, we substitute the form for \mathbf{q}_1 into Equation (5.17), ignoring perturbations of order ϵ^2 or higher:

$$\begin{aligned} \frac{\partial \mathbf{q}}{\partial t} + \frac{\partial}{\partial x^{(a)}} \left\{ \mathbf{f}^{(a)} + \epsilon \frac{\partial \mathbf{f}^{(a)}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left[\frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \left(\mathbf{s}_0 - \frac{\partial \mathbf{f}^{(b)}}{\partial x^{(b)}} \right) + \frac{\partial \bar{\mathbf{f}}_0^{(b)}}{\partial x^{(b)}} \right] \right\} = \mathbf{s}_0 + \\ \epsilon \frac{\partial \mathbf{s}_0}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left[\frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \left(\mathbf{s}_0 - \frac{\partial \mathbf{f}^{(a)}}{\partial x^{(a)}} \right) + \frac{\partial \bar{\mathbf{f}}_0^{(a)}}{\partial x^{(a)}} \right], \end{aligned} \quad (5.22)$$

or,

$$\begin{aligned} \frac{\partial \mathbf{q}}{\partial t} + \frac{\partial}{\partial x^{(a)}} \left[\mathbf{f}^{(a)} + \epsilon \frac{\partial \mathbf{f}^{(a)}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \mathbf{s}_0 \right] &= \mathbf{s}_0 + \epsilon \frac{\partial \mathbf{s}_0}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left\{ \left[\frac{\partial \bar{\mathbf{f}}_0^{(a)}}{\partial x^{(a)}} - \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \frac{\partial \mathbf{f}^{(a)}}{\partial x^{(a)}} \right] + \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \mathbf{s}_0 \right\} \\ &\quad - \epsilon \frac{\partial}{\partial x^{(a)}} \left\{ \frac{\partial \mathbf{f}^{(a)}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left[\frac{\partial \bar{\mathbf{f}}_0^{(b)}}{\partial x^{(b)}} - \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \frac{\partial \mathbf{f}^{(b)}}{\partial x^{(b)}} \right] \right\}. \end{aligned} \quad (5.23)$$

We can now rewrite the perturbed system in a more intuitive and compact form:

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial (\mathbf{f}^{(a)} + \hat{\mathbf{f}}^{(a)})}{\partial x^{(a)}} = \mathbf{s}_0 + \hat{\mathbf{s}} + \frac{\partial \hat{\mathbf{D}}^{(a)}}{\partial x^{(a)}}. \quad (5.24)$$

where we have absorbed the timescale, ϵ , into the definition of the stiff source, $\bar{\mathbf{s}}$.

We define the modified flux term, modified source term, and the diffusion-like term as

$$\hat{\mathbf{f}}^{(a)} = \frac{\partial \mathbf{f}^{(a)}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \mathbf{s}_0 \quad (5.25a)$$

$$\hat{\mathbf{s}} = \frac{\partial \mathbf{s}_0}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left[\left(\frac{\partial \bar{\mathbf{f}}_0^{(a)}}{\partial x^{(a)}} - \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \frac{\partial \mathbf{f}^{(a)}}{\partial x^{(a)}} \right) + \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \mathbf{s}_0 \right] \quad (5.25b)$$

$$\hat{\mathbf{D}}^{(a)} = - \frac{\partial \mathbf{f}^{(a)}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} \left[\frac{\partial \bar{\mathbf{f}}_0^{(b)}}{\partial x^{(b)}} - \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \frac{\partial \mathbf{f}^{(b)}}{\partial x^{(b)}} \right]. \quad (5.25c)$$

The hats in Equations (5.25) are to signify that these terms are perturbations on top of the equilibrium system, coming in at order $\mathcal{O}(\epsilon^1)$.

Looking at the form for the iMHD equations (2.2.2), we can see that there is zero source term when we exclude the unphysical field due to the divergence cleaning method—i.e. $\mathbf{s}(\mathbf{q}) \equiv 0$. As a result, the modified flux and source terms in the previous equations reduce to zero, and we are left with only the diffusion term being non-zero. This means the equations of motion of iMHD are modified only by a single diffusion term¹,

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}^{(a)}}{\partial x^{(a)}} = \frac{\partial \hat{\mathbf{D}}^{(a)}}{\partial x^{(a)}}. \quad (5.26)$$

We can see that in the ideal limit, which corresponds to $\{\sigma, \epsilon\} \rightarrow \{\infty, 0\}$, the fact that $\hat{\mathbf{D}} \propto \epsilon$ means Equation (5.26) reduces to the standard form for iMHD, as expected. Then, as the conductivity becomes finite and smaller, larger corrections are made through the diffusive source term to mimic behaviour that should be present in the rMHD model.

Finally, observe how the source term for rMHD is proportional to σ , but that the REGIME source term scales as $\epsilon \propto \sigma^{-1}$. This means that the two models become stiff in opposing limits—near the ideal limit (large σ) REGIME will be stable as a result of a small source term, and will only become stiff, and potentially unstable, as $\sigma \rightarrow 0$. The big benefit of this behaviour is that near the ideal limit we can confidently evolve REGIME with explicit time integrators, knowing that source contributions will remain small. In contrast, in the event of very low conductivities, $\sigma \sim 0$, it will not be sensible to evolve REGIME using implicit schemes. The

¹We should point out that general relativistic iMHD has a non-zero source term, and so this simplification may not be made—see Section 6.1.

numerical flux function appears in the diffusion vector (i.e. $\partial_{(a)} \mathbf{f}^{(a)}$), and as this has a dependence on neighbouring cells, an implicit integration for REGIME would require solving the system for all cells in the domain at once (i.e. an $N_x \times N_y \times N_z \times N_{cons}$ dimensional root-find, where N_{cons} is the size of the conserved vector). Regardless, near this limit rMHD is likely to be stable using explicit integrators and conventional CFL conditions. We present a hybrid of these two models to take advantage of this feature in Chapter 7.

5.3 Numerical implementation

Now we have the form for the source term, our attention turns to implementing it numerically. The interpretation of a term in a fluid's equations of motion may be intuited by the order of the spatial derivative of the conserved fields. For example, advection is noted by a first order derivative, diffusion by a second order and dispersion via a third order derivative. This is why we call the new piece in Equation (5.26) a diffusion-like term, as it can be re-written to include second order spatial derivatives of the conserved fields:

$$\frac{\partial \mathbf{q}}{\partial x} + \frac{\partial \mathbf{f}^{(a)}}{\partial x^{(a)}} = -\frac{\partial}{\partial x} \left[\frac{\partial \mathbf{f}^{(a)}}{\partial \mathbf{q}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{q}} \right)^{-1} \left(\frac{\partial \bar{\mathbf{f}}_0^{(b)}}{\partial \mathbf{q}} - \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}} \frac{\partial \mathbf{f}^{(b)}}{\partial \mathbf{q}} \right) \frac{\partial \mathbf{q}}{\partial x^{(b)}} \right], \quad (5.27)$$

recalling that both $\hat{\mathbf{f}}^{(a)} = \hat{\mathbf{s}} = \mathbf{0}$ when considering ideal and resistive MHD in the special relativistic limit.

In order to implement this new piece, it is useful to understand more about what it involves. If we look that the diffusion vector given in Equation (5.25c), the first term is the Jacobian of the non-stiff system with respect to the stiff variables, $\partial_{\mathbf{q}} \mathbf{f}^{(a)}$, and is therefore a matrix operation—as, indeed, is the second and fourth term. The third and fifth terms correspond to the spatial derivatives of the stiff and non-stiff flux vectors. In fact, the form of the diffusion term in Equation (5.25c) is preferable to that in Equation (5.27), as any numerical flux function that computes $\mathcal{F}(\mathbf{q}) = -\partial_{(a)} \mathbf{f}^{(a)}(\mathbf{q})$ can be re-used *as is* for the generation of the source term.

All that is required then is to determine the forms of the two matrices,

$$\underline{\underline{M}}_1^{(a)} = \frac{\partial \mathbf{f}^{(a)}}{\partial \mathbf{q}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{q}} \right)^{-1}, \quad (5.28a)$$

$$\underline{\underline{M}}_2 = \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{q}}. \quad (5.28b)$$

Differentiating with respect to the conserved variables is challenging, as we do not know how to express the flux and source vectors in terms of only the conserved quantities. We can, however, express the conserved, flux and source vectors in terms of the primitive variables, and thus we can differentiate using the primitive quantities. For example,

$$\frac{\partial \mathbf{f}}{\partial \mathbf{q}} = \frac{\partial \mathbf{f}}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \mathbf{q}}, \quad (5.29)$$

with \mathbf{w} as the vector of primitive variables. In this way, the new form for the M matrices is

$$\underline{\underline{M}}_1^{(a)} = \frac{\partial \mathbf{f}^{(a)}}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \bar{\mathbf{q}}} \right)^{-1} \quad (5.30a)$$

$$= \frac{\partial \mathbf{f}^{(a)}}{\partial \mathbf{w}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{w}} \right)^+ \quad (5.30b)$$

$$\underline{\underline{M}}_2 = \frac{\partial \bar{\mathbf{q}}_0}{\partial \mathbf{w}} \left(\frac{\partial \mathbf{q}}{\partial \mathbf{w}} \right)^+, \quad (5.30c)$$

Here, we have used the superscript $+$ to denote the Moore-Penrose pseudoinverse (Barata and Hussein, 2012). As the length of the vectors $\bar{\mathbf{s}}_0$ and \mathbf{w} are, in general, not the same, the resulting matrix will not be square, nor have a corresponding inverse. As a result, we use the definition of the right-pseudoinverse of a matrix Q as, $Q^+ = Q^T(QQ^T)^{-1}$. We will use the term inverse to refer to the pseudoinverse henceforth.

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 and hence will slow 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 iMHD with the intention of being faster to evolve than rMHD, it is sensible to adopt the performance gains of a purely symbolic source term, and perhaps sacrifice on accuracy.

In order to generate human readable terms for Equations (5.30), we have had to make a number of assumptions. Firstly, we will assume a low velocity limit, in which terms of $\mathcal{O}(v^2)$ can be ignored. For instance, with this approximation in rMHD, the y -momentum flux in the x -direction reduces to $S_{12} = -E_x E_y - B_x B_y$. Second, we assume that the fluid only couples weakly to the magnetic field, and to enforce this we set the electric and magnetic fields to zero. This is only necessary to reduce the number of terms in the inverted matrices, and so for un-inverted matrices we do not make this assumption. This means that electric and magnetic field terms do still make it through to the final source term. As a check of these assumptions and their effect on the system we also implemented REGIME numerically with no approximations. After doing this, we got virtually no difference in the simulation output for mildly relativistic flows, $v \sim 0.5c$.

Making these simplifications, we can use a symbolic algebra program to compute the form for the source term. For this task, we used Wolfram Mathematica v11.2 (Wolfram Research, 2018). The full notebooks, along with greater detail, have been made available through the GitHub page².

For the systems of equations presented here, namely ideal and resistive MHD in the special relativistic limit and with the above assumptions made, we get the simple result that $\underline{\underline{M}}_2 = 0$. The reason for this being that $\partial \bar{\mathbf{q}}_0 / \partial \mathbf{w}$ is a 3×12 matrix in which only three entries are non-zero—these correspond to the $\partial E_x / \partial E_x$, $\partial E_y / \partial E_y$ and $\partial E_z / \partial E_z$ terms. These three terms are then dotted with $\partial \mathbf{E} / \partial \mathbf{S}$, which on account of the weak coupling approximation are all zero.

²See ‘REGIMEExtensionTerms.nb’ at www.github.com/AlexJamesWright/METHOD/tree/master/PubSrc

With regards to the $\underline{\underline{M}}_1^{(a)}$ matrix, we get

$$\frac{\partial \mathbf{f}^{(1)}}{\partial \mathbf{w}} = \begin{bmatrix} v_x & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -B_x & B_y & B_z & -E_x & E_y & E_z & 0 \\ 0 & 0 & 0 & 0 & 0 & -B_y & -B_x & 0 & -E_y & -E_x & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -B_z & 0 & -B_x & -E_z & 0 & -E_x & 0 \\ 0 & \frac{\Gamma p}{\Gamma-1} & 0 & 0 & \frac{\Gamma v_x}{\Gamma-1} & 0 & -E_z & E_y & 0 & B_z & -B_y & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}, \quad (5.31a)$$

$$\frac{\partial \mathbf{f}^{(2)}}{\partial \mathbf{w}} = \begin{bmatrix} v_y & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -B_y & -B_x & 0 & -E_y & -E_x & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & B_x & -B_y & B_z & E_x & -E_y & E_z & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -B_z & -B_y & 0 & -E_z & -E_y & 0 \\ 0 & 0 & \frac{\Gamma p}{\Gamma-1} & 0 & \frac{\Gamma v_y}{\Gamma-1} & E_z & 0 & -E_x & -B_z & 0 & B_x & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \end{bmatrix}, \quad (5.31b)$$

$$\frac{\partial \mathbf{f}^{(3)}}{\partial \mathbf{w}} = \begin{bmatrix} v_z & 0 & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -B_z & 0 & -B_x & -E_z & 0 & -E_x & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -B_z & -B_y & 0 & -E_z & -E_y & 0 \\ 0 & 0 & 0 & 0 & 1 & B_x & B_y & -B_z & E_x & E_y & -E_z & 0 \\ 0 & 0 & 0 & \frac{\Gamma p}{\Gamma-1} & \frac{\Gamma v_z}{\Gamma-1} & -E_y & E_x & 0 & B_y & -B_x & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad (5.31c)$$

$$\left(\frac{\partial \bar{\mathbf{s}}}{\partial \mathbf{w}} \right)^{+T} = \alpha \begin{bmatrix} 0 & \mathcal{A}_{11} & \mathcal{A}_{21} & \mathcal{A}_{31} & 0 & \mathcal{B}_{11} & \mathcal{B}_{21} & \mathcal{B}_{31} & \mathcal{C}_{11} & \mathcal{C}_{21} & \mathcal{C}_{31} & \mathcal{D}_1 \\ 0 & \mathcal{A}_{12} & \mathcal{A}_{22} & \mathcal{A}_{32} & 0 & \mathcal{B}_{12} & \mathcal{B}_{22} & \mathcal{B}_{32} & \mathcal{C}_{12} & \mathcal{C}_{22} & \mathcal{C}_{32} & \mathcal{D}_2 \\ 0 & \mathcal{A}_{13} & \mathcal{A}_{23} & \mathcal{A}_{33} & 0 & \mathcal{B}_{13} & \mathcal{B}_{23} & \mathcal{B}_{33} & \mathcal{C}_{13} & \mathcal{C}_{23} & \mathcal{C}_{33} & \mathcal{D}_3 \end{bmatrix}, \quad (5.31d)$$

where $(\partial_{\mathbf{w}} \bar{\mathbf{s}})^{+T}$ is the transpose of $(\partial_{\mathbf{w}} \bar{\mathbf{s}})^+$, and we define the elements of the matrices \mathcal{A} , \mathcal{B} and \mathcal{C} as

$$\mathcal{A}_{ij} = -\varrho \sigma^2 B_i B_j + \epsilon_{ijk} B^k (\varrho^2 + \sigma^2) \sigma - \delta_{ij} \varrho (\varrho^2 + \sigma^2), \quad (5.32)$$

$$\mathcal{B}_{ij} = B_j \sigma^3 \epsilon_{ikl} v^k B^l - \epsilon_{ijk} v^k (\varrho^2 + \sigma^2) \sigma, \quad (5.33)$$

$$\mathcal{C}_{ij} = -\sigma^3 B_i B_j - \delta_{ij} \sigma (\varrho^2 + \sigma^2), \quad (5.34)$$

the vector \mathcal{D} as

$$\mathcal{D}_i = -v_i (\varrho^2 + \sigma^2) - \sigma^2 B_i v_k B_k, \quad (5.35)$$

and the pre-factor, α , as

$$\alpha = \frac{1}{(\varrho^2 + \sigma^2)(\varrho^2 + (1 + B^2)\sigma^2)}. \quad (5.36)$$

The pre-factor, α , in Equation (5.31d) now acts in a similar way to the previous timescale, ϵ . That is, in the ideal limit, $\alpha \propto \sigma^{-4} \rightarrow 0$, and the source term tends to zero, recovering iMHD. For large but finite conductivities however, the source term will begin to have more of an effect on the system.

All that remains is to compute the diffusion vector is the $\partial_{(a)} \bar{\mathbf{f}}_0^{(a)}$ term. This can be done using any existing numerical flux function, which in the case of `METHOD` is flux vector splitting using a WENO3 reconstruction (Shu, 1997), see Chapter 3 and Section 2.5. To determine the derivative of the diffusion vector we employ second-order central differencing.

It should also be noted that in Equation (5.31d), the charge density, ϱ , appears a number of times. As this term is not explicitly evolved in iMHD, we compute its value from Gauss' law, $\nabla \cdot \mathbf{E} = \varrho$. In this calculation, we assume the ideal form for the electric fields, namely that $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$, which is the correct form up to $\mathcal{O}(\frac{1}{\sigma}^0)$. As we are only constructing the extension to be accurate up to first order, $\mathcal{O}(\frac{1}{\sigma}^1)$, this is sufficient.

5.4 Stability criterion

The source term we have derived is diffusion-like, containing a second order spatial derivative. Whilst the original equilibrium system is in a strongly hyperbolic form, the addition we have made would be parabolic on its own, i.e. if there were no flux. As a result, we must first understand the limits on the spatio-temporal resolution such that the new system is stable. The following analysis will be limited to one dimension, which we then generalise at the end of this section.

Recall our system has the following form:

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}}{\partial x} = \frac{\partial}{\partial x} \left(-\underline{\underline{M}}_1 \cdot \frac{\partial \bar{\mathbf{f}}_0}{\partial x} \right). \quad (5.37)$$

In order to proceed, we will make the assumption that in a small region of the domain, $\underline{\underline{M}}_1$ is approximately constant. We will also rewrite the stiff flux derivative using the Jacobian, and further assume that it, too, is constant in this region. Finally, if we assume that the resolution constraints set by the hyperbolic system are met, and that we are only interested in the second derivative term, then we can simplify and say

$$\frac{\partial \mathbf{q}}{\partial t} = \underline{\underline{K}} \cdot \frac{\partial^2 \mathbf{q}}{\partial x^2}, \quad (5.38)$$

Here, we make analogy with the scalar case the diffusion equation, $\partial_t q = k \partial_x^2 q$, and the analysis in Section 5.1.1, in which an explicit scheme is generally considered stable if

$$k \frac{\Delta t}{\Delta x^2} \leq \frac{1}{2}, \quad (5.39)$$

in which Δt and Δx are the size of the timestep and spatial resolution for a given simulation.

For the system case, we suggest that k be given by the largest magnitude eigenvalue of $\underline{\underline{K}}$. This, therefore, requires an eigenanalysis of the matrix

$$\underline{\underline{K}} = -\frac{\partial \mathbf{f}}{\partial \mathbf{w}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{w}} \right)^+ \frac{\partial \bar{\mathbf{f}}_0}{\partial \mathbf{w}} \left(\frac{\partial \mathbf{q}}{\partial \mathbf{w}} \right)^+. \quad (5.40)$$

With the help of Mathematica, it can be shown that the largest eigenvalue is given by

$$\lambda_{\max} = \frac{\sigma(\varrho + \sigma^2[1 + B_y^2 + B_z^2])}{(\varrho^2 + \sigma^2)(\varrho^2 + \sigma^2[1 + B^2])} \sim \sigma^{-1}. \quad (5.41)$$

This result suggests that, assuming a CFL constraint given by $\Delta t/\Delta x = \nu$, the spatial resolution of any simulation using this source term must (roughly) satisfy

$$\Delta x \gtrsim \frac{2\delta\nu}{\sigma}, \quad (5.42)$$

where δ is some softening factor which we discuss in Section 5.6. Intuitively, this means there is a limit on how well a simulation may be resolved. For high conductivities, this limit is not an issue as the RHS tends to zero, allowing (almost) arbitrarily fine resolutions. For resistive simulations, however, there is a finest grid resolution for which the system is numerically stable (for a given Courant factor), as the source term begins to become large compared to the hyperbolic part. Of course, one may always reduce the CFL condition and take smaller timesteps in order to achieve the highest resolutions.

Higher dimensions

With regards to the stability analysis in three dimensions, Equation (5.38) becomes,

$$\frac{\partial \mathbf{q}}{\partial t} = \underline{\underline{K}}^{(ab)} \cdot \frac{\partial^2 \mathbf{q}}{\partial x^{(a)} \partial x^{(b)}}, \quad (5.43)$$

where there are now nine matrices,

$$\underline{\underline{K}}^{(ab)} = -\frac{\partial \mathbf{f}^{(a)}}{\partial \mathbf{w}} \left(\frac{\partial \bar{\mathbf{s}}}{\partial \mathbf{w}} \right)^+ \frac{\partial \bar{\mathbf{f}}^{(b)}}{\partial \mathbf{w}} \left(\frac{\partial \mathbf{q}}{\partial \mathbf{w}} \right)^+. \quad (5.44)$$

There are now cross terms in the derivative, but to simplify the analysis we once again take the largest eigenvalue of all K^{ab} to be the value of k as in Equation (5.40). The largest eigenvalue is the same as in the previous section, and so we assume the same form for the resolution criterion, Equation (5.42). As this is an order of magnitude estimation of the stability, we assume that any effects due to the additional terms are captured in δ .

5.5 Convergence with σ

In this section, we will see how the proposed extension behaves for various values of the conductivity. To implement it, we extend the iMHD model in `METHOD` with the source term of Equation (5.26), and use the second order, explicit, operator split, RKSplit integrator mentioned in algorithm A.2, Section A.2.1.

For the convergence plots presented here, we will consider the results of the rMHD model to be the *exact* solution. We now present the results of a number of well-known initial data, and the solutions of rMHD and REGIME are compared and contrasted. For clarity, we have also evolved some simulations using iMHD to highlight its inability in modelling various resistive behaviours, and to show how the REGIME extension to it can capture these behaviours. The initial data for the Brio-Wu shock tube, self-similar current sheet and KHI can be found in Section 3.1.

Brio-Wu shock tube




FIGURE 5.1: Final state of the Brio-Wu test problem for a range of conductivities. Top row is the magnetic field in the y -direction, followed by the hydrodynamic pressure, density, and x -velocity. Columns from left to right correspond to conductivities of $\sigma = (50, 100, 200, 400)$ respectively.




FIGURE 5.2: The L_2 error norm for the Brio-Wu test of iMHD and REGIME, using rMHD as the *exact* solution. The error of both models grows as the system becomes more resistive, however for all but the velocity, the REGIME error grows many factors more slowly. The increased error of REGIME in the velocity plot is relatively small in magnitude (compared to the magnetic field error, for example).

Figure (5.1) shows the final state for iMHD (blue solid), rMHD (black dashed) and REGIME (red solid), for varying conductivities. Clearly for the larger conductivities, there is little difference in the output of the models, which is to be expected as both rMHD and REGIME limit to iMHD. For more resistive simulations (smaller σ), the resistive models differ more greatly from the iMHD results, and REGIME and rMHD results are in excellent agreement.

We now wish to quantify the error between REGIME and rMHD. To do this, we calculate the error as the L_2 norm of the difference between the model output and rMHD, and see how this norm varies with changes in σ , Figure 5.2.

Clearly, by looking at the final magnetic field and pressure, there is a large improvement in the accuracy of REGIME—the diffusion term captures the important features in both. The magnetic field and pressure are the variables for which there is the largest change in behaviour when using a finite conductivity. The electrical resistivity acts to smear out the y -direction magnetic field and the shock present in the pressure at $x \approx 0.6$. The density plot is mostly unchanged, and as a result there is good agreement with all models. The velocity is the only

variable in which the error is greater when using REGIME, although the magnitude of this error is small compared to the improvements in the pressure and magnetic field.

Self-similar current sheet




FIGURE 5.3: Final state comparison of the y -direction magnetic field for the self-similar current sheet problem. Larger resistivities lead to a greater rate of magnetic diffusion and thus greater smearing out compared to the iMHD solution. The REGIME source term captures this diffusion exceptionally well, even for very small conductivities.

As the form of the proposed source term looks much like a diffusion term, we now present a problem in which the solution obeys the diffusion equation. An ideal simulation of the current sheet problem should not evolve from the initial condition, which given an infinite resolution would be a step function. Simulations with greater resistivities should expect more diffusion, and so the solution should appear more smeared out. Clearly from Figure 5.3, both rMHD and REGIME capture this process, with excellent agreement between both down to conductivities of only $\sigma = 20$. The error plot, Figure 5.4, shows only minimal error growth between rMHD and REGIME.

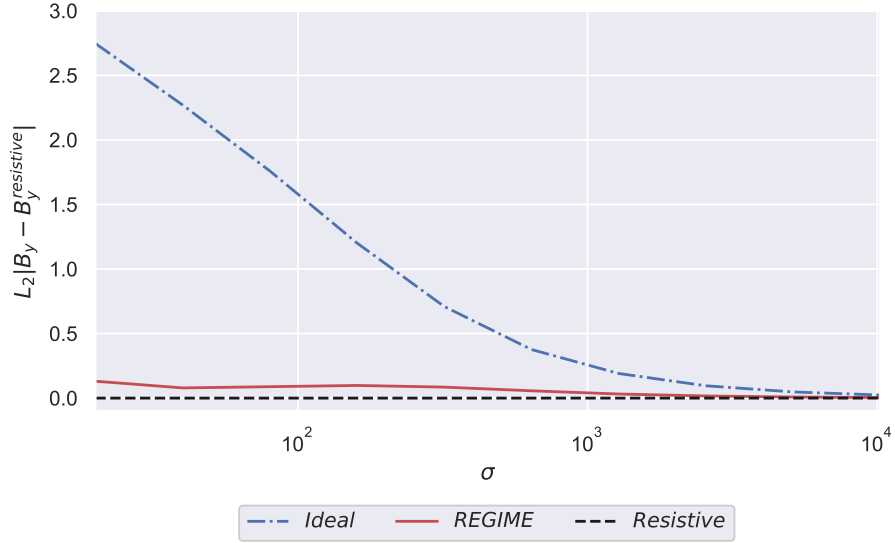


FIGURE 5.4: Error growth for the self-similar current sheet problem. The error is calculated as the L_2 -norm of the difference of the y -direction magnetic fields, using rMHD as the *exact* solution. The difference between rMHD and REGIME is marginal down to $\sigma = 20$, whereas iMHD quickly deviates from the exact solution due to the lack of any resistive diffusion.

Resistive magnetic reconnection

Next, to demonstrate the effectiveness of REGIME in higher dimensions, we present the 2D magnetic reconnection problem. Similarly to the current sheet problem, magnetic reconnection is a purely resistive phenomenon—the lack of electrical resistivity in iMHD should mean that there is no evolution from the initial condition (excluding effects due to numerical diffusion). Furthermore, this kind of phenomenon should occur extensively during simulations of neutron star mergers, due to the complex interaction between the progenitors’ intense magnetic fields, and so makes a useful test of REGIME in its applicability to these simulations.

Reconnection occurs when two magnetic fields whose vectors are not parallel intersect, forming a new magnetic vector and releasing stored magnetic energy into thermal and kinetic energy. There are multiple configurations of magnetic fields that will lead to some form of reconnection—indeed, understanding magnetic reconnection is so vital for solar physics that there is extensive literature and multiple models to choose from. As we are only interested in demonstrating the effect, we will adopt the first and simplest attempt of understanding reconnection, the Sweet-Parker model (Sweet, 1958; Parker, 1957). One should note, however, that the reconnection rates that the Sweet-Parker model predicts are slow compared to alternative models and experimental data. In fact, Loureiro and Uzdensky (2015) claim that “unsatisfactory predictions are obtained for reconnection events in almost all plasmas that one cares to examine.”

In the Sweet-Parker model, the reconnection layer is defined to have a thickness of 2λ , and a length of $2L$. The initial magnetic field strength is given by B_0 , with the corresponding Alfvén velocity as V_A . With this, it is possible to estimate the rate of magnetic reconnection, $\mathcal{E} \equiv u_{\text{in}}/u_{\text{out}}$. By conservation-of-mass arguments, we have the relation

$$\frac{u_{\text{in}}}{u_{\text{out}}} = \frac{\lambda}{L}, \quad (5.45)$$

where $u_{\text{in/out}}$ is the velocity of the fluid flowing in/out of the boundary layer. Using Ohm's law, one can show that the rate of flow in to the boundary layer can be expressed as $u_{\text{in}} = \eta/\lambda$, in which the resistivity is $\eta = \sigma^{-1}$. The Lundquist number, S , is defined via $S \equiv LV_A/\eta$ and so $\eta/V_A = L/S$. Using the fact that $u_{\text{out}} = V_A$, one can show that,

$$\mathcal{E} \equiv \frac{u_{\text{in}}}{u_{\text{out}}} = \frac{\lambda}{L} \propto \eta^{1/2}. \quad (5.46)$$

The domain, where $x \in [-12.8, 12.8]$ and $y \in [-6.4, 6.4]$, is set up as follows:

$$p = 0.5 \quad (5.47)$$

$$\rho = \rho_\infty + \rho_0 \cosh^2(y/\lambda) \quad (5.48)$$

$$B_x = B_0 \tanh(y/\lambda), \quad (5.49)$$

with all other variables set to zero. The magnetic fields are then perturbed by,

$$\delta B_x = -\frac{\pi\psi_0}{L_y} \sin(\pi y/L_y) \cos(2\pi x/L_x) \quad (5.50)$$

$$\delta B_y = \frac{2\pi\psi_0}{L_x} \sin(2\pi x/L_x) \cos(\pi y/L_y). \quad (5.51)$$

Here, we have $(L_x, L_y) = (25.6, 12.8)$, with the initial thickness of the boundary layer $\lambda = 0.5$. The background density is $\rho_\infty = 0.2$ with $\rho_0 = 1.0$, a magnetic field strength of $B_0 = 1.0$, and a perturbation size of $\psi_0 = 0.1$. The simulation is then run using 1024×512 cells until $T = 100$, with the adiabatic index as $\Gamma = 2$. We use periodic boundaries along the x -axis and outflow boundaries along the y -axis.

In Figure 5.5, we show the final state of the hydrodynamic pressure for a range of conductivities for both rMHD and REGIME. The hotspots in the pressure coincide with the formation of two magnetic islands. Clearly, for the simulations with greater conductivity, there is significantly less diffusion occurring—this is seen as a more distinct separation of the two magnetic islands. As the resistivity of the simulations is increased, this separation reduces more quickly.

The same behaviour can be asserted from Figures 5.6, where we instead plot snapshots of the time-series of magnetic field strength for a low, $\sigma = 100$ and a high, $\sigma = 1000$, conductivity. A comparison of the evolution of these two simulations confirms that the magnetic fields, too, show more diffusion on account of the greater resistivity.

With these plots, one can see no differences in the output of rMHD and REGIME. The proposed source term of REGIME captures the diffusive behaviour of the full model exceptionally well.

Next, to determine the rate of diffusion for a given resistivity for both models, we compute the current density from Ampere's law, $\mathbf{J} = \nabla \times \mathbf{B}$, at the line $x = 0$. As the rate of reconnection can be written $\mathcal{E} = \lambda/L$, where $L = L_x$ is the length of the domain in the x -direction, we compute the width of the boundary layer, λ , as the width of the Gaussian profile that best approximates \mathbf{J} at $T = 100.0$.

As we can see from the Gaussian fits, Figure 5.7, the fit for $\sigma = 100$ is poor, and so has been excluded from the reconnection rate analysis. The results of this analysis are shown in Figure 5.8

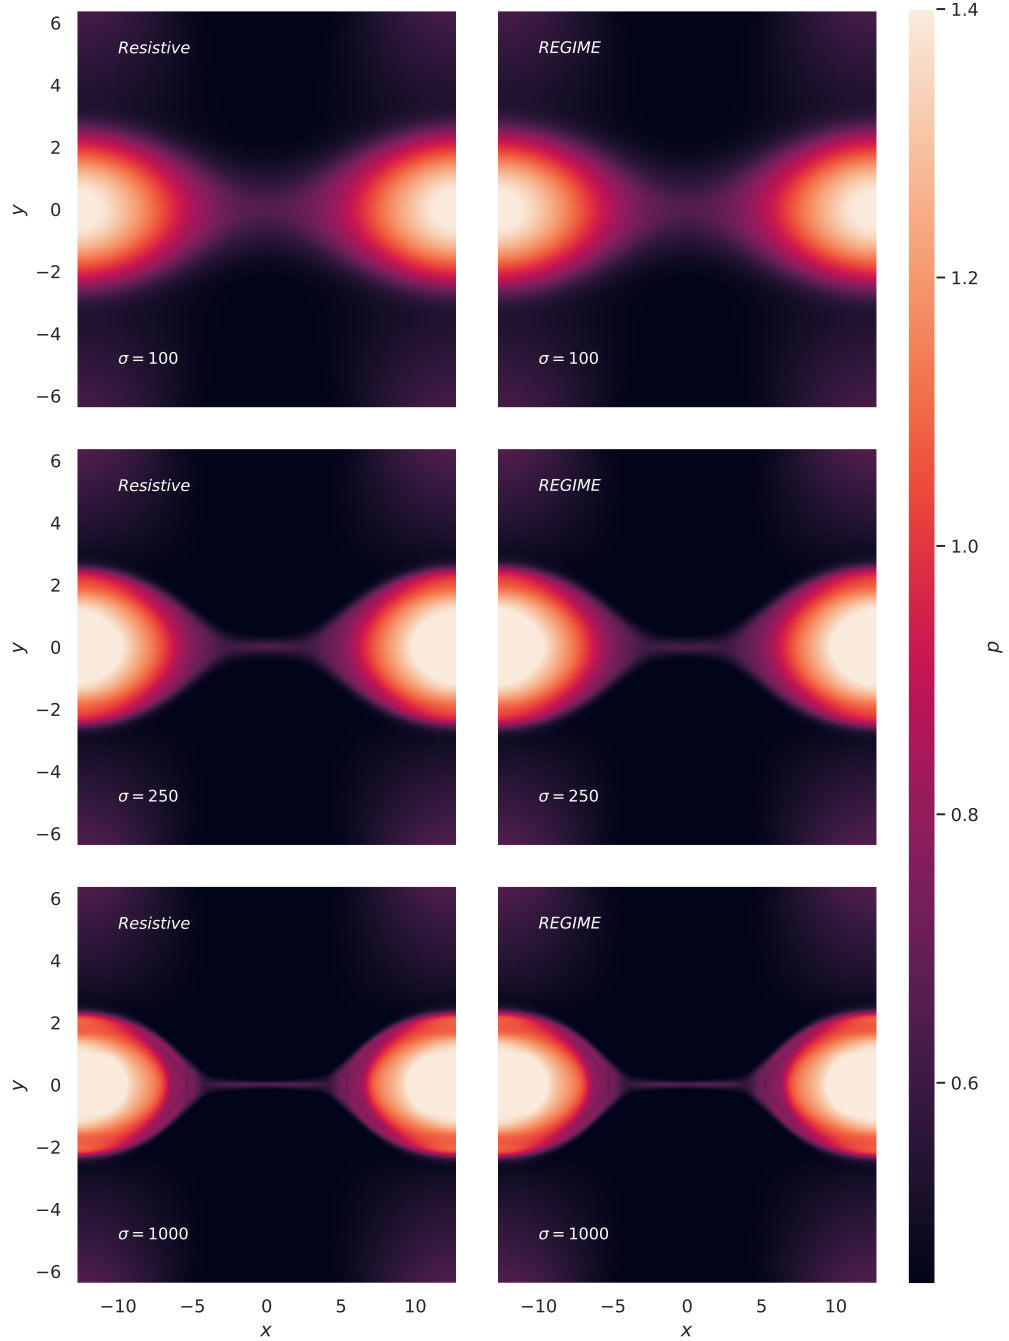



FIGURE 5.5: Final state, at $T = 100$, of the pressure for the magnetic reconnection problem and a range of conductivities. The bright spots correspond to the location of the magnetic islands. We can see clearly that the more resistive systems exhibit significantly more diffusion.

Once again, rMHD and REGIME show indistinguishable results.

for a range of conductivities, and shows a scaling slightly greater than what is expected from the Sweet-Parker analysis. Clearly REGIME captures this event, and as a result the reconnection rate of rMHD and REGIME agree well.



(a) Time-series magnetic field for $\sigma = 10^3$. Different snapshots are taken (top to bottom) and shown for rMHD (left) and REGIME (right).


(b) Same as the above figure, but for $\sigma = 10^2$.

FIGURE 5.6: Strength of the magnetic field, B^2 , for the magnetic reconnection problem at increasing time (moving downwards). The more conductive system, 5.6(a), has a very slow rate of reconnection compared to a more resistive set-up, 5.6(b). In either case, the results of rMHD and REGIME are in agreement. For each sub-figure, the left column corresponds to rMHD and the right to REGIME.

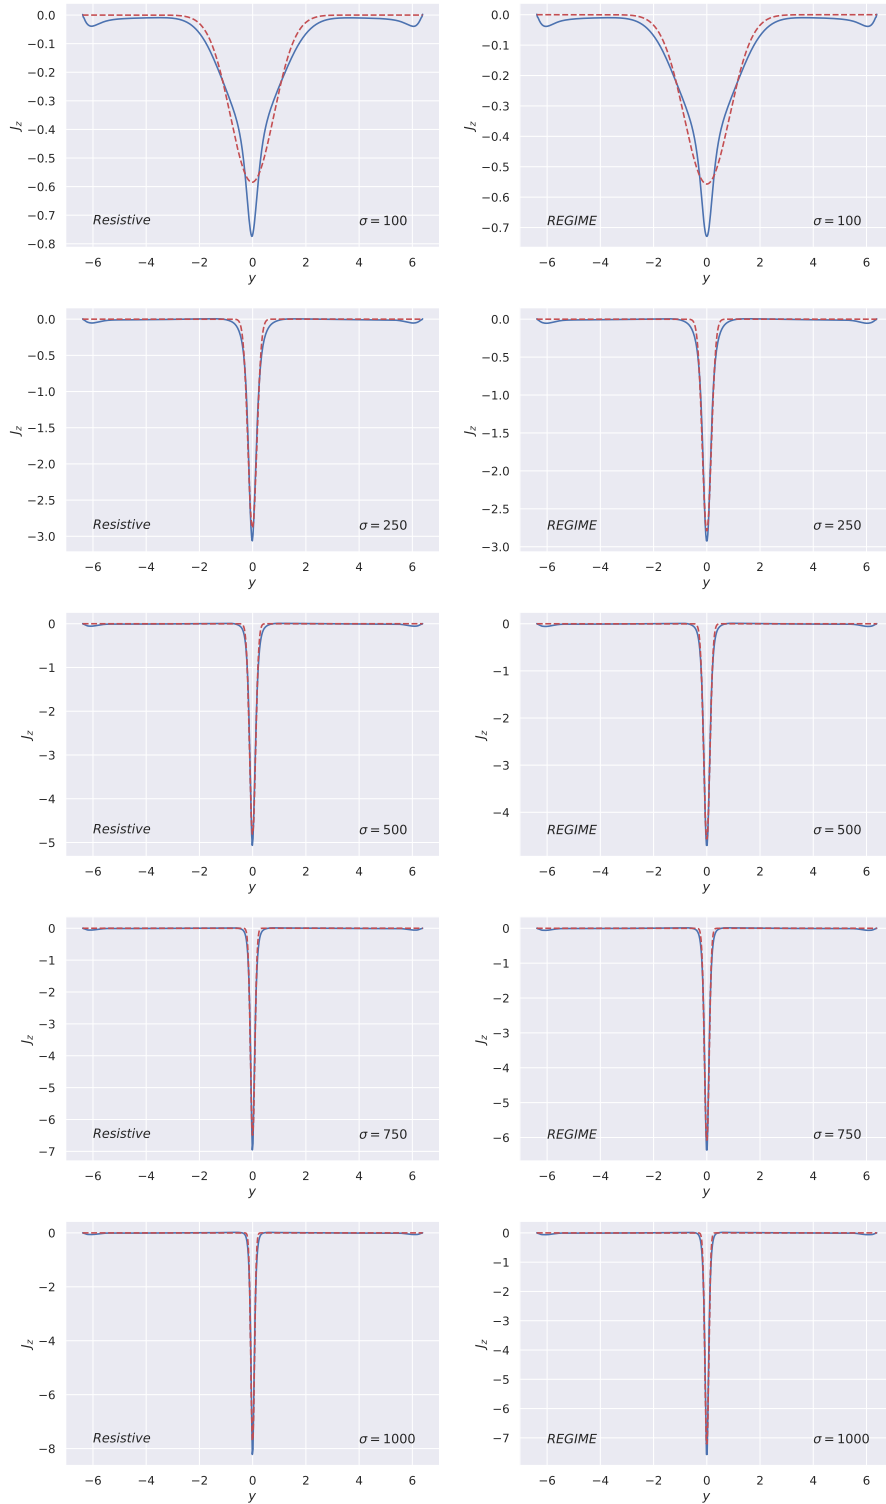


FIGURE 5.7: Plots of the current density at $T = 100$ for both models and a range of conductivities. The solid line is the numerical solution, and the dashed line is the best fit Gaussian. The width of the Gaussian profile gives 2λ .




FIGURE 5.8: Using the width of the best fit Gaussian, we plot here the reconnection rate for a range of resistivities for the rMHD model (left) and REGIME (right). The gradient is highlighted in the legend, and shows the models' reconnection rates agree well with each other and the Sweet-Parker prediction.

Kelvin-Helmholtz instability

We now perform the same KHI simulations as in Section 3.1.5. We use a conductivity of $\sigma = 10^2$ for these tests such that there will be a difference between ideal and resistive solutions. This is a more demanding test of REGIME on account of the range of behaviour that occurs across all scales of the simulation. For these simulations we use a resolution of 512×1024 .




FIGURE 5.9: Final state, $T = 6.0$, of the density for the Kelvin-Helmholtz instability for each model. Visually, there is almost no difference between the models, which is to be expected due to the weak magnetic field.

From Figure 5.9 we can see that there are no visual differences between the models' densities, as expected. If instead we turn our attention to the evolution of the magnetic fields, Figure 5.10, and perform the same analysis as in Section 3.1.5.2 (note: a different conductivity to Section 3.1.5.2 is being used), we can see that for the large scale behaviour, REGIME captures the magnetic evolution strikingly well. During the linear growth of the instability, both the maximum magnetic field strength and the average of the magnetic energy density of REGIME and the full resistive MHD model correlate closely.




FIGURE 5.10: Evolution of the magnetic fields for the KHI. Top row is the average magnetic energy density, bottom row is the peak magnetic field strength.

To understand how REGIME behaves across all scales of the simulation, we apply the power spectrum analysis used in Section 3.1.5.2. Here is where we begin to see some noticeable differences between the output of rMHD and REGIME. Figure 5.11 shows the kinetic energy density, rest-mass density and magnetic energy density power spectrum for the three models. Results for the kinetic energy density and rest-mass density power spectra are virtually indistinguishable as was the case in Section 3.1.5. The spectrum for the magnetic energy density, however, begins to show differences for high modes.

The cause for this disparity on small scales is ultimately unknown, however we can say that the behaviour of REGIME in this test case is sensitive to the choice of numerical methods used. When taking the derivative of the diffusion vector, Equation (5.26), we use a second order central differencing scheme. In previous iterations, however, MINMOD slope limiting (Roe, 1986) was employed with the hope of minimising the onset of Gibbs oscillations near discontinuous data and improving stability. Whilst the stability was improved for simple behaviour, such as Brio-Wu shock tubes, it led to a very large amplification of the magnetic fields in the KHI. Whilst the effects of alternative slope limiting methods have not yet been investigated, the results that are generated using the central differencing scheme are sufficient for our purposes.




FIGURE 5.11: Results for the KHI power spectrum at $T=3$. Top row is the kinetic energy density power, middle is rest-mass density power, and bottom is the magnetic energy density power. REGIME has some trouble accurately emulating rMHD on very small scales, but approximates it on large scales.

5.6 Stability

We now turn our attention to the stability of the new model. In Section 5.4, we argue that there should be some relationship between the resolution and the system's stability, and that the result should be some maximum possible spatial resolution whilst still giving stable evolutions.

Returning to the Brio-Wu shock tube test, Section 3.1.2, we can see how quickly the simulation can move from a stable region to an unstable region. Figure 5.12 shows the final state of the density for the Brio-Wu shock tube, where $\sigma = 1000$, for four different resolutions. The first two (top row) use $N_x = 1000$ and 1498 cells, leading to stable simulations. Adding only




FIGURE 5.12: Final state of the density for the Brio-Wu shock-tube problem with increasing resolution. For simulations with $N_x > 1498$ the stability criterion is violated and the evolution becomes unstable.

two additional cells, however, results in an unstable simulation, as can be seen by the spikes present in the data at $x \approx 0.2$. If we continue to use finer resolutions, the effect becomes more pronounced, and additional artefacts appear.

When performing the stability analysis, we assumed that in moving to higher dimensions it was the same largest eigenvalue that is the main factor in determining stability. To test this, we perform a number of simulations in all dimensions. The initial data for the Orszag-Tang vortex is taken from [Beckwith and Stone \(2011\)](#) for 2D, with the addition of a z -perturbation to the fluid velocity of the form $0.2\sin(2\pi z)$ for the three dimensional case.

Figure 5.13 shows the relationship between the finest possible resolution for a range of conductivities. The one-dimensional Brio-Wu and current sheet problems produce an average $\delta \approx 0.92$ and 0.93 , respectively. When we move to higher dimensions, the δ value for all problems drops to a value of $\delta \approx 0.3 - 0.4$, corresponding to a less restrictive stability criterion.

Whilst the reasoning behind this drop in δ is not fully understood, we believe it comes from the numerics. The fifth plot (bottom left) shows the data for a 2D Brio-Wu test, in which the standard 1D data is rotated about the z -axis by $\pi/4$. If the stability was set by the 1D nature of the initial data, a rotation about the z -axis should not change the resolution that is possible along the propagating shock wave, and we would expect the same softening factor. In reality, the rotated problem allows for higher resolutions to be achieved, possibly due to the effects of the cross derivatives in the REGIME source term.




FIGURE 5.13: Stability criterion for 1D, 2D and 3D test problems. Points represent the resolution of the last stable simulation for the indicated initial data, with the gradient highlighted in the legend, and the corresponding value for the softening parameter, δ . This suggests that the criterion, Equation 5.42, implies stability when using a softening factor of $\delta = 1$.

Whatever the reason, the effect is to allow a greater resolution in the multi-dimensional case, tending towards a factor $3\times$ for large conductivities. If we therefore take a conservative value for the softening factor of $\delta = 1$, we can determine whether the types of astrophysical simulations we are interested in are possible with REGIME. The conductivities one would expect in a warm neutron star crust have values of $\sigma \gtrsim 10^{16}\text{s}^{-1}$ (Harutyunyan and Sedrakian, 2016). Calculations of the minimum conductivity that would be possible with typical resolutions used in merger simulations give $\sigma_{\min} \sim 10$ in *Whisky* code units, or $2 \cdot 10^6\text{s}^{-1}$. As a result, it appears that REGIME will be suitable for realistic neutron star simulations, and that for physical magnitudes of the conductivity it should be stable, even for the finest resolutions to date.

5.7 Performance

Knowing that REGIME gives sensible results for a range of simulations, and agrees well with rMHD in various regimes, we now turn our attention to the relative performance of the two models. The motivation for the model is that rMHD simulations generally require implicit time-integration schemes to ensure stable evolutions, and as a result suffer from slow evolution especially when near the iMHD limit (large σ). By instead using the REGIME source term, which is small in this limit compared to the rMHD source term, it is possible to perform the evolution using explicit time integrators, and ultimately simulations should benefit in terms of performance.

To determine this, we once again return to the Brio-Wu test, and to the magnetic reconnection problem. The Brio-Wu shock-tube is run using $N_x = 400$ cells until $T = 0.4$. The 2D reconnection problem uses $(N_x, N_y) = (128, 64)$ and run until $T = 2.0$. Due to time restraints, the 3D reconnection problem uses $(N_x, N_y, N_z) = (128, 64, 64)$, and is run until $T = 0.5$ ³. The

³It may be possible that these simulations are ultimately unstable at later times than are evolved here, and so the optimum execution time may change, however this would be true of all models and so should still be representative of the true, relative performance of the models




FIGURE 5.14: The optimum execution time for the Brio-Wu and magnetic reconnection (2D and 3D) test problems for both models. The red line corresponds to REGIME using the RK split integrator, and the blue and green line represent rMHD with the RK split and IMEX integrators, respectively. REGIME negates the use of the implicit integrators for rMHD.

problems are evolved using both models for a range of conductivities, and the optimum Courant factor giving the fastest execution, whilst generating indistinguishable results, is used. For the resistive model, we present the execution time using both the operator-split RK2 method and the SSP2(222) IMEX integrator for comparison. Conductivities that have no corresponding execution time failed to run even with unreasonably small Courant factors.

Results are given in Figure 5.14 that demonstrate the performance benefits of REGIME over rMHD. For the majority of the range of σ that we present here, REGIME is the most efficient model to evolve for both problems, being only second to the resistive model (using explicit integrators) for the most resistive set-ups.

With these results, we can see that for any given resistivity there is no need to employ implicit integrators to evolve the system—either the system has a small enough σ to allow use of the explicit integrators to evolve rMHD, or REGIME will provide a faster execution than rMHD when implicit schemes would be required. Expected speed-ups are dependant upon the magnitude of the resistivity, but for 1D, 2D and 3D simulations, one can expect speed-ups of around $10\times$, $4\times$ and $2.5\times$, respectively.

5.8 Discussion

The work of this chapter has been motivated by the relative success of model extensions, in the form of additional source terms, in capturing additional physics. Capturing a greater range of physics in a simulation invariably requires greater computation, and thus time is the limiting factor in performing more realistic simulations of mergers.

In this chapter we are encouraged by three models: The LES subgrid model of Radice et al. (2018a) has been used to show that viscosity plays an important role in the outflow dynamics of unequal mass binaries; the source of Giacomazzo et al. (2015) can be used to generate strong magnetic fields at the point of merger; and Dionysopoulou et al. (2015) has demonstrated how electrical resistivity may impact the collapse of the post-merger remnant. The two former required very few changes to existing code bases and models. In contrast, the work of Dionysopoulou et al. (2015) requires changing the conserved, and flux vectors, including additional scalar fields to satisfy Gauss' law, and requires the ability to handle large source terms with a class of semi-implicit time integrators. Previous sections have demonstrated the impact that Dionysopoulou's set-up has on the performance of a simulation—a significant amount of additional work is required at each timestep, not to mention the difficulties one encounters when solving a non-linear system of over ten variables implicitly.

Section 5.2 culminates in an addition to the standard equations of special relativistic iMHD, which capture the additional dynamics of the rMHD models used by Dionysopoulou, but in a scheme as simple to implement as a sub-grid model. The addition depends inversely upon the magnitude of the conductivity, and as a result is only small when near the ideal regime. This is a nice result, as for a large number of astrophysical systems the resistivity tends to be non-zero, and very small. For example, we believe REGIME would be suitable for evolution with standard, explicit, RK-like integrators, not only for merger simulations (Chapter 6), but also for simulations of magneto-rotational instabilities, magnetic reconnection, and jet launching from BH accretion.

In Section 5.5, we demonstrate how the model converges appropriately with the conductivity over a number of orders of magnitude—values which are realistic for a number of interesting astrophysical applications. The effect REGIME has on the solution for the magnetic fields is pronounced, and in all simulations presented agrees well with the solution of the full rMHD

model. Furthermore, as the extension may be evolved using explicit integrators for a large range of conductivities, REGIME also benefits from accelerated evolution—one can typically expect speed-ups of many factors.

6

General relativistic REGIME

No matter how hard you try to teach your cat general relativity, you're going to fail.

– Brian Greene

The work of Chapter 5 produced some encouraging results. The model developed there shows the expected convergence with conductivity, and provides performance improvements against an implicitly evolved resistive model. The aim of this chapter is to extend that work such that REGIME can be used in merger simulations, allowing us to investigate the effects of a realistic resistivity in a BNS merger simulation. With this, we will be able to ascertain whether NS resistivity is an important factor to include in future BNS investigations.

For REGIME to work successfully in a merger simulation, we must extend it to deal with dynamic spacetimes. The introduction of GR into REGIME begins with the hyperbolic conservation laws describing GRMHD. We then use a similar analysis to the SR case in deriving the resulting model. In Section 6.1 we begin with the equations of motion of resistive GRMHD, and perform the Chapman-Enskog expansion on the fast variables, i.e. the electric fields. Section 6.2 describes a number of details that are useful when going from the higher level mathematical description to implementing the resulting terms into a GRMHD code, including how to efficiently take the matrix derivatives and matrix inversions, and some examples of how to compute the matrix elements. Following this, in Section 6.3 we provide some tests to validate that GR REGIME (GRR) gives the expected behaviour with conductivity, before moving on to a number of BNS simulations.

Previous resistive merger simulations have either treated the NS matter as ideal (Lehner et al., 2012), or as unrealistically resistive Ponce et al. (2014); Dionysopoulou et al. (2015) to deal with stiffness issues with the electric field evolution. The BNS analyses in the final section of this chapter correspond to the first merger simulations that use realistic magnitude for the resistivity of NS matter, something that is not possible without GRR.

6.1 Derivation

The mathematics for the expansion in general relativity is mostly similar to the special relativistic case—one difference, however, lies in the equations we start with. The equations of motion for the matter and electromagnetic fields are the following (Palenzuela et al., 2009; Etienne et al., 2010, 2012; Baumgarte and Shapiro, 2003):

$$\partial_t \begin{pmatrix} \sqrt{\gamma} D \\ \sqrt{\gamma} S_j \\ \sqrt{\gamma} \tau \\ A_i \\ \sqrt{\gamma} \Phi \\ \sqrt{\gamma} E^j \end{pmatrix} + \partial_i \begin{pmatrix} \sqrt{-g} \left(\frac{-\beta^i}{\alpha} D + v^i D \right) \\ \sqrt{-g} \left(\frac{-\beta^i}{\alpha} S_j + S_j^i \right) \\ \sqrt{-g} \left(\frac{-\beta^i}{\alpha} \tau + S^i - v^i D \right) \\ \alpha \Phi - A_j \beta^j \\ \sqrt{-g} \left(\frac{-\beta^i}{\alpha} \Phi + A^i \right) \\ \sqrt{-g} \left(\frac{-\beta^i}{\alpha} E^j + \epsilon^{ijk} B_k \right) \end{pmatrix} = \begin{pmatrix} 0 \\ \sqrt{-g} \left(\frac{S^{lm}}{2} \partial_j \gamma_{lm} + \frac{S_k}{\alpha} \partial_j \beta^k - \frac{\tau - D}{\alpha} \partial_j \alpha \right) \\ \sqrt{-g} (S^{lm} K_{lm} - \frac{S^k}{\alpha} \partial_k \alpha) \\ -\alpha E_i - \epsilon_{ijk} \beta^j B^k \\ -\xi \sqrt{-g} \Phi \\ \sqrt{-g} \left(\frac{-E^k}{\alpha} \partial_k \beta^j - J^j \right) \end{pmatrix} \quad (6.1)$$

where A and Φ are the electromagnetic vector and scalar potential, $\xi = 1.5/\Delta t_{\max}$ where Δt_{\max} is the timestep for the coarsest refinement level, and the rest have their standard definitions as given in Equations (2.25) and Section 2.3. In contrast to Equations (5.13), we now identify three separate classes of equation in this system, which we write explicitly as

$$\partial_t(\sqrt{\gamma} \mathbf{q}) + \partial_i[\sqrt{-g} \mathbf{f}^i(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}})] = \sqrt{-g} \mathbf{s}(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}}), \quad (6.2a)$$

$$\partial_t \tilde{\mathbf{q}} + \partial_i \tilde{\mathbf{f}}^i(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}}) = \tilde{\mathbf{s}}(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}}), \quad (6.2b)$$

$$\partial_t(\sqrt{\gamma} \bar{\mathbf{q}}) + \partial_i[\sqrt{-g} \bar{\mathbf{f}}^i(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}})] = \sqrt{-g} \frac{\bar{\mathbf{s}}(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}})}{\epsilon}. \quad (6.2c)$$

Due to the time dependence of the metric components, we must treat the densitised evolution equations, \mathbf{q} and $\bar{\mathbf{q}}$, separately to the undensitised ones, $\tilde{\mathbf{q}}$. This appears as **Whisky** evolves the vector potential, in contrast to the magnetic fields, to reduce the onset of divergence-free violations of the magnetic fields at mesh refinement boundaries (Etienne et al., 2010; Giacomazzo et al., 2011).

As before, we have identified the stiff source to act on some small timescale, ϵ . As $\epsilon \rightarrow 0$, the stiff variables assume their equilibrium solution, which we identify as $\lim_{\epsilon \rightarrow 0} \bar{\mathbf{q}} = \bar{\mathbf{q}}_0$. In this limit, the stiff source is $\bar{\mathbf{s}}(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}}_0) = 0$. As the timescale grows, the stiff solution will be perturbed about its equilibrium. As in the SR case, we can therefore perform a Chapman-Enskog expansion for the stiff variables around this equilibrium, such that

$$\bar{\mathbf{q}} = \bar{\mathbf{q}}_0 + \epsilon \bar{\mathbf{q}}_1 + \mathcal{O}(\epsilon^2). \quad (6.3)$$

We can substitute this into the equations for the non-stiff variables, \mathbf{q} and $\tilde{\mathbf{q}}$, and keep terms of order $\mathcal{O}(\epsilon^1)$,

$$\partial_t(\sqrt{\gamma} \mathbf{q}) + \partial_i \left(\sqrt{-g} \left[\mathbf{f}_0^i + \epsilon \frac{\partial \mathbf{f}_0^i}{\partial \bar{\mathbf{q}}} \bar{\mathbf{q}}_1 \right] \right) = \sqrt{-g} \left[\mathbf{s}_0 + \epsilon \frac{\partial \mathbf{s}_0}{\partial \bar{\mathbf{q}}} \bar{\mathbf{q}}_1 \right], \quad (6.4a)$$

$$\partial_t \tilde{\mathbf{q}} + \partial_i \left(\tilde{\mathbf{f}}_0^i + \epsilon \frac{\partial \tilde{\mathbf{f}}_0^i}{\partial \bar{\mathbf{q}}} \bar{\mathbf{q}}_1 \right) = \tilde{\mathbf{s}}_0 + \epsilon \frac{\partial \tilde{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \bar{\mathbf{q}}_1, \quad (6.4b)$$

where we are making the definition that for some vector, \mathbf{V} , that depends on the conserved vectors, $\mathbf{V}(\mathbf{q}, \tilde{\mathbf{q}}, \bar{\mathbf{q}}_0) \equiv \mathbf{V}_0$. Doing the same for the stiff variables, but only keeping terms up to order $\mathcal{O}(\epsilon^0)$, gives

$$\partial_t(\sqrt{\gamma}\bar{\mathbf{q}}_0) + \partial_i(\sqrt{-g}\bar{\mathbf{f}}_0^i) = \sqrt{-g}\frac{\partial\bar{s}_0}{\partial\bar{\mathbf{q}}}\bar{\mathbf{q}}_1, \quad (6.5)$$

which gives us a form for the perturbation,

$$\bar{\mathbf{q}}_1 = \left(\sqrt{-g}\frac{\partial\bar{s}_0}{\partial\bar{\mathbf{q}}}\right)^{-1} \left[\partial_t(\sqrt{\gamma}\bar{\mathbf{q}}_0) + \partial_i(\sqrt{-g}\bar{\mathbf{f}}_0^i)\right]. \quad (6.6)$$

To remove the time dependence (of the stiff variables, we will not be able to completely remove the time dependence of the spatial metric), we assume that the equilibrium values of the stiff variables are only functions of the non-stiff variables, that is, $\bar{\mathbf{q}}_0 = \bar{\mathbf{q}}_0(\mathbf{q}, \tilde{\mathbf{q}})$, and thus

$$\begin{aligned} \partial_t(\sqrt{\gamma}\bar{\mathbf{q}}_0) &= \bar{\mathbf{q}}_0\partial_t\sqrt{\gamma} + \sqrt{\gamma}\partial_t\bar{\mathbf{q}}_0 \\ &= \bar{\mathbf{q}}_0\partial_t\sqrt{\gamma} + \sqrt{\gamma}\frac{\partial\bar{\mathbf{q}}_0}{\partial\mathbf{q}}\partial_t\mathbf{q} + \sqrt{\gamma}\frac{\partial\bar{\mathbf{q}}_0}{\partial\tilde{\mathbf{q}}}\partial_t\tilde{\mathbf{q}}. \end{aligned} \quad (6.7)$$

Equations (6.4) can be rewritten as

$$\sqrt{\gamma}\partial_t\mathbf{q} = \sqrt{-g}\mathbf{s}_0 - \partial_i(\sqrt{-g}\mathbf{f}_0^i) - \mathbf{q}\partial_t\sqrt{\gamma}, \quad (6.8a)$$

$$\partial_t\tilde{\mathbf{q}} = \tilde{\mathbf{s}}_0 - \partial_i\tilde{\mathbf{f}}_0^i, \quad (6.8b)$$

up to $\mathcal{O}(\epsilon^0)$, which upon substitution into (6.7) gives

$$\partial_t(\sqrt{\gamma}\bar{\mathbf{q}}_0) = \bar{\mathbf{q}}_0\partial_t\sqrt{\gamma} + \frac{\partial\bar{\mathbf{q}}_0}{\partial\mathbf{q}} \left[\sqrt{-g}\mathbf{s}_0 - \partial_i(\sqrt{-g}\mathbf{f}_0^i) - \mathbf{q}\partial_t\sqrt{\gamma}\right] + \sqrt{\gamma}\frac{\partial\bar{\mathbf{q}}_0}{\partial\tilde{\mathbf{q}}} \left[\tilde{\mathbf{s}}_0 - \partial_i\tilde{\mathbf{f}}_0^i\right]. \quad (6.9)$$

We can now use this, along with Equation (6.6), to get the perturbation,

$$\begin{aligned} \bar{\mathbf{q}}_1 &= \left(\sqrt{-g}\frac{\partial\bar{s}_0}{\partial\bar{\mathbf{q}}}\right)^{-1} \left\{ \frac{\partial\bar{\mathbf{q}}_0}{\partial\mathbf{q}} \left[\sqrt{-g}\mathbf{s}_0 - \partial_i(\sqrt{-g}\mathbf{f}_0^i)\right] + \sqrt{\gamma}\frac{\partial\bar{\mathbf{q}}_0}{\partial\tilde{\mathbf{q}}} \left[\tilde{\mathbf{s}}_0 - \partial_i\tilde{\mathbf{f}}_0^i\right] + \right. \\ &\quad \left. \partial_i(\sqrt{-g}\bar{\mathbf{f}}_0^i) + \partial_t\sqrt{\gamma} \left[\bar{\mathbf{q}}_0 - \frac{\partial\bar{\mathbf{q}}_0}{\partial\mathbf{q}}\mathbf{q}\right] \right\}. \end{aligned} \quad (6.10)$$

A quick check shows this is consistent with the perturbation for SR, Equation (5.21), as $\sqrt{-g} = \sqrt{\gamma} = 1$ and $\tilde{\mathbf{s}} = \tilde{\mathbf{f}} = 0$ for that system.

Substituting the above form for the perturbation into Equations (6.4) gives the new evolution equations for the non-stiff variables,

$$\partial_t(\sqrt{\gamma}\mathbf{q}) + \partial_i(\sqrt{-g}\mathbf{f}_0^i + \mathbf{F}^i) = \sqrt{-g}\mathbf{s}_0 + \mathbf{S} + \partial_i\mathbf{D}^i, \quad (6.11a)$$

$$\partial_t\tilde{\mathbf{q}} + \partial_i(\tilde{\mathbf{f}}_0^i + \tilde{\mathbf{F}}^i) = \tilde{\mathbf{s}}_0 + \tilde{\mathbf{S}} + \partial_i\tilde{\mathbf{D}}^i, \quad (6.11b)$$

with

$$\mathbf{F}^i = \frac{\partial \mathbf{f}_0^i}{\partial \bar{q}} \mathbf{l} \quad \text{and} \quad \tilde{\mathbf{F}}^i = \sqrt{-g}^{-1} \frac{\partial \tilde{\mathbf{f}}_0^i}{\partial \bar{q}} \mathbf{l}, \quad (6.12a)$$

$$\mathbf{D}^i = \frac{\partial \mathbf{f}_0^i}{\partial \bar{q}} \mathbf{m} \quad \text{and} \quad \tilde{\mathbf{D}}^i = \sqrt{-g}^{-1} \frac{\partial \tilde{\mathbf{f}}_0^i}{\partial \bar{q}} \mathbf{m}, \quad (6.12b)$$

$$\mathbf{S} = \frac{\partial s_0}{\partial \bar{q}} \mathbf{n} \quad \text{and} \quad \tilde{\mathbf{S}} = \sqrt{-g}^{-1} \frac{\partial \tilde{s}_0}{\partial \bar{q}} \mathbf{n}, \quad (6.12c)$$

and

$$\mathbf{l} = \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} \left[\sqrt{-g} \frac{\partial \bar{q}_0}{\partial \bar{q}} s_0 + \sqrt{\gamma} \frac{\partial \bar{q}_0}{\partial \bar{q}} \tilde{s}_0 + \left(\bar{q}_0 - \frac{\partial \bar{q}_0}{\partial \bar{q}} \mathbf{q} \right) \partial_t \sqrt{\gamma} \right], \quad (6.13a)$$

$$\mathbf{m} = - \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} \left[\partial_j (\sqrt{-g} \bar{\mathbf{f}}_0^j) - \frac{\partial \bar{q}_0}{\partial \bar{q}} \partial_j (\sqrt{-g} \mathbf{f}_0^j) - \sqrt{\gamma} \frac{\partial \bar{q}_0}{\partial \bar{q}} \partial_j \tilde{\mathbf{f}}_0^j \right], \quad (6.13b)$$

$$\begin{aligned} \mathbf{n} = \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} & \left[\frac{\partial \bar{q}_0}{\partial \bar{q}} \left(\sqrt{-g} s_0 - \partial_j (\sqrt{-g} \mathbf{f}_0^j) \right) + \sqrt{\gamma} \frac{\partial \bar{q}_0}{\partial \bar{q}} \left(\tilde{s}_0 - \partial_j \tilde{\mathbf{f}}_0^j \right) + \right. \\ & \left. \partial_j (\sqrt{-g} \bar{\mathbf{f}}_0^j) + \left(\bar{q}_0 - \frac{\partial \bar{q}_0}{\partial \bar{q}} \mathbf{q} \right) \partial_t \sqrt{\gamma} \right], \end{aligned} \quad (6.13c)$$

where we have absorbed the timescale ϵ into the definition of the stiff source, $\bar{\mathbf{s}}$.

These represent the full set of equations in general, however, some further simplifications can be made to reduce the number of terms. First, we note that, in general,

$$\frac{\partial \bar{q}_0}{\partial \bar{q}} = \frac{\partial \bar{q}_0}{\partial \tilde{q}} = \mathbf{0}, \quad (6.14)$$

as was the case for SR (Section 5.3). By inspection of Equations (6.1), we can also see that $\partial_{\bar{q}} \tilde{\mathbf{f}}_0 = \mathbf{0}$.

Our system, therefore, looks like

$$\partial_t (\sqrt{\gamma} \mathbf{q}) + \partial_i (\sqrt{-g} \mathbf{f}_0^i + \mathbf{F}^i) = \sqrt{-g} s_0 + \mathbf{S} + \partial_i \mathbf{D}^i, \quad (6.15a)$$

$$\partial_t \tilde{\mathbf{q}} + \partial_i \tilde{\mathbf{f}}_0^i = \tilde{s}_0 + \tilde{\mathbf{S}}, \quad (6.15b)$$

with the modifications as

$$\mathbf{F}^i = \frac{\partial \mathbf{f}_0^i}{\partial \bar{q}} \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} \bar{\mathbf{q}}_0 \partial_t \sqrt{\gamma}, \quad (6.16a)$$

$$\mathbf{D}^i = - \frac{\partial \mathbf{f}_0^i}{\partial \bar{q}} \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} \partial_j (\sqrt{-g} \bar{\mathbf{f}}_0^j), \quad (6.16b)$$

$$\mathbf{S} = \frac{\partial s_0}{\partial \bar{q}} \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} \left[\partial_j (\sqrt{-g} \bar{\mathbf{f}}_0^j) + \bar{\mathbf{q}}_0 \partial_t \sqrt{\gamma} \right], \quad (6.16c)$$

$$\tilde{\mathbf{S}} = \sqrt{-g}^{-1} \frac{\partial \tilde{s}_0}{\partial \bar{q}} \left(\frac{\partial \bar{s}_0}{\partial \bar{q}} \right)^{-1} \left[\partial_j (\sqrt{-g} \bar{\mathbf{f}}_0^j) + \bar{\mathbf{q}}_0 \partial_t \sqrt{\gamma} \right], \quad (6.16d)$$

and with $\partial_t \sqrt{\gamma} = \frac{-\sqrt{\gamma}}{\alpha} \partial_t \alpha$ (see Section 2.3.3).

6.2 Implementation

In this section we outline some details about implementing GRR as an extension to the `Whisky` code. In contrast to Section 5.3, where in the SR case we assumed low velocities and weak coupling between electromagnetism and hydrodynamics, in the GR case we do not attempt to use symbolic algebra tools to derive the terms. Rather, we derive by hand the necessary matrix elements and numerically compute the terms that result.

6.2.1 Matrix derivatives

Initial attempts to implement Equations (6.16), where we assumed $\bar{\mathbf{q}} \equiv \mathbf{E}$, were unsuccessful. One reason for this may be that $\bar{\mathbf{q}}$ represents conserved quantities in the equations of motion that act on a fast timescale. Looking at the full set of equations of motion, Equation (6.1), we can see that there is also a dependence of the momentum and energy terms on the electric fields and thus they are implicitly fast variables. This information is lost if we only consider the electric fields to be fast.

Alternatively, we may use the chain rule to perform these derivatives, noting that

$$\frac{\partial \mathbf{s}_0}{\partial \bar{\mathbf{q}}} = \frac{\partial \mathbf{s}_0}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \bar{\mathbf{q}}}. \quad (6.17)$$

Thus, an expression such as in Equation (6.16c) becomes

$$\frac{\partial \mathbf{s}_0}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \bar{\mathbf{q}}} \right)^{-1} = \frac{\partial \mathbf{s}_0}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \bar{\mathbf{q}}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{w}} \frac{\partial \mathbf{w}}{\partial \bar{\mathbf{q}}} \right)^{-1} \quad (6.18)$$

$$= \frac{\partial \mathbf{s}_0}{\partial \mathbf{w}} \left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{w}} \right)^+, \quad (6.19)$$

where $(\cdot)^+$ denotes the pseudoinverse, see Section 6.2.2.

This approach is similar to the method used in the SR implementation. In this case, however, we will let the vector of auxiliary variables be $\mathbf{w} = (v^i, B^i, E^i)$. This turns out to be sufficient, and we need not consider other primitive variables such as pressure or density. This is a result of all terms in Equation (6.16) containing a $\partial \bar{\mathbf{s}}_0 / \partial \mathbf{w}$ term, and by inspecting the form of the stiff source term we see that $\bar{\mathbf{s}}_0 = \bar{\mathbf{s}}_0(\mathbf{v}, \mathbf{B}, \mathbf{E})$ only.

6.2.2 Inverting numerically

Within Equation (6.19) lies the pseudoinverse of a 3×9 matrix:

$$\left(\frac{\partial \bar{\mathbf{s}}_0}{\partial \mathbf{w}} \right)^+. \quad (6.20)$$

Here, we have used the definition of the right pseudoinverse as $\underline{\underline{A}}^+ \equiv \underline{\underline{A}}^T (\underline{\underline{A}} \underline{\underline{A}}^T)^{-1}$, such that $\underline{\underline{A}} \underline{\underline{A}}^+ = I$. In contrast to the SR case, we have decided not to invert this matrix symbolically. Symbolic inversion, whilst possible, leads to a very large number of terms and increases the

likelihood of errors when implementing it in the code. In addition, numerically inverting a matrix is considered bad practice and can lead to large errors for ill-conditioned matrices.

Instead, what is important here is the solution to the system $\underline{\underline{A}}\mathbf{x} = \mathbf{b}$, where $\underline{\underline{A}}$ is some non-singular matrix, and \mathbf{b} and \mathbf{x} are some known and unknown vectors, respectively. For Equations (6.16), this corresponds to

$$\frac{\partial \bar{s}_0}{\partial \mathbf{w}} \mathbf{X} = \bar{q}_0 \partial_t \sqrt{\gamma} \quad (6.21a)$$

$$\frac{\partial \bar{s}_0}{\partial \mathbf{w}} \mathbf{Y} = \partial_j (\sqrt{-g} \mathbf{f}_o^j), \quad (6.21b)$$

where the solutions are the vectors \mathbf{X} and \mathbf{Y} :

$$\mathbf{X} = \left(\frac{\partial \bar{s}_0}{\partial \mathbf{w}} \right)^+ \bar{q}_0 \partial_t \sqrt{\gamma} \quad (6.22a)$$

$$\mathbf{Y} = \left(\frac{\partial \bar{s}_0}{\partial \mathbf{w}} \right)^+ \partial_j (\sqrt{-g} \mathbf{f}_o^j). \quad (6.22b)$$

As the matrices $\partial \bar{s}_0 / \partial \mathbf{w}^+$ are non-square, the system we solve is instead

$$\underline{\underline{T}} \mathbf{X} = \bar{q}_0 \partial_t \sqrt{\gamma} \quad (6.23a)$$

$$\underline{\underline{T}} \mathbf{Y} = \partial_j (\sqrt{-g} \mathbf{f}_o^j), \quad (6.23b)$$

where $\underline{\underline{T}}$ is the following square matrix,

$$\underline{\underline{T}} = \frac{\partial \bar{s}}{\partial \mathbf{w}} \frac{\partial \bar{s}}{\partial \mathbf{w}}^T, \quad (6.24)$$

and after which we define the new vectors

$$\mathbf{V} = \frac{\partial \bar{s}}{\partial \mathbf{w}}^T \mathbf{X} \quad (6.25a)$$

$$\mathbf{W} = \frac{\partial \bar{s}}{\partial \mathbf{w}}^T \mathbf{Y}. \quad (6.25b)$$

With these definitions, the modifications of Equation (6.16) are now

$$\mathbf{F}^i = \frac{\partial \mathbf{f}_0^i}{\partial \mathbf{w}} \mathbf{V}, \quad (6.26a)$$

$$\mathbf{D}^i = -\frac{\partial \mathbf{f}_0^i}{\partial \mathbf{w}} \mathbf{W}, \quad (6.26b)$$

$$\mathbf{S} = \frac{\partial s_0}{\partial \mathbf{w}} (\mathbf{V} + \mathbf{W}), \quad (6.26c)$$

$$\tilde{\mathbf{S}} = \frac{\partial \tilde{s}_0}{\partial \mathbf{w}} (\mathbf{V} + \mathbf{W}) / \sqrt{-g}, \quad (6.26d)$$

This method allows us to avoid inverting matrices explicitly, something which can be more inaccurate than the solution of a system of Equations (6.23). To solve this system, we use the well-tested and trusted linear algebra package LAPACK.

6.2.3 Matrix terms

We will now briefly illustrate how to compute the terms given in Equation (6.26). In the following, all indices run over the spatial dimensions. Taking the form for the stiff source,

$$\bar{s}_0^j = \frac{-E^k}{\alpha} \partial_k \beta^j - \varrho v^j - W \sigma \left[E^j + \epsilon^{jkl} v_k B_l - v_k E^k v^j \right], \quad (6.27)$$

we take the derivatives with respect to the auxiliary variables, $\mathbf{w} = (v^i, B^i, E^i)$. These are, respectively,

$$\frac{\partial \bar{s}_0^r}{\partial v^s} = -\varrho \frac{\partial v^r}{\partial v^s} - \sigma \left(\frac{\partial W}{\partial v^s} \left[E^r + \epsilon^{rnp} v_n B_p - v_k E^k v^r \right] + W \left[\epsilon^{rnp} \frac{\partial v_n}{\partial v^s} B_p - \frac{\partial v_k}{\partial v^s} E^k v^r - v_k E^k \frac{\partial v^r}{\partial v^s} \right] \right) \quad (6.28)$$

or,

$$\frac{\partial \bar{s}_0^r}{\partial v^s} = -\varrho \delta_s^r - \sigma \left(v_s W^3 \left[E^r + \epsilon^{rnp} v_n B_p - v_k E^k v^r \right] + W \left[\epsilon^{rnp} \gamma_{ns} B_p - \gamma_{ks} E^k v^r - v_k E^k \delta_s^r \right] \right) \quad (6.29)$$

where we have used $\partial v^r / \partial v^s = \delta_s^r$, the Kronecker delta, $\partial v_r / \partial v^s = \gamma_{rs}$, the induced spatial metric, and $\partial W / \partial v^s = W^3 v_s$. Similarly, we have

$$\frac{\partial \bar{s}_0^r}{\partial B^s} = -\epsilon^{rnp} v_n \gamma_{ps}, \quad (6.30)$$

$$\frac{\partial \bar{s}_0^r}{\partial E^s} = -\frac{\partial_s \beta^r}{\alpha} - W \sigma \delta_s^r + W \sigma v^r v_s, \quad (6.31)$$

where $\partial_s \beta^r$ is computed via second order central differencing¹.

These relations give us the elements of the matrix $\partial \bar{s}_0^r / \partial w^p = (\partial \bar{s}_0 / \partial w)^r_p$. From here, we construct the following matrix,

$$T_p^r = (\partial \bar{s}_0 / \partial w)^r_s (\partial \bar{s}_0 / \partial w^T)_p^s, \quad (6.32)$$

which allows us to solve Equations (6.23). To calculate the right-hand side, we assume that $\bar{q}_0^i \equiv E^i = -\epsilon^{ijk} v_j B_k$, which is the correct form of the equilibrium electric fields, and use the relation from Section 2.3.3 for the time derivative of the induced metric determinant, $\partial_t \sqrt{\gamma}$. In order to evaluate $\partial_j (\sqrt{-g} \bar{f}_0^{ij})$, where the index j indicates the flux direction for $j \in [x, y, z]$, and i indicates the i^{th} component of the stiff flux vector, we use the third order WENO reconstruction with a local Lax-Friedrichs flux splitting. That is, we deconstruct the flux into purely upwind and downwind contributions via

$$\bar{f}_0^{ij(\pm)} = \bar{f}_0^{ij} \pm v_{sig} \bar{q}_0^i, \quad (6.33)$$

reconstruct the up- and downwind fluxes to the cell edges using the WENO reconstruction, and then apply flux differencing to get the numerical flux for each cell. Generally, v_{sig} is the maximum signal speed as computed by the eigenvalues of the Jacobian of the system, which as

¹Choosing fourth order central differencing would have required additional MPI calls, increasing communication time and decreasing performance improvements for large simulations. This was deemed more important than the benefits of a higher order method for this term.

we are including electromagnetic fields, we set equal to unity. We discuss this method when applied to the modified flux vector at the end of this section.

We now have all the terms in Equation (6.23) to solve for the \mathbf{X} and \mathbf{Y} vectors via the LAPACK routine DGESV. From here, it is simple to calculate the \mathbf{V} and \mathbf{W} vectors, Equations (6.25). All terms in Equations (6.26) can then be computed, using the same methods to calculate matrix terms as earlier in this section.

Once we have the terms from Equation (6.26) computed, all that remains is to determine the derivatives of the modified flux and diffusion vectors in Equations (6.15). For the diffusion vector, this is done via a fourth order central differencing scheme. The modified flux vector requires a little more attention, however. As we are interpreting it as a flux, it is sensible to compute it via some high-resolution shock capturing scheme, for which we chose flux vector splitting and a third order WENO reconstruction. As in Section 2.5, we use the local Lax-Friedrichs flux splitting, however in this case the maximum signal velocity is not unity. To determine the signal speed, we take a simplified, hyperbolic, conservation-law system in one dimension:

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{q})}{\partial x} = 0. \quad (6.34)$$

In quasi-linear form, this is

$$\frac{\partial \mathbf{q}}{\partial t} + A \frac{\partial \mathbf{q}}{\partial x} = 0, \quad (6.35)$$

with $A = \partial_{\mathbf{q}} \mathbf{f}(\mathbf{q})$ the Jacobian of the system. The maximum signal speed is then $v_{sig} = \max |\lambda_i|$, where λ_i is the i^{th} eigenvalue of A . Typically, when considering MHD, this will be the speed of light, however, the flux vector given by Equation (6.16a) has an inverse dependence on the conductivity, and so we include this in the approximation of the signal speed, $v_{sig} \sim \sigma^{-1}$. As a result, the formula for the Lax-Friedrichs flux splitting for the modified flux vector takes the form

$$F^{ij(\pm)} = F^{ij} \pm \frac{q^i}{\sigma}. \quad (6.36)$$

The rest of the reconstruction is as described in Section 2.5.

6.2.4 A small hack

Here, we must mention a fix we have implemented to improve the quality of the solutions our model produces. In index form, the derivative of the energy flux with respect to the electric fields, and velocity derivative of the stiff source, are respectively

$$\frac{\partial f_{\tau}^r}{\partial E^p} = \frac{-\beta^r}{\alpha} E_p + \gamma^{rs} \epsilon_{spu} B^u, \quad (6.37)$$

and

$$\begin{aligned} \frac{\partial \bar{s}^r}{\partial v^s} = & -\varrho \delta_s^r - v_s W^3 \sigma (E^r + \epsilon^{rnp} v_n B_p - v_n E^n v^r) \\ & - W \sigma (\epsilon^{rnp} \gamma_{ns} B_p - \gamma_{ns} E^n v^r - v_n E^n \delta_s^r), \end{aligned} \quad (6.38)$$

where f_{τ}^r is the flux of the conserved energy, τ , in the r -direction, and E^p , \bar{s}^p , and v^p , are the p^{th} components of the electric field, stiff source and velocity vector, respectively.

In the production code, we have opted to remove the contributions from the following terms: $\gamma^{rs}\epsilon_{spu}B^u$ in the flux derivative; and $\epsilon^{rnp}\gamma_{ns}B_p$ in the stiff source derivative². The reason for this is accuracy. With these two cross-product- B terms included, the diffusion of the magnetic field in simple problems was clearly unphysical. In contrast, once they were removed from the extension, the solutions matched both exact solutions in the current sheet test (see Section 6.3.2) and have near identical results to the METHOD code. The reason why these terms diminish the quality of the solution is not fully clear. It may be that the cause of the error is numerical—i.e. that the two terms should cancel exactly but due to finite precision a residual remains that spoils the solution—but this has not been proven. Nevertheless, we are confident in the solution without these terms as in all test cases performed, we get the expected behaviour and convergence, and simulations are stable in the long term.

6.2.5 Stability

Similarly to REGIME in SR, we must infer the stability properties of GRR and estimate when the model should be stable. The analysis in Section 5.4 was predicated on the fact that the source term is diffusive. This was the case in SR, as the modified source term and flux term are both uniquely zero. For GRR, however, the system is more complicated as we have non-zero source and flux vectors due to the effects of spacetime curvature. In addition, the vector potential evolution equation contains a non-zero source. This is, in fact, the source that acts as the diffusion vector for the magnetic fields, as it is a first order derivative in space, and picks up the second order derivative when transforming to the magnetic fields via $\mathbf{B} = \nabla \times \mathbf{A}$. We can see this by taking the evolution equation for the vector potential,

$$\partial_t \mathbf{A} + \nabla f_A = \mathbf{s}_{A0} + \tilde{\mathbf{S}}, \quad (6.39)$$

with the flux given by f_A , equilibrium source vector given by \mathbf{s}_{A0} , and the modified source given by $\tilde{\mathbf{S}}$. Computing its curl gives

$$\partial_t \mathbf{B} + \partial_i \mathbf{f}_B^i = \mathbf{s}_{B0} + \nabla \times \tilde{\mathbf{S}}, \quad (6.40)$$

where now the flux and source vector for the magnetic fields is \mathbf{f}_B and \mathbf{s}_{B0} . When we compare this to the SR version, Equation (5.26), we see that the modified source for the vector potential is related to the magnetic field diffusion vector through

$$\epsilon^{ijk}\partial_j \tilde{S}_k = \partial_k D^{ki}, \quad (6.41)$$

where D^{ki} represents the i^{th} component of the diffusion vector in the k^{th} direction. We therefore see how $\tilde{\mathbf{S}}$ and \mathbf{D}^k perform very similar roles in diffusing the magnetic fields.

Furthermore, we performed the same stability analyses on GRR as for REGIME in Section 5.6. The results showed that despite the differences in the vector potential formulation, the two models showed signs of instabilities at the same point and in accordance with the stability criterion, Equation 5.42. As a result, we will once again assume a conservative softening factor of $\delta = 1$.

²These terms are included in unit tests, however.

Of course, this assumption has been made based on results performed in special relativity, and there may be contributions to the loss of stability due to general relativistic effects. However, we note that the source of the instability comes from terms that produce diffusive behaviour, namely from $\partial_i \mathbf{D}^i$ for the non-stiff densitised evolution equation and $\tilde{\mathbf{S}}$ for the undensitised one. Neither of these terms include contributions from the curvature of the spacetime, rather they depend upon the choice of gauge, α and β^i .

For the SR stability criterion to apply, we would need two things: the GRR modifications to be gauge independent; and for the cause of the diffusive instability to be gauge independent. If both these are true, then the analysis should hold and the stability criterion should be accurate. Although we lack a formal proof either way, it will be clear very quickly if we have misjudged the stability properties of GRR in curved spacetimes. Indeed, at the resistivities we have used in later BNS simulations, Section 6.4, there were no signs of any diffusive instabilities forming.

6.3 Testing

In this section, we show some tests of our implementation of GRR into the `Whisky` code. Where possible, we compare results with a full resistive model, namely the special relativistic, resistive MHD model of Palenzuela et al. (2009) implemented in the `METHOD` code (Wright, 2018), as verified in Wright and Hawke (2019a,b)³. For error convergence with the conductivity, we use the numerical solution of this resistive model as the *exact* solution.

6.3.1 Unit testing and test driven development

```
INFO (REGIME_GR): REGIME test: Testing whisky functions...
INFO (REGIME_GR): REGIME test: Whisky functions passed!
INFO (REGIME_GR): REGIME test: Testing electric fields...
INFO (REGIME_GR): REGIME test: electric fields passed!
INFO (REGIME_GR): REGIME test: Testing fbar...
INFO (REGIME_GR): REGIME test: fbar passed!
INFO (REGIME_GR): REGIME test: Testing dsbardwT...
INFO (REGIME_GR): REGIME test: dsbardwT passed!
INFO (REGIME_GR): REGIME test: Testing Tmatrix...
INFO (REGIME_GR): REGIME test: Tmatrix passed!
INFO (REGIME_GR): REGIME test: Testing dfdw...
INFO (REGIME_GR): REGIME test: dfdw passed!
INFO (REGIME_GR): REGIME test: Testing dstildBE...
INFO (REGIME_GR): REGIME test: dstildBE passed!
INFO (REGIME_GR): REGIME test: Testing XandYvec...
INFO (REGIME_GR): REGIME test: XandYvec passed!
INFO (REGIME_GR): REGIME test: Testing VandWvec...
INFO (REGIME_GR): REGIME test: VandWvec passed!
INFO (REGIME_GR): REGIME test: Testing Smodtil...
INFO (REGIME_GR): REGIME test: Smodtil passed!
INFO (REGIME_GR): REGIME test: Testing FmodXYZ...
INFO (REGIME_GR): REGIME test: FmodXYZ passed!
WARNING level 0 from host alex process 0
while executing schedule bin REGIME_TEST_END, routine REGIME_GR::killSim
in thorn REGIME_GR, file /home/alex/Cactus/configs/sim/build/REGIME_GR/testScripts/simKiller.cc:14:
->
Testing completed successfully!
```

FIGURE 6.1: Confirmation of tests passing in the output file. Due to the non-physical initial data, we have to kill the simulation after the tests are completed.

³It is worth noting that the results for REGIME in `METHOD` are better now than when initially published due to a bug fix in the time integration scheme.

A testing paradigm that is ubiquitous in software development is the notion of unit testing. In this style of testing, one breaks down a program into the smallest operable chunks and explicitly tests the output of these methods with some given data.

It is virtually certain that when coding some new physics thorn, which may include thousands or tens of thousands of lines of code, that there will be unintentional bugs introduced through typos, missed terms, etcetera. In order to check that our implementation of the routines is correct, therefore, we have automated some unit tests.

We have broken down equations (6.26) into what we call the *vars* and *terms*. *Vars* represent the matrix terms that we have in analytical form, including $\partial_{\mathbf{w}} \mathbf{f}_0^i$, $\partial_{\mathbf{w}} \mathbf{s}_0$, $\partial_{\mathbf{w}} \tilde{\mathbf{s}}_0$, etcetera, whereas the *terms* are those that we compute using the vars, such as \underline{T} , \mathbf{V} and \mathbf{W} .

Some specific initial data is chosen (which is not necessarily physical), and the *vars* and *terms* are then computed and compared to their expected values. The expected values are computed using a Python script which uses `numpy`'s `einsum` function to reduce the amount of typing, and thus lower the chance of errors.

Crucially, these tests are written and implemented before the *vars* and *terms* routines are written. This methodology, known as test driven development, forces the programmer to consider the design of the module first, and removes the temptation to ignore tests that do not pass, or to write tests that are designed to pass. The resulting output of the unit tests implemented in GRR are shown in Figure 6.1.

6.3.2 Tests problems

To ensure that the system as a whole works correctly, we now present some simple tests with known solutions. The solutions of these tests can be compared against the verified output of the special relativistic REGIME model, Section 5.5.

Current sheet

A resistive problem that can be used to show convergence of GRR with the conductivity is the self-similar current sheet (Komissarov, 2007), Section 3.1.3. In Figure 6.2, we show the numerical solution at the final time, $T = 8$, for the current sheet problem using the `Whisky` code, the full resistive model in `METHOD`, and REGIME as implemented in both `METHOD` (special relativistic) and `Whisky` (general relativistic). As expected, the error growth for REGIME is significantly reduced as the system becomes resistive. The larger error for `Whisky` + REGIME is due to the magnetic fields being stored at cell faces, compared to cell centres for `METHOD`.

Brio-Wu shock tube

Another simple test to assess the implementation of REGIME is the shock tube test (Brio and Wu, 1988), Section 3.1.2, the results of which can be seen in Figure 6.3 for a range of conductivities. Whilst the error for GRR is typically much worse in this case than for the current sheet problem, this is mainly due to the non-staggered vector potential in `Whisky` producing

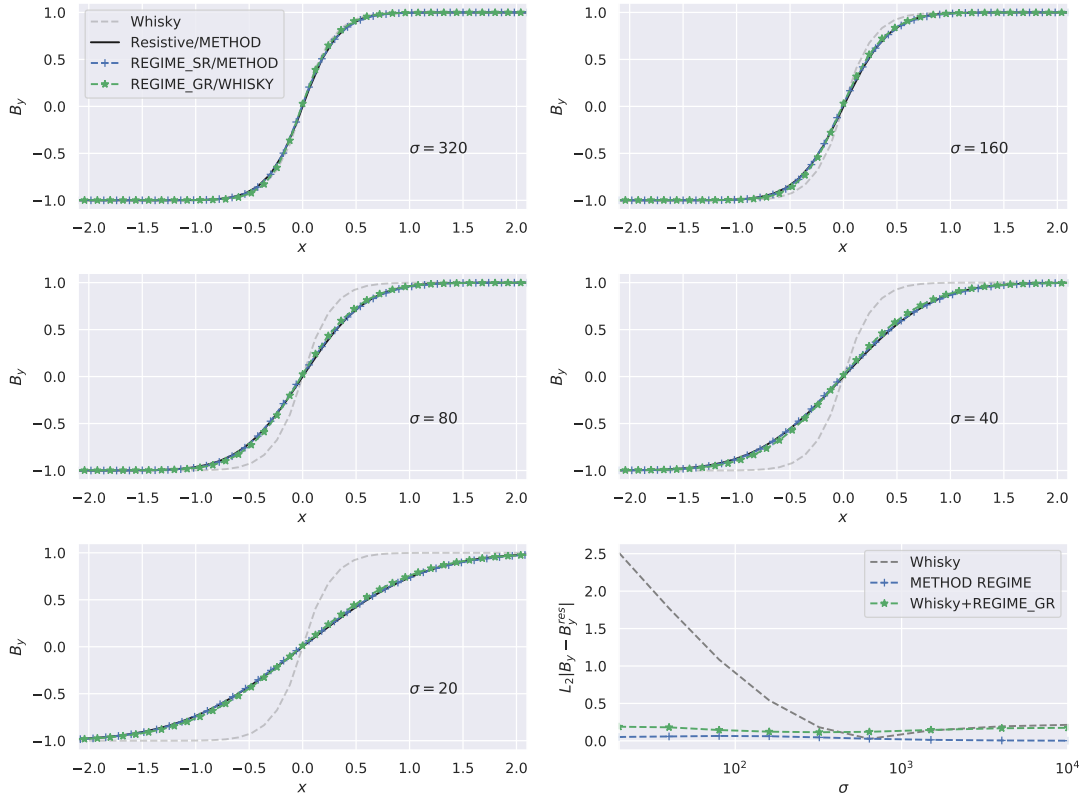


FIGURE 6.2: Convergence of the current sheet with the conductivity, σ , at a final time of $T = 8$. The convergence of GRR and REGIME with the conductivity is very similar.

numerical oscillations about the shock, as discussed in the **Spritz** paper (Cipolletta et al., 2019)⁴. Whilst the error values are less useful here, we can still see that the diffusion of the magnetic field is captured well by GRR in all cases.

6.4 Binary neutron star mergers

As we have demonstrated that GRR produces the expected output in a number of scenarios, we may now turn our attention to BNS mergers. Our aim in the remaining sections is to determine whether or not modelling NS matter with a realistic size resistivity will have a meaningful impact in the output of merger simulations, and what that impact might be.

6.4.1 Simulation set-up

We will now give an overview of the set-up and thorns we use for the BNS simulations. As mentioned in Section 2.3.5, we utilise the open source *Einstein Toolkit*⁵, built on top of the **Cactus** computational toolkit. More specifically, for the evolution of the spacetime, we make use of the McLachlan⁶ thorn, which implements the BSSNOK (Shibata and Nakamura, 1995;

⁴GRR has been developed in such a way as to have minimal dependencies within the ET, and so in future work, it should require little effort to couple GRR to **Spritz** to remove this artefact.

⁵www.einsteintoolkit.org

⁶www.cct.lsu.edu/~eschnett/McLachlan/

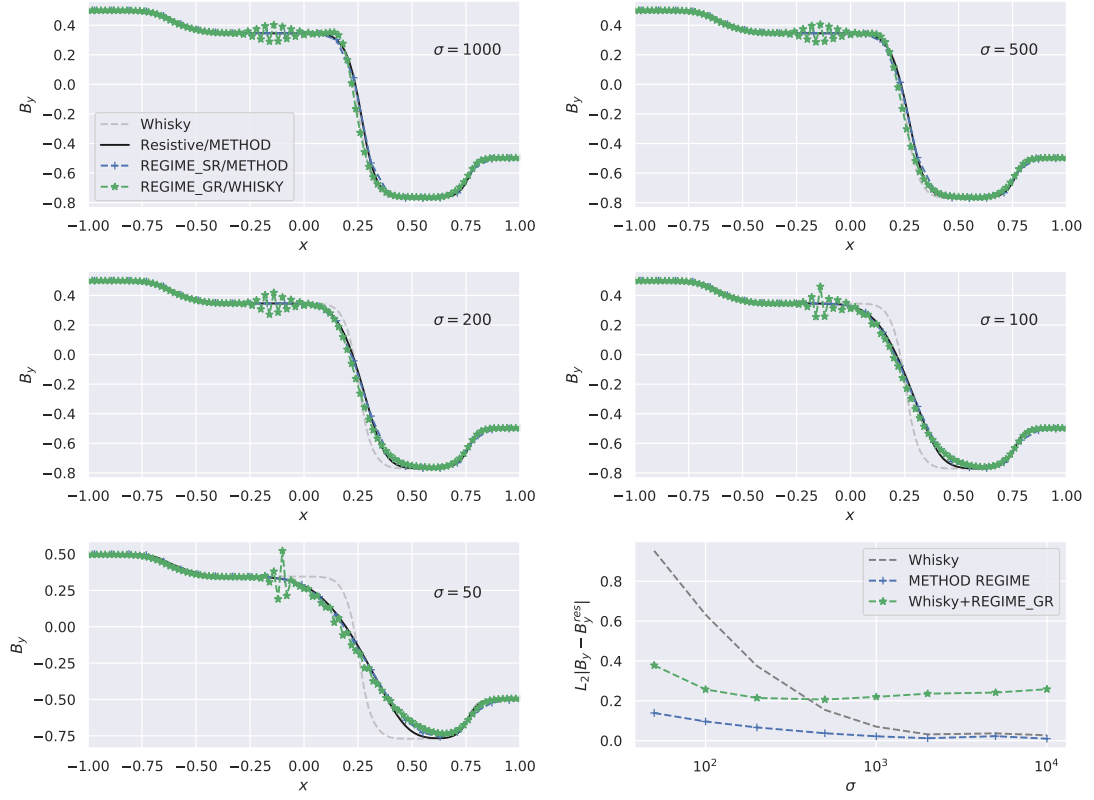


FIGURE 6.3: Convergence of the Brio-Wu shock tube with the conductivity, σ , at a final time of $T = 0.8$. The oscillations left of the shock are caused by the non-staggered vector potential formulation in **Whisky**. Clearly, GRR successfully diffusing the magnetic field, with similar results to the **METHOD** code.

Baumgarte and Shapiro, 1998; Nakamura et al., 1987) formulation of the ADM equations with $1 + \log$ slicing and the Gamma-driver shift condition, allowing for long term, stable simulations of dynamic spacetimes. Integration of the equations of motion are performed using the method of lines, and a fourth order Runge-Kutta method.

For the magnetohydrodynamics, we make use of the **Whisky** code (Giacomazzo and Rezzolla, 2007) with the vector potential formulation for the magnetic fields (Giacomazzo et al., 2011; Etienne et al., 2010) as the resulting magnetic fields will be divergence-free by construction. To minimise any divergence-free violations at the mesh boundaries, **Whisky** uses the generalised Lorenz gauge (Etienne et al., 2012; Farris et al., 2012). The equations of motion are in the conservative, *Valencia* formulation, with the numerical flux at cell interfaces computed using the HLLE approximate Riemann solver (Harten et al., 1983) and the primitive variables reconstructed using the piecewise-parabolic method (Colella and Woodward, 1984). Finally, to reduce numerical issues with the primitive transformation, when the rest-mass density in a computational cell drops below a specified floor value, $\rho_{\text{floor}} = 10^{-13} \approx 6.2 \cdot 10^4 \text{gcm}^{-3}$, **Whisky** treats that cell as the atmosphere, setting the density to the floor value and any velocities to zero.

The simulations we perform in this section are inspired by the ideal MHD simulations of Kawamura et al. (2016), such that we evolve the same initial data with identical numerical methods. The initial data is generated using the **Lorene**⁷ library, and consists of two equal-mass

⁷www.lorene.obspm.fr

NSs, each with a gravitational mass of $1.5M_{\odot}$, as originally reported in [Rezzolla et al. \(2011\)](#). We set up the stars using an ideal-fluid EOS, initially using the polytropic EOS with polytropic index $\Gamma = 2$ and polytropic constant $K = 100$. When evolving the system, we use the Γ -law EOS with $\Gamma = 2$, as done in [Kawamura et al. \(2016\)](#). To save on computation, we also employ a reflection symmetry across the orbital plane.

The magnetic fields are initially poloidal, aligned along the orbital rotation axis, and confined to the interior of the NSs. Confining the magnetic field in this way is done to avoid numerical problems in the atmosphere and has no impact on important phenomena such as the formation of a jet ([Ruiz et al., 2016](#)), and typically a few ms after merger the toroidal component is of the same strength as the poloidal field ([Giacomazzo et al., 2011](#)). The magnetic fields are added *a posteriori* to the initial data produced by *Lorene*, using

$$A_{\phi} = \bar{\omega} A_b \max(p - p_{cut}, 0)^{n_s} \quad (6.42)$$

where $\bar{\omega}$ is the coordinate distance from the NS spin axis, p_{cut} is 4% of the maximum pressure and determines where the magnetic fields go to zero, $A_b = 2.2$ and n_s is the degree of differentiability of the magnetic field strength ([Giacomazzo et al., 2011](#)). Adding the magnetic field in this way perturbs the NSs such that they are out of equilibrium at the beginning of the simulation, however the size of this perturbation is small enough so as not to have a significant effect on the dynamics of the simulation ([Dionysopoulou et al., 2015](#)), and the system typically reaches quasi-equilibrium within one orbit.

The initial magnetic fields described here will result in post merger magnetic field strengths many orders of magnitude weaker would be expected if the simulations were better resolved ([Kiuchi et al., 2015, 2017](#)). This can be accounted for by using stronger initial magnetic field configurations such as in [Ciolfi et al. \(2017\)](#), resulting in more realistic magnetic behaviour post-merger. Preliminary results of a highly magnetised merger with the same conductivity is presented at the end of this chapter. However, we are only interested in a rough comparison of ideal and resistive simulations at this stage, rather than a rigorous investigation of realistic physics. These magnetic configurations should allow us to determine the scale of the impact of a somewhat realistic electrical resistivity. Positive results will require further investigation to confidently extract the effect this resistivity will have on observables.

Including resistivity

The resistive merger simulations performed by [Ponce et al. \(2014\)](#) and [Dionysopoulou et al. \(2015\)](#) both used third order accurate IMEX integrators ([Pareschi and Russo, 2005](#); [Palenzuela, 2013](#)) to evolve the full set of resistive GRMHD equations. This set-up allows them to evolve the force-free and electro-vacuum equations, respectively, in the atmosphere via profiles for the conductivity and Ohm's law, effectively decoupling the electromagnetic fields and the matter. However, due to problems related to stiffness in the source of the electric fields, the highly conducting core is unrealistically resistive, with a conductivity of $\sigma = 2 \cdot 10^{11} \text{s}^{-1}$. The significant increase in computation time as a result of the added stiffness (rooted in the necessity for implicit time-integration) forced the authors to do some numerical gymnastics in [Dionysopoulou et al. \(2015\)](#). In order to get the simulations to finish in sensible times, an additional velocity boost

of 20% the orbital speed kicked the NSs towards the centre of the binary. The authors claim that these grazing collisions lead to qualitatively similar results as collisions from quasi-circular orbits, however this remains to be seen as one would expect differently excited modes in the merger remnant (Chirenti et al., 2017) and distinctive GW and EM observables from eccentric binaries (East et al., 2016).

In our simulations, we wish to model more physical set-ups, particularly regarding value of the resistivity in the NS interior, and thus we refer to Harutyunyan and Sedrakian (2016) on typical NS conductivities. In their discussion, Harutyunyan and Sedrakian (2016) consider the electrical conductivity tensor as a function of local temperature and density for matter composed of carbon and iron, assuming a warm NS star crust ($T \lesssim 10\text{MeV}$). These temperatures are greater than would be expected for the NSs pre-merger, however as the resistive diffusion timescales for smooth and weak magnetic fields are orders of magnitude longer than the dynamical timescale of the orbital period (Harutyunyan et al., 2018), this disparity should not be important during the inspiral. Post merger, however, shocks will be present and temperatures are expected with approximately this magnitude, thus the calculations of conductivities in Harutyunyan and Sedrakian (2016) should better apply.

With this in mind, for our simulations we choose a homogeneous conductivity of $\sigma = 10^{12} \approx 2 \cdot 10^{17}\text{s}^{-1}$. This corresponds to the smallest values of the first two diagonal terms in the conductivity tensor, σ_0 , of Harutyunyan and Sedrakian (2016), for a magnetic field of $B = 10^{12}\text{G}$, a temperature of $T = 1\text{MeV}$ and a density of $\rho \sim 10^6\text{g/cm}^3$. Of course the inner regions of the NSs will be much more highly conducting than we are modelling them with this prescription, however, if we can not demonstrate that there is a significant change in the simulations with values such as these, we would be able to discount resistivity as an important feature in future merger simulations. Our aim here is to determine whether a realistic order of magnitude resistivity can impact the results of merger simulations, and if so what those effects might generally be. We save yet more realistic resistivities for future work.

We run simulations of two different resolutions corresponding to the low and medium resolution ideal runs of Kawamura et al. (2016). The resolutions correspond to a finest grid spacing of 277m and 222m respectively. A convergence study in Kawamura et al. (2016) demonstrates that the solution, particularly the remnant collapse time and magnetic field amplification, is highly sensitive to the resolution. Despite this, the authors of Kawamura et al. (2016) are confident enough that a resolution of 222m is sufficient to describe the bulk motion, general EM structure and the gravitational wave signals that $\Delta x_{\min} = 222\text{m}$ is their default resolution. Similarly, we argue that we are interested in determining the general features that including resistivity will result in, and that for this, this medium resolution is sufficient.

The initial data we are using has been studied in a number of other papers (Rezzolla et al., 2011; Kawamura et al., 2016). Using GRR as opposed to the full resistive model means we can evolve the system using conventional, explicit time integrators, one result of which is that there is no need to modify the initial data in order to keep runtimes manageable. The results of this section, therefore, represent the most realistic, resistive description of NS matter in BNS merger simulations to date.

6.4.2 Results

In the following section, we report the first results of the GRR extension applied to merger simulations. Despite the lower resolutions we employ here than in previous resistive merger simulations (Dionysopoulou, 2016), we demonstrate that a more realistic resistivity of NS matter, a factor 10^6 smaller in this case, can still meaningfully affect the merger dynamics. Throughout this section, lines labelled with REGIME12 correspond to the low resolution run, and MedREGIME12 to the medium resolution counterpart. We will also refer to the REGIME BNS simulations as rBNS, iBNS for the ideal simulations, and the BMP in the ideal and resistive simulations as iBMP and rBMP respectively. Much of the plotting has been done using a closed-source version of the PyCactus⁸ library, in addition to some routines borrowed from Kawamura et al. (2016); Ciolfi et al. (2017). We also define $t = 0$ as the time of maximum gravitational wave strain, as is common.


6.4.3 Low resolution simulation

First, we present the results from the lower resolution run, with a finest grid spacing of 277m. At this resolution, we expect there to be a significant amount of numerical diffusion, but it still allows us to assess the stability of REGIME, particularly in the presence of shocks, and determine general features present in BNS merger simulations. An overview of the simulations is given in Figures 6.4(a) and 6.4(b), where we show the rest-mass density at three different times on the $y = 0$ and $z = 0$ planes, respectively. The total mass of the binary is within the hypermassive regime for an isolated, spinning NS, and so the BMP is a hypermassive neutron star (HMNS), lasting just under 7ms before collapsing to form a black-hole.


Throughout the evolution, from the inspiral and point of merger through to the delayed collapse to a black-hole, the resistive model is stable and handles discontinuities well. As a result, the behaviour of the matter is qualitatively what one would expect from a low resolution ideal MHD simulation. This can be seen more quantitatively in the plot of the maximum rest-mass density, where we compare the rBNS and iBNS simulations, Figure 6.5(a). For the first 5ms after merger, the two simulations are indistinguishable in this respect. The reason for such small variations in the matter is that resistive diffusion primarily affects the magnetic fields, which at the start of these merger simulations are contained within the NSs and do not interact with one another. The small changes in the magnetic fields at this stage, therefore, do not induce significant changes in the behaviour of the matter. Only just before the black-hole formation, and after the magnetic fields have had time to interact, are there any noticeable differences in the matter and the maximum rest-mass density evolution, resulting in the resistive simulation collapsing slightly before the ideal simulation.

If instead we look at the evolution of the magnetic fields, Figures 6.5(b), 6.5(c) and 6.5(d), we see larger differences between the two simulations at earlier times. The maximum of the magnetic field for the rBNS occasionally peaks above the ideal, a behaviour which is reflected in both the mean magnetic field and total magnetic energy plots—for both these quantities, their size is larger than in the ideal MHD case. Another common feature is the drop in the magnetic quantities for the resistive simulations after black-hole formation.

⁸www.github.com/wokast/PyCactus



(a) Density cut of the ideal (top) and resistive (bottom) BNS simulations for the rest-mass density in the x - z plane, for the times shown where $t = 0$ is the time of maximum GW strain and t_{BH} is the time of black-hole formation. The horizon is indicated by the red circle in the last column.




(b) Same as Figure 6.4(a) for the x - y plane. The results are qualitatively the same, with very few differences between the models in the matter.


FIGURE 6.4: Density slices in the x - z (top) and x - y (bottom) planes

In a resistive simulation, the magnetic fields are not locked to the fluid motion as they are in the ideal case (see the discussion on magnetic freezing, Section A.1). As a result, energy in the magnetic fields that is produced through magnetic winding within the BMP can not be as easily transported away from the central regions of the simulation by the fluid flow. This should mean that initially the magnetic energy would be more confined to the central regions, leading to higher energy densities, and greater maximum and mean magnetic field strengths. After collapse, however, a larger portion of the total magnetic energy is near the centre and swallowed by the BH horizon. The remaining fraction would be smaller than in an ideal simulation, leading to larger drops in the magnetic quantities for the resistive simulation. This is exactly the behaviour we see.


Generally, both models demonstrate the same bulk features expected in merger simulations. The seed magnetic fields are amplified exponentially after merger, a result of the vortical fluid




(a) Timeseries of maximum density for the low resolution merger.



(b) Timeseries of maximum magnetic field strength.



(c) Timeseries of mean magnetic field.



(d) Timeseries of total magnetic energy.

FIGURE 6.5: Evolution of various quantities for the low resolution ideal and resistive merger simulations. Circles indicate the time the horizon is first detected and the vertical yellow line shows the time of merger. Solutions are qualitatively the same, with somewhat larger differences between the models in the magnetic fields compared to the matter.




FIGURE 6.6: Evolution of the poloidal and toroidal components for the iBNS (left) and rBNS (right).

motion due to the KHI winding up the magnetic field lines. This continues for a few milliseconds, then ceases due to the low resolution filtering out higher frequency modes that are essential for continued amplification. It is a general feature of these types of simulations that the finer the grid spacing, the greater the effect of the amplification of the seed magnetic fields (Price and Rosswog, 2006; Zhang et al., 2009; Obergaulinger et al., 2010; Kiuchi et al., 2014, 2015; Giacomazzo et al., 2015; Kawamura et al., 2016; Ciolfi et al., 2017).

In addition to amplification, magnetic winding also results in the initially poloidal seed magnetic fields developing a significant toroidal component, as can be seen in Figure 6.6. Shortly after merger, the poloidal and toroidal components of the magnetic field reach equipartition despite the initial configuration being entirely poloidal. This is a completely general feature of BNS simulations, that initially poloidal magnetic fields quickly develop strong toroidal components after merger.

With regards to the rest-mass density, the binary merger product is initially supported from collapse via differential rotation. During this stage, we see the rest-mass density oscillate around $6 \cdot 10^{14} \text{gcm}^{-3}$. Angular momentum is transported away from the core, and energy is released through gravitational radiation, until the core is no longer supported from collapse and forms a BH. This is followed by a sharp drop in the maximum rest-mass density and the magnetic quantities as any regions that fall within the BH horizon are essentially lost from the simulation.

Post-collapse, the system comprises a black-hole, surrounded by a low density accretion torus. At this point, one would expect the magnetic behaviour to be governed by the magneto-rotational instability. The MRI requires a minimum of ~ 10 grid cells covering the wavelength of the fast growing mode in order to be fully resolved (Kawamura et al., 2016; Siegel et al., 2013), however at the low resolution presented here this is not achieved, and as such the amplification is not present and the strength of the magnetic fields continues to fall for both simulations.

6.4.4 Medium resolution simulations

For the remainder of this chapter, we refer to the solutions of the medium resolution simulations. As numerical diffusion will be reduced by a significant factor, the effects of resistive diffusion are





FIGURE 6.7: Slices of the density in the x - z (top) and x - y plane (bottom) for the medium resolution merger simulation. Filled red circles in the right column indicate the BH horizon. Qualitatively there are significant differences between the ideal (upper panels) and resistive (lower panels) simulations, especially post-collapse.


more pronounced, and hence indicate the kinds of behaviours that arise from the resistive description. Typically, for higher resolution merger simulations, one should expect a more delayed collapse and more greatly amplified magnetic fields via Kelvin-Helmholtz and magneto-rotational instabilities (Kiuchi et al., 2014; Kawamura et al., 2016).

6.4.4.1 Overview of evolution


First, we describe the changes in the evolution of the matter as a result of including resistivity. Figures 6.7 give an overview of the matter distribution via slices in the density in the x - z (top) and x - y planes (bottom). General features between the simulations are comparable, just as in the low resolution case.



(a) Evolution of the maximum rest-mass density where filled circles represent time the BH horizon is first detected. The collapse of the BMP in the resistive simulation is delayed by ≈ 2 ms compared to the ideal run.



(b) Total ejected matter at 295km. The delayed collapse of rBNS results in a greater mass of ejecta.



(c) Gravitational wave strain (top) and frequency (bottom) plots.

FIGURE 6.8: The delayed collapse of the rBNS BMP induces changes in the total amount of ejected matter and the gravitational wave signal.

We can be more quantitative by looking at the evolution of the maximum rest-mass density, Figure 6.8(a), where we see the rBMP completes an additional oscillation before collapsing to a BH, $\approx 2\text{ms}$ after the iBMP collapses. This result agrees with the findings of [Dionysopoulou et al. \(2015\)](#), demonstrating that, even with a significantly reduced magnitude, including resistivity can have a major impact on the matter in the system. In addition, the fundamental mode frequency of the rBMP is slightly greater than for the iBMP, indicating that a change in the resistivity can effectively modify the equation of state via the interaction of the matter with the EM fields—reconfiguring the magnetic fields through resistive effects will lead to changes in matter interactions and, ultimately, the oscillation of the remnant.


A consequence of the delayed collapse for rBNS can be seen in the outflow plot, Figure 6.8(b)—the extended lifetime of the BMP results in an increase in the total ejected mass, which may impact on the later formation of a jet. For more concrete statements, however, longer term simulations are required ([Ciolfi et al., 2019](#); [Ciolfi and Kalinani, 2020](#)).

In addition to the total ejecta mass, the delayed collapse has an effect on the resulting gravitational wave signal, Figure 6.8(c). Whilst the waveforms of the two simulations share most of their features, they begin to diverge a few ms after merger, with rBNS showing an extended ringdown. The GW strain, top of Figure 6.8(c), also gives information regarding the oscillation modes of the remnant. The fundamental mode of the BMP is seen as the rapid oscillation, 3-4kHz, during the ringdown. There is also a second mode modulating this high frequency signal at around 1kHz. This mode is present for both simulations, but appears to decay more rapidly in the resistive simulation. This suggests that some process is removing energy from this mode of oscillation and transferring it elsewhere.


There are a number of ideas about where this energy is transferred to. First, it may be the case that the energy is transferred to different modes of oscillations within the BMP. However, the tendency for energy to be filtered towards more fundamental modes makes this seem unlikely, as there appears to be little change in the oscillation of the fundamental mode. Second, the energy of the secondary mode may be being dissipated via some resistive effect, producing entropy and heat. Finally, and similar to the production of heat, resistive processes may be transferring the energy into the magnetic fields, changing their properties and configuration.

It may be the case that a consequence of the latter two is an increase in the total pressure in some regions within the BMP, which in effect is resisting collapse in the resistive case. With the current methods and data available, however, confidently pinning down any of these pathways is challenging. We shall investigate some of these pathways in the next section.


In Figure 6.9 we present the evolution of the magnetic fields. In this medium resolution run, we see fewer signs of the magnetic energy remaining near the core in the initial plots than we did for the lower resolution case, suggesting that this is not a general feature of resistivity. Whereas we previously saw a transition from higher magnetic energies pre-collapse to lower magnetic energies post-collapse for rBNS, at this resolution we instead see differences mainly due to the delayed collapse. The milliseconds between the collapse of the iBMP and rBMP lead to continued growth of the magnetic quantities in the resistive simulation, followed by comparatively similar drops to the ideal simulation once the horizon is formed.



(a) Maximum magnetic field strength evolution.



(b) Mean magnetic field evolution.



(c) Total magnetic energy evolution.

FIGURE 6.9: Evolution of various quantities for the medium resolution ideal and resistive merger simulations. Circles indicate the time the horizon is first detected. Solutions between the models are qualitatively the same, with somewhat larger differences between the models in the magnetic fields compared to the matter.

At first glance, it appears that the delayed collapse of rBMP is the most significant difference between the two simulations pre-merger. The delay results in an extra oscillation of the rBMP that slightly extends GW ringdown and allows, for a short period of time, continued growth of the magnetic quantities. In the next section, we analyse some potential causes of the delayed collapse of rBMP.

Finally, as the impact of REGIME is generally largest for the magnetic fields, and that resistivity is closely linked to the electromagnetic properties of the system, the effects of including resistivity in merger simulations should grow with increasing magnetic field strength. It is

well known that numerical simulations of mergers can not generate sufficiently strong magnetic fields from realistic initial data (Kiuchi et al., 2014)—we expect the magnetic fields to grow more rapidly in nature. As a result, we should also expect simulations of resistive mergers to demonstrate larger differences for finer grid spacings, or for simulations in which higher strength seed fields are imposed. We investigate the latter in Section 6.4.5

6.4.4.2 Delayed rBNS collapse


To understand the origin of the delayed rBMP collapse, we will need to consider the distribution of various quantities throughout the BMPs, with the differences between the distributions hinting at possible causes for the delay. We will focus on snapshots at three times after merger, $t \approx 5.65$, 7.6 and 9.55ms, corresponding to the last three times that the peaks of maximum rest-mass density of the two simulations coincide before collapse (see Figure 6.8(a)). For the final time shown, there should be a significant difference between the two solutions that results in the delayed collapse of the rBMP.

Part of the following analysis will require looking at angle-averaged quantities in the x - z plane. For this analysis to be accurate, we need to be confident that the data is approximately symmetric about the z -axis. Figures 6.10 show the three snapshots of the rest-mass density at four heights, $z = 0, 15, 30$ and 45km, for the two simulations. Similarly, Figures 6.11 show the magnetic norm, $|\mathbf{B}|$, at the same snapshots and heights. We can see that there is no sign of any large scale, asymmetric features such as spiral arms, and so can confidently assume that the data is approximately symmetric about the z -axis for both simulations.


Next, we investigate more quantitatively the distribution of various quantities that may be related to the delayed collapse of rBMP. All of the quantities we discuss in the following are the angle-averaged, and a function of distance from the centre of the BMP in the x - z plane, Figure 6.12. These plots describe what the average of a quantity is within a certain radius in each slice. We have not shown the corresponding slices in the x - y plane as almost all the quantities of interest have values at least an order of magnitude less than in the vertical slices. Equivalent slices in y - z plane have also been omitted as they show the same features as the x - z slices.

First, we consider the matter distribution. From the top row of Figure 6.12 we can see that whilst there is some small variation, the matter is distributed nearly identically in both simulations. The similarities are a result of the time at which the snapshots are taken, and should mean that differences in other interesting quantities, such as internal energy density for example, are not a result of variations in the matter distribution but rather some other process, likely due to the inclusion of resistivity.

The second row shows the internal energy density distribution, e . If some resistive process was more efficiently transferring energy into the fluid or increasing entropy, this would be shown as a rise in e . There are some differences at $t = 7.6$ ms, with a larger internal energy densities for iBNS at all radii, but this feature seems to be insignificant with variations of $\lesssim 10\%$. Furthermore, the differences seem not to persist, and are not present in the later snapshot. We can conclude from this that the damping of the oscillations for the rBMP is not caused by transfer of this energy into heat.




(a) Rest-mass density slices for iBNS at the heights and times shown.




(b) Rest-mass density slices for rBNS at the heights and times shown.

FIGURE 6.10: Rest-mass density slice snapshots at $t = 5.65$, 7.6 and 9.55 ms for iBNS (top) and rBNS (bottom). Slices have been taken at heights $z = 0, 15, 30$ and 45 km. Clearly there are no large scale asymmetric features such spiral arms, showing that the matter is distributed roughly symmetrically about the z -axis.

In the final two rows we plot the distribution of the magnetic norm and the corresponding magnetic energy density. These plots show significantly larger differences between the two simulations. The ideal simulation shows a general trend in which the magnetic quantities begin large and evolve to lower values. This is in contrast the resistive simulation, in which we see the opposite effect, culminating in a stronger and more energetic magnetic fields by the time the iBMP collapses. These results do not explain why the rBMP collapses later than the iBMP, they do suggest, however, that the magnetic fields may be responsible in some capacity. To a degree,



(a) Magnetic norm slices for iBNS at the heights and times shown.



(b) Magnetic norm slices for rBNS at the heights and times shown.

FIGURE 6.11: Magnetic norm slice snapshots at $t \approx 5.65$, 7.6 and 9.55ms for iBNS (top) and rBNS (bottom). Slices have been taken at heights $z = 0, 15, 30$ and 45km, and show that the magnetic norm is also symmetric about the z -axis.

we should expect this, as the effects of resistivity tend to be more pronounced on the magnetic fields.

Fundamentally, the delayed collapse of rBMP must be a result of some additional outward net force, or pressure gradient, that is present in the resistive simulation but absent from the ideal. One hypothesis is that the pressure gradient associated with the magnetic fields is greater in rBMP, and that this small change to the total pressure is enough to prevent the collapse of the critical HMNS for a short period of time. To understand the interplay of the different sources

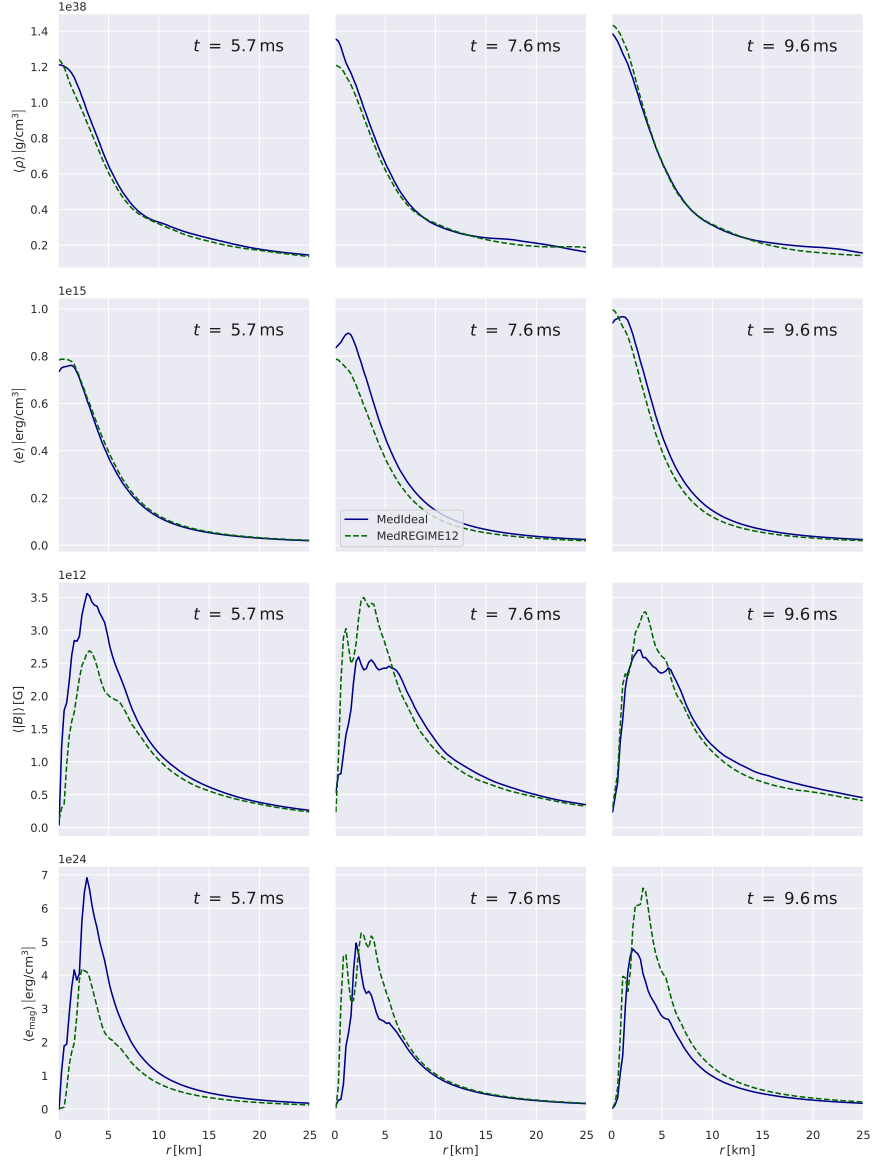


FIGURE 6.12: Angle-average quantities as a function of distance from the BMP centre in the x - z plane. The time of each slice is shown. These plots show what the average of a quantity is within a certain radius in the x - z plane.

of pressure in the system, we now present the angle-averaged pressure distributions and their corresponding gradients, Figure 6.13, at the final snapshot where $t = 9.55$ ms.

The first row shows the hydrodynamic pressure and its associated gradient. Generally, the two simulations show rather similar profiles in this respect, with negative pressure gradients providing resistance against collapse. Quantitatively, the profiles are essentially identical for $r \gtrsim 3$ km, with the rBMP showing a larger negative ∇p_{hydro} for $r \lesssim 3$ km, and therefore resisting collapse more strongly.

The next row shows the pressure and associated pressure gradient of the magnetic fields. We noted from Figure 6.12 that there is a significant difference in the strength and energy contained within the magnetic fields, raising the question as to whether an increased magnetic pressure (gradient) is responsible for the delay. From the plot of the magnetic pressure, it is clear that

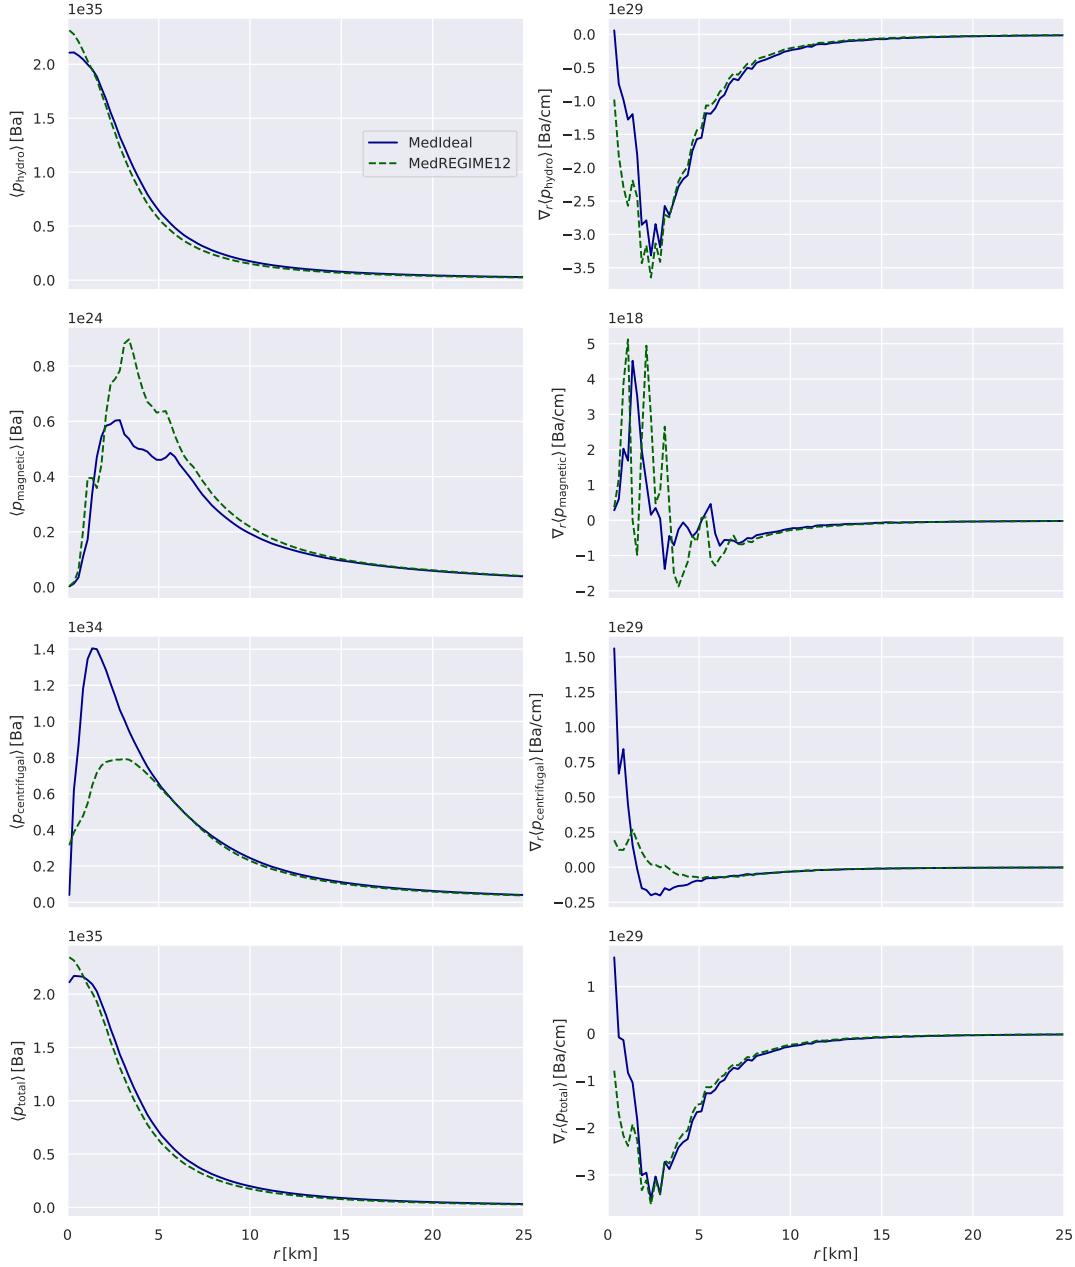


FIGURE 6.13: Angle-averaged pressure distribution in the x - z plane at $t \approx 9.55$ ms. Plots on the left correspond to the different pressure sources (top to bottom: hydro pressure; magnetic pressure; centrifugal pressure; total pressure), with their corresponding pressure gradients on the right hand side. Positive values in the plots on the right favour an early collapse.

whilst there are substantial changes in the profiles, the magnitude is far too small to affect the collapse time of rBMP (when compared to the magnitude of the hydro pressure). Of course, this effect may grow and be a contributing factor for simulations of larger strength initial fields, but for these simulations can be discounted.

In the third row, we show the effective centrifugal pressure, $p_{\text{centrifugal}}$, within the two BMPs and calculated as the centrifugal force per area, and the corresponding effective pressure gradient, $\nabla p_{\text{centrifugal}}$. Here, we can see magnitudes that are comparable to those of the hydrodynamic pressure, and differences between the ideal and resistive simulations that are significant. In the

core region, $r \lesssim 3\text{km}$, differences in the pressure gradient profiles between the simulations are an order of magnitude, with the iBMP showing a substantial *positive* pressure gradient. A positive effective pressure gradient would act to suck matter towards the core, promoting collapse rather than resisting it. We can see the effect of this in the final row, where we plot the sum of all pressure sources,

$$p_{\text{total}} = p_{\text{hydro}} + p_{\text{magnetic}} + p_{\text{centrifugal}}. \quad (6.43)$$

The addition of the effective centrifugal pressure has a large effect on the total pressure gradient in the core, with the iBMP showing signs of a pressure profile that will promote collapse in contrast with the profile of the rBMP which is opposing the collapse strongly. It is likely that the size of this effect is large enough to be a contributing factor to the subsequent collapse of the iBMP, and the delay in the collapse of rBMP.

Centrifugal effects

As we have identified differences in the effective centrifugal pressure as a likely contributing factor for the delayed collapse of rBMP, we will now analyse it in more detail. Figure 6.14 shows the distribution of the effective centrifugal pressure gradient in the x - z slice for the ideal (top row) and resistive (bottom row), at each snapshot in time. In the initial snapshot ($t \approx 5.7\text{ms}$) the profiles are very similar, with no substantial differences between the simulations. By the second and third snapshots, however, there is a large, positive effective pressure gradient that has developed in the ideal simulation that is not present in the resistive simulation. Furthermore, the effect is highly concentrated in the region where the rest-mass density is greatest, and where the horizon is formed—i.e. the core.




FIGURE 6.14: Evolution of the effective centrifugal pressure gradient in the x - z plane. The profile of the centrifugal pressure gradient is similar at $t \approx 5.7\text{ms}$ for both simulations, however, at later times a strong positive gradient develops in the core of iBMP which is absent from rBMP.

The relative contribution of the effective centrifugal pressure can be seen in Figure 6.15, where we show slices in the x - z plane at $t \approx 9.55\text{ms}$ of the gradients of the centrifugal pressure (first column), the magnetohydrodynamic pressure (second column) and the total pressure (final column). The total pressure profile in the core of the iBMP is significantly altered by the effective centrifugal pressure, with a large increase in the size of the region of positive gradient.




FIGURE 6.15: Slices in the x - z plane of the contributions of the different pressures (first two columns) and their sum (final column), at the final snapshot, $t \approx 9.55$ ms, for the ideal (top row) and resistive (bottom row) simulations. Including the effective centrifugal pressure gradient has a large impact on the total pressure gradient profile in the core of the BMPs.

This suggests that this effect can have a large impact on the stability of the HMNS remnant, providing a mechanism to speed up or slow down eventual collapse. A similar region exists in the core of the rBMP, but is significantly reduced in both extent and magnitude.

Ultimately, the effective centrifugal pressure is a result of the mass distribution and the profile of the velocity that is perpendicular to the radius. We have seen that the matter is distributed in the same way in both simulations, Figure 6.12, as thus can conclude that the velocity profiles must be different in the BMP cores. Differences in the velocity profiles will be due to the braking of the flux-freezing condition of ideal MHD in the resistive simulation, Section A.1. In the ideal simulation, the fluid cannot cross the magnetic field lines, resulting in a force that resists the fluid motion, changing the fluid’s velocity vector and strongly reconfiguring the magnetic fields. This will not occur to the same extent in the resistive simulation, whereby the fluid may cross field lines to some degree.




FIGURE 6.16: Slices in the x - z plane of the toroidal (first column) and poloidal (second column) components of the magnetic field for the ideal (top row) and resistive (bottom row) simulation.

Slice is taken at $t \approx 9.55$ ms. The poloidal component is noticeably stronger in the rBMP.




FIGURE 6.17: Time evolution of the poloidal and toroidal components of the magnetic fields for the ideal (left) and resistive (right) simulations, where the yellow vertical line is the time of merger and the black vertical line indicates the formation of the horizon.

One would expect this interplay of the fluid motion and magnetic field lines to result in changes in the magnetic configuration and the fluid velocity vector for the ideal and resistive simulations. The extra freedom of the magnetic fields in the resistive simulation to move in contrast to the fluid flow should result in differences in the growth of toroidal and poloidal components. There are hints of this in Figure 6.16, where we can see a stronger poloidal component for the resistive simulation.

This is made more quantitative in Figure 6.17, where we show the individual components of the magnetic field and their sum. Despite the fields initially being totally poloidal, a toroidal component quickly develops and reaches equipartition with the poloidal field after merger, subsequently becoming dominant. The growth rate of the toroidal field appears to be the dominating factor for the growth of the mean magnetic field strength post-merger. We can see that this rate of growth is suppressed more in the resistive simulation.

With the current data and methods of analysis, it is difficult to pin down the exact causal chain leading to the delayed collapse of the rBMP. It seems likely that it is, at least in part, due to a change in the velocity profile within the core of the BMP. Changes in the fluid flow result in smaller, positive pressure gradients in the resistive simulation, which ultimately prevent collapse compared to the ideal simulation. The mechanism for resistivity altering the fluid velocity will most likely be related to the fluid’s interaction with the magnetic field—resistivity has a large effect on the magnetic field configuration, and leads to less coupling between the motion of the fluid and electromagnetic fields than in the ideal simulations. Further investigation requires better methods for understanding the configuration of the magnetic fields and fluid trajectories.

6.4.4.3 Post-collapse dynamics

So far, the discussion has been focussed on the behaviour of the two models from the point of merger leading up to the collapse. In this section, we will turn our attention to the dynamics after the BH horizon has formed. Important features to study for the post-collapse dynamics are generally related to the formation of a magnetic funnel and the right conditions for an incipient jet, bound and unbound ejecta outflow and the composition of the matter for the purpose of

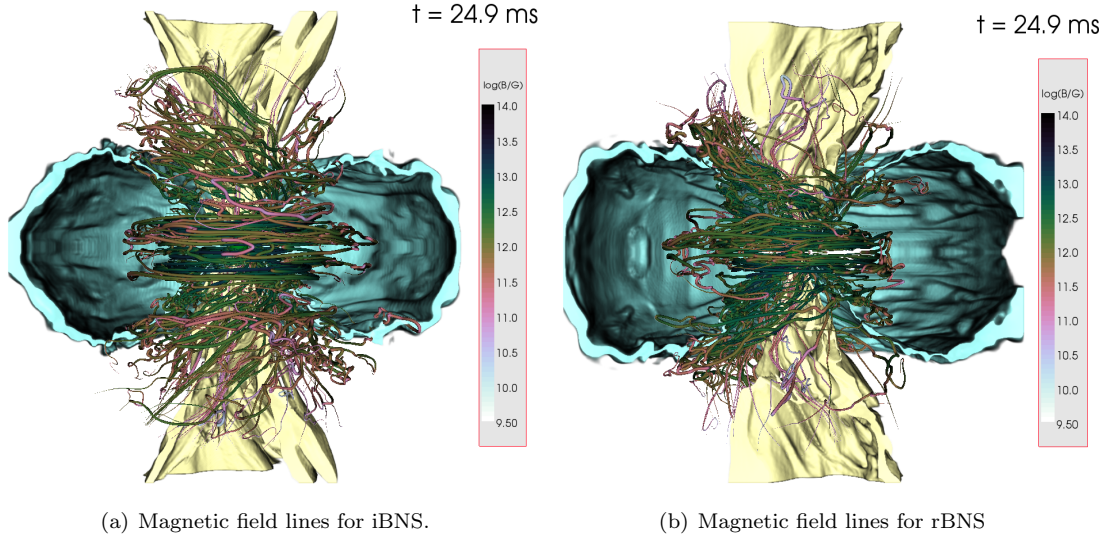


FIGURE 6.18: Magnetic field line configuration for the ideal (left) and resistive (right) simulations ≈ 24.9 ms after merger. Two surfaces of constant density are shown for 10^8 (yellow) and 10^{10} g/cm³ (cyan).

kilonova signatures. Regarding the development of a kilonova, the simulations presented here have little to say. Due to the simplicity of the equation of state, and the fact that we have not evolved the composition of the matter (via the electron fraction, Y_e), it is difficult to assert how the ejecta properties will have been affected by the inclusion of resistivity.

More can be said regarding the formation of a jet, however. Despite our simulations only running for a short period of time after merger (≈ 25 ms), we can see the initial stages of the formation of a magnetic funnel in Figure 6.18 in both simulations. The magnetic field lines are being wrapped around the axis of rotation, forming a helical pattern, whilst in the disc the magnetic field is clearly becoming toroidally dominated. Furthermore, the region above the poles comprises of a poloidally dominated magnetic field situated in a lower density atmosphere. Eventual infall of matter in this region, combined with further magnetic amplification by magnetohydrodynamic instabilities, such as the MRI, should lead to this region becoming magnetically dominated, $\beta = p_{\text{magnetic}}/p_{\text{hydro}} > 1$, and favourable configurations for a collimated, mildly relativistic outflow (Ruiz et al., 2016; Kawamura et al., 2016). The state of the configuration for the resistive simulations is comparable to that of the ideal simulation, and matches those at a similar stage in Kawamura et al. (2016). Longer term simulations will show that the winding of the magnetic fields about the rotation axis becomes more pronounced and ordered.

To make the formation of a jet more quantitative, we now employ the same analysis as in the pre-collapse case by looking at the angle-averaged quantities in a vertical slice. Figure 6.19 shows the radial dependence of the mean magnetic field (top row), rest-mass density (middle row) and magnetic energy density (bottom row) for both simulations in the x - z plane at three times after the horizon is formed, $t - t_{\text{BH}} \approx 4.0, 8.6$ and 13.2 ms, where t_{BH} is the time of formation of the horizon.

Conditions for a jet require a magnetically dominated region above the poles. We can see in Figure 6.19 that the resistive simulation possess stronger magnetic fields, particularly within $r < 15$ km from the BH centre (top row). This should lead to a higher likelihood of a relativistic

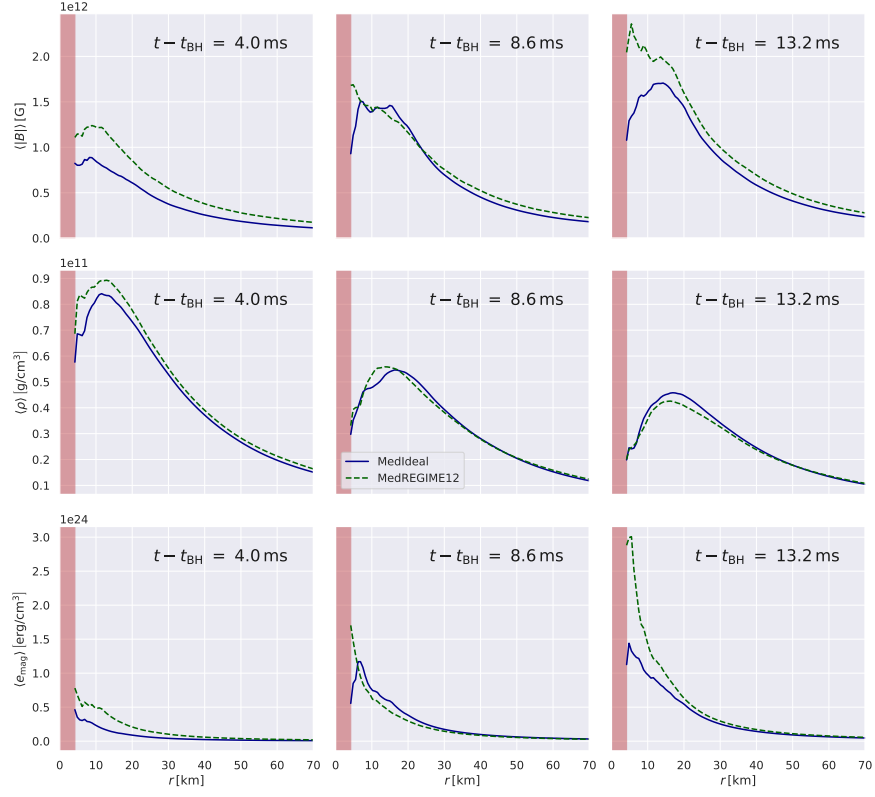


FIGURE 6.19: Angle-averaged quantities in the x - z plane at three times after the horizon is formed (indicated). The vertical red bar shows the approximate position of the BH horizon. Note the r -axis now spans a much larger distance from the centre of the BH than in Figures 6.12 and 6.13.




FIGURE 6.20: Horizontal magnetic field slices at heights $z = 0, 15, 30$ and $45 km$ and times $t - t_{BH} \approx 4.0, 8.6$ and $13.2 ms$ for the ideal (left) and resistive (right) simulations. The resistive simulation possess stronger magnetic fields out of the plane compared to the ideal simulation.

outflow for the same mass distribution (middle row), although at this stage the difference is only minor. Regardless, the significant increase in the magnetic energy density in the resistive simulation (final row) can only aid jet formation.

Finally, Figures 6.20 show the magnetic field norm in horizontal slices at a number of different heights, $z = 0, 15, 30$ and $45 km$. It can be seen from these figures that the strength of the

magnetic field out of the x - y plane is greater in the resistive simulation, seen as brighter spots in the right hand plots compared to the left for $z \neq 0$. The regions of stronger magnetic fields correspond to where the magnetic funnel is forming above the poles. As mentioned, stronger magnetic fields in this region should lead more likely to the formation of an incipient jet.

Other features that are important in the post-collapse dynamics are magnetohydrodynamic instabilities such as the MRI. Unfortunately, due to the insufficient resolution and relatively weak magnetic fields in these simulations, we have not fully resolved the fastest growing mode of the MRI. As a result, we should expect that the amplification of the magnetic field in the disc and the corresponding outflow to represent a lower bound to what is expected.

6.4.5 Highly magnetised case

A measure that characterises information regarding the geometry of the magnetic fields is the magnetic helicity. Roughly speaking, the helicity of the magnetic field lines, H , tells us how tangled they are and the degree of linking between toroidal and poloidal components (Berger and Hornig, 2018). The change in the helicity within some volume, \mathcal{V} , enclosed by a surface, \mathcal{S} , is given by (Berger and Hornig, 2018)

$$\frac{dH_{\mathcal{V}}}{dt} = -2 \int_{\mathcal{V}} \frac{\mathbf{J} \cdot \mathbf{B}}{\sigma} d^3x + \oint_{\mathcal{S}} \tilde{\mathbf{A}} \times \mathbf{E} \cdot \hat{\mathbf{n}} d^2x, \quad (6.44)$$


where $\hat{\mathbf{n}}$ is the unit normal to the surface, \mathcal{S} , and $\tilde{\mathbf{A}}$ is the component of the vector potential parallel to the surface \mathcal{S} . Using our definition of the current density, Equation (2.25), and considering only the volume integral, this derivative reduces to

$$\frac{dH_{\mathcal{V}}}{dt}^{(1)} = -2 \int_{\mathcal{V}} \frac{q\mathbf{v} \cdot \mathbf{B}}{\sigma} d^3x. \quad (6.45)$$


Intuitively, there is no change in the helicity for iMHD ($\sigma \rightarrow \infty$). Finite conductivities, however, alter the helicity proportionally to the strength of the magnetic field. We should therefore expect that the differences we have seen in this chapter by including resistivity in merger simulations should grow with the magnetic field strength. This is an important point, as we know that all merger simulations to date have had insufficient resolution to accurately capture the growth of realistic magnetic fields to equipartition (Kiuchi et al., 2014). One should expect, therefore, that as simulations reach higher resolutions and magnetic fields are amplified at a greater rate, the effects of resistivity should become more pronounced.

In the two merger simulations presented in this chapter, we have used initial magnetic field strengths of $B_{\text{max}} \approx 2 \cdot 10^{12} \text{G}$. This is roughly the magnetic field strength expected for old neutron stars that merge via typical pathways. However, investigations that require magnetar strength magnetic fields post-merger, $B_{\text{max}} \sim 10^{16} \text{G}$, but have only limited computational resources often impose significantly stronger initial seed fields to mimic the growth of the magnetic fields by the KHI (Kiuchi et al., 2017; Ciolfi et al., 2019; Ruiz et al., 2020).

We can present, here, some preliminary results whereby initial magnetic fields of $B_{\text{max}} \approx 2 \cdot 10^{15} \text{G}$ have been imposed on the same initial data as the medium resolution simulations of Section 6.4.4. The evolution of the maximum rest-mass density is shown in Figure 6.21(a) for



(a) Evolution of the maximum rest-mass density for the highly magnetised (HM) case. Blue and green vertical lines indicate the times the snapshots for later figures have been taken for the ideal and resistive simulation, respectively. The effects of resistivity are greatly amplified with an increase in the initial magnetic field strength, delaying the collapse of the rBMP by ≈ 14 ms, more than doubling the lifetime of the HMNS.




(b) Slices in the x - z plane of the gradient of pressure contributions for the highly magnetised (HM) case. Slices are taken at the peak of the density oscillations before the iBMP collapses (see vertical lines of Figure 6.21(a)). Similar to Figure 6.15, there is clearly a larger positive effective centrifugal pressure gradient in the ideal case, promoting collapse.


FIGURE 6.21: Evolution of the highly magnetised BNS merger. The effects of resistivity clearly grow with the strength of the magnetic fields, with a large increase in the difference between the collapse times compared to the lower magnetised case.

both models. The increased magnetic field strength has had a dramatic impact on the lifetime of the rBMP, as should be expected from Equation (6.45), more than doubling the time until collapse compared to the ideal simulation. Furthermore, we can see in Figure 6.21(b) that the arguments made regarding the centrifugal pressure in the previous section seem to carry over to the highly magnetised case. The positive gradient of the effective centrifugal pressure is larger in the ideal simulation, although by a lesser degree than in the lower magnetisation case. It is not fully clear yet whether this is the primary cause for the delayed collapse in resistive simulations in general, but it is clear that the cause is closely linked to the magnetic fields and their structure—stronger fields result in larger changes in the magnetic helicity/configuration, leading to dramatic changes in the behaviour of the remnant.

To highlight the differences in the configurations of the magnetic fields, in Figures 6.22(a)



(a) Evolution of the magnetic field components for the highly magnetised case. The extended lifetime of the rBNS leads to greater magnetic amplification. The magnetic field seems to be contained more within the centre of the resistive simulation, as the drop in the mean values of the components is larger for rBNS when the horizon is formed.



(b) Snapshots of the toroidal and poloidal components of the magnetic field taken at the peak of the density oscillations before the iBMP collapses (see vertical lines of Figure 6.21(a)).

FIGURE 6.22: Magnetic components for the highly magnetised BNS. There are clear differences in the configuration of the magnetic fields in the two simulations just before the iBMP collapses.

and 6.22(b) we show plots of the toroidal and poloidal components of the magnetic field. The much extended lifetime of the rBMP results in greater amplification of the magnetic fields, as one would expect, with the toroidal field becoming the dominant component with the fastest rate of growth in both simulations. The slow growth of the poloidal component appears to be suppressed more in the resistive simulation. Interestingly, the mean magnetic field strength in the resistive simulation clearly goes from larger values than the ideal simulation pre-collapse, to lesser values post-collapse, as in the low resolution case, Figure 6.5(c). We interpret this again as a redistribution of the magnetic energy, with a greater magnetic energy density near the core of the remnant in the resistive simulation. Once the horizon is formed and this energy is effectively lost from the simulation, this appears as a larger drop in the mean magnetic field strength.

Finally, turning to Figure 6.22(b), we see similar behaviour to the previous, lower magnetisation case, Figure 6.16. The magnitude of the toroidal component in both simulations is roughly comparable, while the poloidal component within the rBMP is noticeably larger.

6.5 Discussion

In this chapter, we have extended the resistive model of Chapter 5 to be fit for applications in general relativistic simulations with arbitrary, dynamic spacetimes. Our implementation is limited via our choice for Ohm’s law and the form of the conserved vector used in *Whisky*, however, the method is general and extensible to systems of a similar form with the only free parameter being the conductivity. Validation has been demonstrated through unit-tests and a number of simulations that show GRR converges as expected with conductivity. Furthermore, the extension has been written as a thorn for the Einstein Toolkit, and in the future may be adapted to work with *Spritz* (Cipolletta et al., 2019), a soon-to-be open-source addition to the ET. Upon publication, we will release the code for GRR.

Due to the inverse dependence on the conductivity compared to the conventional resistive GRMHD model, we have been able to use GRR to evolve an equal mass BNS merger with a realistic magnitude of conductivity and typical initial data. Previous resistive simulations of mergers (Ponce et al., 2014; Dionysopoulou et al., 2015) used a conductivity of $\sigma_{\max} \sim 10^{11} \text{s}^{-1}$, which is at least six orders of magnitude more resistive than would be expected in reality according to Harutyunyan and Sedrakian (2016)—these simulations were also thought to be conductive enough to be considered as the ideal MHD limit. It was, therefore, unclear whether some of the resistive effects seen in those simulations would be present with a more physically motivated value. Furthermore, the initial data of Dionysopoulou et al. (2015) does not accurately represent expected pre-merger conditions due to a need to accelerate the inspiral, something which we have been able to overcome as a result of a more computationally efficient numerical methods.

With this in mind, we have performed BNS merger simulations using realistic initial data, treating the matter as resistive and with a magnitude that is likely to occur in NS interiors, $\sigma_{\max} \approx 2 \cdot 10^{17} \text{s}^{-1}$. The cost of modelling resistivity in this way adds around 56% in terms of the total execution time (for the medium resolution including data I/O etcetera)—a significant improvement compared to at least $5\times$ that is required when using IMEX schemes and the full resistive equations.

Our simulations have shown that even at this magnitude, resistivity can have a significant impact on the dynamics of a merger simulation, with the most notable difference between simulations being a delayed collapse in the resistive case. Further differences between the simulations include changes in the gravitational wave signals, magnetic structure and outflow properties of the ejecta.

We can be certain that the changes in the behaviour between the two sets of simulations are a direct result of including electrical resistivity, as initial data and numerical methods are identical throughout. We can also say with a high degree of certainty that the resulting changes in the matter (collapse time, outflow properties, etcetera) are due to a reconfiguration of the magnetic fields as a result of the resistivity, and have demonstrated that these effects grow strongly with both spatial resolution and the magnetic field strength of the remnant.

Noting the dynamical importance of the magnetic fields in the properties of the matter evolution, we argue that the delayed collapse may be the result of an altered velocity profile in the core of the BMP, leading to a greatly reduced, positive centrifugal pressure gradient. This change

in the profile of the total pressure is enough to allow the critical rBMP additional oscillations before collapse. The cause of the change in the velocity profile is ultimately unknown, but likely related to the breaking of the flux-freezing condition of iMHD—the fluid will experience reduced drag caused when moving perpendicular to the magnetic field lines in the resistive simulation. In order to clarify these arguments, better methods are required for analysing the configuration of the magnetic fields and the fluid trajectory within the BMPs.

Our results appear to be in stark contradiction with the predictions of [Harutyunyan et al. \(2018\)](#), in which the authors argue that a realistic resistivity should be unimportant on merger timescales. One reason for this disparity may be related to the formation of shocks within the simulations. Within strong shocks, gradients can become very large and the analysis of [Harutyunyan et al. \(2018\)](#) becomes invalid. Furthermore, their analysis predominantly considers the equations of motion of the magnetic fields in isolation, and yet in simulations of NS mergers there is a strong, non-linear coupling of the equations of motion of the magnetic fields and the fluid in which differences can compound in long term simulations.

As a proof of concept of realistic resistivities, the work of this chapter goes a long way to highlighting the importance of including accurate, finite conductivities in merger simulations. We have shown that the influence of resistivity can produce changes in the dynamics of the simulation that are comparable to those in parameter studies of equations of state, magnetic field orientations, variable mass ratios, etcetera. Similar parameter studies are required, and indeed now possible due to the efficient evolution of GRR, in order to confidently understand the effects resistivity has, and where it fits into this parameter space.

Finally, whilst the scale of the resistivity employed in this simulations is realistic, a homogeneous description will not represent the physical system well. Future simulations will benefit from functional forms for the conductivity, similar in style to those used in previous merger simulations ([Dionysopoulou et al., 2015](#)). To this end, [Harutyunyan and Sedrakian \(2016\)](#) provide a best fit for their model of conductivity for NS matter that would be far more accurate and relatively cheap to implement, although it lacks an analytic dependence on the magnetic field strength.

The results of the simulations provided in this chapter represent the first in a set of future investigations, and demonstrate the significance that resistivity can have in BNS merger simulations. We have demonstrated that resistivity can not be ignored any more in BNS merger simulations, and have provided a formalism in which these effects can be studied efficiently, without incurring significant computational costs. Once the code implementing GRR is released, it will allow for yet another step towards more realistic, non-ideal simulations of neutron star mergers.

7

A hybrid approach to resistive MHD

Balance, that's the secret. Moderate extremism. The best of both worlds.

– Edward Abbey

We mentioned in Section 1.3.3 that a hybrid approach to modelling can be a powerful means to capturing physics in multiple regimes simultaneously. Whilst a single physical model may be accurate in a single domain, hybrid models attempt to combine two or more models along with some suitable boundary conditions (generally for switching between models) in order to achieve superior accuracy, performance, or both.

7.1 Introduction

Modern simulations are required to capture physical phenomena across length scales spanning many orders of magnitude. Consider a merger, for example, where the typical wavelengths of GWs is on the order 10^5m , and contrast this to the scales of the smallest hydrodynamic eddies on the order of 1m (Radice, 2017). Such disparities in typical length or time scales are common in a number of astrophysical systems.

Hybrid methods aim at utilising multiple descriptions, or approximations, for different regions in a simulation, allowing to capture phenomena with the most appropriate model. For example, particle-in-cell (PIC) methods are well suited to resolving plasma behaviours at length scales comparable to the skin depth, but are impractical for global simulations where typical length scales are much larger. In contrast, MHD models are efficient at describing bulk motion of a plasma, but become expensive quickly as resolution increases. Hybrid MHD/PIC codes, such as MHD-EPIC (Daldorff et al., 2014), combine the benefits of both methods to allow for an efficient two-way coupling of large and small scale behaviour. Methods such as this have been successfully used in modelling the interaction of the Jovian winds and the magnetosphere of Jupiter's moon,

Ganymede (Tóth et al., 2016), and for global studies of the Earth’s magnetosphere (Chen et al., 2017).

In addition, there are the themes of Harpole’s thesis (Harpole, 2018), aimed at accurately modelling the turbulent burning on the surface of a neutron star during a Type I X-ray burst, and simultaneously the motion of the shock front on larger scales. The propagation of the burning front is sufficiently slow (compared to acoustic speeds) that the low Mach approximation can be used, reducing the need for small timesteps on the finest scales. In contrast, the large scale dynamics of the shock front will be affected by the rotation of the NS, for example, and need to be captured also.

The scale of the NS ($\sim 10\text{km}$) is much greater than the depth of the ocean in which the burning takes place ($\sim 10\text{m}$). This means the dimensionality of the problem can be reduced by employing a shallow water approximation of the equations of motion. Using a modified AMR algorithm, it was possible to switch between the models dependant upon the refinement level, utilising accurate models when needed whilst benefiting from some reduced computation over a fully compressible code.

Another hybrid approach that utilised multiple approximations of resistive MHD is the model used by Lehner et al. (2012). Working under the assumption that resistive MHD is an accurate description of a NS and its magnetosphere, they proceed to make two approximations. First, that the NS matter is well modelled as a perfect conductor (iMHD); and second, that the exterior of the NS is magnetically dominated (*force-free*), i.e. the fluid’s inertia is much less than the inertia of the EM field. This has the effect of decoupling the motion of the EM fields to the motion of fluid, allowing for a more accurate representation of the EM outbursts that emanate from the collapse of a HMNS.

The two models are evolved in unison, with the solution for the NS providing the boundary conditions for the force-free magnetosphere over a transition region, defined to be where the energy-density is $\sim 10^2 - 10^3 \times$ the atmospheric value. Such studies hint at the importance of correctly modelling the NS exterior, highlighting the role that magnetic reconnection has in the subsequent EM emission.

Clearly, there are benefits to be had by using the correct tool for the job. In the following sections, we will outline two hybrid models that we have tested and discuss their implications for the future of resistive BNS simulations.

7.2 A special relativistic example

For the first example, we recall the equations of rMHD in special relativity—the same logic applies to the models in GR but has been not been explicitly investigated due to time restrictions. As we have explored previously, the reason the rMHD equations are often avoided when they may be relevant to the physics is numerical. Large source terms in highly conducting regions result in numerical instabilities, and solving this problem with implicit integrators can still be computationally impractical. It remains true, however, that capturing regions where $\sigma \sim 0$ may be desirable, especially if trying to determine electromagnetic outbursts consistently with

Maxwell’s vacuum equations. In this section, we describe a method of accurately modelling both limits, $\sigma \rightarrow 0$ and $\sigma \rightarrow \infty$, whilst avoiding any stiffness issues.

7.2.1 Hybrid++

A possible approach to this problem that is capable of approximating the solution in all regions is to simply evolve the rMHD model up to the region where numerical issues begin to arise, and then switch over to an ideal description for conductivities greater than some cut-off value. This approach has the benefits of being accurate in both highly resistive and conducting regions, whilst being numerically suitable for explicit time integration throughout. Of course, for intermediate conductivities, we would expect its solution to deviate somewhat from the true solution.

To further improve the solution in these intermediate regions, we may also employ the REGIME source term. The modified model, which we call Hybrid++, now has the benefit of being able to effectively capture any conductivity. This is something that is not possible for the resistive model even with the IMEX time integrators. This is because for conductivities of $\sigma \gtrsim 10^8$ (in code units), the source term is too large even for implicit integrators to maintain reasonable step sizes. In contrast, the source of Hybrid++ in this region vanishes. Similarly, for $\sigma \sim 1$, where the REGIME source term would become stiff, the resistive source vanishes, and thus the evolution can continue accurately and efficiently with explicit RK methods.

The question now arises, how do we match the models to achieve the best possible accuracy and speed, and to maintain stability? First, we must define some region beyond which rMHD (with explicit RK methods), becomes unstable. Of course, this will depend somewhat on the choice of numerical methods, but the general magnitude of these parameters is not likely to change for typical methods.

As in [Lehner et al. \(2012\)](#), we use a penalty technique when switching between the models, which in our case is dependant upon the local conductivity. The switching between models is defined to occur centred at the *cross-over* conductivity, σ_c , and to vary over some range defined by the conductivity *span*, σ_s . We define our smoothing kernel as

$$P(\sigma) = \frac{\tanh\left(\frac{\sigma - \sigma_c}{\sigma_s/3}\right) + 1}{2}, \quad (7.1)$$

such that for a conductivity of $\sigma_c + \sigma_s$, the source and flux vectors are contributed $\approx 99.8\%$ by iMHD—similarly a conductivity of $\sigma_c - \sigma_s$ is $\approx 99.8\%$ the contribution of rMHD, see [Figure 7.1](#).

Once calculated, the penalty function, $P(\sigma)$, tells us how much of the flux/source vectors should be contributed by the iMHD model, and how much from the rMHD model. These contributions are then summed and used as the corresponding flux/source vector for the Hybrid model. For Hybrid++, we simply calculate the contribution of the source term of iMHD and REGIME combined. When calculating the primitive variables, using the resistive transformation when $\sigma < \sigma_c$, and the ideal transformation otherwise, allows for stable evolutions.

Currently, in the form used to produce these results, all source and flux vectors are computed for every cell in the domain—i.e. both the resistive and ideal contributions. After calculating them, the Hybrid model combines them as defined above. As a result, we are essentially evolving




FIGURE 7.1: Penalty method for the Hybrid model, showing the fraction of the contribution from each model. For $\sigma < \sigma_c - \sigma_s$, the Hybrid model is simply rMHD. For $\sigma > \sigma_c + \sigma_s$ Hybrid is simply iMHD (and REGIME for Hybrid++). In the transition region, the contribution to the source and flux vectors is given by the graph above.

two models, and so should expect the current implementation to be roughly half as quick as a simple iMHD simulation. A more efficient implementation—in which only one model is calculated beyond the bounds of the conductivity span—should take no longer to evolve than simple ideal models, and in fact, due to the cheaper conservative-to-primitive transformation of rMHD, may even be faster.

7.2.2 Results

In this section we present some results of our hybrid models in special relativity. These results have been generated using the `METHOD` code, Chapter 3. For the time integration, we use the second order operator splitting RK method (Algorithm A.2), and refer to the full resistive model using the IMEX scheme as the *exact* solution when comparing error growth.

Current sheet

For our first example, we return to the resistive current sheet problem, Section 3.1.3. Using this example, one can see how the error, with respect to implicitly evolved rMHD, of the explicitly evolved iMHD and rMHD models change with the conductivity. The point at which the error curves of the two models intersect is the location to define the parameters, $(\sigma_c \pm \sigma_s)$, which in the case of the Hybrid model is (400 ± 350) . For the Hybrid++ model, REGIME will drastically improve the accuracy at higher conductivities, pushing down the cross-over conductivity and reducing the necessary span. We found the optimum parameters here to be $(\sigma_c \pm \sigma_s) = (150 \pm 50)$.

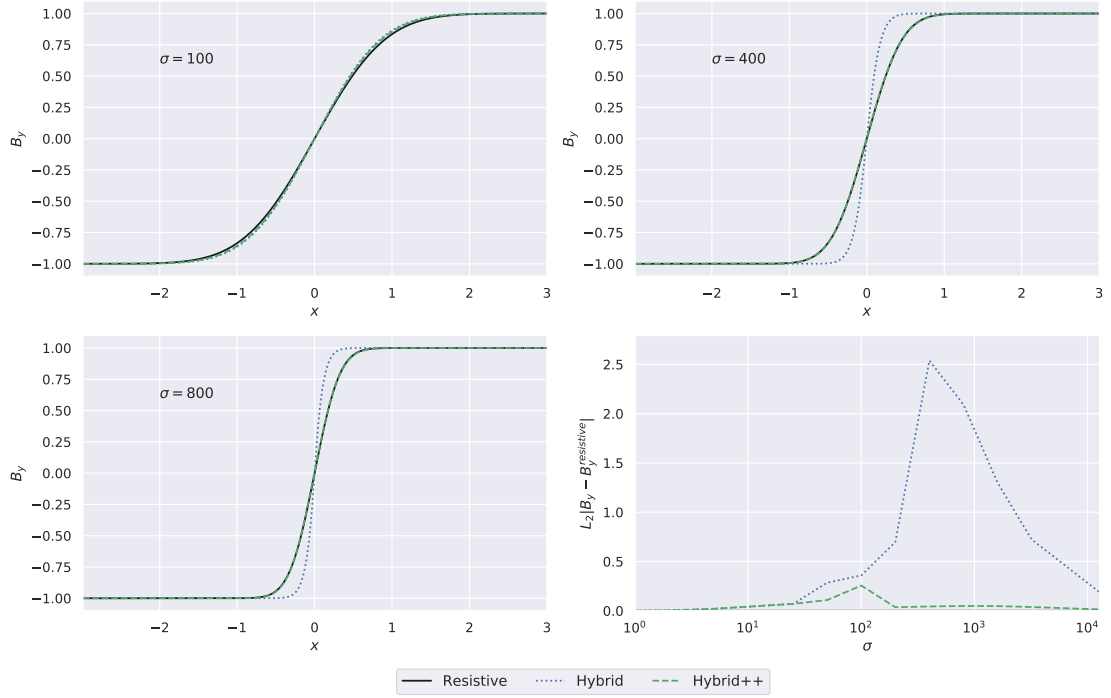


FIGURE 7.2: Error of the Hybrid and Hybrid++ models for the current sheet. The *exact* solution is given by the full resistive model using the second order IMEX scheme. Each model is evolved until $T = 26$, using $N_x = 400$ and a Courant factor of $\nu = 0.2$.

The first three plots of Figure 7.2 show the final state of the y -direction magnetic fields for the Hybrid and Hybrid++ models, with the resistive IMEX model solution for comparison. For the Hybrid model, we can see that at $\sigma = 100$, where the resistive model is the largest contributor, the solution is very close to the exact solution. For the latter two conductivities however, the resistive model contributes less and the model is almost entirely due to iMHD, resulting in much larger errors.

In contrast, the Hybrid++ model shows almost no discernible difference in its solution with the exact solution. We can see this more clearly in the final, bottom-right-most plot, the error for a given conductivity. We can see that the Hybrid++ model has a small peak in error at σ_c , but otherwise varies very little from the exact solution, especially in comparison to the simpler Hybrid model.

One of the benefits of the hybrid models presented here, as can be seen comparing the error plots of Figures 7.2 and 5.4, is the behaviour at very low conductivities. Figure 5.4 does not show conductivities below $\sigma = 20$, because at this resolution and conductivity the REGIME source is not numerically stable. This is a problem if we want to consider functional conductivity profiles which may reach values this low. With the Hybrid++ model, conductivities may be as large or as small as needed, with source terms only becoming smaller at the two extremes.

Functional conductivity

To demonstrate the range of conductivities that can be comfortably handled by the Hybrid++ model, we will now reconsider the Brio-Wu shock tube problem, Section 3.1.2, with a functional

conductivity profile. The profile we use is taken from [Dionysopoulou et al. \(2013\)](#), defined as

$$\sigma(D) = \sigma_0 D^\lambda \quad (7.2)$$

where $D = \rho W$ as usual, and (σ_0, λ) are constants.




FIGURE 7.3: Final state for the density, x -velocity, pressure and y -direction magnetic field for the Brio-Wu shock tube and a functional conductivity using Hybrid++.

In Figure 7.3 we show the final state for this problem for two conductivity profiles. We can see that the Hybrid++ model is able to capture the change in conductivity and avoid spurious oscillations at shocks. Furthermore the conductivity profile is shown at the final time, $T = 0.4$, in Figure 7.4 for a number of different parameters. We can see that Hybrid++ can evolve conductivity profiles that increase/decrease by factors of a googol, or more. In principle, Hybrid++ will become more stable as the range is pushed further—due to the source contributions from the relevant model continuing to decrease—however we have not tried larger ranges and believe a factor of a googol to be sufficient for the purpose of this demonstration.

This impressive stability is encouraging for Hybrid++ and its application to astrophysical simulations. Particularly in BNS simulations, one wants to evolve the highly conducting neutron star matter simultaneously with the near-vacuum exterior. While the functional form or typical magnitude of the conductivity during mergers, or even isolated stars, is not known well, one should expect descriptions of the conductivity in NS interiors and the electrovacuum exterior to vary by very many orders of magnitude. The ability to capture these large ranges at sharp interfaces will be essential in the future of resistive merger simulations.

The final question to address with Hybrid++ is the required computation with respect to the alternative models. As mentioned above, in its current implementation in `METHOD`, Hybrid++ evolves both the ideal and resistive models and combines them later according to the conductivity in a given region. This means that in a cell where $\sigma = 0$, the ideal model and `REGIME` source are both computed along with the resistive model, and then only the resistive model is used in




FIGURE 7.4: Conductivity profile for the final state of the Brio-Wu shock tube problem. The profile defined by $(\sigma_0, \lambda) = (10^6, 12)$, shown in the inset, is the most challenging presented in [Dionysopoulou et al. \(2013\)](#), where as with the Hybrid++ model, we can span arbitrary conductivity profiles varying by over 100 orders of magnitude.

the time integration. Obviously this involves a significant amount of wasted calculation, and more fine-tuned attempts could cut the computational cost almost in half.

With this in mind, in Figure 7.5, we show the minimum runtime for the functional Brio-Wu problem as a function of σ_0 . The exponent is given as $\lambda = \log_{10}(\sigma_0)$, such that for very large, and very small σ_0 , some region in the domain will be highly conducting, and another will be highly resistive. The times shown correspond to the maximum Courant factor, ν , for which the solution shows no signs of instabilities, and we also plot an ideal simulation baseline which, of course, does not depend upon σ_0 .

Current implementations of full rMHD for astrophysical simulations in GR use the IMEX integrators ([Ponce et al., 2014](#); [Dionysopoulou, 2016](#)), and would correspond to the red-dashed line. As we can see, the IMEX schemes allow for a much larger range of σ_0 over the explicit schemes, blue-dashed-dotted line, at the expense of longer runtimes in certain regions. In this case, resistive IMEX implementations seem to take at least an order of magnitude longer than a purely ideal MHD model, in regions where σ_0 does not grow beyond 10^8 or below 10^{-12} . Outside of this range, the rMHD model becomes too stiff for reasonable execution times even with the IMEX integrators, and further changes in the conductivity rapidly reduce the required size of the time step, making simulations intractable.




FIGURE 7.5: Minimum time taken to evolve the Brio-Wu shock tube with a functional conductivity defined by σ_0 as shown, and $\lambda = 2 \log_{10}(\sigma_0)$ using $N_x = 200$ cells. As we can see, the execution times of the non-hybrid models diverge as $\sigma_0 \rightarrow \infty$ and $\sigma \rightarrow 0$, where as Hybrid++ remains efficient and accurate.

In contrast, Hybrid++ takes almost the same amount of time regardless of the conductivity. This equates to around a factor $2\times$ slower than the ideal baseline simulation as expected, recalling that we are evolving both iMHD and rMHD at every cell. It is easy to see the benefits of the Hybrid++ model with these results—superior performance over a significantly increased range of conductivities when compared to the resistive IMEX schemes, and with almost identical solutions. For a more carefully implemented Hybrid++ model, in which less computation is wasted, we would expect the performance to be roughly equivalent to iMHD.

One open question with regards to Hybrid++ for use in merger simulations is the values of the parameters (σ_c, σ_s) . At the resolutions used here, we have had a great deal of success in terms of producing accurate solutions across a large conductivity range efficiently. For very high resolution simulations, however, such as in merger simulations, the REGIME source may become unstable, Equation (5.42). The effect of this is that REGIME can not be computed below some critical conductivity. If this critical value is greater than the conductivity at which rMHD becomes unstable with explicit integrators, then the error during the transition between models will be high. If this gap in stability is large enough, it may even be the case that neither the explicitly evolved resistive model nor REGIME can be computed during the transition. If this is the case, one must either reduce the Courant factor in this region, or revert to the IMEX schemes to maintain stability.

7.3 A general relativistic example

The Hybrid++ model of the previous section shows promise for utilising multiple descriptions of MHD for different regions in a simulation. In this section, we consider an alternative hybrid model but now in the context of fully general relativistic merger simulations.

7.3.1 Ideal atmosphere

This next hybrid model comes from the realisation that, for most merger simulations, the magnetic field is initially confined to within the progenitor stars. Throughout the simulation, the motion of the magnetic field and the fluid will result in messy field lines that extend into the exterior of the NSs. The size of the simulated domain, however, is vast compared to the region of interest, and is comprised mostly of the near-vacuum atmosphere. The rest-mass density of the atmosphere is typically many orders of magnitude lower than the peak (Giacomazzo and Rezzolla, 2007), a factor of 10^{-7} in standard *Whisky* simulations. Importantly, any region that is deemed to be in the atmosphere has its hydrodynamic variables set to prespecified values, and thus is homogeneous.

In the atmosphere regions of a merger simulation therefore, one will not expect REGIME to have much of a contribution. This is because it is diffusive in nature, and all terms are based upon spatial variations in (magneto)hydrodynamic quantities with neighbouring cells. In the atmosphere, all fluid variables are homogeneous, and the magnetic fields are not expected to protrude too far into the atmosphere, at least at early times. As a result, it is a good approximation that REGIME will be zero, or very close to zero, for any computational cell in which $\rho = \rho_{\text{atmo}}$, where ρ_{atmo} is the rest-mass density of the atmosphere. This corresponds to simply evolving GR iMHD in the atmosphere. We refer to this model in which we treat the atmosphere in this way as the ideal atmosphere (IA) model.

7.3.2 Results

The details of the simulations in this section are covered in more detail in Section 6.4. We take the initial data from the low resolution, ideal fluid model of Kawamura et al. (2016), and impose a uniform conductivity of $\sigma = 2 \cdot 10^{17} \text{s}^{-1}$, evolving the system until around 12ms after merger.

In Figure 7.6 we show slices in the $z = 0$ plane of the rest-mass density for a homogeneous resistivity (top) and for IA (bottom) for times 0.1ms and 2.5ms after merger, and 1.4ms after black hole formation. There is very little variation between the two density profiles on large scales at any time. We can see the same result in Figure 7.7 on small scales, where we instead show the maximum rest-mass density in the simulation against time. There is quantitatively very little difference between the two models until moments before collapse, but the differences even then are practically small.

From these two plots of the behaviour of the rest-mass density, it appears that, at least at the resolution performed here, treating the atmosphere as ideal has no discernible effect on the dynamics of the matter. The same can be said of the large scale magnetic field behaviour. Figure 7.8(a) plots the evolution of the total magnetic energy for the two models. Once again, deviations are small until moments before collapse, but the same general trend and magnitude can be seen in both cases.

The largest disparities between the models can be seen when we observe the small scale magnetic behaviour. The maximum magnetic field strength, shown in Figure 7.8(b), contains arguably larger differences immediately before collapse, but once again the variations are so small in magnitude that qualitatively the solutions from both models seem practically equivalent.




FIGURE 7.6: Slices in the x - y plane of the rest-mass density at 0.1ms and 2.5ms after merger, and 1.4ms after black hole formation. Top panels: solution for IA. Bottom panels: solution of homogeneous resistivity. Qualitatively there is very little difference between the two models at the times shown.






FIGURE 7.7: Timeseries of the maximum rest-mass density for the IA model and a homogeneous resistivity. As can be seen, there is minimal difference between the two models.

Although IA appears to give very similar solutions to the homogeneous resistivity case, this may be an artefact of the resolution of the simulation we have chosen. Limitations on computational resources have prevented us from exploring higher resolutions for this set up. It is possible that at these low resolutions, numerical diffusion dominates over resistive diffusion. If true, then at higher resolutions, where we expect the numerical diffusion to be reduced, any diffusion will be dominated by resistivity, and thus the solutions will be more sensitive to changes in the model. We may therefore expect that these differences pre-collapse, and as a result post-collapse, to grow with increasing resolution.

Regardless, the atmospheric treatment as ideal is currently the standard in most simulations



(a) Timeseries of total magnetic energy for IA and a homogeneous resistivity. Although there is arguably a slight change in the magnetic energy a few milliseconds after merger, the general behaviour is the same, especially post-collapse.



(b) Timeseries of maximum magnetic field for IA and a homogeneous resistivity. Similar to the case of the total magnetic energy, there is a slighter stronger magnetic field in the IA simulation that only lasts for a few ms until collapse, after which the solutions of the two models are very similar.

FIGURE 7.8: Evolution of the magnetic quantities for the standard REGIME and IA mergers

of BNS mergers, and so a resistive treatment of just the interior of the NSs represents a step forward in complexity, and physicality.

With regards to the computational benefits of IA, the reduced calculation in the atmosphere results in a decrease in computation time of around 30% when compared to a homogeneous resistivity, as in the previous chapter. Of course, this speed-up is dependant upon the domain, and in simulations of isolated neutron stars, for example, where there is a significantly reduced domain and atmosphere, the benefits of this treatment will diminish. For simulations of mergers, however, if it can be determined that we get the expected results at higher resolutions when using IA, reductions in execution time of this factor can amount to a significant allocation of computer time.

7.4 Discussion

In this chapter, we have taken inspiration from a number of different hybrid models. The first model we have presented utilises the most computationally efficient model of rMHD, either the full resistive model or iMHD with the REGIME extension, dependant upon the local conductivity. By selecting the correct model and removing all stiffness issues that arise through the evolution

of the electric fields, Hybrid++ can be evolved only ever requiring explicit time integrators, thereby saving many factors in execution time over the typical IMEX schemes that have been used for some rMHD simulations (Palenzuela et al., 2009; Dionysopoulou et al., 2013). A further consequence of the model selection is that, in both highly conducting, $\sigma \gtrsim 10^3$, and highly resistive simulations, $\sigma \lesssim 10$, the solution becomes ever more accurate, improving the quality of the solution over rMHD for a huge range of conductivities.

Hybrid++ is currently only implemented in SR, using the `METHOD` code, however the logic does apply in GR settings as well. A realistic conductivity profile used in astrophysical simulations is likely to span many orders of magnitude (Harutyunyan and Sedrakian, 2016), and the additional stability gained from using the hybrid approach is likely to ease constraints set on resistive merger simulations. Although there is still potential for stiffness in the crossover regions (a result of the high resolution making REGIME unstable), there is potentially much to gain from extending Hybrid++ to dynamic spacetimes, and we hope to see projects that use these ideas in the future.

The second hybrid model we have presented in this chapter is the ideal atmosphere model when using REGIME for merger simulations. We have argued that we should expect the REGIME source to be negligible in the atmosphere, at least for short times after merger. This allows us to save a significant amount of computation over a standard REGIME simulation, in which the almost-zero REGIME terms are calculated everywhere. Standard REGIME merger simulations should expect to be around 50% slower than a normal, ideal merger simulation, but when using the ideal atmosphere approximation this is reduced to only 35%.

At least at low resolutions, the IA model produces the same qualitative behaviour as a standard REGIME merger simulation, and for much of the simulation both local and global quantities are similar. This suggests that, indeed, the REGIME source does not significantly change the dynamics of the simulation in the atmosphere, particularly for the features that are interesting in resistive merger simulations (i.e. the magnetic field evolution and collapse to a black hole).

There are questions regarding the long term solutions of IA and REGIME, and whether for longer simulations, $t > 30\text{ms}$, the differences in the evolutions will diverge. Furthermore, it is uncertain whether such similar results will be obtained for higher resolution simulations, and as a result this is currently the priority for future simulations.

We have seen in the literature effective uses of the simultaneous evolution of multiple models for improving accuracy and efficiency in simulations (Lehner et al., 2012; Ponce et al., 2014; Harpole, 2018). This chapter represents an additional direction in which merger simulations can go, whereby the cutting edge models that have been used for today's merger simulations are switched between dynamically, improving the stability of the scheme as a whole, reducing the execution time relative to previous implementations, and improving the overall accuracy of the solutions.

Final thoughts

Humans were always far better at inventing tools than using them wisely.

– Yuval Noah Harari

Neutron star mergers provide a natural laboratory for investigating physics that occurs over a range of scales—from the microphysics acting on atomic scales to the generation and propagation of gravitational waves. In recent years, our capability of extracting information from these events has grown in leaps, with the successful detection of GWs and electromagnetic counterparts from these events becoming a semi-regular occurrence. Now, we are in need of extensive parameter surveys such that we can further understand the origins of the properties of the signals we detect. Until data from next generation detectors becomes available, the necessity to understand these events will continue to grow. For now, the best way to refine our models and better understand the parameter space is through numerical simulations.

8.1 Overview

In Chapter 1, we review the literature regarding what is currently understood about mergers and the progress made in simulations of compact object mergers, and highlighted some of the key limitations with the current state-of-the-art approaches. With this in mind, we described the outline of the work within this thesis—that is, to investigate non-ideal models of MHD for use in merger simulations, and to reduce the computational costs that come with non-ideal descriptions of MHD.

Next, in Chapter 2, we outlined some the mathematical tools that are required when modelling BNS mergers. In particular, we presented the equations describing a number of ideal and non-ideal models of MHD; some numerical tools used in their evolution; and a brief overview of solving Einstein’s equations of GR for dynamical spacetimes.

In Chapter 3, we provided a description and validation of our numerical fluid code, `METHOD`. This code has been used for model development, and provides an easy interface to explore new numerical tools or physics models. `METHOD` has been instrumental in refining the tools used in

later chapters. Furthermore, the code has been made public¹, and is currently being used for new research outside of the work of this thesis.

The first result of this thesis is presented in Chapter 4, whereby we executed the most computationally demanding tasks involved in some non-ideal models of MHD on a graphics processing unit. This corresponds to the numerical flux function and, crucially, the semi-implicit time integration. The large size of the work array required for the IMEX rootfinding procedure limits the potential speed-ups that can be achieved on a GPU—as the work array cannot fit inside shared memory, a large amount of time is spent accessing data that lies in slow, global memory. Despite these limitations, we were still able to produce parallel speed ups of just over $21\times$ when compared to serial execution.

In Chapter 5, we focussed our attention specifically on resistive MHD. Motivated by the requirement of semi-implicit time-integrators to evolve the rMHD equations, we developed a resistive extension to iMHD that approximates the full resistive model. A host of test problems were then presented that demonstrate how the solution converges as expected with the conductivity. We also provided a stability analysis that predicts REGIME’s stability for resistive, BNS merger simulations. We note that the benefit of REGIME comes in the form of an inverted dependence on conductivity compared to the full resistive model. This allows us to use the explicit RK time-integrators commonly used for these types of simulations, resulting in reduced computation compared to the full resistive model when evolved using IMEX schemes.

REGIME is further extended in Chapter 6, where we make it available for use in general relativistic simulations with dynamical spacetimes. This allows for resistive simulations of BNS mergers that use a conductivity that is expected to exist within NS matter. We demonstrate that even with such a small resistivity, large differences in the matter dynamics are produced compared to ideal MHD simulations—an effect that grows with the strength of the magnetic fields. The largest impact currently appears to be the extended lifetime of the remnant in the resistive simulation, which we believe to be a result of an altered pressure profile due to changes in the fluid velocity in the core.

The work of Chapter 6 represents an important milestone for non-ideal descriptions of MHD. In the vast majority of merger simulations, NS matter is assumed to be perfectly conducting, an approximation that was thought to have little effect on the results of the simulations. In the few merger simulations that use rMHD, the size of the resistivity used for the NS interior was many orders of magnitude larger than is expected in nature—values that were chosen to minimise numerical problems when evolving the electric fields. In contrast, GRR lends itself well to highly conductive simulations, with runtimes expecting only a 50% increase.

Finally, the work of Chapter 7 aims to further improve the accuracy and efficiency of resistive merger simulations in light of the successes of previous chapters by presenting two hybrid models. The first uses a similar approach to the force-free hybrid model of [Lehner et al. \(2012\)](#). Utilising the fact that the full rMHD model can be evolved explicitly for small conductivities, and REGIME can be used for larger conductivities, a penalty method is used for switching between the two, resulting not only in consistently fast execution times for any possible conductivity (with the potential to be at least as fast as an ideal simulation), but also improved stability properties and better accuracy over implicitly evolved rMHD. In the future, GR implementations of this

¹www.github.com/AlexJamesWright/METHOD

model would allow for the efficient evolution of both vacuum and perfectly conducting regions simultaneously.

The second hybrid model is an example set in full GR, building upon the work of GRR. By noting that REGIME should have little effect in the atmosphere, at least at early times, we can save significant computation by reducing the model to iMHD in these regions. Comparable results are generated with the ideal atmosphere hybrid model and the full GRR implementation, but it is unclear whether this would be true also for higher resolution simulations.

8.2 Looking forward

The necessity of high resolution, physically accurate numerical models for BNS merger simulations is only going to grow over the coming decades. With better theoretical models and more precise apparatus for observations, the sources of error in our understanding of these events will continue to decrease. It is of great importance that we continue to refine the computational methods and physical models that we use in the next generation of merger simulations.

That said, modern simulations of compact object mergers have already come along way over the last couple of decades. Researchers' ability to evolve stable, long term models incorporating relativity, thermodynamics and magnetohydrodynamics has allowed for more refined analyses of the important interactions at play during mergers. A lot of attention is now being focussed on processes beyond ideal MHD, whether that be via weak interactions (Foucart et al., 2018), multiple fluid models (Amano, 2016; Andersson et al., 2017b), subgrid scale modelling of turbulence (Radice, 2017; Viganò et al., 2020) or of Kelvin-Helmholtz vortices (Giacomazzo et al., 2015), dissipative phenomena such as thermal conductivity or viscosity (Dommes et al., 2020), or electrically resistive descriptions of the atmosphere (Lehner et al., 2012; Ponce et al., 2014).

Inherent in these more complex models is additional computation. Often, the computational cost of describing new phenomena is a limiting factor when applying novel theoretical insights in simulations. By focussing on the case of electrical resistivity, we have, in this thesis, developed a number of methodologies to reduce these overheads. Whether by utilising GPU hardware on the next generation of supercomputers, or by using multiple models simultaneously to benefit from improved accuracy and efficiency in various domains, the tools presented here will allow for more detailed investigation of the physics occurring in mergers.

Furthermore, the work presented here is not only applicable to BNS simulations. Our analysis of the bottlenecks present in non-ideal MHD on GPUs highlights the limitations of high-dimensional rootfinding procedures on this type of hardware, and is not necessarily specific to models of MHD. In addition, the techniques used for developing REGIME should apply for any relaxation system—that is, any system with potentially stiff source terms that relaxes to an equilibrium state. For example, the Chapman-Enskog analysis could see use in models of MHD that include neutrino transfer, with the expansion timescale driven by the opacity. Furthermore, we also see no reason why the expansion must be limited to a single timescale. The two fluid model presented in Chapter 2 possess two timescales, one describing the resistive timescale and one for the plasma frequency. Although more complex, it is conceivable that an extension to iMHD could approximate both the resistive and multi-fluid behaviour of the full model.

The initial investigations using these techniques for merger simulations were fruitful. Including a realistic scale resistivity resulted in clear differences in the behaviour of the merger remnant, differences that directly impact observables that are of key interest. We have also seen that these new behaviours become more pronounced with strong magnetic fields, and will likely be a dominant factor for studies of subsequent jet formation and kilonovae.

The work presented here only scratches the surface, providing a proof-of-concept that GRR is stable and functions in these complex simulations. Now we can be confident that non-ideal models of MHD are required—or more specifically, that resistive NS matter has a large impact in the post-merger dynamics—it is important that work continues to explore the parameter space and understand further the complex nature of physics occurring within neutron star mergers.

Appendix A

A.1 Magnetic freezing

Here, we briefly show why magnetic fields are locked to the fluid motion for iMHD. Consider the magnetic flux, Φ , through some elemental volume's surface:

$$\frac{d\Phi}{dt} = \frac{d}{dt} \int_{\Sigma} \mathbf{B} \cdot \mathbf{n} \, d\Sigma = \int_{\Sigma} \left[\partial_t \mathbf{B} - \nabla \times (\mathbf{v} \times \mathbf{B}) + \mathbf{v}(\nabla \cdot \mathbf{B}) \right] \cdot \mathbf{n} \, d\Sigma, \quad (\text{A.1})$$

where B is the magnetic field. First we use Maxwell's divergence-free equation, $\nabla \cdot \mathbf{B} = 0$, and the induction equation, $\partial \mathbf{B} / \partial t = -\nabla \times \mathbf{E}$, noting that ideal Ohm's law is $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$. With these relations, we can see that the magnetic flux of a fluid element is unchanging in ideal MHD. This property of perfect magnetofluids has the following consequence: the magnetic lines are “frozen” to the fluid, and are simply advected with the fluid velocity. This result is known as Alfvén's theorem, and such behaviour is not present in fluids that exhibit electrical resistance, where the magnetic field lines are free to drift with respect to the fluid velocity.

A.2 Overview of METHOD

This section presents, for the interested reader, an outline of the numerical fluid code `METHOD` produced for this thesis. We describe the design pattern and class structure, and different physics models that `METHOD` implements. Also outlined in later subsections are the testing paradigms used in its development to ensure that all exiting functionality persists as changes to the code are made.

A.2.1 General Design Features

The aim of `METHOD` is to provide a simple interface in which we can implement and test existing or novel MHD models in small to medium sized domains. `METHOD` is being developed in C++, utilizing object-oriented programming to maximize ease of use, limit code duplication and to allow for the easy development and addition of new functionality.

Care has been taken in the design of each class to ensure that any new features can be seamlessly integrated. In order to make `METHOD` easily extensible to new physics models, time

integrators, and flux reconstructions, abstraction is key. For example, the class hierarchy is built in a top-down fashion, an approach initially promoted by IBM researchers in the '70s and still the golden standard of OO design to date. This approach promotes the paradigm of prospective, as opposed to retrospective, development.

Initially, the overarching design of the code base must be understood, taking into account the desired functionality, and any additional potential improvements that should be catered for. Classes that implement general methods, such as generating the flux vector, are abstracted into base classes, and the input/output of these methods are defined from the outset via a set of pure virtual functions such that any implementation of these methods conform to the standard. Any additional member functionality of these functions, e.g. some subroutine, may be added if needed.

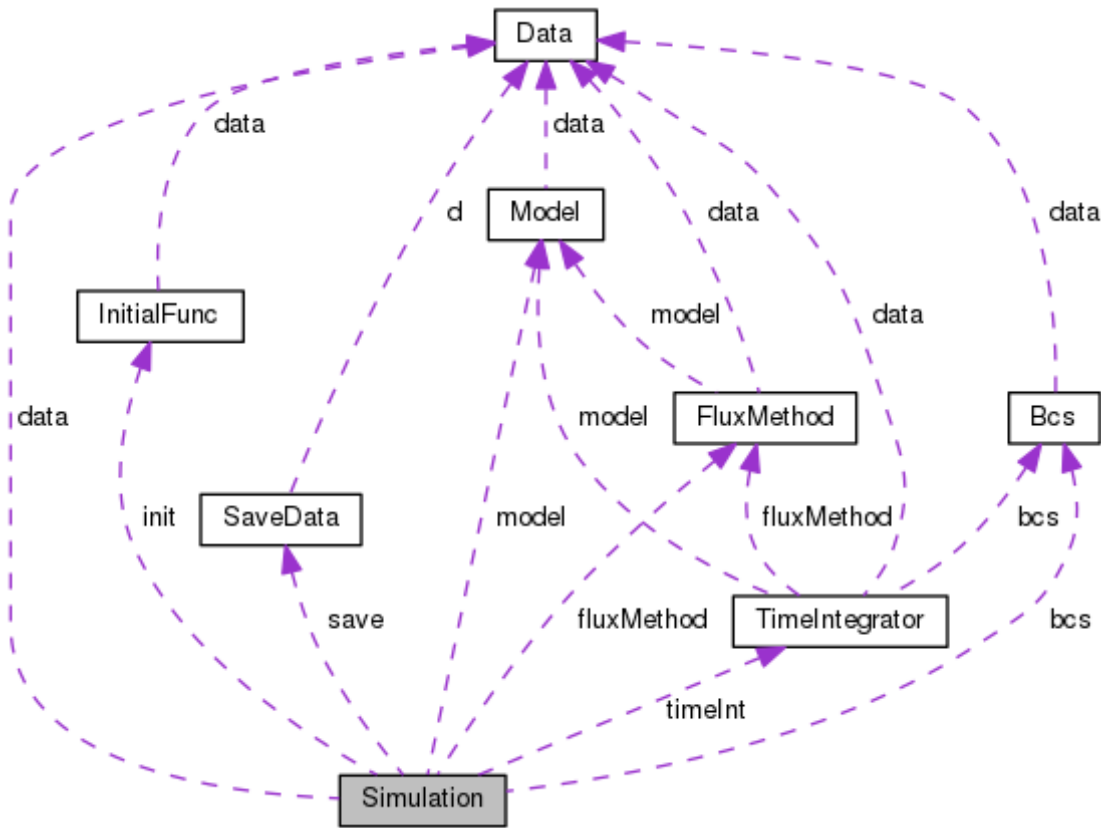


FIGURE A.1: Dependencies of the Simulation class. We see how all class communication is one way—for example, FluxMethod requires knowledge of the physics via Model, but not of the type of TimeIntegrator. This reduces the chances of spaghetti code. Graph generated by Doxygen: www.doxygen.nl.

Communication between classes is one-way. To help in this instance, the most general class contains all relevant simulation data that may be required by any other functional unit. The class *Data* contains arrays holding the current state of the simulation in the form of conserved, primitive, and auxiliary vectors, the current simulation time, and the size and composition of the domain, along with any constants relevant to the simulation, e.g. resistivity. In the parallel edition of METHOD, information regarding GPU specifications is also stored here. Due to the use of abstraction, the *Data* class is entirely general, it does not require knowledge of or data from

any other class. For clarity, there are no instances in which the *Data* class retrieves information from any object that contains it (see figure A.1).

All objects in **METHOD** contain a pointer to the simulation's *Data* class, as is necessary to operate on the current state of the system. As a result, member functions of all other objects have access to *Data*'s work arrays publicly, and are free to alter and update all non-constant data. The *Data* class should not, however, alter any data specific to any other class that contains it. This is a general feature of top-down programming, an object that contains an instance of a class should not have its data altered by the contained class.

To elucidate the use of inheritance throughout **METHOD**, is instructional to look at the implementation of the time integrators (e.g. algorithm A.1¹). Given some current state for the conserved (and primitive and auxiliary) vectors, some physics model to determine the flux vectors, some method for computing the numerical flux from said vector, and some boundary conditions, the time integrator will produce an update on the conserved vector for the next timestep.




FIGURE A.2: The inheritance graph for the RKSplit time integrator. TimeIntegrator is the abstract base class for all integrators that defines the minimum required functionality. RKSplit implements an RK2 class, and then adds the source vector. Graph generated by Doxygen.

ALGORITHM A.1: Definition for the time integrator's step function. This function is pure virtual, meaning it is defined but not implemented. This means that any time integrator inherits this definition and is required to implement their version of the step function.

¹ `virtual void timeInt.step(double *cons, double *prims, double *aux, double dt)=0;`

Using pure virtual functions in this way ensures that the input and output of every implementation of the step function has identical form. This is essential as functions that call *timeInt.step* should be independent of the type of time integrator that is being used.

¹Code snippets may not match the actual implementation—they are simplified and used to reflect the input/output and functionality of a method

We can also see how this technique cuts down on code duplication. If we consider two general physics models, both of which are in conservation form with one containing a source term, we require slightly different methods for computing the next timestep. For the source-less model, a simple Runge-Kutta-type integrator is appropriate. For the sourced model, we would require first an RK step of the same type as for the previous model, followed by the addition of the source vector.

To achieve this so-called operator splitting, it makes sense to inherit the *step* method from the RK integrator in the operator split integrator, *RKSplit2*. The resulting step function is this simplified to the following:

ALGORITHM A.2: Implementation of the second-order, operator split Runge-Kutta integrator, showing the use of inheritance of virtual functions.

```

1  void RKSplit2.step(double *cons, double *prims, double *aux, double dt)
2  {
3      cons += dt / 2 * model->sourceVector(cons, prims, aux);
4      RK2.step(cons, prims, aux, dt);
5      cons += dt / 2 * model->sourceVector(cons, prims, aux);
6  }
```

This has the added benefit that any changes, i.e. bug fixes or optimisations, applied to the RK integrator are implicitly applied to *RKSplit2* also (see also figure A.2).

We should note, however, that we have also implemented methods in which the source term and flux terms are combined into a RHS function. Using this approach, there is no need to treat the two contributions separately, and higher orders may be achieved easily using a standard SSP-RK4 integrator (Gottlieb et al., 2009), for example.

A.2.2 Available Physics Models

There are three models that are currently available for use in **METHOD**, all of which are limited to special relativistic descriptions and can be found in Section 2.2: Ideal MHD; Resistive MHD; Two-fluid EMHD. Some care must be taken in choosing the correct numerical methods and setting the correct constants specific to the model one wishes to use. For example, when using resistive MHD, one should explicitly set the resistivity, σ , or the value will be set to the default, $\sigma = 10^3$. In this case, the user should also implement the semi-implicit SSP2 time integrator to deal with the potential stiffness of the equations. Using the incorrect integrator may cause unstable simulations, or give inaccurate results.

A.2.3 Testing Paradigms

In order to develop a reliable code base, we have implemented some best practices used in the technology sector. We advocate the use of these methods whenever developing some program of substantial size.

A.2.3.1 Test-Driven Development

Testing is essential when developing code. Some code may produce results, but the degree to which those results may be trusted is entirely dependant upon the amount of testing that has been performed on that code. How would one know if the results produce by untested code are genuine and possibly ground breaking, or disastrously wrong?

One may argue that developers should not be trusted to write useful and precise code, that they should instead be required to prove the robustness and accuracy of what they build. If some functionality is built and tested after the fact, there is a real chance that the test may be designed to pass, and that for some other set of arguments the test may fail. If a test is built and passes, we move on!

A far more robust method of testing, and one which forces developers to consider the input/output of a function seriously before blindly writing lines of code, is to build the test first. In this case, one does not worry about the mechanics of the function or how it produces its output until a test is written that defines how the function should behave. The more comprehensive this test is, the more likely the function being tested is error free when that test passes.

To implement this kind of test-driven development, we have used the suite provided by [Google](#), *GoogleTest*, or *gtest*.

A.2.3.2 Regression Testing

When developing software, it can be easy to submit changes to a complex program that change the behaviour of that program somewhere that was not intended. Once some new feature has been added, it is necessary to check that all functionality performs as expected.

To prevent unintentional changes in behaviour from occurring, or rather to spot that they have occurred as soon as possible, it is common practice to implement regression testing. Simply put, regression testing applies all the same tests to the code any time a change is made to the code base, be that whether a method has been altered or the code has been reorganised. As the tests will have been passing before the change, this saves time when debugging issues as the root of the problem can be more easily pinpointed.

As *METHOD* is publicly available, we can offload the regression testing to a continuous integration client, and perform all tests remotely and automatically. We have used Travis² to this end.

A.2.3.3 Code Validation and Unit Testing

Testing provides a method of ensuring that the intended input of some function produces the intended output (and in some cases that erroneous inputs produce well-behaved outputs, such as error messages, preventing those errors from contaminating the results of the simulation). Testing entire code bases all at once, or some large subset of the code base, may be necessary but can be very time consuming—such tests may take a long time to execute. Instead, and

²<https://travis-ci.com/>

especially whilst still developing, it is more common to test the minimum amount of code as is feasible.

A unit is considered to be the smallest part of a program that is testable. Examples include an object’s member function, or in some cases an entire class—anything that has a defined input and output. Often, it is best that the unit being tested be isolated from interactions with any other parts of the code base, although this may not always be possible. We shall now highlight some of the unit tests in `METHOD`.

Any simulation involving a physical domain will require boundary conditions—some rule applied to the outer most grid points that replicate some global behaviour. These may include periodic domains (out one end, and back in the other), or outflow domains. To ensure that our implementation of these types of boundary conditions are correct, small two-dimensional domains of 18^2 cells are created and the conserved vector is initialised in some predictable manner. Once the boundary conditions have been applied, it is simple to check that all values are as expected.

Described in Section 2.7 is the process of transforming the conserved variables into those of the primitive vector, namely (ρ, \mathbf{v}, p) . By choosing some values for the primitive vector, *a priori*, and converting them to the conserved quantities using the various relations for that model, we can ensure that the transformation of the conserved variables back to primitive form should match³ their initial values. Indeed we can go one step further, as the rootfinder requires some initial guess we can vary the accuracy of the guess by a margin that is unlikely to crop up in a simulation (off by $\sim 10\%$). If the resulting primitive variables still match their initial values we can be confident that this process should be accurate during real simulations.

Testing some elements of `METHOD` requires confidence that some function’s dependencies are already working. As an example, to test whether the numerical flux function is producing the correct output, we need to be sure the flux vector is correct also. In order to test a series of methods such as this, some special case tests have been created. One example of these is the equilibrium set up in which there is no variation in any one variable throughout the domain. As a result, we expect that the combination of the time integrator, flux vector and numerical flux function result in no change of the conserved vector up to floating point accuracy.

There are three sub-versions contained within `METHOD`—one serial version, and one parallel. In addition to the unit tests of the type described above, a further set of tests, this time using the pytest framework ([Krekel](#)), compare directly the output of serial and parallel versions of the code. This comparison is particularly useful at highlighting race conditions when parallelising existing code.

³to within the rootfinder’s tolerance

References

- B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):61102, feb 2016. doi:[10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102).
- B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, oct 2017a. doi:[10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101).
- B. P. Abbott et al. Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *The Astrophysical Journal*, 848(2):L13, oct 2017b. doi:[10.3847/2041-8213/aa920c](https://doi.org/10.3847/2041-8213/aa920c).
- B. P. Abbott et al. Multi-messenger Observations of a Binary Neutron Star Merger. *The Astrophysical Journal*, 848(2):L12, oct 2017c. doi:[10.3847/2041-8213/aa91c9](https://doi.org/10.3847/2041-8213/aa91c9).
- B. P. Abbott et al. Low-latency Gravitational-wave Alerts for Multimessenger Astronomy during the Second Advanced {LIGO} and Virgo Observing Run. *The Astrophysical Journal*, 875(2):161, apr 2019a. doi:[10.3847/1538-4357/ab0e8f](https://doi.org/10.3847/1538-4357/ab0e8f).
- B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9(3):31040, sep 2019b. doi:[10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040).
- R. Abbott et al. GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object. *Astrophysical Journal Letters*, 896(2):L44, 2020a. doi:[10.3847/2041-8213/ab960f](https://doi.org/10.3847/2041-8213/ab960f).
- R. Abbott et al. GW190412: Observation of a binary-black-hole coalescence with asymmetric masses. *Phys. Rev. D*, 102(4):43015, 2020b. doi:[10.1103/PhysRevD.102.043015](https://doi.org/10.1103/PhysRevD.102.043015).
- P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein. Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins. *Physical Review D*, 89(8):084041, 2014. doi:[10.1103/PhysRevD.89.084041](https://doi.org/10.1103/PhysRevD.89.084041).
- M. Alcubierre. Introduction to 3+1 relativity. *New York: Oxford University Press*, 2008.
- M. Alcubierre and B. Bru. Simple excision of a black hole in 3+1 numerical relativity. *Physical Review D*, 63(10):104006, 2001. doi:[10.1103/PhysRevD.63.104006](https://doi.org/10.1103/PhysRevD.63.104006).
- D. Alic, C. Bona-casas, C. Bona, L. Rezzolla, and C. Palenzuela. Conformal and covariant formulation of the Z4 system with constraint-violation damping. *Physical Review D*, 85, 2012. doi:[10.1103/PhysRevD.85.064040](https://doi.org/10.1103/PhysRevD.85.064040).

- D. Alic, W. Kastaun, L. Rezzolla, D. Alic, W. Kastaun, and L. Rezzolla. Constraint damping of the conformal and covariant formulation of the Z4 system in simulations of binary neutron stars. *Physical Review D*, 88(6):064049, 2013. doi:[10.1103/PhysRevD.88.064049](https://doi.org/10.1103/PhysRevD.88.064049).
- T. Amano. A Second-Order Divergence-Constrained Multidimensional Numerical Scheme for Relativistic Two-Fluid Electrodynamics. *The Astrophysical Journal*, 831(1):100, 2016. doi:[10.3847/0004-637X/831/1/100](https://doi.org/10.3847/0004-637X/831/1/100).
- J. D. Anderson. *Modern Compressible Flow: With Historical Perspective*. McGraw-Hill Publishing Company, 1990.
- M. Anderson, E. W. Hirschmann, L. Lehner, S. L. Liebling, P. M. Motl, D. Neilsen, C. Palenzuela, and J. E. Tohline. Magnetized Neutron-Star Mergers and Gravitational-Wave Signals. *Physical Review Letters*, 100:191101, 2008. doi:[10.1103/PhysRevLett.100.191101](https://doi.org/10.1103/PhysRevLett.100.191101).
- N. Andersson. Resistive relativistic magnetohydrodynamics from a charged multifluids perspective. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 86(4):043002, 2012. doi:[10.1103/PhysRevD.86.043002](https://doi.org/10.1103/PhysRevD.86.043002).
- N. Andersson, G. L. Comer, and I. Hawke. A variational approach to resistive relativistic plasmas. *Classical and Quantum Gravity*, 34(12):125001, oct 2017a. doi:[10.1088/1361-6382/aa6b37](https://doi.org/10.1088/1361-6382/aa6b37).
- N. Andersson, K. Dionysopoulou, I. Hawke, and G. L. Comer. Beyond ideal magnetohydrodynamics: Resistive, reactive and relativistic plasmas. *Classical and Quantum Gravity*, 34(12):125002, oct 2017b. doi:[10.1088/1361-6382/aa6b3a](https://doi.org/10.1088/1361-6382/aa6b3a).
- N. Andersson, I. Hawke, K. Dionysopoulou, and G. L. Comer. Beyond ideal magnetohydrodynamics: From fibration to 3+1 foliation. *Classical and Quantum Gravity*, 34(12):125003, oct 2017c. doi:[doi:10.1088/1361-6382/aa6b39](https://doi.org/10.1088/1361-6382/aa6b39).
- L. Antón, J. A. Miralles, J. M. Martí, J. M. Ibáñez, M. A. Aloy, and P. Mimica. Relativistic Magnetohydrodynamics: Renormalized Eigenvectors and Full Wave Decomposition Riemann Solver. *The Astrophysical Journal Supplement Series*, 188(1):1–31, 2010. doi:[10.1088/0067-0049/188/1/1](https://doi.org/10.1088/0067-0049/188/1/1).
- R. Arnowitt, S. Deser, and C. W. Misner. Republication of: The dynamics of general relativity. *General Relativity and Gravitation*, (40):1997–2027, 2008. doi:[10.1007/s10714-008-0661-1](https://doi.org/10.1007/s10714-008-0661-1).
- L. Baiotti and L. Rezzolla. Binary neutron star mergers: a review of Einstein’s richest laboratory. *Reports on Progress in Physics*, 80(6):096901, 2017. doi:[10.1088/1361-6633/aa67bb](https://doi.org/10.1088/1361-6633/aa67bb).
- D. S. Balsara and J. Kim. A Comparison between Divergence-Cleaning and Staggered-Mesh Formulations for Numerical Magnetohydrodynamics. *The Astrophysical Journal*, 602(2):1079–1090, 2004. doi:[10.1086/381051](https://doi.org/10.1086/381051).
- D. S. Balsara and D. S. Spicer. A Staggered Mesh Algorithm Using High Order Godunov Fluxes to Ensure Solenoidal Magnetic Fields in Magnetohydrodynamic Simulations. *Journal of Computational Physics*, 149(2):270–292, 1999. doi:[10.1006/jcph.1998.6153](https://doi.org/10.1006/jcph.1998.6153).
- J. C. A. Barata and M. S. Hussein. The Moore-Penrose Pseudoinverse. A Tutorial Review of the Theory. *Brazilian Journal of Physics*, 42(1):146–165, 2012. doi:[10.1007/s13538-011-0052-z](https://doi.org/10.1007/s13538-011-0052-z).

- M. Bauer, H. Cook, and B. Khailany. CudaDMA: Optimizing GPU Memory Bandwidth via Warp Specialization. *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis on - SC '11*, pages 1–11, 2011. doi:[10.1145/2063384.2063400](https://doi.org/10.1145/2063384.2063400).
- T. W. Baumgarte and S. L. Shapiro. Numerical integration of Einstein’s field equations. *Physical Review D*, 59(2):1–7, 1998. doi:[10.1103/PhysRevD.59.024007](https://doi.org/10.1103/PhysRevD.59.024007).
- T. W. Baumgarte and S. L. Shapiro. Numerical Relativity and Compact Binaries. *Physics Reports*, 376(2):41–131, 2002. doi:[10.1016/S0370-1573\(02\)00537-9](https://doi.org/10.1016/S0370-1573(02)00537-9).
- T. W. Baumgarte and S. L. Shapiro. Collapse of a magnetized star to a black hole. *The Astrophysical Journal*, 585(2):930–947, 2003. doi:[10.1086/346104](https://doi.org/10.1086/346104).
- T. W. Baumgarte and S. L. Shapiro. Numerical relativity: Solving Einstein’s equations on the computer. *Oxford University Press*, 2010. doi:[10.1017/CBO9781139193344](https://doi.org/10.1017/CBO9781139193344).
- G. Baym, C. Pethick, and D. Pines. Superfluidity in Neutron Stars. *Nature*, 224(5220):673–674, 1969. doi:[10.1038/224673a0](https://doi.org/10.1038/224673a0).
- D. Bazow, U. Heinz, and M. Strickland. Massively parallel simulations of relativistic fluid dynamics on graphics processing units with CUDA. *Computer Physics Communications*, 225(1): 92–113, 2016. doi:[10.1016/j.cpc.2017.01.015](https://doi.org/10.1016/j.cpc.2017.01.015).
- K. Beckwith and J. M. Stone. A second-order Godunov method for multi-dimensional relativistic magnetohydrodynamics. *Astrophysical Journal, Supplement Series*, 193(1):6, 2011. doi:[10.1088/0067-0049/193/1/6](https://doi.org/10.1088/0067-0049/193/1/6).
- M. A. Berger and G. Hornig. A generalized poloidaltoroidal decomposition and an absolute measure of helicity. *Journal of Physics A: Mathematical and Theoretical*, 51(49):495501, 2018. doi:[10.1088/1751-8121/aaca88](https://doi.org/10.1088/1751-8121/aaca88).
- M. J. Berger and J. Olinger. Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations. *Journal of Computational Physics*, 53(3):484–512, 1984. doi:[10.1016/0021-9991\(84\)90073-1](https://doi.org/10.1016/0021-9991(84)90073-1).
- R. D. Blandford and R. L. Znajek. Electromagnetic extraction of energy from Kerr black holes. *Monthly Notices of the Royal Astronomical Society*, 179(3):433–456, 1977. doi:[10.1093/mnras/179.3.433](https://doi.org/10.1093/mnras/179.3.433).
- C. Bona, J. Masso, E. Seidel, and J. Stela. New Formalism for Numerical Relativity. *American Physical Society*, 75(4):600–603, 1995. doi:[10.1103/PhysRevLett.75.600](https://doi.org/10.1103/PhysRevLett.75.600).
- M. Brio and C. C. Wu. An upwind differencing scheme for the equations of ideal magnetohydrodynamics. *Journal of Computational Physics*, 75(2):400–422, 1988. doi:[10.1016/0021-9991\(88\)90120-9](https://doi.org/10.1016/0021-9991(88)90120-9).
- A. Burrows. Supernova explosions in the Universe. *Nature*, 403(6771):727–733, 2000. doi:[10.1038/35001501](https://doi.org/10.1038/35001501).
- J. C. Butcher. A history of Runge-Kutta methods. *Applied Numerical Mathematics*, 20(3): 247–260, 1996. doi:[10.1016/0168-9274\(95\)00108-5](https://doi.org/10.1016/0168-9274(95)00108-5).

- M. Campanelli, C. O. Lousto, and Y. Zlochower. Spinning-black-hole binaries: The orbital hang-up. *Physical Review D*, 74(4):041501, 2006. doi:[10.1103/PhysRevD.74.041501](https://doi.org/10.1103/PhysRevD.74.041501).
- F. Carrasco, D. Vigan, C. Palenzuela, and J. A. Pons. Triggering magnetar outbursts in 3D force-free simulations. *Monthly Notices of the Royal Astronomical Society*, 484(1):L124–L129, 2019. doi:[10.1093/mnrasl/slz016](https://doi.org/10.1093/mnrasl/slz016).
- C. Cercignani and G. Kremer. *The Relativistic Boltzmann Equation: Theory and Applications*, volume 22. 2002. doi:[10.1007/978-3-0348-8165-4](https://doi.org/10.1007/978-3-0348-8165-4).
- S. Chawla, M. Anderson, M. Besselman, L. Lehner, S. L. Liebling, P. M. Motl, and D. Neilsen. Mergers of Magnetized Neutron Stars with Spinning Black Holes: Disruption, Accretion, and Fallback. *Physical Review Letters*, 105(11):111101, 2010. doi:[10.1103/PhysRevLett.105.111101](https://doi.org/10.1103/PhysRevLett.105.111101).
- Y. Chen, G. Tóth, P. Cassak, X. Jia, T. I. Gombosi, J. A. Slavin, S. Markidis, I. B. Peng, V. K. Jordanova, and M. G. Henderson. Global Three-Dimensional Simulation of Earth’s Dayside Reconnection Using a Two-Way Coupled Magnetohydrodynamics With Embedded Particle-in-Cell Model: Initial Results. *Journal of Geophysical Research : Space Physics*, 122(10):10318–10335, 2017. doi:[10.1002/2017JA024186](https://doi.org/10.1002/2017JA024186).
- C. Chirenti, R. Gold, and M. C. Miller. Gravitational Waves from F-modes Excited by the Inspiral of Highly Eccentric Neutron Star Binaries. *The Astrophysical Journal*, 837(1):67, 2017. doi:[10.3847/1538-4357/aa5ebb](https://doi.org/10.3847/1538-4357/aa5ebb).
- R. Ciolfi. Short gamma-ray burst central engines. *International Journal of Modern Physics D*, 27(13), 2018. doi:[10.1142/S021827181842004X](https://doi.org/10.1142/S021827181842004X).
- R. Ciolfi. Binary neutron star mergers after GW170817. *Frontiers in Astronomy and Space Sciences*, 7(27):1–5, 2020. doi:[10.3389/fspas.2020.00027](https://doi.org/10.3389/fspas.2020.00027).
- R. Ciolfi and J. V. Kalinani. Magnetically driven baryon winds from binary neutron star merger remnants and the blue kilonova of August 2017. *American Astronomical Society*, 900(2):L35, 2020. doi:[10.3847/2041-8213/abb240](https://doi.org/10.3847/2041-8213/abb240).
- R. Ciolfi, W. Kastaun, B. Giacomazzo, A. Endrizzi, D. M. Siegel, and R. Perna. General relativistic magnetohydrodynamic simulations of binary neutron star mergers forming a long-lived neutron star. *Physical Review D*, 95(6):063016, 2017. doi:[10.1103/PhysRevD.95.063016](https://doi.org/10.1103/PhysRevD.95.063016).
- R. Ciolfi, W. Kastaun, J. V. Kalinani, and B. Giacomazzo. First 100 ms of a long-lived magnetized neutron star formed in a binary neutron star merger. *Physical Review D*, 100(2):23005, 2019. doi:[10.1103/PhysRevD.100.023005](https://doi.org/10.1103/PhysRevD.100.023005).
- F. Ciolletta, J. V. Kalinani, and B. Giacomazzo. Spritz: a new fully general-relativistic magnetohydrodynamic code. *Classical and Quantum Gravity*, 37(13):135010, 2019. doi:[10.1088/1361-6382/ab8be8](https://doi.org/10.1088/1361-6382/ab8be8).
- P. Colella and P. R. Woodward. The Piecewise Parabolic Method for Gas-Dynamical Systems. *Journal of Computational Physics*, 54(1):174–201, 1984. doi:[10.1016/0021-9991\(84\)90143-8](https://doi.org/10.1016/0021-9991(84)90143-8).
- I. Contopoulos, D. Kazanas, and C. Fendt. The axisymmetric pulsar magnetosphere. *Astrophysical Journal*, 511(1):351–358, 1999. doi:[10.1086/306652](https://doi.org/10.1086/306652).

- D. F. Cowen. The Dawn of Multimessenger Astronomy. *Journal of Physics: Conference Series*, 1342(1):9–12, 2020. doi:[10.1088/1742-6596/1342/1/012001](https://doi.org/10.1088/1742-6596/1342/1/012001).
- L. K. S. Daldorff, G. Tóth, T. I. Gombosi, G. Lapenta, J. Amaya, S. Markidis, and J. U. Brackbill. Two-way coupling of a global Hall magnetohydrodynamics model with a local implicit particle-in-cell model. *Journal of Computational Physics*, 268(1):236–254, 2014. doi:[10.1016/j.jcp.2014.03.009](https://doi.org/10.1016/j.jcp.2014.03.009).
- M. B. Davies and W. Benz. Merging neutron stars I: Initial results for coalescence of non-corotating systems. *Astrophysical Journal*, 431(1):742, 1994. doi:[10.1086/174525](https://doi.org/10.1086/174525).
- S. De, D. Finstad, J. M. Lattimer, D. A. Brown, E. Berger, and C. M. Biwer. Tidal Deformabilities and Radii of Neutron Stars from the Observation of GW170817. *Physical Review Letters*, 121(9):91102, 2018. doi:[10.1103/PhysRevLett.121.091102](https://doi.org/10.1103/PhysRevLett.121.091102).
- R. De Pietri, A. Feo, F. Maione, and F. Löffler. Modeling equal and unequal mass binary neutron star mergers using public codes. *Physical Review D*, 93(6), 2016. doi:[10.1103/PhysRevD.93.064047](https://doi.org/10.1103/PhysRevD.93.064047).
- V. Decoene, C. Guépin, K. Fang, K. Kotera, and B. Metzger. High-energy neutrinos from fallback accretion of binary neutron star merger remnants. *Journal of Cosmology and Astroparticle Physics*, 2020(4):045–045, 2020. doi:[10.1088/1475-7516/2020/04/045](https://doi.org/10.1088/1475-7516/2020/04/045).
- A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg. Hyperbolic Divergence Cleaning for the MHD Equations. *Journal of Computational Physics*, 175(2): 645–673, 2002. doi:[10.1006/jcph.2001.6961](https://doi.org/10.1006/jcph.2001.6961).
- D. Desai, B. D. Metzger, and F. Foucart. Imprints of r-process heating on fall-back accretion: distinguishing black hole-neutron star from double neutron star mergers. *Monthly Notices of the Royal Astronomical Society*, 485(3):4404–4412, 2019. doi:[10.1093/mnras/stz644](https://doi.org/10.1093/mnras/stz644).
- T. Dietrich, M. Ujevic, W. Tichy, S. Bernuzzi, and B. Brügmann. Gravitational waves and mass ejecta from binary neutron star mergers: Effect of the mass ratio. *Physical Review D*, 95(2): 024029, 2017. doi:[10.1103/PhysRevD.95.024029](https://doi.org/10.1103/PhysRevD.95.024029).
- T. Dietrich, D. Radice, S. Bernuzzi, F. Zappa, A. Perego, B. Brügmann, S. V. Chaurasia, R. Dudi, W. Tichy, and M. Ujevic. CoRe database of binary neutron star merger waveforms. *Classical and Quantum Gravity*, 35(24):24LT01, 2018. doi:[10.1088/1361-6382/aaebc0](https://doi.org/10.1088/1361-6382/aaebc0).
- K. Dionysopoulou. *General-Relativistic Magnetohydrodynamics in Compact Objects*. PhD thesis, Universität Potsdam, 2016.
- K. Dionysopoulou, D. Alic, C. Palenzuela, L. Rezzolla, and B. Giacomazzo. General-relativistic resistive magnetohydrodynamics in three dimensions: Formulation and tests. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 88(4):044020, 2013. doi:[10.1103/PhysRevD.88.044020](https://doi.org/10.1103/PhysRevD.88.044020).
- K. Dionysopoulou, D. Alic, and L. Rezzolla. General-relativistic resistive-magnetohydrodynamic simulations of binary neutron stars. *Physical Review D*, 92(8):084064, 2015. doi:[10.1103/PhysRevD.92.084064](https://doi.org/10.1103/PhysRevD.92.084064).

- V. A. Dommès, M. E. Gusakov, and P. S. Shternin. Dissipative relativistic magnetohydrodynamics of a multicomponent mixture and its application to neutron stars. *Physical Review D*, 101(10):103020, 2020. doi:[10.1103/PhysRevD.101.103020](https://doi.org/10.1103/PhysRevD.101.103020).
- W. E. East, V. Paschalidis, F. Pretorius, and S. L. Shapiro. Relativistic simulations of eccentric binary neutron star mergers: One-arm spiral instability and effects of neutron star spin. *Physical Review D*, 93(2):024011, 2016. doi:[10.1103/PhysRevD.93.024011](https://doi.org/10.1103/PhysRevD.93.024011).
- D. Eichler, M. Liviot, T. Piran, and D. N. Schramm. Nucleosynthesis, neutrino bursts and gamma-rays from coalescing neutron stars. *Nature*, 340(6229):126–128, 1989. doi:[10.1038/340126a0](https://doi.org/10.1038/340126a0).
- B. Einfeldt. On Godunov-type methods for gas. *SIAM Journal on Numerical Analysis*, 25(2):294–318, 1988. doi:[10.1137/0725021](https://doi.org/10.1137/0725021).
- A. Einstein. Approximate integration of the field equations of gravitation. *Prussian Academy of Sciences*, pages 688–696, 1916. URL <https://einsteinpapers.press.princeton.edu/vol6-trans/213>.
- K. R. Eppley. *The numerical evolution of the collision of two black holes*. PhD thesis, Princeton University, 1975.
- Z. B. Etienne, Y. T. Liu, and S. L. Shapiro. Relativistic magnetohydrodynamics in dynamical spacetimes: A new AMR implementation. *Physical Review D*, 82(8):084031, 2010. doi:[10.1103/PhysRevD.82.084031](https://doi.org/10.1103/PhysRevD.82.084031).
- Z. B. Etienne, V. Paschalidis, Y. T. Liu, and S. L. Shapiro. Relativistic MHD in dynamical spacetimes: Improved EM gauge condition for AMR grids. *Physical Review D*, 85(2):024013, 2012. doi:[10.1103/PhysRevD.85.024013](https://doi.org/10.1103/PhysRevD.85.024013).
- C. R. Evans and J. F. Hawley. Simulation of Magnetohydrodynamic Flows:a Constrained Transport Method. *The Astrophysical Journal*, 332(1):659–677, 1988. doi:[10.1086/166684](https://doi.org/10.1086/166684).
- J. a. Faber and F. a. Rasio. Binary Neutron Star Mergers. *Living Reviews in Relativity*, 15(1):8, 2012. doi:[10.12942/lrr-2012-8](https://doi.org/10.12942/lrr-2012-8).
- B. D. Farris, R. Gold, V. Paschalidis, Z. B. Etienne, and S. L. Shapiro. Binary black-hole mergers in magnetized disks: Simulations in full general relativity. *Physical Review Letters*, 109(22):221102, 2012. doi:[10.1103/PhysRevLett.109.221102](https://doi.org/10.1103/PhysRevLett.109.221102).
- R. Fernández, A. Tchekhovskoy, E. Quataert, F. Foucart, and D. Kasen. Long-term GRMHD simulations of neutron star merger accretion discs: implications for electromagnetic counterparts. *Monthly Notices of the Royal Astronomical Society*, 482(3):3373–3393, 2019. doi:[10.1093/mnras/sty2932](https://doi.org/10.1093/mnras/sty2932).
- T. Flåtten and H. Lund. Relaxation two-phase flow models and the subcharacteristic condition. *Mathematical Models and Methods in Applied Sciences*, 21(12):2379–2407, 2011. doi:[10.1142/S0218202511005775](https://doi.org/10.1142/S0218202511005775).
- J. A. Font. Numerical Hydrodynamics in General Relativity. *Living Reviews in Relativity*, 3(2), 2000. doi:[10.12942/lrr-2000-2](https://doi.org/10.12942/lrr-2000-2).

- J. A. Font. Numerical hydrodynamics and magnetohydrodynamics in general relativity. *Living Reviews in Relativity*, 11(7), 2008. doi:[10.12942/lrr-2008-7](https://doi.org/10.12942/lrr-2008-7).
- F. Foucart, M. B. Deaton, M. D. Duez, E. O. Connor, C. D. Ott, R. Haas, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilagyi. Neutron star-black hole mergers with a nuclear equation of state and neutrino cooling: Dependence in the binary parameters. *Physical Review D*, 90(2):024026, 2014. doi:[10.1103/PhysRevD.90.024026](https://doi.org/10.1103/PhysRevD.90.024026).
- F. Foucart, M. D. Duez, L. E. Kidder, R. Nguyen, H. P. Pfeiffer, and M. A. Scheel. Evaluating radiation transport errors in merger simulations using a Monte Carlo algorithm. *Physical Review D*, 98(6):063007, 2018. doi:[10.1103/PhysRevD.98.063007](https://doi.org/10.1103/PhysRevD.98.063007).
- F. Foucart, M. D. Duez, L. E. Kidder, S. M. Nissanke, H. P. Pfeiffer, and M. A. Scheel. Numerical simulations of neutron star-black hole binaries in the near-equal-mass regime. *Physical Review D*, 99(10):103025, 2019. doi:[10.1103/PhysRevD.99.103025](https://doi.org/10.1103/PhysRevD.99.103025).
- C. Freiburghaus, S. Rosswog, and F.-K. Thielemann. r-process in neutron star mergers. *Astrophysical Journal*, 525(2):121–124, 1999. doi:[10.1086/312343](https://doi.org/10.1086/312343).
- H. Friedrich. On the regular and the asymptotic characteristic initial value problem for Einstein’s vacuum field equations. *Proceedings of the National Academy of Sciences of the United States of America*, 375(1761):169–184, 1981. doi:[10.1098/rspa.1981.0045](https://doi.org/10.1098/rspa.1981.0045).
- E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R. H. Castain, D. J. Daniel, R. L. Graham, and T. S. Woodall. Open MPI: Goals, concept, and design of a next generation MPI implementation. *11th European PVM/MPI Users’ Group Meeting*, 3241, 2004. doi:[10.1007/978-3-540-30218-6_19](https://doi.org/10.1007/978-3-540-30218-6_19).
- T. A. Gardiner and J. M. Stone. An unsplit Godunov method for ideal MHD via constrained transport in three dimensions. *Journal of Computational Physics*, 227(8):4123–4141, 2008. doi:[10.1016/j.jcp.2007.12.017](https://doi.org/10.1016/j.jcp.2007.12.017).
- E. Gawehn, J. A. Hiss, J. B. Brown, G. Schneider, E. Gawehn, J. A. Hiss, J. B. Brown, G. S. Advancing, E. Gawehn, J. A. Hiss, J. B. Brown, and G. Schneider. Advancing drug discovery via GPU-based deep learning. *Expert Opinion on Drug Discovery*, 13(7):579–582, 2018. doi:[10.1080/17460441.2018.1465407](https://doi.org/10.1080/17460441.2018.1465407).
- B. Giacomazzo. Exact Riemann Solver. URL http://www.brunogiacomazzo.org/old_{_}web_{_}page/exactriemannsolv.html.
- B. Giacomazzo and L. Rezzolla. WhiskyMHD: a new numerical code for general relativistic magnetohydrodynamics. *Classical and Quantum Gravity*, 24(12):S235–S258, 2007. doi:[10.1088/0264-9381/24/12/S16](https://doi.org/10.1088/0264-9381/24/12/S16). URL <http://stacks.iop.org/0264-9381/24/i=12/a=S16>.
- B. Giacomazzo, L. Rezzolla, and L. Baiotti. Can magnetic fields be detected during the inspiral of binary neutron stars? *Monthly Notices of the Royal Astronomical Society*, 399(1):164–168, 2009. doi:[10.1111/j.1745-3933.2009.00745.x](https://doi.org/10.1111/j.1745-3933.2009.00745.x).

- B. Giacomazzo, L. Rezzolla, and L. Baiotti. Accurate evolutions of inspiralling and magnetized neutron stars : Equal-mass binaries. *Phys. Rev. D*, 044014(4):044014, 2011. doi:[10.1103/PhysRevD.83.044014](https://doi.org/10.1103/PhysRevD.83.044014).
- B. Giacomazzo, J. Zrake, P. C. Duffell, A. I. Macfadyen, and R. Perna. Producing Magnetar Magnetic Fields in the Merger of Binary Neutron Stars. *Astrophysical Journal*, 809(1):39, 2015. doi:[10.1088/0004-637X/809/1/39](https://doi.org/10.1088/0004-637X/809/1/39).
- R. Gill, A. Nathanail, and L. Rezzolla. When did the remnant of GW170817 collapse to a black hole? *The Astrophysical Journal*, 876(2):139, 2019. doi:[10.3847/1538-4357/ab16da](https://doi.org/10.3847/1538-4357/ab16da).
- D. A. Godzieba, D. Radice, and S. Bernuzzi. On the maximum mass of neutron stars and GW190814. *arXiv e-prints*, 2020. URL [arxiv:2007.10999v1](https://arxiv.org/abs/2007.10999v1).
- J. A. González, U. Sperhake, B. Brügmann, M. Hannam, and S. Husa. Maximum Kick from Nonspinning Black-Hole Binary Inspiral. *Physical Review Letters*, 98(9):091101, 2007. doi:[10.1103/PhysRevLett.98.091101](https://doi.org/10.1103/PhysRevLett.98.091101).
- Google. Google Test. URL <https://github.com/google/googletest>.
- D. Gottlieb and S. Gottlieb. Spectral methods for discontinuous problems. *D.F. Griffiths, G.A. Watson (Eds.), Proceedings of the 20th Biennial Conference on Numerical Analysis, University of Dundee*, pages 65–71, 2003.
- S. Gottlieb and C.-W. Shu. Total variation diminishing Runge-Kutta schemes. *Mathematics of Computation of the American Mathematical Society*, 67(221):73–85, 1998. doi:[10.1090/S0025-5718-98-00913-2](https://doi.org/10.1090/S0025-5718-98-00913-2).
- S. Gottlieb, D. I. Ketcheson, and C.-W. Shu. High order strong stability preserving time discretizations. *Journal of Scientific Computing*, 38(3):251–289, 2009. doi:[10.1007/s10915-008-9239-z](https://doi.org/10.1007/s10915-008-9239-z).
- A. Gruzinov. Pulsar emission and force-free electrodynamics. *Astrophysical Journal*, 677(1):L69—L71, 2007. doi:[10.1086/519839](https://doi.org/10.1086/519839).
- J. Guilet, R. Raynaud, and M. Bugli. A global model of the magnetorotational instability in protoneutron stars. *arXiv e-prints*, 2020. URL [arxiv:2005.03567v1](https://arxiv.org/abs/2005.03567v1).
- M. Hanauske, J. Steinheimer, A. Motornenko, V. Vovchenko, L. Bovard, E. R. Most, L. J. Papenfort, S. Schramm, and H. Stöcker. Neutron star mergers: Probing the EoS of hot, dense matter by gravitational waves. *Particles*, 2(1):44–56, 2019. doi:[10.3390/particles2010004](https://doi.org/10.3390/particles2010004).
- A. Harpole. *Multiscale Modelling Of Neutron Star Oceans*. PhD thesis, University of Southampton, 2018.
- A. Harten, P. D. Lax, and B. van Leer. On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Review*, 25(1):35–61, 1983. doi:[10.1137/1025002](https://doi.org/10.1137/1025002).
- A. Harutyunyan and A. Sedrakian. Electrical conductivity of a warm neutron star crust in magnetic fields. *Physical Review C*, 94(2):025805, 2016. doi:[10.1103/PhysRevC.94.025805](https://doi.org/10.1103/PhysRevC.94.025805).
- A. Harutyunyan, A. Nathanail, L. Rezzolla, and A. Sedrakian. Electrical resistivity and Hall effect in binary neutron star merger. *The European Physical Journal A*, 54(11):191, 2018. doi:[10.1140/epja/i2018-12624-1](https://doi.org/10.1140/epja/i2018-12624-1).

- D. Hilditch, A. Schoepe, and F.-s.-u. Jena. Hyperbolicity of divergence cleaning and vector potential formulations of general relativistic magnetohydrodynamics. *Physical Review D*, 99(10):104034, 2019. doi:[10.1103/PhysRevD.99.104034](https://doi.org/10.1103/PhysRevD.99.104034).
- V. Honkkila and P. Janhunen. HLLC solver for ideal relativistic MHD. *Journal of Computational Physics*, 223(2):643–656, 2007. doi:[10.1016/j.jcp.2006.09.027](https://doi.org/10.1016/j.jcp.2006.09.027).
- J. D. Hunter. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*, 9(90), 2007. doi:[10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).
- O. Just, M. Obergaulinger, H. Janka, A. Bauswein, and N. Schwarz. Neutron-star merger ejecta as obstacles to neutrino-powered jets of gamma-ray bursts. *The Astrophysical Journal Letters*, 816(2):L30, 2016. doi:[10.3847/2041-8205/816/2/L30](https://doi.org/10.3847/2041-8205/816/2/L30).
- D. Kasen, B. Metzger, J. Barnes, E. Quataert, and E. Ramirez-Ruiz. Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event. *Nature*, 551(7678):80–84, 2017. doi:[10.1038/nature24453](https://doi.org/10.1038/nature24453).
- V. M. Kaspi and A. M. Beloborodov. Magnetars. *Annual Review of Astronomy and Astrophysics*, 55(1):261–302, 2017. doi:[10.1146/annurev-astro-081915-023329](https://doi.org/10.1146/annurev-astro-081915-023329).
- T. Kawamura, B. Giacomazzo, W. Kastaun, R. Ciolfi, A. Endrizzi, and L. Baiotti. Binary neutron star mergers and short gamma-ray bursts: Effects of magnetic field orientation, equation of state, and mass ratio. *Physical Review D*, 94(6):064012, 2016. doi:[10.1103/PhysRevD.94.064012](https://doi.org/10.1103/PhysRevD.94.064012).
- B. Kelly, P. Laguna, K. Lockitch, J. Pullin, E. Schnetter, D. Shoemaker, and M. Tiglio. Cure for unstable numerical evolutions of single black holes: Adjusting the standard ADM equations in the spherically symmetric case. *Physical Review D*, 64(8):084013, 2001. doi:[10.1103/PhysRevD.64.084013](https://doi.org/10.1103/PhysRevD.64.084013).
- K. Kiuchi, K. Kyutoku, Y. Sekiguchi, M. Shibata, and T. Wada. High resolution numerical relativity simulations for the merger of binary magnetized neutron stars. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 90(4):041502, 2014. doi:[10.1103/PhysRevD.90.041502](https://doi.org/10.1103/PhysRevD.90.041502).
- K. Kiuchi, P. Cerdá-Durán, K. Kyutoku, Y. Sekiguchi, and M. Shibata. Efficient magnetic-field amplification due to the Kelvin-Helmholtz instability in binary neutron star mergers. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 92(12):124034, 2015. doi:[10.1103/PhysRevD.92.124034](https://doi.org/10.1103/PhysRevD.92.124034).
- K. Kiuchi, K. Kyutoku, Y. Sekiguchi, and M. Shibata. Global simulations of strongly magnetized remnant massive neutron stars formed in binary neutron star mergers. *Physical Review D*, 97(12):124039, 2017. doi:[10.1103/PhysRevD.97.124039](https://doi.org/10.1103/PhysRevD.97.124039).
- S. S. Komissarov. Multi-dimensional Numerical Scheme for Resistive Relativistic MHD. *Monthly Notices of the Royal Astronomical Society*, 382(3):995–1004, 2007. doi:[10.1111/j.1365-2966.2007.12448.x](https://doi.org/10.1111/j.1365-2966.2007.12448.x).
- C. Kong. *Comparison of Approximate Riemann Solvers*. PhD thesis, University of Reading, 2011.

- H. Krekel. pytest. URL <https://docs.pytest.org/en/latest/index.html>.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, 12(141):04015009, 2012. doi:[10.1061/\(ASCE\)GT.1943-5606.0001284](https://doi.org/10.1061/(ASCE)GT.1943-5606.0001284).
- C. B. Laney. *Computational Gasdynamics*. Cambridge: Cambridge University Press., 1998. doi:[10.1017/CBO9780511605604](https://doi.org/10.1017/CBO9780511605604).
- J. M. Lattimer and M. Prakash. The Physics of Neutron Stars. *Science*, 304(5670):536–542, 2004. doi:[10.1126/science.1090720](https://doi.org/10.1126/science.1090720).
- J. M. Lattimer and F. D. Swesty. A generalized equation of state for hot, dense matter. *Nuclear Physics A*, 535:331–376, 1991. doi:[10.1016/0375-9474\(91\)90452-C](https://doi.org/10.1016/0375-9474(91)90452-C).
- D. Lazzati, R. Ciolfi, and R. Perna. Intrinsic properties of the engine and jet that powered the short gamma-ray burst associated with GW170817. *American Astronomical Society*, 898(1): 59, 2020. doi:[10.3847/1538-4357/ab9a44](https://doi.org/10.3847/1538-4357/ab9a44).
- D. Lee, I. Dinov, B. Dong, B. Gutman, I. Yanovsky, and A. W. Toga. CUDA Optimization Strategies for Compute- and Memory-Bound Neuroimaging Algorithms. *Computer Methods and Programs in Biomedicine*, 106(3):175–187, 2012. doi:[10.1016/j.cmpb.2010.10.013](https://doi.org/10.1016/j.cmpb.2010.10.013).
- L. Lehner and F. Pretorius. Numerical Relativity and Astrophysics. *Annual Review of Astronomy and Astrophysics*, 52(1):661–694, 2014. doi:[10.1146/annurev-astro-081913-040031](https://doi.org/10.1146/annurev-astro-081913-040031).
- L. Lehner, C. Palenzuela, S. L. Liebling, C. Thompson, and C. Hanna. Intense electromagnetic outbursts from collapsing hypermassive neutron stars. *Physical Review D*, 86(10):104035, 2012. doi:[10.1103/PhysRevD.86.104035](https://doi.org/10.1103/PhysRevD.86.104035).
- L. Lehner, S. L. Liebling, C. Palenzuela, O. L. Caballero, E. O. Connor, M. Anderson, and D. Neilsen. Unequal mass binary neutron star mergers and multimessenger signals. *Classical and Quantum Gravity*, 33(18):184002, 2016. doi:[10.1088/0264-9381/33/18/184002](https://doi.org/10.1088/0264-9381/33/18/184002).
- R. J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, 2002.
- A. G. M. Lewis and H. P. Pfeiffer. GPU-Accelerated Simulations of Isolated Black Holes. *Classical and Quantum Gravity*, 35(9):095017, 2018. doi:[10.1088/1361-6382aab256](https://doi.org/10.1088/1361-6382aab256).
- L.-X. Li and B. Paczynski. Transient events from neutron star mergers. *Astrophysical Journal*, 507(1):L59—L62, 1998. doi:[10.1086/311680](https://doi.org/10.1086/311680).
- S. L. Liebling. Singularity threshold of the non-linear sigma model using 3D adaptive mesh refinement. *Physical Review D*, 66(4):041703, 2002. doi:[10.1103/PhysRevD.66.041703](https://doi.org/10.1103/PhysRevD.66.041703).
- S. Liu, J. Tang, C. Wang, Q. Wang, and J.-L. Gaudiot. A Unified Cloud Platform for Autonomous Driving. *IEEE*, 50(12):42–49, 2017. doi:[10.1109/MC.2017.4451224](https://doi.org/10.1109/MC.2017.4451224).
- T.-p. Liu. Hyperbolic Conservation Laws with Relaxation. *Communication Mathematical Physics*, 108(1):153–175, 1987. doi:[10.1007/BF01210707](https://doi.org/10.1007/BF01210707).

- Y. T. Liu, S. L. Shapiro, Z. B. Etienne, and K. Taniguchi. General relativistic simulations of magnetized binary neutron star mergers. *Physical Review D*, 78(2):024012, 2008. doi:[10.1103/PhysRevD.78.024012](https://doi.org/10.1103/PhysRevD.78.024012).
- F. Löffler, J. Faber, E. Bentivegna, T. Bode, P. Diener, R. Haas, I. Hinder, B. C. Mundim, C. D. Ott, E. Schnetter, G. Allen, M. Campanelli, and P. Laguna. The Einstein Toolkit: a community computational infrastructure for relativistic astrophysics. *Classical and Quantum Gravity*, 29(11):115001, 2012. doi:[10.1088/0264-9381/29/11/115001](https://doi.org/10.1088/0264-9381/29/11/115001).
- P. Londrillo and L. Del Zanna. On the Divergence-Free Condition in Godunov-Type Schemes for Ideal Magnetohydrodynamics: the Upwind Constrained Transport Method. *Journal of Computational Physics*, 195(1):17–48, 2004. doi:[10.1016/j.jcp.2003.09.016](https://doi.org/10.1016/j.jcp.2003.09.016).
- D. R. Lorimer. Binary and Millisecond Pulsars. *Living Reviews in Relativity*, 11(1):8, 2008. doi:[10.12942/lrr-2008-8](https://doi.org/10.12942/lrr-2008-8).
- N. F. Loureiro and D. A. Uzdensky. Magnetic reconnection: from the Sweet Parker model to stochastic plasmoid chains. *Plasma Physics and Controlled Fusion*, 58(1):014021, 2015. doi:[10.1088/0741-3335/58/1/014021](https://doi.org/10.1088/0741-3335/58/1/014021).
- P. MacNeice, K. M. Olson, C. Mobarry, R. D. Fainchtein, and C. Packer. PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126(3):330–354, 2000. doi:[10.1016/S0010-4655\(99\)00501-9](https://doi.org/10.1016/S0010-4655(99)00501-9).
- B. Margalit and B. D. Metzger. Constraining the Maximum Mass of Neutron Stars from Multimessenger Observations of GW170817. *The Astrophysical Journal Letters*, 850(2):L19, 2017. doi:[10.3847/2041-8213/aa991c](https://doi.org/10.3847/2041-8213/aa991c).
- J. M. Martí and E. Müller. Numerical hydrodynamics in special relativity. *Living Reviews in Relativity*, 2(1):3, 1999. doi:[10.12942/lrr-2003-7](https://doi.org/10.12942/lrr-2003-7).
- J. M. Martí and E. Müller. Grid-based Methods in Relativistic Hydrodynamics and Magnetohydrodynamics. *Living Reviews in Computational Astrophysics*, 1(1):3, 2015. doi:[10.1007/lrca-2015-3](https://doi.org/10.1007/lrca-2015-3).
- J. M. Martí, J. M. Ibanez, and J. A. Miralles. Numerical relativistic hydrodynamics: Local characteristic approach. *Physical Review D*, 43(12):1, 1991. doi:[10.1103/PhysRevD.43.3794](https://doi.org/10.1103/PhysRevD.43.3794).
- J. C. McKinney. General relativistic force-free electrodynamics: a new code and applications to black hole magnetospheres. *Monthly Notices of the Royal Astronomical Society*, 367(4):1797–1807, 2006. doi:[j.1365-2966.2006.10087.x](https://doi.org/10.1111/j.1365-2966.2006.10087.x).
- B. D. Metzger, A. Arcones, E. Quataert, and G. Martínez-Pinedo. The effects of r-process heating on fallback accretion in compact object mergers. *Monthly Notices of the Royal Astronomical Society*, 402(4):2771–2777, 2010a. doi:[10.1111/j.1365-2966.2009.16107.x](https://doi.org/10.1111/j.1365-2966.2009.16107.x).
- B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *Monthly Notices of the Royal Astronomical Society*, 406(4):2650–2662, 2010b. doi:[10.1111/j.1365-2966.2010.16864.x](https://doi.org/10.1111/j.1365-2966.2010.16864.x).

- A. Mignone and J. C. McKinney. Equation of state in relativistic magnetohydrodynamics: Variable versus constant adiabatic index. *Monthly Notices of the Royal Astronomical Society*, 378(3):1118–1130, 2007. doi:[10.1111/j.1365-2966.2007.11849.x](https://doi.org/10.1111/j.1365-2966.2007.11849.x).
- A. Mignone, M. Ugliano, and G. Bodo. A five-wave Harten-Lax-van Leer Riemann solver for relativistic magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 393(4):1141–1156, 2009. doi:[10.1111/j.1365-2966.2008.14221.x](https://doi.org/10.1111/j.1365-2966.2008.14221.x).
- C. W. Misner, K. Thorne, and J. Wheeler. Gravitation. *W.H. Freeman and Co., San Francisco*, 1974.
- J. Muddle. *Advanced numerical methods for neutron star interfaces*. PhD thesis, University of Southampton, 2015. URL <http://eprints.soton.ac.uk/375551/>.
- A. Murguia-Berthier, G. Montes, E. Ramirez-Ruiz, F. D. Colle, and W. H. Lee. Necessary conditions for short gamma-ray burst production in binary neutron star mergers. *The Astrophysical Journal Letters*, 788(1):L8, 2014. doi:[10.1088/2041-8205/788/1/L8](https://doi.org/10.1088/2041-8205/788/1/L8).
- A. Murguia-berthier, E. Ramirez-ruiz, G. Montes, F. D. Colle, L. Rezzolla, S. Rosswog, K. Takami, A. Perego, and W. H. Lee. The Properties of Short Gamma-Ray Burst Jets Triggered by Neutron Star Mergers. *The Astrophysical Journal Letters*, 835(2):L34, 2017. doi:[10.3847/2041-8213/aa5b9e](https://doi.org/10.3847/2041-8213/aa5b9e).
- A. Murguia-Berthier, E. Ramirez-Ruiz, F. D. Colle, A. Janiuk, S. Rosswog, and W. H. Lee. The fate of the merger remnant in GW170817 and its imprint on the jet structure. *arXiv e-prints*, 2020. URL <https://arxiv.org/abs/2007.12245>.
- T. Nakamura, K. Oohara, and Y. Kojima. General relativistic collapse to black holes and gravitational waves from black holes. *Progress of Theoretical Physics Supplement*, 90(1):1–218, 1987. doi:[10.1143/PTPS.90.1](https://doi.org/10.1143/PTPS.90.1).
- E. Nakar. Short-hard gamma-ray bursts. *Physics Reports*, 442(1):166–236, 2007. doi:[10.1016/j.physrep.2007.02.005](https://doi.org/10.1016/j.physrep.2007.02.005).
- R. Narayan, B. Paczynski, and T. Piran. Gamma-ray bursts as the death throes of massive binary stars. *The Astrophysical Journal*, 395(1):L83, 1992. doi:[10.1086/186493](https://doi.org/10.1086/186493).
- A. Nathanail, R. Gill, O. Porth, C. M. Fromm, and L. Rezzolla. On the opening angle of magnetized jets from neutron-star mergers: the case of GRB170817a. *Monthly Notices of the Royal Astronomical Society*, 495(4):3780–3787, 2020. doi:[10.1093/mnras/staa1454](https://doi.org/10.1093/mnras/staa1454).
- M. Nicholl, E. Berger, D. Kasen, B. D. Metzger, J. Elias, C. Briceno, K. D. Alexander, P. K. Blanchard, R. Chornock, P. S. Cowperthwaite, T. Eftekhari, W. Fong, R. Margutti, V. A. Villar, P. K. G. Williams, W. Brown, J. Annis, A. Bahramian, D. Brout, D. A. Brown, H. Y. Chen, J. C. Clemens, E. Denny, B. Dunlap, D. E. Holz, E. Marchesini, F. Massaro, N. Moskowitz, I. Pelisoli, A. Rest, F. Ricci, M. Sako, M. Soares-Santos, and J. Strader. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. III. Optical and UV Spectra of a Blue Kilonova From Fast Polar Ejecta. *American Astronomical Society*, 848(2):L18, 2017. doi:[10.3847/2041-8213/aa9029](https://doi.org/10.3847/2041-8213/aa9029).
- J. Nickolls, I. Buck, M. Garland, and K. Skadron. Scalable parallel programming with CUDA. *Association for Computing Machinery*, 16(1):14, 2008. doi:[10.1145/1401132.1401152](https://doi.org/10.1145/1401132.1401152).

- Nvidia. Nvidia GeForce GTX 1080. *Whitepaper*, page 52, 2016. URL http://international.download.nvidia.com/geforce-com/international/pdfs/GeForce_GTX_1080_Whitepaper_FINAL.pdf.
- Nvidia. Nvidia Tesla V100 GPU Architecture. *White Paper*, page 53, 2017. URL <http://www.nvidia.com/content/gated-pdfs/Volta-Architecture-Whitepaper-v1.1.pdf>.
- M. Obergaulinger, M. A. Aloy, and E. Müller. Local simulations of the magnetized Kelvin-Helmholtz instability in neutron-star mergers. *Astronomy & Astrophysics*, 515(1):A30, 2010. doi:[10.1051/0004-6361/200913386](https://doi.org/10.1051/0004-6361/200913386).
- T. Oliphant. *A guide to NumPy*. CreateSpace Independent Publishing Platform, 2006.
- OpenMP Architecture Review Board. The OpenMP API specification for parallel programming. 2015. URL <http://www.openmp.org/wp-content/uploads/openmp-4.5.pdf>.
- S. A. Orszag and C. M. Tang. Small-Scale Structure of 2-Dimensional Magnetohydrodynamic Turbulence. *J. Fluid Mech.*, 90(1):129–143, 1979. doi:[10.1017/S002211207900210X](https://doi.org/10.1017/S002211207900210X).
- B. Paczynski. Gamma-ray bursters at cosmological distances. *The Astrophysical Journal*, 308(1):43–46, 1986. doi:[10.1086/184740](https://doi.org/10.1086/184740).
- C. Palenzuela. Modelling magnetized neutron stars using resistive magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 431(2):1853–1865, 2013. doi:[10.1093/mnras/stt311](https://doi.org/10.1093/mnras/stt311).
- C. Palenzuela, L. Lehner, O. Reula, and L. Rezzolla. Beyond ideal MHD: Towards a more realistic modelling of relativistic astrophysical plasmas. *Monthly Notices of the Royal Astronomical Society*, 394(4):1727–1740, 2009. doi:[10.1111/j.1365-2966.2009.14454.x](https://doi.org/10.1111/j.1365-2966.2009.14454.x).
- C. Palenzuela, T. Garrett, L. Lehner, and S. L. Liebling. Magnetospheres of black hole systems in force-free plasma. *Physical Review D*, 82(4):044045, 2010. doi:[10.1103/PhysRevD.82.044045](https://doi.org/10.1103/PhysRevD.82.044045).
- C. Palenzuela, S. L. Liebling, D. Neilsen, L. Lehner, O. L. Caballero, E. O’Connor, and M. Anderson. Effects of the microphysical equation of state in the mergers of magnetized neutron stars with neutrino cooling. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 92(4):044045, 2015. doi:[10.1103/PhysRevD.92.044045](https://doi.org/10.1103/PhysRevD.92.044045).
- L. Pareschi and G. Russo. Implicit-Explicit Runge-Kutta Schemes and Applications to Hyperbolic Systems with Relaxation. *Journal of Scientific Computing*, 25(1):129–155, 2005. doi:[10.1007/s10915-004-4636-4](https://doi.org/10.1007/s10915-004-4636-4).
- E. N. Parker. Sweet’s mechanism for merging magnetic fields in conducting fluids. *Journal of Geophysical Research*, 62(4):509–520, 1957. doi:[10.1029/JZ062i004p00509](https://doi.org/10.1029/JZ062i004p00509).
- A. Perego, H. Yasin, and A. Arcones. Neutrino pair annihilation above merger remnants: implications of a long-lived massive neutron star. *Journal of Physics G: Nuclear and Particle Physics*, 44(8):084007, 2017. doi:[10.1088/1361-6471/aa7bdc](https://doi.org/10.1088/1361-6471/aa7bdc).
- J. Pétri. General-relativistic force-free pulsar magnetospheres. *Monthly Notices of the Royal Astronomical Society*, 455(4):3779–3805, 2016. doi:[10.1093/mnras/stv2613](https://doi.org/10.1093/mnras/stv2613).

- M. Ponce, C. Palenzuela, L. Lehner, and S. L. Liebling. Interaction of misaligned magnetospheres in the coalescence of binary neutron stars. *Physical Review D*, 90(4):044007, 2014. doi:[10.1103/PhysRevD.90.044007](https://doi.org/10.1103/PhysRevD.90.044007).
- F. Pretorius. Evolution of Binary Black-Hole Spacetimes. *Physical Review Letters*, 95(12):121101, 2005. doi:[10.1103/PhysRevLett.95.121101](https://doi.org/10.1103/PhysRevLett.95.121101).
- D. J. Price and S. Rosswog. Producing ultrastrong magnetic fields in neutron star mergers. *Science*, 312(5774):719–722, 2006. doi:[10.1126/science.1125201](https://doi.org/10.1126/science.1125201).
- M. Punturo, M. Abernathy, F. Acernese, B. Allen, N. Andersson, K. Arun, F. Barone, B. Barr, M. Barsuglia, M. Beker, and N. Beveridge. The Einstein Telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity*, 27(19):194002, 2010. doi:[10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002).
- M. Punturo, M. Abernathy, F. Acernese, B. Allen, N. Andersson, K. Arun, B. Barr, M. Barsuglia, M. Beker, M. Punturo, M. Abernathy, F. Acernese, B. Allen, and N. Andersson. The third generation of gravitational wave observatories and their science reach. *Classical and Quantum Gravity*, 27(8):084007, 2011. doi:[10.1088/0264-9381/27/8/084007](https://doi.org/10.1088/0264-9381/27/8/084007).
- J. Qian. Generalization of the Kolmogorov $-5/3$ law of turbulence. *Physical Review E*, 50(1):611–613, 1994. doi:[10.1103/PhysRevE.50.611](https://doi.org/10.1103/PhysRevE.50.611).
- Q. Qian, C. Fendt, S. Noble, and M. Bugli. rHARM: Accretion and Ejection in Resistive GR-MHD. *The Astrophysical Journal*, 834(1), 2016. ISSN 1538-4357. doi:[10.3847/1538-4357/834/1/29](https://doi.org/10.3847/1538-4357/834/1/29). URL <http://arxiv.org/abs/1610.04445>{%}0Ahttp://dx.doi.org/10.3847/1538-4357/834/1/29.
- Q. Qian, C. Fendt, and C. Vourellis. Jet Launching in Resistive GR-MHD Black Hole-Accretion Disk Systems. *The Astrophysical Journal*, 859(1):28, 2018. doi:[10.3847/1538-4357/aabd36](https://doi.org/10.3847/1538-4357/aabd36).
- D. Radice. *Advanced Numerical Approaches in the Dynamics of Relativistic Flows*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, 2013.
- D. Radice. General-Relativistic Large-Eddy Simulations of Binary Neutron Star Mergers. *The Astrophysical Journal Letters*, 838(1):L2, 2017. doi:[10.3847/2041-8213/aa6483](https://doi.org/10.3847/2041-8213/aa6483).
- D. Radice, A. Perego, K. Hotokezaka, S. Bernuzzi, S. A. Fromm, and L. F. Roberts. Viscous-dynamical ejecta from binary neutron star mergers. *American Astronomical Society*, 869(2):L35, 2018a. doi:[10.3847/2041-8213/aaf053](https://doi.org/10.3847/2041-8213/aaf053).
- D. Radice, A. Perego, F. Zappa, and S. Bernuzzi. GW170817: Joint Constraint on the Neutron Star Equation of State from Multimessenger Observations. *The Astrophysical Journal Letters*, 852(2):L29, 2018b. doi:[10.3847/2041-8213/aaa402](https://doi.org/10.3847/2041-8213/aaa402).
- L. Rezzolla and O. Zanotti. Relativistic hydrodynamics. *Oxford University Press*, 2013.
- L. Rezzolla, B. Giacomazzo, L. Baiotti, J. Granot, C. Kouveliotou, and M. A. Aloy. The missing link: Merging neutron stars naturally produce jet-like structures and can power short-gamma ray bursts. *The Astrophysical Journal Letters*, 732(1):L6, 2011. doi:[10.1088/2041-8205/732/1/L6](https://doi.org/10.1088/2041-8205/732/1/L6).

- P. L. Roe. Approximate Riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics*, 43(2):357–372, 1981. doi:[10.1016/0021-9991\(81\)90128-5](https://doi.org/10.1016/0021-9991(81)90128-5).
- P. L. Roe. Characteristic-Based Schemes for the Euler Equations. *Annual Review of Fluid Mechanics*, 18(1):337–365, 1986. doi:[10.1146/annurev.fl.18.010186.002005](https://doi.org/10.1146/annurev.fl.18.010186.002005).
- S. Rosswog. Relativistic smooth particle hydrodynamics on a given background spacetime. *Classical and Quantum Gravity*, 27(11):114108, 2010. doi:[10.1088/0264-9381/27/11/114108](https://doi.org/10.1088/0264-9381/27/11/114108).
- S. Rosswog. SPH Methods in the Modelling of Compact Objects. *Living Reviews in Computational Astrophysics*, 1(1):1–109, 2014. doi:[10.1007/lrca-2015-1](https://doi.org/10.1007/lrca-2015-1).
- S. Rosswog and M. B. Davies. High-resolution calculations of merging neutron stars I. Model description and hydrodynamic evolution. *Monthly Notices of the Royal Astronomical Society*, 334(3):481–497, 2002. doi:[10.1046/j.1365-8711.2002.05409.x](https://doi.org/10.1046/j.1365-8711.2002.05409.x).
- S. Rosswog and M. Liebendörfer. High-resolution calculations of merging neutron stars II. Neutrino emission. *Monthly Notices of the Royal Astronomical Society*, 342(3):673–689, 2003. doi:[10.1046/j.1365-8711.2003.06579.x](https://doi.org/10.1046/j.1365-8711.2003.06579.x).
- S. Rosswog, M. Liebendörfer, F.-K. Thielemann, M. Davies, W. Benz, and T. Piran. Mass ejection in neutron star mergers. *Astronomy & Astrophysics*, 341(1):499–526, 1999. URL <https://ui.adsabs.harvard.edu/abs/1999A&A...341..499R>.
- S. Rosswog, E. Ramirez-Ruiz, W. R. Hix, and M. Dan. Simulating black hole white dwarf encounters. *Computer Physics Communications*, 179(1):184–189, 2008. doi:[10.1016/j.cpc.2008.01.031](https://doi.org/10.1016/j.cpc.2008.01.031).
- S. Rosswog, J. Sollerman, U. Feindt, A. Goobar, O. Korobkin, R. Wollaeger, C. Fremling, and M. M. Kasliwal. The first direct double neutron star merger detection: Implications for cosmic nucleosynthesis. *Astronomy & Astrophysics*, 615(1):A132, 2018. doi:[10.1051/0004-6361/201732117](https://doi.org/10.1051/0004-6361/201732117).
- M. Ruiz, R. N. Lang, V. Paschalidis, and S. L. Shapiro. Binary neutron star mergers: A jet engine for short gamma-ray bursts. *The Astrophysical Journal Letters*, 824(1):1–5, 2016. doi:[10.3847/2041-8205/824/1/L6](https://doi.org/10.3847/2041-8205/824/1/L6).
- M. Ruiz, A. Tsokaros, V. Paschalidis, and S. L. Shapiro. Effects of spin on magnetized binary neutron star mergers and jet launching. *Physical Review D*, 99(8):84032, 2019. doi:[10.1103/PhysRevD.99.084032](https://doi.org/10.1103/PhysRevD.99.084032).
- M. Ruiz, A. Tsokaros, and S. L. Shapiro. Magnetohydrodynamic simulations of binary neutron star mergers in general relativity: Effects of magnetic field orientation on jet launching. *Physical Review D*, 101(6):64042, 2020. doi:[10.1103/PhysRevD.101.064042](https://doi.org/10.1103/PhysRevD.101.064042).
- C. Runge. Über die numerische Auflöfung von Differentialgleichungen. *Mathematische Annalen*, 46(1):167–178, 1895. doi:[10.1016/j.compfluid.2018.07.001](https://doi.org/10.1016/j.compfluid.2018.07.001).
- S. Sangeetha and J. C. Mandal. A Cure for numerical shock instability in HLLC Riemann solver using antidiffusion control. pages 1–37, 2018. URL <http://arxiv.org/abs/1803.04954>.
- J. P. Sauppe and C. R. Sovinec. Two-fluid and finite Larmor radius effects on helicity evolution in a plasma pinch. *Physics of Plasmas*, 23(3):032303, 2016. doi:[10.1063/1.4942761](https://doi.org/10.1063/1.4942761).

- E. Schnetter, S. H. Hawley, and I. Hawke. Evolutions in 3D numerical relativity using fixed mesh refinement. *Classical and Quantum Gravity*, 21(6):1465–1488, 2004. doi:[10.1088/0264-9381/21/6/014](https://doi.org/10.1088/0264-9381/21/6/014).
- Y. Sekiguchi, K. Kiuchi, K. Kyutoku, M. Shibata, and K. Taniguchi. Dynamical mass ejection from the merger of asymmetric binary neutron stars: Radiation-hydrodynamics study in general relativity. *Physical Review D*, 93(12):124046, 2016. doi:[10.1103/PhysRevD.93.124046](https://doi.org/10.1103/PhysRevD.93.124046).
- M. Shibata. Fully general relativistic simulation of coalescing binary neutron stars: Preparatory tests. *Physical Review D*, 60(10):104052, 1999. doi:[10.1103/PhysRevD.60.104052](https://doi.org/10.1103/PhysRevD.60.104052).
- M. Shibata. Simulation of merging binary neutron stars in full general relativity: $\Gamma=2$ case. *Physical Review D*, 61(6):064001, 2000. doi:[10.1103/PhysRevD.61.064001](https://doi.org/10.1103/PhysRevD.61.064001).
- M. Shibata and T. Nakamura. Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Physical Review D*, 52(10):5428—5444, 1995. doi:[10.1103/PhysRevD.52.5428](https://doi.org/10.1103/PhysRevD.52.5428).
- H.-a. Shinkai and G. Yoneda. Hyperbolic formulations and numerical relativity: experiments using Ashtekar’s connection variables. *Classical and Quantum Gravity*, 17(23):4799—4822, 2000. doi:[10.1088/0264-9381/17/23/302](https://doi.org/10.1088/0264-9381/17/23/302).
- C.-w. Shu. Essentially Non-Oscillatory and Weighted Essentially Non-Oscillatory Schemes for Hyperbolic Conservation Laws Operated by Universities Space Research Association. In *ICASE Report*, volume 1697, pages 325—432. 1997. doi:[10.1007/BFb0096355](https://doi.org/10.1007/BFb0096355).
- D. M. Siegel and B. D. Metzger. Three-Dimensional General-Relativistic Magnetohydrodynamic Simulations of Remnant Accretion Disks from Neutron Star Mergers: Outflows and r -Process Nucleosynthesis. *Physical Review Letters*, 119(23):231102, 2017. doi:[10.1103/PhysRevLett.119.231102](https://doi.org/10.1103/PhysRevLett.119.231102).
- D. M. Siegel and B. D. Metzger. Three-dimensional GRMHD Simulations of Neutrino-cooled Accretion Disks from Neutron Star Mergers. *The Astrophysical Journal*, 858(1):52, 2018. doi:[10.3847/1538-4357/aabaec](https://doi.org/10.3847/1538-4357/aabaec).
- D. M. Siegel, R. Ciolfi, A. I. Harte, and L. Rezzolla. Magnetorotational instability in relativistic hypermassive neutron stars. *Physical Review D*, 87(12):121302, 2013. doi:[10.1103/PhysRevD.87.121302](https://doi.org/10.1103/PhysRevD.87.121302).
- J. Smagorinsky. General circulation experiments with the primitive equations. *Monthly Weather Review*, 91(3):99–164, 1963. doi:[10.1175/1520-0493\(1963\)091<0099:GCEWTP>2.3.CO;2](https://doi.org/10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2).
- A. Spitkovsky. Time-dependent force-free pulsar magnetospheres: axisymmetric and oblique rotators. *Astrophysical Journal*, 648(1):L51—L54, 2006. doi:[10.1086/507518](https://doi.org/10.1086/507518).
- H. C. Spruit. Origin of neutron star magnetic fields. *AIP Conference Proceedings*, 983(1):391–398, 2008. doi:[10.1063/1.2900262](https://doi.org/10.1063/1.2900262).
- D. Strigl, K. Kofler, and S. Podlipnig. Performance and scalability of GPU-based convolutional neural networks. *Proceedings of the 18th Euromicro Conference on Parallel, Distributed and Network-Based Processing, PDP 2010*, pages 317–324, 2010. doi:[10.1109/PDP.2010.43](https://doi.org/10.1109/PDP.2010.43).
- E. Süli and D. F. Mayers. *An Introduction to Numerical Analysis*. Cambridge University Press, 2003. doi:[10.1017/CBO9780511801181](https://doi.org/10.1017/CBO9780511801181).

- P. A. Sweet. The neutral point theory of solar flares. *Electromagnetic Phenomena in Cosmical Physics*, 6(1):123, 1958. URL <https://ui.adsabs.harvard.edu/abs/1958IAUS....6.123S>.
- N. R. Tanvir, A. J. Levan, A. S. Fruchter, J. Hjorth, R. A. Hounsell, K. Wiersema, and R. L. Tunnicliffe. A ‘kilonova’ associated with the short-duration γ -ray burst GRB 130603B. *Nature*, 500(7464):547–549, 2013. doi:[10.1038/nature12505](https://doi.org/10.1038/nature12505).
- M. Thierfelder, S. Bernuzzi, D. Hilditch, B. Bru, and L. Rezzolla. Trumpet solution from spherical gravitational collapse with puncture gauges. *Physical Review D*, 83(6):064022, 2011. doi:[10.1103/PhysRevD.83.064022](https://doi.org/10.1103/PhysRevD.83.064022).
- E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4(1):25–34, 1994. doi:[10.1007/BF01414629](https://doi.org/10.1007/BF01414629).
- G. Tóth. The $B=0$ Constraint in Shock-Capturing Magnetohydrodynamics Codes. *Journal of Computational Physics*, 121(2):1273–1293, 2000. doi:[10.1002/2015JA021997](https://doi.org/10.1002/2015JA021997).
- G. Tóth, X. Jia, S. Markidis, I. B. Peng, Y. Chen, L. K. S. Daldorff, V. M. Tenishev, D. Borovikov, J. D. Haiducek, T. I. Gombosi, A. Gloer, and J. C. Dorelli. Extended magnetohydrodynamics with embedded particle-in-cell simulation of Ganymede’s magnetosphere. *Journal of Geophysical Research: Space Physics*, 121:1273–1293, 2016. doi:[10.1002/2015JA021997](https://doi.org/10.1002/2015JA021997).Received.
- S. Valenti, J. Sand, S. Yang, E. Cappellaro, L. Tartaglia, A. Corsi, S. W. Jha, D. E. Reichart, J. Haislip, and V. Kouprianov. The Discovery of the Electromagnetic Counterpart of GW170817: Kilonova AT 2017gfo/DLT17ck. *The Astrophysical Journal Letters*, 848(2):L24, 2017. doi:[10.3847/2041-8213/aa8edf](https://doi.org/10.3847/2041-8213/aa8edf).
- D. Viganò, R. Aguilera-Miret, and C. Palenzuela. Extension of the sub-grid-scale gradient model for compressible magnetohydrodynamics turbulent instabilities. *Physics of Fluids*, 31(10):105102, 2019. doi:[10.1063/1.5121546](https://doi.org/10.1063/1.5121546).
- D. Viganò, R. Aguilera-Miret, F. Carrasco, B. Miñano, and C. Palenzuela. GRMHD large eddy simulations with gradient subgrid-scale model. *Phys. Rev. D*, 101(12):123019, 2020.
- P. Virtanen, R. E. Gommers, T. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. J. Bright, S. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, I. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, and P. van Mulbregt. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17(3):261—272, 2020. doi:[10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- J. Von Neumann and H. H. Goldstine. Numerical inverting of matrices of high order. *Bulletin of the American Mathematical Society*, 53(11):1021–1099, 1947. doi:[10.1090/S0002-9904-1947-08909-6](https://doi.org/10.1090/S0002-9904-1947-08909-6).
- J. Wallis. *A treatise of algebra, both historical and practical: shewing the original, progress, and advancement thereof, from time to time, and by what steps it hath attained to the heighth at which now it is; with some additional treatises*. R. Davis, 1685. URL <https://lib.ugent.be/catalog/rug01:001602641>.

- M. Waskom et al. Seaborn. *Zenodo*, 2014. doi:[doi:10.5281/zenodo.12710](https://doi.org/10.5281/zenodo.12710).
- I. Wolfram Research. Mathematica, 2018.
- H. C. Wong, U. H. Wong, X. Feng, and Z. Tang. Efficient magnetohydrodynamic simulations on graphics processing units with CUDA. *Computer Physics Communications*, 182(10):2132–2160, 2011. doi:[10.1016/j.cpc.2011.05.011](https://doi.org/10.1016/j.cpc.2011.05.011).
- A. Wright. METHOD: Initial Public Release. 2018. doi:[10.5281/zenodo.1404697](https://doi.org/10.5281/zenodo.1404697).
- A. J. Wright and I. Hawke. Resistive and Multi-fluid RMHD on Graphics Processing Units. *The Astrophysical Journal Supplement Series*, 240(1):8, 2019a. doi:[10.3847/1538-4365/aaf1b0](https://doi.org/10.3847/1538-4365/aaf1b0).
- A. J. Wright and I. Hawke. A resistive extension for ideal MHD. *Monthly Notices of the Royal Astronomical Society*, 491(4):5510–5523, 2019b. doi:[10.1093/mnras/stz2779](https://doi.org/10.1093/mnras/stz2779).
- U.-H. Zhang, H.-Y. Schive, and T. Chiueh. Magnetohydrodynamics with GAMER. *The Astrophysical Journal Supplement Series*, 236(2):50, 2018. doi:[10.3847/1538-4365/aac49e](https://doi.org/10.3847/1538-4365/aac49e).
- W. Zhang, A. MacFadyen, and P. Wang. Three-dimensional relativistic magnetohydrodynamic simulations of the kelvin-helmholtz instability: Magnetic field amplification by a turbulent dynamo. *Astrophysical Journal*, 692(1):L40—L44, 2009. doi:[10.1088/0004-637X/692/1/L40](https://doi.org/10.1088/0004-637X/692/1/L40).
- W. Zhang, A. Almgren, V. Beckner, J. Bell, C. Chan, M. Day, B. Friesen, K. Gott, D. Graves, M. P. Katz, A. Myers, T. Nguyen, M. Rosso, S. Williams, and M. Zingale. AMReX: a framework for block-structured adaptive mesh refinement. *The Journal of Open Source Software*, 4(37):1370, 2019. doi:[10.21105/joss.01370](https://doi.org/10.21105/joss.01370).
- X. Zhuge, J. M. Centrella, and S. L. Mcmillan. Gravitational radiation from coalescing binary neutron stars. *Physical Review D*, 50(10):6247—6261, 1994. doi:[10.1103/PhysRevD.50.6247](https://doi.org/10.1103/PhysRevD.50.6247).
- B. Zink. HORIZON: Accelerated General Relativistic Magnetohydrodynamics. *arXiv e-prints*, 2011. URL <http://arxiv.org/abs/1102.5202>.
- J. Zrake, A. I. Macfadyen, and S. Bat. Magnetic energy production by turbulence in binary neutron star mergers. *The Astrophysical Journal Letters*, 769(2):L29, 2013. doi:[10.1088/2041-8205/769/2/L29](https://doi.org/10.1088/2041-8205/769/2/L29).