

University of Southampton Research Repository
ePrints Soton

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

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON
FACULTY OF ENGINEERING, SCIENCE AND
MATHEMATICS

School of Mathematics

**Interfaces in Numerical Relativistic
Hydrodynamics**

by

Stephen Timothy Millmore

Thesis for the degree of Doctor of Philosophy

April 2010

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF MATHEMATICS

Doctor of Philosophy

INTERFACES IN NUMERICAL RELATIVISTIC HYDRODYNAMICS

by Stephen Timothy Millmore

This thesis investigates numerical techniques for modelling sharp interfaces between relativistic fluids. The motivation for this work lies in obtaining accurate models of neutron star interiors for use in multidimensional simulations in general relativity. The interior structure of a neutron star is believed to contain several regions, often separated by sharp transition layers. These layers are too thin to be explicitly incorporated in a numerical simulation of the entire star. We investigate how techniques can be developed to model these layers as sharp interfaces, across which the matter model can change, with the microphysical behaviour of the transition layer described through some appropriate boundary conditions.

The physical situations in which strong, detectable, gravitational waves are produced are, by their nature, violent events. As a result, we expect that large non-linear features, such as shock waves, will be formed. Therefore it is essential that the techniques developed to incorporate these sharp interfaces allow for their interaction with non-linear features in a stable manner numerically.

The techniques required for modelling sharp interfaces between two fluid components has not previously been considered in relativity. However, in Newtonian computational fluid dynamics, the boundary conditions required for stable, accurate behaviour across a sharp interface between two fluids, modelled using level set methods, have been developed. These techniques lend themselves naturally to an extension to the relativistic situations we wish to consider. In this thesis we start from the Ghost Fluid Method of Fedkiw et al. We first investigate whether it can be extended to simple relativistic situations, hence use special relativity in 1+1 dimensions. In order to use this method in neutron star simulations, however, full general relativity is required. We therefore extend these initial results to a spherically symmetric self-gravitating body in 1+1 dimensional general relativity. Finally, since gravitational wave production requires a fully asymmetric system, we show that our method extends to multidimensional relativistic situations. To this end, the final chapter presents results using 2+1 dimensional special relativistic simulations.

Contents

List of figures	vi
Acknowledgements	x
1 Introduction	1
1.1 Neutron stars and general relativity	2
1.2 Neutron star formation and structure	3
1.3 Gravitational waves from neutron stars	12
1.3.1 Gravitational wave detection	16
1.4 Modelling a neutron star	19
1.4.1 Current neutron star simulations	21
1.4.2 Modelling interfaces in a neutron star	23
1.5 Notation and structure	25
2 Theoretical basis	28
2.1 Newtonian fluid evolution equations	28
2.1.1 Spherically symmetric fluid evolution equations	33
2.2 Relativistic fluid evolution equations	34
2.2.1 3+1 decomposition of the Einstein equations	38
2.2.2 The special relativistic fluid evolution equations	43
2.2.3 The general relativistic fluid evolution equations in spherical symmetry	46
2.3 Equation of state	53
2.4 The Riemann problem	57
2.4.1 Wave types in the Newtonian Riemann problem	60
2.4.2 Solving the Newtonian Riemann problem	65
2.4.3 Wave types in the special relativistic Riemann problem .	68
2.4.4 Solving the special relativistic Riemann problem	71

2.4.5	The special relativistic Riemann problem with non-zero tangential velocities	72
2.5	Approximate Riemann solvers	75
2.5.1	Newtonian HLLE solver	76
2.5.2	Relativistic HLLE solver	77
2.5.3	Roe and Marquina solvers	78
2.6	The TOV equations	80
2.6.1	The TOV equations for multicomponent stars	82
2.7	Recovering the primitive variables	84
3	Numerical methods	87
3.1	Method of lines	91
3.2	Runge-Kutta methods	92
3.3	The Lax-Wendroff Theorem and convergence	94
3.3.1	Measuring convergence	98
3.4	Reconstruction	102
3.4.1	Minmod slope limiter	103
3.4.2	MC limiter	105
3.4.3	PPM	106
3.4.4	Alternative reconstruction methods	107
3.5	Boundary conditions	107
3.5.1	Exact boundary conditions	109
3.5.2	Periodic boundary conditions	110
3.5.3	Solid boundary conditions	110
3.5.4	Far-field boundary conditions	111
3.5.5	Vacuum boundary conditions	112
3.6	Source terms	113
3.7	Multicomponent fluids	113
3.7.1	Level set methods	117
3.7.2	The level set function	119
3.7.3	Numerical methods for the level set function	122
3.7.4	Ghost Fluid Method	124
3.7.5	Alternative methods	130
3.8	Atmosphere treatment	131

4 Results	133
4.1 Preliminary results for Newtonian hydrodynamics	133
4.1.1 Motivation for use of the Ghost Fluid Method	134
4.1.2 The Newtonian Ghost Fluid Method	138
4.2 The Ghost Fluid Method in special relativity	145
4.2.1 Comparison of reconstruction methods for the Ghost Fluid Method in special relativity	167
4.3 The Ghost Fluid Method in general relativity	170
4.4 Summary	188
5 2+1 dimensional Ghost Fluid Method	190
5.1 Numerical techniques for solving multidimensional equations . .	190
5.2 Multicomponent fluids in more than one dimension	192
5.2.1 Level set methods	192
5.2.2 Ghost Fluid Method	192
5.2.3 Relativistic effects in the Ghost Fluid Method	194
5.3 Higher order methods for the level set evolution	194
5.4 Results for the Newtonian multidimensional Ghost Fluid Method	197
5.5 Results for the special relativistic multidimensional Ghost Fluid Method	198
5.5.1 Comparing the velocity extrapolation techniques	198
5.5.2 Bubble tests for the Ghost Fluid Method in special relativity	202
5.5.3 Summary	208
6 Conclusions	212
6.1 Future work	215
A Calculations for relativistic equations	217
A.1 Derivation of the 3+1 decomposition of the Einstein equations .	217
A.1.1 Projection operators	217
A.1.2 Extrinsic curvature tensor	221
A.1.3 Preliminary results for the 3+1 decomposition of the Einstein equations	224
A.1.4 The 3+1 split of the Einstein equations	231
A.1.5 The Lorentz factor	236

A.2	Derivation of the spacetime and fluid evolution equations in spherical symmetry	238
A.2.1	Spacetime quantities	238
A.2.2	Hamiltonian constraint	239
A.2.3	Momentum constraint	240
A.2.4	Slicing condition	242
A.2.5	Conservation of rest mass density	244
A.2.6	Conservation of total energy density	246
A.2.7	Conservation of momentum	248
A.3	Derivation of the TOV equations	249
B	Calculations for numerical techniques	252
B.1	PPM	252

List of Figures

1.1	Maximum mass of a neutron star	4
1.2	Neutron star interior	9
1.3	Inner core possibilities	10
2.1	The 3+1 split	41
2.2	Possible form of the Riemann problem	59
2.3	Rarefaction fan	60
2.4	Rarefaction shock	61
2.5	Shock wave characteristics	63
3.1	Conservative vs. non-conservative methods	96
3.2	Minmod slope limiter	104
3.3	MC limiter	105
3.4	Issues at the boundaries	108
3.5	The level set function	118
3.6	Motivation for the Ghost Fluid boundary conditions	126
3.7	The Ghost Fluid Method	129
4.1	The Newtonian Sod test	135
4.2	The failure of the colour function	137
4.3	The Newtonian Sod test with an artificial interface	139
4.4	A shock wave hitting an interface in Newtonian physics	141
4.5	A shock wave hitting a slab of material in Newtonian physics	143
4.6	A low resolution version of a shock wave hitting a slab of material in Newtonian physics	144
4.7	A Sod-type test in special relativity	146
4.8	A moving contact discontinuity in special relativity	148

4.9	A zoom of the moving contact discontinuity test	149
4.10	Two rarefaction waves in special relativity	151
4.11	A zoom of the contact discontinuity in the two rarefaction test .	152
4.12	Two shock waves in special relativity	153
4.13	A zoom of the contact discontinuity in the two shock case . .	154
4.14	A zoom of the shock in the two shock case	155
4.15	A shock hitting a contact discontinuity in special relativity . .	157
4.16	A zoom of contact discontinuity after it has been hit by a shock	158
4.17	A zoom of various features after a shock has hit a contact discontinuity	159
4.18	A highly relativistic shock hitting a contact discontinuity in special relativity	161
4.19	A shock hitting a slab of material in special relativity	164
4.20	A zoom of a slab hit by a shock	165
4.21	A low resolution version of a shock hitting a slab of material in special relativity	166
4.22	An advected sine wave in special relativity: 200 cells	168
4.23	An advected sine wave in special relativity: 800 cells	169
4.24	A shock hitting a perturbed density profile in special relativity: 200 cells	171
4.25	A shock hitting a perturbed density profile in special relativity: 800 cells	172
4.26	A stable single component star in general relativity	173
4.27	Convergence of a stable single component star	175
4.28	A stable single component star with an artificial interface in general relativity	176
4.29	Convergence of a stable single component star with an artificial interface	177
4.30	Further examination of a single component star with an artificial interface	178
4.31	A two component star with $\gamma_{\text{outer}} = 1.9$ in general relativity . .	180
4.32	Convergence for a two component star with $\gamma_{\text{outer}} = 1.9$	181
4.33	A two component star with $\gamma_{\text{outer}} = \frac{5}{3}$ in general relativity . .	182
4.34	Convergence for a two component star with $\gamma_{\text{outer}} = \frac{5}{3}$	183
4.35	A perturbed two component star in general relativity	185

4.36	Convergence of a perturbed two component star	186
4.37	Velocity and interface movement in a perturbed two component star	187
5.1	Shock hitting a bubble in Newtonian physics	199
5.2	Shock hitting a slab in 2+1 dimensional special relativity	201
5.3	Shock hitting a bubble in special relativity using MC limiter	204
5.4	Shock hitting a bubble in special relativity using PPM	205
5.5	Schlieren comparison of shock-bubble interaction	206
5.6	Moderate shock hitting a bubble in special relativity	207
5.7	Large shock hitting a bubble in special relativity	209
5.8	Labelling features from a shock hitting a bubble	210

DECLARATION OF AUTHORSHIP

I, Stephen Timothy Millmore, declare that the thesis entitled “Interfaces in Numerical Relativistic Hydrodynamics” and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- 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;
- where I have consulted the published work of others, this is always clearly attributed;
- 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;
- I have acknowledged all main sources of help;
- 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;
- parts of this work have been published as:
Millmore S. T., Hawke I., “Numerical simulations of interfaces in relativistic hydrodynamics”, 2010, *Classical and Quantum Gravity*, 27, 015007.

Signed:

Date:

Acknowledgements

I would first like to thank my supervisor Ian Hawke for all his help and guidance that lead to the production of this thesis. I have enjoyed working with him immensely, and have learnt to appreciate the intricate workings of numerical methods.

I am incredibly grateful for the support provided throughout my education by my parents, Mike and Alice Millmore, and to my sister, Jennifer, for teaching (forcing me to learn) how to argue my point.

My time working in Southampton has been enjoyable, and I would like to thank those in the general relativity group for their support and company. Notably I thank Sam Lander for coping with my deluded ramblings for three and a half years, as well as being a continuous source of interesting music.

Throughout the time spent getting code to produce the correct results, and the final stages of this thesis, I thank Katie Ng for the love, support and cakes.

The task of proof reading my thesis cannot have been an easy one, and I therefore owe gratitude to those who have dedicated their time to this: Leonie Hicks, Mike Hogg and, again, Ian.

I feel I should offer a special mention to my housemates, who have caused me (almost) no problems, and have kept the house in such a state that turning up to work in the mornings was a much more enjoyable experience than it might otherwise have been.

Finally I would like to thank everyone who has contributed to my collection of penguins, both in the office and at home.

Chapter 1

Introduction

Astronomical objects have long been identified as valuable test-beds for many areas of physics since they are subject to conditions that cannot be replicated in a laboratory situation. Neutron stars (NSs) are a classic example of this, with a typical mass of $\sim 1.4M_{\odot}$ and radius of $\sim 10\text{km}$, they can achieve a central density in excess of nuclear density ($2 \times 10^{14}\text{g cm}^{-3}$) making them one of the densest objects known [119]. They provide the means to allow us to probe the extremes of compactness, magnetic field strength and rotation and provide an environment which may lead to the creation of exotic particles. It is useful to consider just how extreme these conditions are in a quantifiable manner. We quantify compactness by determining just how ‘close’ a NS is to being a black hole (BH) (i.e. how close it is to the Schwarzschild radius). Specifically if we take $2GM/(Rc^2)$ where G is the gravitational constant, M the total mass, R radius and c the speed of light, then we find the compactness of a typical NS is $\sim 10^{-1}$ of that for a BH. The magnetic field of a NS can exceed 10^{15}G [60], a factor of $\sim 10^{15}$ times greater than that of the Earth. The fastest observed rotation rate of a NS is $\sim 1\text{ms}$ [120], hence the surface of this star will be rotating $\sim 0.1c$. The densities within a NS mean that matter cannot remain bound within nuclei. It is possible that at these densities we may find strange matter (baryons which contain at least one strange quark), which would be unstable in terrestrial situations [209]. It is perhaps not surprising that with behaviour so far removed from that found on Earth, to model a NS accurately we would expect to use general relativity (GR), magnetohydrodynamics and nuclear physics. One result of these conditions, in particular the

extreme densities within a NS, is that this makes these objects perhaps the most promising source for detectable gravitational wave (GW) emission. Only by carefully matching observational data with accurate theoretical models can we hope to probe the physics of these objects.

The first theoretical predictions of the existence of NSs were made by Baade and Zwicky in 1934 [21], just two years after the discovery of the neutron [42]. The actual discovery of NSs came serendipitously in 1967 when a pulsating radio source was discovered by Bell and Hewish [96]. Shortly after this discovery, the source was identified as a NS [77]. Since then many of these rapidly pulsating radio sources (*pulsars*) have been identified and represent one of the most commonly observed forms of NS. Several more classes of NS have also been identified, including those with x-ray and gamma-ray emissions in addition to further radio sources, see e.g. [78, 215].

1.1 Neutron stars and general relativity

Einstein's theory of general relativity, first published in 1915, is a geometric theory of gravitation [62]. The theory states that matter and energy distort the geometry of space and time (spacetime), and that the geometry of spacetime in turn affects how matter and energy behave. In a weak gravitational field, the differences between GR and the Newtonian theory of gravity are small. In a strong field, Newtonian gravity can, however, no longer offer a reasonable description of how matter moves. It is therefore unsurprising that NSs, which fall into this strong gravitational regime, require GR to accurately describe their behaviour.

Observational evidence also suggests that an accurate model of a NS must use GR. From the observations, information about the mass range of NSs can be inferred. Current estimates of the maximum mass are in the range $1.5 M_{\odot}$ to $3 M_{\odot}$. Observational difficulties and model-based dependence of the interpretation of these observations may lead to the uncertainties in the maximum mass estimates, but for individual models, there is always a clear maximum [87]. In Newtonian physics, there is no limiting constraint that can explain this maximum mass. In GR, by contrast we do have a limit above which a NS becomes unstable, at which point there is no known mechanism

which can prevent collapse to a BH. The actual value of this maximum is dependent on the model used to describe the interior of the NS, its rotation and its magnetic field. For all realistic models in GR, we find the maximum mass is consistent with that provided by the observations [86, 136, 177, 210]. The maximum masses of a variety of models are shown in figure 1.1.

1.2 Neutron star formation and structure

In order to model NSs we need to understand both how they are formed and their composition. Even in the first predictions of NSs, Baade and Zwicky suggested that NSs could be formed during a supernova (SN) [21]. Since this initial prediction, observational evidence has been identified in support of this. The Crab pulsar was discovered in the SN remnant SN1054, and the SN itself was observed and documented at the time [164, 165]. Similarly the Vela pulsar is also associated with a SN remnant [117]. It is now apparent that most NSs are formed from type I Ib SN, the collapse of the core of a massive star. Progenitor stars for NS production are thought to be those with masses in the range $8M_{\odot} \lesssim M \lesssim 100M_{\odot}$ [113, 119].

Core collapse SNe are just one of the possible outcomes that arise at the end of a star’s ‘life’. The majority of the life of a star consists of the fusion of hydrogen into helium within its core. This provides radiative pressure to prevent gravitational collapse, sustaining the star for $\sim 10^6 - 10^8$ years for a NS progenitor, or $\sim 10^{10}$ years for a solar mass star [107]. Once approximately 10% of the hydrogen in the core of the star has been used, the pressure provided by this fusion is no longer sufficient to support the star. At this point the core of the star starts to contract, increasing pressure and temperature in this region. If the star is sufficiently massive, this can initiate the fusion of helium in addition to the remaining hydrogen, and preventing further contraction of the core. The quantity of helium is significantly less than that of hydrogen, hence this helium fusion cannot be sustained nearly as long $\sim 10^5 - 10^6$ years [107]. If the star is massive enough, fusion of heavier elements may continue until iron is produced. For elements lighter than iron, the fusion of two nuclei generates energy, but for fusion to continue beyond this point energy must be added to the system, therefore radiative pressure no longer has the capability to halt




Figure 1.1: Illustrating the maximum masses of a NS based on the choice of model. The central density (ρ_c) is plotted against total gravitational mass (M). It is clear that for each model a maximum is reached (above which increasing the central density appears to reduce the total mass, suggesting the NS is unstable). The shaded region represents the range of known (accurately measured) NS masses. Source: [86].

matter falling towards the core [177]. In order for a star to generate iron in the core (and hence later undergo a type IIb SN) it must be massive enough for the contraction of the core to provide temperature and pressure required for the fusion of all nuclei up until this point. It is this that provides the $\sim 8M_{\odot}$ minimum mass constraint for NS formation.

The unavoidable gravitational collapse that occurs at this point that triggers the SN explosion. Once this collapse has started, the pressure (and temperature) within the core continue to rise as a result of the infalling matter. This continues until the density of the core is sufficient for the Pauli exclusion principle to apply. The general case of this principle states that no two fermions (quarks or leptons) may share the same quantum state. Consequently the fermions in sufficiently dense regions effectively cannot be too close together and, in the core, they experience a pressure known as *degeneracy pressure*. This pressure prevents further collapse of the material at the centre of the core and effectively provides a solid boundary against which the remaining infalling matter will impact [103]. In these conditions subatomic interactions favour isolated neutron production, hence a neutron-rich core is formed, referred to as a proto NS [103].

This impact of the infalling matter onto the degenerate core is known as *core bounce*. This sends a shock wave propagating outwards through the infalling matter. The remaining infalling matter is then prevented from reaching the proto NS as it hits this outgoing shock wave. The lack of matter reaching the proto NS itself means that aside from some accretion of the post-shocked matter, it does not grow significantly. It then can cool rapidly, in ~ 100 s, to form a (cold) NS ($< 10^9$ K) [150, 151]. It would seem that this outgoing shock can be linked to the explosion seen during a SN. Current theory, however, suggests that the energy within the shock is rapidly lost through photodissociation of atomic nuclei into nucleons, and through neutrino losses in the post shock region, causing the shock to stall. Since type IIb SN are well documented, there is clearly some mechanism through which the shock wave is revived. The causes of this are not yet completely understood. For a review the overall SN procedure see e.g. [103] and the possible mechanisms by which the shock is revived are found in e.g. [151].

For a sufficiently massive progenitor star, the accretion onto the proto NS can reach a point where even neutron degeneracy pressure is no longer suf-

ficient to prevent further collapse. At this point it is believed that a BH is formed. This can either occur as a result of mass accreted onto the proto NS, or, in the most extreme cases, directly as a result of the collapse. In these cases a SN is still observed, since the infalling matter will still bounce off the core when it has become degenerate, but before the BH is formed. We expect this to happen for stars with masses above ~ 25 to $\sim 100M_{\odot}$, with the composition of the progenitor star in this mass range playing an important role in determining whether a BH is formed [111, 177].

In this thesis we focus on NSs in the absence of e.g. further accretion from a partner star in a binary system, and hence can be considered to have reached some equilibrium configuration. In order to understand the behaviour of these objects we must consider the internal structure and common features of such NSs. The precise composition of a NS is then dependent on both the conditions in which it was formed, e.g. the mass and the composition of the progenitor star and the model assumed for the interior matter. Typically we expect that NSs in some equilibrium configuration will have several distinct regions: the inner core, outer core, crust, envelope (or ocean) and the atmosphere [87, 119].

The outer layers

The two outer layers - the atmosphere and the envelope - are very thin, with depths of $\sim 1\text{cm}$ and $\sim 10 - 100\text{m}$ respectively. The atmosphere of a NS can be directly observed through x-ray emission. It is expected that the atmosphere is composed of the lightest element available (typically hydrogen), which can be accreted from the interstellar medium or from the SN remnant. Heavier elements then ‘sink’ into the deeper layers in the envelope [8]. This lightest element then forms either a true (gaseous) atmosphere, or a solid or liquid surface, dependent on the temperature of the NS [218]. It is this region which dictates the emission spectrum from the actual NS, although there is a further contribution from the magnetic field. As a result, the atmosphere provides much of the information about the NS. The temperature of the surface is directly calculated through this emission, and further information can be inferred e.g. mass and radius [87]. The internal temperature profile of the

NS can also be inferred from these observations if we understand the processes through which a NS cools and either know, or can infer, its age.

In the envelope, the density is sufficient that electrons can no longer remain bound to nuclei. Here a fluid region is formed (hence the alternative name, the ocean) consisting of these nuclei and electrons. The envelope then plays a major role in determining how a NS cools [84].

The crust

The crust of a neutron star has a depth of $\sim 1\text{km}$, and is often split into subregions, namely the inner and outer crust. A common feature of these subregions is the presence of neutron rich nuclei bound in a crystal lattice structure [29]. In the outer crust it is this lattice which, along with free electrons, forms the primary component of the crust. At a depth of $\sim 300\text{m}$, and a density $\sim 10^{11}\text{g cm}^{-3}$, we move to the inner crust [87]. At these densities, we approach *neutron drip density*, where neutrons start to dissociate from nuclei [32]. The lattice structure is then permeated by a neutron fluid, which for a typical cold NS (with temperature $\lesssim 10^9\text{K}$), can form a superfluid [173]. We believe that the solid lattice structure of the crust results in it displaying elastic properties [41, 106].

At the base of the crust, at densities of $\sim 10^{14}\text{g cm}^{-3}$, neutrons and protons can no longer be contained within nuclei. A transition to a uniform neutron fluid then occurs. Some models include an extended transition region through which the nuclei move through a variety of structures, becoming increasingly elongated with greater density, known as the *pasta phase* [153, 156]. The width of this phase is then equally model-dependent, but could be present in as much as a third of the crust.

The inner and outer core

The two core regions of a NS make up the majority of the volume, comprising all but the last kilometre of the star, and containing up to 99% of the mass. The outer core consists primarily of a neutron fluid, with protons, electrons and possibly muons, making up a few percent of the matter [87]. The presence of this neutron fluid means that, as with the crust, it is believed this region

has superfluid properties, see e.g. [17, 174]. The presence of the electrons in this region can also form a superconducting fluid in the outermost regions of the core [163].

At densities of $\sim 5 \times 10^{14} \text{ g cm}^{-3}$ our intuition as to the nature of matter breaks down. In this region, termed the inner core, there is great uncertainty about the composition of a NS. At these densities it is possible for particles to exist that would usually be unstable in free space, or within a nucleus. We expect that even in the outer core the density may be sufficient for muons to exist, and there are many further models that have been proposed for the inner core, see e.g. [30, 209]. One of the most favoured possibilities for exotic matter is the existence of hyperons. These are hadrons which contain at least one strange quark (instead of being composed entirely of up and down quarks). These particles (denoted Λ , Σ and Ξ in order of ascending mass) would lead to a more compact NS if they are present, and hence accurate measurements of mass and radius, as well as GW information (see section 1.3) could lead to evidence for such a configuration. Other possibilities include a meson core (particles composed of two quarks), either pions (π^-) or kaons (K^-), or deconfined quarks. We currently believe that such meson cores (especially pions) are unlikely [146]. Finally, it is possible that the presence of quark matter in the inner core could then trigger a phase transition throughout much of the star, causing the entire core to be predominantly composed of such matter [120].

The internal regions of a NS, along with some of the structure within these regions, are illustrated in figure 1.2. The various transition structures of the pasta phase are illustrated, showing the transition from the crystal lattice of nuclei to the homogeneous neutron fluid. Also highlighted are vortices within the inner crust and outer core, expected to exist in the neutron superfluid [17, 186]. Figure 1.3 shows some of the theorised models for the inner core of a NS. Various models for a quark core star and a strange star (for which a phase transition has resulted in quark matter throughout the star) are listed.




Figure 1.2: Illustrating the possible regions present within a NS. Some additional features of the NS are included. The crust lattice permeated by superfluid (the lower left inset), the pasta phase transition between the crust (the upper inset) and the core and the superfluid/superconducting outer core region are also shown (the lower right inset). Source: <http://www.astroscu.unam.mx/neutrones/dany.html>




Figure 1.3: A figure detailing some of the models for the inner core of a NS. Source: <http://www-rohan.sdsu.edu/~fweber>.

External features

In addition to information about the interior of NSs, observations reveal much about the general external behaviour of these objects. From the onset of NS observations it was apparent that these objects rotate very rapidly, with the first measured pulsar having a period of 1.337s [96]. Since these initial measurements, rotational periods from ~ 10 s to ~ 1 ms have been observed [120]. For all but the most rapidly rotating NSs, this rotation rate arises directly from the conservation of angular momentum in the contracting core of the progenitor star during SN. This is not the case for the rapidly rotating *millisecond pulsars*. It is expected that they are in some accreting binary system, and this results in *spinning up* of the pulsar, see e.g. [213].

It has also been observed that NSs have very strong magnetic fields, with surface fields of $\sim 10^{12}$ G. These fields can be inferred from measurements of the

rotational period, and the rate of decrease of the period (a rotating magnetic dipole has been shown experimentally to be subject to loss of energy). It should be noted that the estimates for magnetic fields are subject to large (orders of magnitude) uncertainties based on poorly known quantities, such as the age of the NS. It is this magnetic field that leads to the electromagnetically ‘beamed’ emission that we detect as pulsars, through acceleration of protons and electrons on the NSs surface [108]. A class of NSs with even stronger magnetic fields, up to $\sim 10^{15}$ G, termed *magnetars*, have also been identified [60, 152]. The magnetic field of NSs are still far from fully understood with many open questions, such as whether it penetrates the core. Therefore the theory of these fields, and observations of phenomena associated with magnetic fields, are very active fields of research.

A complete model

It is clear that an accurate model of a NS will need to take many (or all) of the features discussed in this section into account. Practically this is not a simple task since neither the full behaviour, nor the composition, of these regions is completely understood. We can describe the regions within a NS with some accuracy through an *equation of state* (EOS) which effectively describes the microphysical aspects of the matter being modelled in a macrophysical manner. This would then allow us to investigate behaviour of the entire star. We would expect, therefore, that this EOS (in some manner) represents the various regions at the appropriate densities. The choice of EOS that we use in this thesis is discussed in section 2.3. Ideally we need to include rotation and magnetic fields within the model too. Qualitatively many of the properties of a non-rotating NS are also present for rotating models (e.g. maximum mass and internal structure). We therefore treat NSs in this thesis as non-rotating (and hence spherically-symmetric) objects. It has also been shown that, with the exception of very strong magnetic fields, there are many cases in which they do not have a large effect on NS behaviour (see e.g. [72]), therefore we will model NSs without magnetic fields in this thesis.

1.3 Gravitational waves from neutron stars

One of the key predictions of GR is that asymmetric bodies will radiate gravitationally. These GWs are exceptionally weak. The effect of a wave is expected to have a *maximum* effect of approximately 1 part in 10^{21} . As a result no direct GW detection has yet been made. The greatest chances of detectable GWs arise from situations in which the largest GWs will be produced, i.e. in the strongest gravitational fields. It is not surprising that NSs fall into this category and the (comparatively) common nature of these stars, compared to the even more compact BHs, makes them very promising sources for the first direct GW detection. Any scenario which leads to an asymmetric structure, and hence a mass quadrupole in a NS, or in the external structure of a system, will release (potentially detectable) GWs. Here we consider some of the key GW release mechanisms, more details can be found in e.g. [15].

Mountains

The rapid rotation of NSs discussed in the previous section results in these objects being close to perfectly axisymmetric. If a rotating NS has even a slight asymmetric deformation, however, then this will result in the release of GWs. These deformations are typically referred to as *mountains*, though the size of NSs means that for a mountain to have a noticeable GW signature, it may need to be only a few millimetres high. Such mountains can be formed through strains within the crust, as a result of accretion, or even due to effects of the magnetic field. The magnitude of the GWs will obviously be dependent on the size of these mountains, and this in turn is dependent on the composition of the NS. For example, a mountain on the crust is dependent on the maximum strain the crust can sustain before it ‘breaks’ and the star returns to a more axisymmetric configuration. For these deformations it has been suggested that the maximum supported mountain will result in the NS having an ellipticity of $\sim 10^{-5}$ [93, 97]. Although the NS can support this maximum, there is still an issue as to whether mechanisms that lead to the formation of mountains could achieve ellipticities this large. It may be that accretion, which would lead to obvious deformations, is the best candidate for this, but the complex nature of accretion means there is currently a lack of satisfactory models quantifying these effects [15].

Neutron star instabilities

Any oscillation of a NS interior which causes the structure to deviate from axisymmetry will lead to some GW production. If there is some unstable behaviour which excites these oscillations significantly then there is a possibility that the GW signal will be significantly large that we have a chance of detecting it. NSs can experience a wide range of instabilities, many of which are associated with unstable modes of oscillation: a comprehensive description of these instabilities is presented in [13]. Modelling these instabilities requires an accurate description of the interior structure of a NS. Currently there is a number of promising mechanisms for producing GWs, but further work is required to quantify the expected wavesignal [15].

Detection of GWs from these oscillations will allow insight into the interior structure of NSs. Since the structure of a NS will provide constraints on these oscillations, or when instabilities become apparent, this *GW asteroseismology* will allow for the testing of various models [16, 30]. In order for the NS to actually emit any GWs, some behaviour which excites these instabilities must occur. Various methods by which this can happen are proposed based on observational data, with one prominent example being NS glitches. Observations of the rotational rate of pulsars show an expected slow decrease in angular frequency due to loss of angular momentum through e.g. magnetic torque and maybe GWs. Some observations have also shown that occasionally the frequency suddenly jumps; this phenomena is generally referred to as a *glitch* [112, 132]. It is thought that glitches occur when the angular momentum from vortices within the superfluid region of a NS, shown in figure 1.2, is transferred to the entire star. Mechanisms through which this transfer occurs are described in e.g. [124, 170, 208].

Precession

If the rotational axis of a NS is not aligned with the axis of symmetry (assuming an otherwise axisymmetric NS), then this can be seen as a ‘wobble’ known as *free precession*. This occurs in order for the star to conserve angular momentum [105]. The energy of this precession will be dissipated within the star as thermal energy, with some of the damping also manifesting itself as GWs [105, 219]. Evidence for precessing NSs may exist, [189], although the

number of precessing stars is not as high as may be expected. It has been suggested that this could be due to the interior behaviour of a NS as detailed in e.g. [74, 75]. The actual mechanisms which will lead to precession are also not fully understood. It is however believed that behaviour such as glitches may be one such method by which precession can be induced [105].

Core collapse supernova

The formation of a NS through a core collapse SN is, by its nature, a violent event involving matter at very high densities, and therefore a promising source of GW emission. There are many stages of a SN that can lead to the asymmetric behaviour required for GW emission. One such situation is the revival of the stalled shock. It is expected that this will be a highly non-linear event, since it is believed (and observed in computational simulations) that the stalled shock experiences an instability, the *standing accretion shock instability* (SASI) [151]. In addition to this, GWs may be produced from asymmetries in rotational instabilities during collapse, core bounce, convection behind the shock or from pulsations in the proto NS [111].

In order to gain a more complete insight into the processes that occur during SN, and understand how e.g. the stalled shock is revived, it would be desirable to obtain observational data from within the collapsing core. Although observations of the progenitor star and the SN remnant yield some useful information, the opacity of the matter surrounding the SN means that electromagnetic observations are ruled out as a means of probing the core behaviour. GWs and neutrinos can, however, pass through matter with almost no interaction, and therefore they could reveal information about the conditions within the core of a star during SN. In a SN as much as 99% of the energy is carried away through neutrino losses, with these neutrinos actually observed in SN1987a [33, 203]. At this time, the detections were not significant enough to determine the behaviour in the core of this SN.

The signature of detected GWs and neutrinos is dependent on the model used to model the behaviour that occurs during SN [55]. As an example, there are currently three favoured models which could explain the revival of the stalled shock. Firstly it is possible that there is neutrino heating of the shock [103]. Alternatively magnetic fields could transfer rotational energy

from the proto NS to the stalled shock [144]. Finally an acoustic method has been proposed where shocked sound waves from g -mode oscillations in the core transfer energy to the stalled shock [36].

Binary inspiral

When any two bodies are in orbit there is a very small fraction of the angular momentum of the system lost continuously through GWs. In the case of compact objects (white dwarfs, NSs and BHs), this loss can be sufficient for these objects to inspiral and eventually merge on a timescale such that we could expect to observe these events. It is believed that short gamma-ray bursts are such observations of (distant) compact binary mergers [139].

In the case of white dwarf mergers, it is the GW signal from the final stages of inspiral that is most prominent [100]. For binaries involving BHs or NSs, we would, however, not only expect strong GW emission from the late stages of the inspiral but also from the merger itself and possibly from the post-merger behaviour, the *ringdown*. Compared to other sources of GWs, the emission mechanism for compact binaries is well understood, since there is significantly less dependence on the internal composition of the objects. It is therefore possible, for the inspiral phase in particular, for accurate waveforms to be obtained.

For binary BH systems, the full waveforms for the merger and ringdown phases can also be calculated. The violent non-linear behaviour of the system at merger makes computer simulations a necessity. Through these simulations it is possible to obtain exact waveforms to very high accuracy. These simulations, pioneered in [160], can now cover a wide parameter space for comparable mass (a mass ratio $M_1/M_2 \leq 10$) BH binaries [79, 88]. Alternatively, at the centres of galaxies, *extreme mass ratio inspirals* exist, in which a small stellar BH is inspiralling into a supermassive galactic BH. In such cases we typically have a mass ratio of $M_1/M_2 \geq 10^5$. Here GW calculations rely on a perturbative approach, with a review given in [27]. For mergers that fall within these two mass ranges, perturbative methods are generally unable to provide the necessary accuracy, whilst there is insufficient computational power to provide simulations in full GR, see e.g. [11].

For binaries involving NSs (either BHNS or NSNS) the waveforms are de-

pendent on the parameters of the NSs (e.g. mass and composition) [24]. Exact calculations are therefore not feasible, but once a model for the NSs has been selected, computer simulations can still be undertaken. For the BHNS case, it is not surprising that many of the features of the merger are found to be similar to those in the BHBH case. The merger results in almost all of the NS matter accreted onto the BH, with maybe a small fraction remaining in some disk structure around the BH [64, 183]. In the NSNS case the final outcome of the merger is again believed to be a BH in almost all cases. In this case, however, the collapse of the merged object to this BH is not necessarily immediate. For some NSs, a massive (unstable) merged object can exist temporarily before a delayed collapse occurs [23]. The current status of the field of NSNS inspiral simulations is detailed in section 1.4.1, where we consider effects of different parameters in the NS models.

1.3.1 Gravitational wave detection

The inherently weak nature of GWs means that they currently have not yet been directly detected. In contrast, indirect evidence for GWs has been observed. As detailed above, compact binaries inspiral due to angular momentum losses through GWs. This decreases the orbital period of the binary system, an effect which can be measured. The first such measurement came from the binary pulsar PSR1913+16 (the Hulse-Taylor pulsar), showing an excellent agreement with the predictions of GR [197]. This has since been augmented with measurements of the inspiral in other binary systems.

Despite this indirect evidence for GWs, direct detection is desirable for a full confirmation of the predictions of GR. In the 1960s with Weber made the first efforts to observe GWs using the first bar detectors [211]. Despite early optimism ([212]), these bar detectors could not reach sufficient sensitivity for any reliable GW detection. Instead it is now possible to exceed the sensitivity of bar detectors by many orders of magnitude using laser interferometers, and as such, almost all current (and future) efforts to detect GWs are concentrated on such interferometers.

Attempts to detect GWs are conducted by measuring how they move suspended test masses, insulated (as well as possible) against other interference.

Laser interferometers consist of two perpendicular arms, with lengths of up to 4km and with two mirrors (the test masses) suspended at either end. A laser beam is then fired between the mirrors, and the effects of a passing GW on the spacetime will be apparent through a phase difference between the two laser beams when they arrive at a final detector.

Current laser interferometers are the LIGO detectors (two 4km and one 2km detectors) [7], the 3km VIRGO detector [40], and the smaller 600m GEO and 300m TAMA detectors [130, 204]. The larger detectors (LIGO and VIRGO) can detect finer movement of the test masses, and hence smaller effects of GWs and are the most likely to yield detection results. The smaller GEO and TAMA detectors are primarily test-beds for technological developments, but can also be used when the larger detectors are offline providing cover primarily in the case of a large galactic event occurring. Both the LIGO and VIRGO detectors have achieved their design sensitivities (each at a level of 10^{-19} m/ $\sqrt{\text{Hz}}$) [2, 4, 5], and upgrades to even more sensitive detectors, advanced LIGO and VIRGO+, are currently under construction (*second generation* ground-based detectors) aiming for an order of magnitude improvement in the sensitivity [172]. Beyond that a *third generation* ground based detector, the Einstein Telescope (ET) has also been proposed with a further one to two orders of magnitude increase in sensitivity [1].

Even with the sensitivity of current detectors, the direct detection of GWs remains elusive. We are, however, still able to make some predictions as to how likely it is that these detectors (and the future generations) will make a detection. We can measure this likelihood based on, for a specific type of event, the maximum distance for which we would expect to see the GWs of the event (those listed at the beginning of this section). Considering the amount of such events we would expect to see within this distance, we can then obtain an *event rate* for a given detector. Many of the phenomena associated with the interiors of NSs (e.g. for waves from NS glitches) are too small to allow for GW detection even with third generation detectors [187]. SN and particularly binary merger, however, offer more promising sources of GWs.

Galactic SN would offer an excellent potential source of GWs, but with only one such event every $\sim 30 - 100$ years, a lot of luck is required for such an observation. With the third generation detectors, the observable distance

makes for more reliable rates, though even these may not be more than 0.5 per year [18, 150].

By far the most promising source of GWs, with achievable detection rates for at least the second generation detectors, are NSNS binary inspirals. Here, we might expect $\sim 10 - 100$ detections per year [15, 110]. BHNS and BHBS inspirals also offer strong potentially detectable signals, however, these events are significantly rarer than NSNS binaries (due to the high mass progenitors for BHs being rarer themselves), and as a result the event rates are significantly lower.

So far we have only considered the mechanisms by which GWs will reach the detectors. In addition to this, the actual task of extracting GW signatures from the data collected by the interferometer data is by no means trivial. The weak nature of these signals results in them only being discernible amongst the detector noise and other terrestrial noise sources through careful statistical analysis. The quantity of data that can be collected, combined with the frequency range of expected GW signals, mean that for practical purposes, this analysis must be as efficient as possible. The best chance of successful detection relies on the availability of accurate templates of the expected waveforms. The interferometer data can then be searched for these waveforms, a technique known as *matched filtering* [95]. Alternative searches for the data are available: e.g. direct analysis of the waveforms for large amplitude signals. For more detailed description of data analysis see e.g. [172].

The lack of GW signals means that if we desire to test whether these data analysis techniques can detect waveforms we must use artificially created data. This has been the goal of the *Numerical INjection Analysis* (NINJA) project [20, 38]. This project initially used the expected signals for BH binary inspirals, injected into (unrealistic) detector noise, for use in the data analysis pipelines. Attempts to use more realistic detector noise, and to use NSNS inspiral data are underway.

The success of data analysis pipelines, such as the matched filtering techniques, requires accurate templates of the expected GW signal to be available. We must therefore have some method to obtain this signal from the GW release mechanisms detailed at the beginning of this section. More explicitly, we should solve the Einstein equations and matter evolution equations for these

scenarios and extract the GWs from some region suitably far from the source. This is not a trivial task; the equations form a coupled system in which non-linear behaviour is often expected to form. In order to obtain an accurate solution to such a system we require computational simulations (such as the inspiral simulations mentioned at the start of this section). We must therefore be able to express the Einstein equations, and the relativistic matter evolution equations, in a manner that allows these simulations to be carried out, this field is known as *numerical relativity* (NR). Since the exact details of a NS interior are unknown we must also use some approximate model which contains (to some degree of accuracy) all of the features we wish to use in the simulation. We now consider how such a model can be obtained. It should be noted that full GR is not always necessary for an accurate waveform, in some cases, where non-linear behaviour is not expected; the post-Newtonian (PN) approximation can give suitable GW data.

1.4 Modelling a neutron star

Before undertaking any numerical simulation of a NS, we must make choices (and assumptions) about how the star is modelled. These choices are, in part, based on the behaviour we wish to model (e.g. fast rotation rates, strong magnetic fields or hyperons in the inner core) and on unknown factors in the models (e.g. the extent of a nuclear pasta region or magnetic fields within the core). These choices will also be heavily influenced by the computational demands of the problem being considered. Effectively there must be some ‘trade-off’ between the accuracy of the model and the computational resources required for the simulation.

Whenever any properties of a physical system are evaluated numerically, the system must be discretised to some extent (both in space and time for a dynamical system). The smaller the length of the discretisation, the better the accuracy obtained from the simulation. Note that increases in accuracy are not automatically guaranteed from decreasing the length of the discretisation. The methods used for the simulation must be chosen carefully to ensure that this is possible (see chapter 3 for further information). However, any decrease in the size of the discretisation does lead directly to an increase in

the computational cost of the simulation. For a system to emit GWs it must be asymmetric, and therefore the computational simulation must be carried out in three spatial dimensions. This means that by halving the level of discretisation, the computational requirements for the simulation increase by a factor of 2^3 . Additionally the discretisation time must also be decreased by this same factor, leading to a total factor of 2^4 increase in the computational requirements. Therefore for practical situations, the microphysics of the different regions of the NS, and the behaviour at the transitions between these regions cannot be incorporated into a system that requires the evolution of the entire NS over timescales such as those required for merger simulations.

It should be noted that it is not always necessary for the discretisation of a system to be uniform. We might consider decreasing the level of discretisation only where some feature changes on small lengthscales. This can even be achieved dynamically (if such features are expected to move during the simulation) using techniques such as adaptive mesh refinement [31]. Even using these advanced techniques, the computational requirements will still increase dramatically as the accuracy of the solutions obtained is increased.

We now consider some of the assumptions that go into modelling an entire NS for numerical simulations. We need some system of equations which describes how the matter within the star evolves over time. The extreme densities within a NS mean that for all but the finest scales, it is a valid assumption to treat the matter as a continuum. Furthermore, as described in section 1.2, it is believed that much of the NS interior is a fluid. We assume, therefore, that the NS can be modelled using some system of hydrodynamical evolution equations. If necessary, such a model can be extended to include e.g. an elastic crust or solid core.

If we were to model a hydrodynamical system using Newtonian physics then the Navier-Stokes equations would be an obvious choice for the evolution equations. These are a system of non-linear evolution equations for a viscous fluid allowing for the inclusion of external forces, such as gravity. For modelling a NS, however, we must use GR, and ideally we need an extension of the Navier-Stokes equations to incorporate relativistic effects. Attempts at such an extension have been made, [101, 114, 155], but currently there is no single generally accepted form. The issue with these extensions is there is a possible

violation of causality in the propagation of viscous forces and heat flow [214]. It is, however, often assumed that the viscous effects in a NS are less important in determining short term behaviour than e.g. pressure gradients. As a result these terms can be neglected in many cases such as NSNS mergers. In Newtonian physics, for this inviscid, or *perfect*, fluid the Navier-Stokes equations reduce to the Euler equations. These equations can be extended to relativistic hydrodynamics and as such are the natural choice for use in NSs [201]. In addition to this, the effects of magnetic fields can also be included within the model, see [69] for a review.

The Euler equations, both in their original Newtonian form and under the relativistic extension, do not form a complete set of equations. There is a degree of freedom remaining, and this allows the type of fluid being modelled to be constrained: e.g. we need to be able to differentiate whether water or a neutron fluid is being used. This is achieved through the EOS which, as mentioned in section 1.2, effectively encompasses the microphysical behaviour of the fluid as a global property. The requirements of an EOS, and some specific choices are considered in detail in section 2.3.

1.4.1 Current neutron star simulations

Computational simulations have been carried out for many of the mechanisms for producing GWs listed in section 1.3. These simulations have considered oscillation modes [115, 154], glitches [187], core collapse SNe [150] and NSNS binary merger [23, 24, 65]. In this thesis we are interested in developing numerical techniques for simulations involving NSs initially in, or close to, some equilibrium configuration and where the entire star is modelled. One major area for which these techniques are applicable is NSNS binary mergers. We therefore consider the current status of these simulations in more detail.

For the majority of a NSNS inspiral, the system is in a quasi-equilibrium state with angular momentum being lost through GW radiation, this can be modelled satisfactorily through a PN approximation [35]. Once the GW timescale becomes comparable to the dynamical timescale of the system, the PN approximation is no longer sufficient. Beyond this point a numerical solution in

full GR is required for an accurate description of the behaviour of the two objects [65]. It is this stage of the event that is most likely to produce detectable GWs, and therefore must be modelled accurately to obtain a suitable wave-signal. In general we require some simulation using NR to describe *at least* the last few orbits of the inspiral, the merger itself and the eventual collapse to BH. If there are sufficient orbits included in the NR simulation, then there will be some overlap with the PN quasi-equilibrium solution in the region where this approximation is still valid. This allows the GW signals from the two solutions to be matched in this region, giving a full description of the waveform of the merger.

In order to undertake a NR simulation of a merger, some initial data for the two NSs in quasi-equilibrium orbits must be provided. This is not a trivial task, see e.g. [47, 83] which describes how this initial data can be formulated, and e.g. the Meudon group’s work, in which this is implemented using the **Lorene** code [3, 194]. To this end, a variety of initial configurations for merger simulations have been formulated, with varying initial separations, masses and EOSs. At the beginning of any simulation using such quasi-equilibrium initial data, there will be some artificial error as it ‘settles’ into the inspiral orbits. This is then recorded as a burst of *junk radiation* in the GW signal from the simulation. The matching of the signal to the PN waveform can again be used here, allowing this junk radiation to be identified and removed.

It is, in fact, only comparatively recently that the computational requirements became available for a NR simulation of a few inspiral orbits, merger and ringdown. The first NSNS inspiral simulations were produced by Shibata et al. [182]. These, along with other early simulations ([12, 23, 216]) used equal mass NSs with a simple EOS, the polytropic EOS, described in section 2.3. Later simulations considered unequal mass NSNS binaries [23, 180]. The results of these simulations show that the larger the mass of the system, the more rapidly the merged object collapses to a BH. As a result, the low-mass systems produce a longer, noisier GW signal from the merger, and this should be obvious in any detections of these events. It is clearly apparent, therefore, how successful GW detection from mergers can give insight into the properties of the NSs undergoing such events. Further differences arise if the EOS of the NSs in the simulation varies, as shown by [181, 23]. Here it is seen that if the

EOS is changed, allowing the temperature within the NSs to change, then the delayed collapse of the merger object to a BH is observed even for higher mass systems. EOSs based on nuclear physics have been used in [181, 179], and some of the most realistic, EOSs currently only available in Newtonian physics have been implemented through an ‘approximative treatment’ to GR [145]. The simulation of NSNS mergers with magnetic fields has also been considered [72, 161]. Here it is found that with the exception of very large magnetic fields (which are not expected to exist during the final merger stages), there is little difference compared to the non-magnetised cases.

In addition to these scenarios which have currently been investigated, there is a wide parameter space which may need to be considered in further NSNS merger simulations. It is likely that energy transport through neutrinos must be included, with particular relevance for the post-merger object [59]. Additionally all simulations so far have used only a single component fluid i.e. the star is described entirely by a single EOS. As we discussed in section 1.2 we know, however, that NSs have several different interior regions. In order to develop an accurate model of a NS for use in simulations, we need it to be able to include these differing regions. It is believed that these regions are described by different EOSs, therefore we would want our model to also reflect this. The techniques required to include multiple regions within a NS has been the focus of this thesis. We present a treatment that allows these features to be considered, and hence lead to an accurate model of a NS.

1.4.2 Modelling interfaces in a neutron star

Including different regions of a NS in a NR simulation of e.g. a NSNS inspiral and merger is not a trivial task. Ideally we would want to consider a full model in which each region (e.g. crust and outer core) is separated by some transition layer. In this transition layer we would then need to include the microphysics that describes how the matter changes, such as that shown in the pasta phase. In many cases, these transition layers will be very thin (e.g. the Ekman layer which may exist at the crust/core boundary has a width $\sim 1 - 10$ cm [14]). At the beginning of this section we highlighted some of the issues associated with numerical simulations when dealing with features which change on small lengthscales. The length of the discretisation required for these features is

in many cases incompatible with the requirements of the overall simulation. These transition layers are an example of this: in section 3.7.1 we give a detailed quantitative presentation of the difficulties associated with including a transition layer within a NS in the context of binary inspiral.

An alternative method for incorporating these different regions into a NS is to neglect the transition layers and assume there is a genuinely sharp interface between regions (i.e. the interface has zero width). This does not necessarily mean that the effects these transition layers have on the matter within the NS cannot be included in the model. By some suitable choice of boundary conditions at these sharp interfaces we may hope to reproduce accurately the behaviour here. In addition to correctly providing the physical conditions at the interface, we must also take great care with how these are implemented numerically. Since we are assuming a sharp interface between regions of the NS, this means that at least some of the variables in the evolution equations will be discontinuous at this point. Such discontinuities often present difficulties in numerical simulations, with unphysical errors, such as oscillations, being introduced. For example, in section 4.1.1 we show some of the issues a naive implementation of such conditions can have on a numerical simulation. This implies that we must be very careful with how any boundary conditions are introduced to model behaviour across interfaces.

Techniques to implement such boundary conditions in Newtonian computational fluid dynamics CFD have been considered, see e.g. [148]. One such successfully implemented method is to treat the system in a *multicomponent* manner, in which each fluid region (each region which obeys a different EOS) is treated individually. This allows a final solution to be ‘pieced’ together from these components. Typically the interface treatment in such multicomponent models can be considered as being composed of two techniques; one to identify the location of the interface and one to apply appropriate boundary conditions.

When the location of some sharp feature, such as interfaces in this treatment, is required, then *level set* methods are frequently used. These have been applied to such features in both Newtonian CFD [148, 149, 176] and relativistic situations (such as BH event horizons [53, 199]). Level set methods can accurately track the location of a sharp feature, and deal naturally with topological changes in behaviour. Work on boundary conditions in a multicomponent fluid system has been considered, however, in this case only

in Newtonian CFD [67, 129]. We therefore offer the first analysis of the viability of extending these methods into relativistic situations. To achieve this we consider the Ghost Fluid Method, devised for Newtonian CFD by Fedkiw et al. [67]. This technique has been selected for both comparative simplicity and the ease in which it can be extended to multidimensional situations with general EOSs.

Techniques such as the Ghost Fluid Method are required to make some assumptions of the fluid behaviour around the interface in order to provide boundary conditions. It is possible that there will be cases in which these are not applicable to NSs. Firstly, the evolution of the level set function assumes that there is no mass transfer across the interface. In a long-scale evolution of a NS, this may not go to be accurate, matter is expected to move between regions, particularly in the case of the NS returning to an equilibrium configuration after e.g. additional matter is added to the star. However, in many of the cases we consider, the effects of matter transfer between regions will occur on a longer timescale than for the dynamic behaviour which is being simulated, such as if merger is followed by prompt collapse to BH. Therefore we expect such methods to be useful in a wide range of NS simulations. The Ghost Fluid Method itself also relies on a constant extrapolation of entropy across the interface. This may also not be the case in some NS simulations, and hence in these cases we may not obtain the correct behaviour at the interface. However, if this problem arises, methods, such as those described in section 3.7.5, deal with cases in which these assumptions are not applicable. The success of the Ghost Fluid Method in Newtonian CFD, even within complex situations, suggests that the relativistic extension will be of value to studying interfaces within NSs.

1.5 Notation and structure

In this work, geometric units ($c = G = 1$) and the renormalisation $M_\odot = 1$ are used throughout (unless otherwise stated) and we use the metric signature $(-1, 1, 1, 1)$. In tensor components, Greek indices refer to all spacetime components ($\mu, \nu, \dots \in \{0, 1, 2, 3\}$) and Latin indices refer to only the spatial components ($i, j, \dots \in \{1, 2, 3\}$). The Einstein summation convention applies

to all tensor indices except those in reference to a particular coordinate system i.e. $\{t, x, y, z\}$ or $\{t, r, \theta, \phi\}$ where indices represent a specific tensor component.

Approximated quantities, such as those used in approximate Riemann solvers, or in discretised quantities, are written with a circumflex e.g. \hat{u} . Discretised quantities are specified through subscripts i, j, k, \dots , referring to grid location, and through a superscript n referring to the time at which this discretisation is made e.g. \hat{u}_i^n . In chapter 3 we deal with *reconstructed* quantities, which are notated with a bar, e.g. $\bar{u}_{i+1/2}^n$.

The work within this thesis is carried out in the Eulerian frame. There are several advantages to this, the most important of which is that many of the techniques devised for modelling sharp interfaces are described in this frame. Furthermore, when considering multidimensional situations, if a Lagrangian frame is used, then it is possible for the coordinate grid to become ‘twisted’ due to e.g. rotational behaviour. Clearly the fixed grid implemented for an Eulerian implementation of the hydrodynamics evolution equations offers advantages.

The terminology used within this thesis follows many of the standard conventions of the relativistic literature, and certain phrases may differ from those used in e.g. Newtonian CFD. Therefore it is advantageous to define a few terms in more detail. When referring to multicomponent fluids, this describes any case in which an interface is present (i.e. level set methods are being used). The two (or more) regions which arise based on the zeroes of the level set function are each termed a component. Across these interfaces we will also make reference to a ‘jump in the EOS’. This incorporates any case in which the properties of the fluid either side of the interface change, including if the parametrisation of the EOS changes. We therefore consider the situation in which each component of the system is described by the ideal EOS, but each region has a different value of the parameter γ (see section 2.3), as containing a jump in the EOS.

The structure of this thesis is as follows. In chapter 2 we describe the theoretical background required to model fluids in the simplified model of a NS set out in section 1.4.2. We also consider how initial data for various problems can be set up. In particular, we show how a stable multicomponent star can be initialised in section 2.6.1. In chapter 3 we detail the numerical methods used to solve the fluid (and spacetime) evolution equations. We give special

attention to the methods for evolving multicomponent fluids in section 3.7. The results for simulating multicomponent fluids are presented in chapter 4. Here we demonstrate our relativistic extension of the Ghost Fluid Method introduced in section 1.4.2. We first investigate the method in 1+1 dimensional special relativity, and then extend this to general relativity. Work on our relativistic extension to the Ghost Fluid Method and the results in this section have been published in [135]. Finally we give some preliminary results for our multidimensional relativistic Ghost Fluid Method. Conclusions and extensions to this work are described in chapter 6.

Chapter 2

Theoretical basis

In order to simulate a physical system, we must describe the behaviour at any point in spacetime through a set of evolution equations, one for each independent quantity. In this section, we outline the derivation of such equations for a perfect fluid and the spacetime containing the fluid. The techniques used throughout this thesis originate in Newtonian CFD. As a result, we gain an advantage from first considering the Newtonian evolution equations from which we construct reference models using these currently existing techniques. The effects of extending these techniques to relativistic situations will then be clearer.

2.1 Newtonian fluid evolution equations

In Newtonian hydrodynamics, the Euler equations for the evolution of a perfect fluid arise naturally in integral form by considering the conservation of mass, M , momenta, S (for which there is a component in each coordinate direction), and energy, E , [116, 123]. We detail the derivation for a single spatial dimension (i.e. we have a single momentum component), but these techniques extend to multiple dimensions. For simplicity we use Cartesian coordinates for this calculation, but the results can be obtained in any coordinate system by use of the appropriate transformations. It is initially assumed that there are no additional forces acting on the fluid, e.g. no gravity or heat conduction. As we shall show, these can be included in a more general set of equations in which these forces balance the conserved quantities exactly.

We first consider conservation of mass. This conservation law states the change in the mass of fluid within a domain $x_0 < x < x_1$ over time $t_0 < t < t_1$ is equal to the total mass of fluid which passes through the boundaries of this domain over the interval $\Delta t = t_1 - t_0$. We assume that there is no mass created or destroyed within the interval $\Delta x = x_1 - x_0$ and that the flow across the boundaries is measured in the positive x -direction. This can be expressed in integral form as

$$\int_{x_0}^{x_1} [\rho_0(x, t_1) - \rho_0(x, t_0)] dx = - \int_{t_0}^{t_1} [\rho_0(x_1, t) v(x_1, t) - \rho_0(x_0, t) v(x_0, t)] dt \quad (2.1.1)$$

where ρ_0 is the rest mass density, or mass per unit volume, and v the velocity.

Conservation of momentum can be expressed as the change in the total momentum over Δx being equal to the total momentum flow through the boundaries of Δx during Δt and the change of momentum at these boundaries due to pressure, p , in this same time interval. In integral form, this conservation law is

$$\begin{aligned} \int_{x_0}^{x_1} [S(x, t_1) - S(x, t_0)] dx = \\ - \int_{t_0}^{t_1} [S(x_1, t) v(x_1, t) - S(x_0, t) v(x_0, t)] dt \\ - \int_{t_0}^{t_1} [p(x_1, t) - p(x_0, t)] dt. \end{aligned} \quad (2.1.2)$$

Finally, conservation of energy states that the total change of energy over Δx is equal to the energy flow though the boundaries of Δx plus the overall energy change due to pressure during Δt . This has integral form

$$\begin{aligned} \int_{x_0}^{x_1} [E(x, t_1) - E(x, t_0)] dx = \\ - \int_{t_0}^{t_1} [E(x_1, t) v(x_1, t) - E(x_0, t) v(x_0, t)] dt \\ - \int_{t_0}^{t_1} [p(x_1, t) v(x_1, t) - p(x_0, t) v(x_0, t)] dt. \end{aligned} \quad (2.1.3)$$

The three conservation equations, (2.1.1), (2.1.2) and (2.1.3) form a coupled

system of equations. We can express this system in terms of a state vector, \mathbf{q} , and flux vector \mathbf{f} . This allows the integral conservation equations to be written as

$$\int_{x_0}^{x_1} [\mathbf{q}(x, t_1) - \mathbf{q}(x, t_0)] dx = - \int_{t_0}^{t_1} [\mathbf{f}(x_1, t) - \mathbf{f}(x_0, t)] dt. \quad (2.1.4)$$

The two vectors introduced above are then given by

$$\mathbf{q} = \begin{pmatrix} \rho_0 \\ S \\ E \end{pmatrix} \quad (2.1.5)$$

and

$$\mathbf{f} = \begin{pmatrix} S \\ Sv + p \\ (E + p)v \end{pmatrix}. \quad (2.1.6)$$

These *conserved* variables, \mathbf{q} , couple through the so-called *primitive* variables ρ_0 , v , p and specific internal energy ε , where the momentum and energy are given by

$$S = \rho_0 v, \quad E = \rho_0 \left(\varepsilon + \frac{v^2}{2} \right). \quad (2.1.7)$$

We refer to these as primitive variables since they are often those that are the easiest to measure directly in a physical system. Additionally they allow for a more intuitive understanding of the behaviour of a fluid system.

Currently the conservation equations, (2.1.4), do not form a closed set. There are currently three equations, but four independent quantities, i.e. the primitive variables ρ_0 , v , p and ε . The system is closed by an additional equation which describes one of these primitive quantities in terms of the others, and effectively constrains the type of fluid being considered. This is an EOS, typically of the form

$$p = p(\rho_0, \varepsilon). \quad (2.1.8)$$

The EOS is dependent on the fluid being modelled. Further details as to the choice of EOS in this work are in given section 2.3. Once we have an EOS, we

can then define a vector of primitive variables, e.g.

$$\mathbf{w} = \begin{pmatrix} \rho_0 \\ v \\ \varepsilon \end{pmatrix}. \quad (2.1.9)$$

The most common form in which the conservation equations, (2.1.4), are written is not in integral form, but rather as a system of partial differential equations (PDEs). To see how we can write our integral equations in this form, we first note that if \mathbf{q} is differentiable in time, we have

$$\mathbf{q}(x, t_1) - \mathbf{q}(x, t_0) = \int_{t_0}^{t_1} \frac{\partial \mathbf{q}}{\partial t} dt. \quad (2.1.10)$$

A similar result is apparent if \mathbf{f} is differentiable in space,

$$\mathbf{f}(x_1, t) - \mathbf{f}(x_0, t) = \int_{x_0}^{x_1} \frac{\partial \mathbf{f}}{\partial x} dx. \quad (2.1.11)$$

These two relations can be inserted into the integral equation, (2.1.4), to give

$$\int_{x_0}^{x_1} \int_{t_0}^{t_1} \frac{\partial \mathbf{q}}{\partial t} dt dx = - \int_{t_0}^{t_1} \int_{x_0}^{x_1} \frac{\partial \mathbf{f}}{\partial x} dx dt. \quad (2.1.12)$$

Assuming that the order of integration is reversible, this can then be simplified to

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}}{\partial x} = 0. \quad (2.1.13)$$

Note that the constants of integration that arise from this simplification must be zero due to the limits of integration being arbitrary, i.e. the relationship in (2.1.12) must hold for all values of these limits. This result, for the vectors \mathbf{q} and \mathbf{f} given by (2.1.5) and (2.1.6) respectively, is the standard differential equation form of the Newtonian 1+1 dimensional Cartesian Euler equations in conservation form.

The generalisation of this equation, (2.1.13), to multidimensional (Cartesian) systems with non-trivial body forces (forces which act on the fluid, e.g. gravitational potential) can be derived in a similar manner. Such a calculation

yields

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}^i}{\partial x^i} = \mathbf{s} \quad (2.1.14)$$

where $x^i = \{x, y, z\}$. The state vector, \mathbf{q} , is similar to (2.1.5), but now must include momentum components for each direction x^i , i.e.

$$\mathbf{q} = \begin{pmatrix} \rho_0 \\ S_x \\ S_y \\ S_z \\ E \end{pmatrix} \quad (2.1.15)$$

with

$$S_i = \rho_0 v_i, \quad E = \rho_0 \left(\varepsilon + \frac{v_i v^i}{2} \right), \quad (2.1.16)$$

where subscript or superscript i represents the momentum or velocity component in the direction x^i . There is now also a flux vector \mathbf{f} for each direction x^i , given by

$$\mathbf{f}^i = \begin{pmatrix} S^i \\ S_x v^i + \delta_x^i p \\ S_y v^i + \delta_y^i p \\ S_z v^i + \delta_z^i p \\ (E + p) v^i \end{pmatrix}, \quad (2.1.17)$$

where the Kronecker delta,

$$\delta_j^i = \begin{cases} 1 & i = j \\ 0 & \text{otherwise} \end{cases}, \quad (2.1.18)$$

has been used. All body forces are included in the source vector \mathbf{s} . Equation (2.1.14) then gives the full three dimensional Euler equations in Cartesian coordinates. If $\mathbf{s} = 0$ then (2.1.14) is in conservation form. When non-zero sources are present, provided they contain no derivatives of the conserved quantities, these equations can be thought of as being in *balance law* form, since it is these forces that exactly balance the conserved quantities.

Aside from originating from the intuitive description of fluid behaviour, the importance in writing the evolution equations in conservation (or balance law)

form arises when non-linear behaviour is present. The differential equation form was dependent on the solution being differentiable, which is clearly not the case for shocks, where the fluid variables jump. This is explored in more detail in section 3.3.

2.1.1 Spherically symmetric fluid evolution equations

We have so far derived the fluid evolution equations in Cartesian coordinates; however, in many practical cases this will not be the natural coordinate system to use. One case of particular importance to us is the modelling of stars (e.g. NSs), for which spherical coordinates are a more natural choice. Here we shall consider the spherically symmetric evolution equations (i.e. in 1+1 dimension with only the radial derivatives non-vanishing). These can be written in the same form as (2.1.13), but the coordinate transformation results in an additional source term, (see e.g. [201]),

$$\frac{d}{dt} \begin{pmatrix} \rho_0 \\ S \\ E \end{pmatrix} + \frac{d}{dr} \begin{pmatrix} S \\ Sv + p \\ (E + p)v \end{pmatrix} = \begin{pmatrix} -\frac{2}{r}\rho_0 v \\ -\frac{2}{r}\rho_0 v^2 \\ -\frac{2}{r}v(E + p) \end{pmatrix}. \quad (2.1.19)$$

This form of the equations is not necessarily the most ‘natural’ form, particularly for accurate numerical solutions, since the presence of the $1/r$ term in \mathbf{s} leads to singular behaviour at the origin. The source terms can, however, be incorporated into \mathbf{f} , giving

$$\frac{d}{dt} \begin{pmatrix} \rho_0 \\ S \\ E \end{pmatrix} + \frac{1}{r^2} \frac{d}{dr} \begin{pmatrix} r^2 S \\ r^2 Sv \\ r^2(E + p)v \end{pmatrix} + \frac{d}{dr} \begin{pmatrix} 0 \\ p \\ 0 \end{pmatrix} = \mathbf{0}. \quad (2.1.20)$$

This has removed the source terms, and hence, the singular behaviour at $r = 0$ that would arise from such terms. There is now, however, a $1/r^2$ factor in front of the flux vector term. The relationship

$$\frac{1}{r^2} \frac{d}{dr} = 3 \frac{d}{d(r^3)}. \quad (2.1.21)$$

then allows this additional singular term to be removed, and hence we obtain the so-called ‘natural’ form of the spherically symmetric evolution equations,

$$\frac{d}{dt} \begin{pmatrix} \rho_0 \\ S \\ E \end{pmatrix} + 3 \frac{d}{d(r^3)} \begin{pmatrix} r^2 S \\ r^2 S v \\ r^2 (E + p) v \end{pmatrix} + \frac{d}{dr} \begin{pmatrix} 0 \\ p \\ 0 \end{pmatrix} = \mathbf{0}. \quad (2.1.22)$$

The equations above are still expressed in the absence of body forces, and, therefore, are not yet suitable for use in describing the evolution of matter within a star. A source term which incorporates the effects of the gravitational potential can be added to this natural form, [201],

$$\mathbf{s} = \begin{pmatrix} 0 \\ -\rho_0 \nabla \Phi \\ -\rho_0 v \nabla \Phi \end{pmatrix}. \quad (2.1.23)$$

Therefore the evolution equations for a self-gravitating spherically symmetric fluid in Newtonian physics can be written as

$$\frac{d}{dt} \begin{pmatrix} \rho_0 \\ S \\ E \end{pmatrix} + 3 \frac{d}{d(r^3)} \begin{pmatrix} r^2 S \\ r^2 S v \\ r^2 (E + p) v \end{pmatrix} + \frac{d}{dr} \begin{pmatrix} 0 \\ p \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ -\rho_0 \nabla \Phi \\ -\rho_0 v \nabla \Phi \end{pmatrix}. \quad (2.1.24)$$

Note that in these spherically symmetric coordinates, we have $\nabla \Phi = \frac{d\Phi}{dr}$.

2.2 Relativistic fluid evolution equations

In relativistic situations there is a coupling between space and time, and this must be taken into account when considering the equations that govern a fluid’s behaviour. We must also consider the effects of Einstein’s General Theory of Relativity, which describes gravity as the curvature of spacetime. In GR, matter and energy distort the geometry of spacetime, whilst this distortion then affects how matter moves. The geometry of the spacetime is described by a 4-dimensional manifold, \mathcal{M} (see e.g. [9, 136, 205]). On \mathcal{M} , the invariant distance between two events is given by a metric, $g_{\mu\nu}$, which is a symmetric tensor. The structure of this manifold would then be expected to enter equations for the

evolution of the mass-energy of the system. For simplicity, this mass-energy is typically referred to as matter. Similarly, the evolution of the spacetime manifold would then be also expected to be dependent on the matter.

In GR the density and flux of matter is described through the *stress-energy tensor* $T_{\mu\nu}$. Since in this thesis we consider only perfect fluids, we will deal with a stress-energy tensor of the form

$$T_{\mu\nu} = \rho_0 h u_\mu u_\nu + p g_{\mu\nu} \quad (2.2.1)$$

where ρ_0 is the rest-mass density of the fluid, h the specific enthalpy, u_μ the covariant 4-velocity and p the pressure. The 4-velocity of a particle in GR is defined by

$$u^\mu = \frac{dx^\mu}{d\tau} \quad (2.2.2)$$

where τ is proper time and x^μ is the coordinate position of the particle. The specific enthalpy for a relativistic fluid is given by

$$h = 1 + \varepsilon + \frac{p}{\rho_0}, \quad (2.2.3)$$

where (as in Newtonian physics) ε is the specific internal energy of the fluid.

We also have a covariant derivative, ∇_μ associated with the metric $g_{\mu\nu}$, such that $\nabla_\alpha g_{\mu\nu} = 0$. For an arbitrary tensor $S^{\mu_1 \dots \mu_k}{}_{\nu_1 \dots \nu_l}$ the covariant derivative is given by

$$\begin{aligned} \nabla_\alpha S^{\mu_1 \dots \mu_k}{}_{\nu_1 \dots \nu_l} &= \partial_\alpha S^{\mu_1 \dots \mu_k}{}_{\nu_1 \dots \nu_l} + \Gamma_{\beta\alpha}^{\mu_1} S^{\beta \dots \mu_k}{}_{\nu_1 \dots \nu_l} + \dots \\ &\quad + \Gamma_{\beta\alpha}^{\mu_k} S^{\mu_1 \dots \mu_{k-1}\beta}{}_{\nu_1 \dots \nu_l} - \Gamma_{\nu_1\alpha}^\beta S^{\mu_1 \dots \mu_k}{}_{\beta \dots \nu_l} \\ &\quad - \dots - \Gamma_{\nu_l\alpha}^\beta S^{\mu_1 \dots \mu_k}{}_{\nu_1 \dots \nu_{l-1}\beta}. \end{aligned} \quad (2.2.4)$$

Here ∂_μ is a partial derivative and $\Gamma_{\nu\alpha}^\mu$ are connection coefficients, or Christoffel symbols, which are given by

$$\Gamma_{\nu\alpha}^\mu = \frac{1}{2} g^{\mu\beta} (\partial_\alpha g_{\nu\beta} + \partial_\nu g_{\alpha\beta} - \partial_\beta g_{\nu\alpha}). \quad (2.2.5)$$

The presence of the metric terms in the covariant derivative mean that the curvature of spacetime will effect the evolution of a fluid. In addition to this it is already evident that the spacetime will enter the description of matter

through the stress-energy tensor (2.2.1).

The curvature of an arbitrary manifold is given by the Riemann tensor, $R^\mu_{\nu\alpha\beta}$, and can be given in terms of the Christoffel symbols through

$$R^\mu_{\nu\alpha\beta} = \partial_\alpha \Gamma^\mu_{\nu\beta} - \partial_\beta \Gamma^\mu_{\nu\alpha} + \Gamma^\mu_{\gamma\alpha} \Gamma^\gamma_{\nu\beta} - \Gamma^\mu_{\gamma\beta} \Gamma^\gamma_{\nu\alpha}. \quad (2.2.6)$$

Contractions of the Riemann tensor can also be defined, with the Ricci tensor given by

$$R_{\mu\nu} = R^\alpha_{\mu\alpha\nu} \quad (2.2.7)$$

and the Ricci scalar by

$$R = R^\mu_\mu. \quad (2.2.8)$$

The curvature of spacetime (which is subject to various simplifications from a general 4-dimensional manifold, see e.g. [136, 177]) can then be described through the Einstein tensor, $G_{\mu\nu}$. This is given by

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R. \quad (2.2.9)$$

The manner in which the matter distorts the spacetime is then given by the Einstein equations,

$$G_{\mu\nu} = 8\pi T_{\mu\nu}. \quad (2.2.10)$$

Due to the symmetric nature of both the Einstein tensor and the perfect fluid stress-energy tensor, this is a system of 10 coupled PDEs.

As with Newtonian physics, the matter evolution equations are based on locally conserved quantities. In GR these quantities are baryon number, the flow of which is given by $\rho_0 u^\mu$, and stress-energy, $T^{\mu\nu}$. These have analogies in a Newtonian fluid; the former of these can be compared to conservation of mass and the latter incorporates the conservation of momentum and energy. For a vector, \mathbf{F} , which represents the flux of some conserved quantity through the boundaries of a hypersurface Σ , this can be represented in integral form as

$$\oint_{\Sigma} \mathbf{F} \cdot \mathbf{n} \, d\Sigma = 0, \quad (2.2.11)$$

where \mathbf{n} is the unit-normal of the surface Σ . This result also holds for each free index of a conserved tensor quantity (e.g. $T^{\mu\nu}$). Effectively we have stated that there is no transfer of matter through the boundaries of Σ . We can then apply Gauss' theorem to this result to give

$$\int_V (\nabla \cdot \mathbf{F}) \, dV = 0 \quad (2.2.12)$$

where V is the spacetime volume bounded by Σ . This is a generalisation of the result given in (2.1.13) where the integral equations for a vector of conserved quantities, \mathbf{u} , can now be written as a system of PDEs. In this general case we have

$$\nabla_\mu \mathbf{F}^\mu = 0 \quad (2.2.13)$$

where, for a relativistic fluid we have the conserved quantities as described above, $\rho_0 u^\mu$ and $T^{\mu\nu}$. Note that the constants of integration will also vanish in this case, as described in section 2.1. The conservation equations are then often written as the conservation of baryon number,

$$\nabla_\mu (\rho_0 u^\mu) = 0, \quad (2.2.14)$$

and conservation of stress-energy,

$$\nabla_\mu T^{\mu\nu} = 0. \quad (2.2.15)$$

The coupled nature of space and time means interpreting any solutions of the conservation and Einstein equations is not necessarily intuitive. Furthermore the ‘natural’ 4-dimensional geometric form of the equations is often not suitable for use in a numerical simulations of relativistic fluid evolution. Since most physically relevant situations do not have analytic solutions, these numerical simulations are often essential. It is, therefore, useful to consider a decomposition of the 4-dimensional manifold, \mathcal{M} , and hence of the Einstein and evolution equations. Typically a ‘3+1’ formalism used, a decomposition of \mathcal{M} into 3 ‘spatial’ dimensions and one ‘time’ dimension. The first such 3+1 formalism was the Arnowitt-Deser-Misner (ADM) formalism [19], and

the derivation given here is based on the reformulation of York [217], with the applications to fluid evolution equations are based on e.g. [44, 140, 142].

2.2.1 3+1 decomposition of the Einstein equations

Here we outline some of the major features and results from making a 3+1 decomposition of the Einstein equations. Full details of this calculation are in appendix A.1. To make a 3+1 decomposition of the spacetime, \mathcal{M} , we foliate it into a set of spacelike hypersurfaces, Σ_τ , which are level surfaces of a scalar field, τ , interpreted as some ‘time’ function. These hypersurfaces are then described locally by a closed dual-vector field, Ω_μ , where

$$\Omega_\mu = \nabla_\mu \tau. \quad (2.2.16)$$

The norm of this dual-vector field is then defined to be

$$g^{\mu\nu} \Omega_\mu \Omega_\nu = -\alpha^{-2}, \quad (2.2.17)$$

where α is a scalar function, known as the lapse function. The lapse can be thought of as a measure of the perpendicular distance between spacetime slices Σ_τ . From this, a unit-normal dual-vector field, n_μ , can be constructed,

$$n_\mu = -\alpha \Omega_\mu. \quad (2.2.18)$$

The sign of this dual-vector is chosen such that the normal vector n^μ is future pointing. This normal vector then acts as a projection operator so that the contraction of a tensor with n^μ will give its timelike components.

We can also define a projection operator which, when used in the contraction of a tensor, gives the spacelike components. This projection tensor is given by

$$\perp^\mu_\nu \equiv \delta^\mu_\nu + n^\mu n_\nu. \quad (2.2.19)$$

If we have an arbitrary tensor $S^{\mu_1\mu_2\dots\mu_k}_{\nu_1\nu_2\dots\nu_l}$, then we define \perp to be the

spatial projection of every index of this tensor,

$$\perp S^{\mu_1\mu_2\ldots\mu_k}_{\nu_1\nu_2\ldots\nu_l} = \perp^{\mu_1}_{\alpha_1} \perp^{\mu_2}_{\alpha_2} \ldots \perp^{\mu_k}_{\alpha_k} \perp^{\beta_1}_{\nu_1} \perp^{\beta_2}_{\nu_2} \ldots \perp^{\beta_l}_{\nu_l} S^{\alpha_1\alpha_2\ldots\alpha_k}_{\beta_1\beta_2\ldots\beta_l}. \quad (2.2.20)$$

A tensor projected in such a manner is referred to as a *spatial tensor*.

From these projections we can define the metric on the hypersurfaces, the *spatial metric*, $\gamma_{\mu\nu}$, through the projection of the full metric $\perp g_{\mu\nu}$,

$$\gamma_{\mu\nu} = g_{\mu\nu} + n_\mu n_\nu. \quad (2.2.21)$$

We also have the contravariant form of this metric,

$$\gamma^{\mu\nu} = g^{\mu\nu} + n^\mu n^\nu. \quad (2.2.22)$$

Similarly we can define the spatial covariant derivative operator, D_μ , in the obvious manner

$$D_\mu \equiv \perp \nabla_\mu. \quad (2.2.23)$$

The intrinsic curvature of any spacelike slice can be defined through the spatial Riemann tensor, ${}^{(3)}R^\mu_{\nu\alpha\beta}$, which is the Riemann tensor associated with $\gamma_{\mu\nu}$. This can be defined in an analogous manner to the full Riemann tensor, through its action on an arbitrary *spatial* dual-vector ω_μ ,

$$(D_\mu D_\nu - D_\nu D_\mu) w_\alpha = {}^{(3)}R_{\mu\nu\alpha}{}^\beta w_\beta. \quad (2.2.24)$$

As with its 4-dimensional counterpart, we can define a 3-dimensional Ricci tensor and Ricci scalar,

$${}^{(3)}R_{\mu\nu} = {}^{(3)}R^\alpha_{\mu\alpha\nu}, \quad {}^{(3)}R = {}^{(3)}R^\mu_{\mu}. \quad (2.2.25)$$

The intrinsic curvature gives a description of the curvature within a slice Σ_τ , but we also need to define how these slices are embedded in the 4-dimensional spacetime. This is done through the extrinsic curvature tensor, a symmetric

tensor given by

$$K_{\mu\nu} = -\perp \nabla_\mu n_\nu. \quad (2.2.26)$$

The extrinsic curvature can also be defined through a Lie-derivative of the spatial metric along the normal vector field,

$$K_{\mu\nu} = -\frac{1}{2} \mathcal{L}_n \gamma_{\mu\nu} = -\frac{1}{2} \perp \mathcal{L}_n g_{\mu\nu}. \quad (2.2.27)$$

The 3+1 decomposition of a tensor effectively splits it into timelike and space-like components. For reference, the 3+1 decomposition of an arbitrary symmetric tensor $\sigma_{\mu\nu}$ is

$$\sigma_{\mu\nu} = \perp \sigma_{\mu\nu} - 2n_{(\mu} \perp \sigma_{\nu)\hat{n}} + n_\mu n_\nu \sigma_{\hat{n}\hat{n}} \quad (2.2.28)$$

where an index \hat{n} represents a contraction with n^μ . It is convention that a contraction with n_μ introduces a minus sign, i.e. for arbitrary vector ω^μ , $\omega^{\hat{n}} = -\omega^\mu n_\mu$.

When considering the 3+1 decomposition of the stress-energy tensor, $T_{\mu\nu}$, the projections are defined as

$$\varrho \equiv T_{\hat{n}\hat{n}} = T_{\mu\nu} n^\mu n^\nu, \quad (2.2.29a)$$

$$j^\mu \equiv \perp T^{\mu\hat{n}} = -\perp (T^{\mu\nu} n_\nu), \quad (2.2.29b)$$

$$\mathfrak{E}_{\mu\nu} \equiv \perp T_{\mu\nu}. \quad (2.2.29c)$$

These projections then have the physical interpretation of the local energy density, ϱ , and momentum density, j^μ , and the local stress-energy tensor measured by observers moving along n^μ , $\mathfrak{E}_{\mu\nu}$.

In order to obtain a 3+1 form of the Einstein equations, since we wish to express the derivatives of quantities in terms of spatial and time derivatives, it is necessary to specify a vector along which we compute these time derivatives. In section (A.1.3) we justify the natural choice of vector for this is

$$t^\mu = N^\mu + \beta^\mu, \quad (2.2.30)$$




Figure 2.1: Showing how the vector t^μ describes the movement between two spacetime slices. The normal, αn^μ , and tangential, β^μ , components are also shown.

where $N^\mu = \alpha n^\mu$ and β^μ is an arbitrary spatial vector, the *shift vector*. This generalised vector, and how it relates to movement between the spacetime slices, is shown in figure 2.1.

From the generalised vector field, t^μ , the full line element for the 3+1 decomposition, as used in [9, 142], can be given as,

$$ds^2 = - (N^2 - \beta_i \beta^i) dt^2 + 2\beta_i dx^i dt + \gamma_{ij} dx^i dx^j. \quad (2.2.31)$$

and in this case, we have normal vector

$$n^\mu = \left(\frac{1}{\alpha}, -\frac{\beta^i}{\alpha} \right)^T \quad (2.2.32)$$

and its dual-vector

$$n_\mu = (-\alpha, 0, 0, 0). \quad (2.2.33)$$

The contractions involved in making the 3+1 decomposition of the Einstein equations, (2.2.10), along with the definition of the extrinsic curvature tensor, (2.2.27), allow for a full description of the behaviour of the spacetime within this framework. This description provides two constraint equations for the spacetime and two evolution equations, one for the intrinsic spacetime of

Σ_τ and one for the extrinsic curvature. These equations are derived in section A.1.4, and only the results are given here.

Contracting both indices of the Einstein equation with n^μ gives the Hamiltonian constraint,

$$\mathcal{R} + K^2 - K_\nu^\mu K_\mu^\nu = 16\pi\varrho. \quad (2.2.34)$$

The second constraint equation, the momentum constraint, is achieved by contracting the contravariant form of the Einstein equations with $-n_\mu$ and then making the spatial projection. This gives the momentum constraint,

$$D_\nu K^{\mu\nu} - D^\mu K = 8\pi j^\mu. \quad (2.2.35)$$

These two equations are known as the constraint equations since they contain only spatial tensors and their derivatives. They must hold for $\gamma_{\mu\nu}$ and $K_{\mu\nu}$ for all time when solving the Einstein equations.

From the definition of the extrinsic curvature tensor, (2.2.27),

$$K_{\mu\nu} = -\frac{1}{2}\mathcal{L}_n\gamma_{\mu\nu}, \quad (2.2.36)$$

we obtain the evolution equation for the spatial metric $\gamma_{\mu\nu}$,

$$\mathcal{L}_t\gamma_{\mu\nu} = -2\alpha K_{\mu\nu} + \mathcal{L}_\beta\gamma_{\mu\nu}. \quad (2.2.37)$$

Finally the evolution equation for the extrinsic curvature tensor arises from the full spatial contraction of the Einstein equation, which can be written as

$$\mathcal{L}_t K_\nu^\mu = \mathcal{L}_\beta K_\nu^\mu - D^\mu D_\nu \alpha + \alpha \left\{ \mathcal{R}_\nu^\mu + K K_\nu^\mu + 8\pi \left[\frac{1}{2} \perp_\nu^\mu (\mathfrak{E} - \varrho) - \mathfrak{E}_\nu^\mu \right] \right\} \quad (2.2.38)$$

where $\mathfrak{E} = \mathfrak{E}_\mu^\mu$.

We have now given the 3+1 evolution equations for a general spacetime, characterised by the metric $g_{\mu\nu}$. The fluid evolution equations, (2.2.14) and (2.2.15) can also be written in this form using the metric for the 3+1 decomposition, (2.2.31). In order to obtain evolution equations that can be used in computational simulations, we need to specify the coordinates used on the

spacetime, γ_{ij} , as well as the lapse, α , and shift, β .

When making a numerical simulation, it is more natural in many ways to work with the 3-velocity of the fluid. This is defined in terms of the 4-velocity, u^μ , by

$$v^i = \frac{u^i}{\alpha u^t} + \frac{\beta^i}{\alpha}. \quad (2.2.39)$$

We also have a covariant form given by a contraction with γ_{ij} ,

$$v_i = \frac{u_i}{W}, \quad (2.2.40)$$

where we have the scalar parameter

$$W = -u^\mu n_\mu = \alpha u^t. \quad (2.2.41)$$

Using the normalisation $u^\mu u_\mu = -1$, we show, in appendix A.1.5, that this parameter is simply the Lorentz factor as measured by Eulerian observers,

$$W = \frac{1}{\sqrt{1 - v_i v^i}}. \quad (2.2.42)$$

We are then able to use the Lorentz factor to replace u^t terms in evolution equations for the fluid.

2.2.2 The special relativistic fluid evolution equations

Here we consider the equations for the evolution of matter in special relativity (SR). This is a special case for the full GR scenario described above in which gravitational effects are neglected, i.e. we have a static spacetime with no curvature. We therefore work in Minkowski spacetime, which, in Cartesian coordinates, has the metric

$$g_{\mu\nu} = \eta_{\mu\nu} = \text{diag}[-1, 1, 1, 1]. \quad (2.2.43)$$

In SR, the flat and static spacetime leads to all Christoffel symbols $\Gamma_{\nu\alpha}^\mu$ vanishing, hence covariant derivatives become equivalent to partial derivatives. In terms of the 3+1 split, we can choose that the lapse and shift take their

‘trivial’ values, $\alpha = 1$ and $\beta^i = 0$. Therefore we have 4-velocities, obtained using (2.2.39), (2.2.41), (A.1.121) and (A.1.124),

$$u^\mu = W (1, v^i)^T, \quad u_\mu = W (-1, v_i). \quad (2.2.44)$$

We now consider the conservation laws, (2.2.14) and (2.2.15). We start with the conservation of mass, (2.2.14),

$$\nabla_\mu (\rho_0 u^\mu) = 0. \quad (2.2.45)$$

Considering the timelike and spacelike components, and using (2.2.44) to write the 4-velocity component in terms of the 3-velocity and the Lorentz factor, we get

$$\partial_t (\rho_0 W) + \partial_i (\rho_0 W v^i) = 0. \quad (2.2.46)$$

We then define the quantity $D = \rho_0 W$, which is the rest mass density seen from the Eulerian frame, giving

$$\partial_t D + \partial_i (D v^i) = 0. \quad (2.2.47)$$

We also consider the conservation of stress-energy, (2.2.15),

$$\nabla_\mu T^\mu_\nu = 0. \quad (2.2.48)$$

In SR, the components of the stress-energy tensor are given by

$$T^t_t = \rho_0 h u^t u_t - p \delta^t_t = \rho_0 h W^2 - p, \quad (2.2.49)$$

$$T^t_i = \rho_0 h u^t u_i + p \delta^t_i = \rho_0 h W^2 v_i, \quad (2.2.50)$$

$$T^i_t = \rho_0 h u^i u_t + p \delta^i_t = \rho_0 h W^2 v^i, \quad (2.2.51)$$

$$T^i_j = \rho_0 h u^i u_j + p \delta^i_j = \rho_0 h W^2 v^i v_j + p \delta^i_j. \quad (2.2.52)$$

These components, as described in section 2.2, give the total energy density measured in the Eulerian frame, $E = \rho_0 h W^2 - p$, and the momentum $S^i = \rho_0 h W^2 v^i$. These definitions then yield a useful relation,

$$S^i = (E + p) v^i. \quad (2.2.53)$$

Considering the t component of the conservation of stress-energy we have

$$\partial_\mu T^\mu_t = \partial_t T^t_t + \partial_i T^i_t = 0, \quad (2.2.54)$$

which can then be written

$$\partial_t E + \partial_i [(E + p) v^i] = 0. \quad (2.2.55)$$

The total energy density comprises the various energies present in the system; e.g. the self-energy, characterised by the ρ_0 component of E , and the kinetic energy. At low velocities the self energy of the system is much greater than the kinetic energy and dominates the energy term. As a result, the evolution equation for E becomes approximately equal to that for D , (2.2.47). When solving this equation numerically, this can lead to errors in obtaining the correct solution for E . To avoid these problems, the variable τ is defined, such that

$$\tau \equiv E - D \quad (2.2.56)$$

and, by subtracting (2.2.47) from (2.2.55), we obtain

$$\partial_t \tau + \partial_i [(\tau + p) v^i] = 0. \quad (2.2.57)$$

The spatial components of (2.2.48) are given by

$$\partial_\mu T^\mu_i = \partial_t T^t_i + \partial_j T^j_i = 0, \quad (2.2.58)$$

which, using (2.2.52) can be written

$$\partial_t S_i + \partial_j (S_i v^j + p \delta_i^j) = 0. \quad (2.2.59)$$

The three conservation laws derived, (2.2.47), (2.2.59) and (2.2.57) can be written in vector form

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}^i(\mathbf{q})}{\partial x^i} = 0 \quad (2.2.60)$$

with state vector

$$\mathbf{q} = \begin{pmatrix} D \\ S_x \\ S_y \\ S_z \\ \tau \end{pmatrix} \quad (2.2.61)$$

and flux vector

$$\mathbf{f}^i = \begin{pmatrix} Dv^i \\ S_xv^i + \delta_x^i p \\ S_yv^i + \delta_y^i p \\ S_zv^i + \delta_z^i p \\ (\tau + p)v^i \end{pmatrix}. \quad (2.2.62)$$

We can see that the conserved variables in SR, D , S_i and τ , lead to the same form of the state vector, \mathbf{q} , and flux vector, \mathbf{f}^i as in the Newtonian case, (2.1.14).

2.2.3 The general relativistic fluid evolution equations in spherical symmetry

Before we consider the specific case of the GR fluid evolution equations in spherical symmetry, it is worth considering the case in full 3+1 GR. This formalism of the evolution equations is derived in [26] and used in e.g. [25, 70], the equations shown in this section follow from these results. As with the special relativistic case, we have the conserved variables

$$D = \rho_0 W, \quad (2.2.63a)$$

$$S_i = \rho_0 h W^2 v_i, \quad (2.2.63b)$$

$$\tau = \rho_0 h W^2 - p - D. \quad (2.2.63c)$$

The 3+1 form of the conservation equations can then be written

$$\frac{1}{\sqrt{-g}} [\partial_t (\sqrt{\gamma} \mathbf{q}) + \partial_i (\sqrt{-g} \mathbf{f}^i)] = \mathbf{s} \quad (2.2.64)$$

where $g \equiv \det(g_{\mu\nu})$ and $\gamma \equiv \det(\gamma_{i,j})$. The state vector is then given by

$$\mathbf{q} = \begin{pmatrix} D \\ S_j \\ p \end{pmatrix}, \quad (2.2.65)$$

the flux vector by

$$\mathbf{f}^i = \begin{pmatrix} D \left(v^i - \frac{\beta^i}{\alpha} \right) \\ S_j \left(v^i - \frac{\beta^i}{\alpha} \right) + \delta_j^i p \\ \tau \left(v^i - \frac{\beta^i}{\alpha} \right) + p v^i \end{pmatrix}, \quad (2.2.66)$$

and source vector by

$$\mathbf{s} = \begin{pmatrix} 0 \\ T^{\mu\nu} (\partial_\mu g_{\nu j} - \Gamma_{\nu\mu}^\delta g_{\delta j}) \\ \alpha [T^{\mu t} \partial_\mu (\ln \alpha) - T^{\mu\nu} \Gamma_{\nu\mu}^t] \end{pmatrix}. \quad (2.2.67)$$

Polar-areal coordinates

In this work, since we restrict ourselves to spherical symmetry, meaning only radial derivatives are non-vanishing. The most general spherically symmetric 3+1 line element can be written as

$$ds^2 = (-\alpha^2 + a^2 (\beta^r)^2) dt^2 + 2a^2 \beta^r dt dr + a^2 dr^2 + r^2 b^2 d\Omega^2 \quad (2.2.68)$$

where $d\Omega^2 = d\theta^2 + \sin^2\theta d\phi^2$. In spherical symmetry all but the r component of β^i vanish due to spherical symmetry and a and b are functions of r and t only. We then need to make a choice of gauge conditions to specify within which coordinate system we are working. Here we follow e.g. [140, 142] by using the polar slicing condition and the areal or radial condition [9, 28, 45]. The form of the equations presented in this thesis does differ from these sources however, since we have matched the form used in many 3+1 dimensional NS simulations e.g. [25, 26]. The polar slicing condition is a condition such that the only non-zero component of the extrinsic curvature tensor $K_{\mu\nu}$ is K_{rr} . The areal condition chooses the radial coordinate such that the proper area of a

two-sphere is $4\pi r^2$, which gives $b = 1$. When combined, these gauge conditions are result in the shift vector vanishing everywhere i.e. $\beta^r = 0$.

The resulting coordinate system, known as polar-areal coordinates, then has line element

$$ds^2 = -\alpha^2 dt^2 + a^2 dr^2 + r^2 d\Omega^2 \quad (2.2.69)$$

with the metric given by

$$g_{\mu\nu} = \text{diag} [-\alpha^2, a^2, r^2, r^2 \sin^2 \theta] \quad (2.2.70)$$

which has contravariant form

$$g^{\mu\nu} = \text{diag} [-\alpha^{-2}, a^{-2}, r^{-2}, r^{-2} \sin^{-2} \theta] . \quad (2.2.71)$$

For this coordinate choice, we have normal vector n^μ , (2.2.32), and its dual n_μ , (2.2.33), given by

$$n^\mu = \left(\frac{1}{\alpha}, 0, 0, 0 \right)^T, \quad n_\mu = (-\alpha, 0, 0, 0) . \quad (2.2.72)$$

In this section we will outline the derivation of the fluid evolution equations and relevant constraint equations in polar-areal coordinates. Here we use the form of [26], i.e. the evolution equations are obtained from (2.2.64). An alternative derivation of the equations directly from the conservation laws is shown in appendix A.2.

One useful quantity we can define is the *mass aspect function*, $m(t, r)$ which can be interpreted as the gravitational mass within a radius r . We obtain this function by comparison of the polar-areal line element (2.2.69) with the vacuum Schwarzschild line element

$$ds^2 = - \left(1 - \frac{2M}{r} \right) dt^2 + \left(1 - \frac{2M}{r} \right)^{-1} dr^2 + r^2 d\Omega^2 \quad (2.2.73)$$

where M is the (constant) gravitational mass of a body. Outside a star, the polar-areal and Schwarzschild metrics are equivalent, as stated by Birkhoff's

theorem ([34]), hence we have

$$a^2(r) = \left(1 - \frac{2M}{r}\right)^{-1} \quad (2.2.74)$$

This leads naturally to the definition of the mass aspect function within the star,

$$m(t, r) \equiv \frac{r}{2} \left(1 - \frac{1}{a(t, r)^2}\right). \quad (2.2.75)$$

In section 2.6 we will justify why this function can be interpreted as the gravitational mass within r .

We now need to compute the components of the perfect fluid stress-energy tensor, $T^{\mu\nu} = \rho_0 h u^\mu u^\nu + p g^{\mu\nu}$, in polar-areal coordinates. We first need the components of the 4-velocity u^μ which are given by (2.2.41), $u^t = W/\alpha$, and (2.2.39), $u^i = u^t (\alpha v^i - \beta^i)$, hence we have

$$u^\mu = W(\alpha^{-1}, v^r, 0, 0). \quad (2.2.76)$$

The non-vanishing components of the stress-energy tensor in polar areal coordinates are then given by

$$\begin{aligned} T^{tt} &= \frac{\rho_0 h W^2}{\alpha^2} - \frac{p}{\alpha^2} &= \frac{E}{\alpha^2}, \\ T^{tr} &= \frac{\rho_0 h W^2 v^r}{\alpha} &= \frac{S^r}{\alpha}, \\ T^{rr} &= \rho_0 h W^2 v^r v^r + \frac{p}{a^2} &= S^r v^r + \frac{p}{a^2}, \\ T^{\theta\theta} &= \frac{p}{r^2}, \\ T^{\phi\phi} &= \frac{p}{r^2 \sin^2 \theta}. \end{aligned} \quad (2.2.77)$$

The use of the perfect fluid stress energy tensor also allows us to write the constraint and evolution equations for the spacetime variables α and a explicitly. To achieve this it is useful to consider how the contractions of the stress-energy tensor, given in (2.2.29), relate to the conserved variables. The local energy density, (2.2.29a), is

$$\varrho = T_{\hat{n}\hat{n}} = T_{\mu\nu} n^\mu n^\nu = T_{tt} n^t n^t = E \quad (2.2.78)$$

where we have used the fact that the only non-zero component of the normal

vector is n^t . The momentum density, (2.2.29b) can be written

$$\begin{aligned}
 j^\mu &= \perp T^{\mu\hat{n}} = -\perp (T^{\mu\nu} n_\nu) \\
 &= -(\delta_\alpha^\mu + n_\alpha n^\mu) T^{\alpha\nu} n_\nu \\
 &= -(\delta_\alpha^\mu + n_\alpha n^\mu) T^{\alpha t} n_t \\
 &= -(\delta_t^\mu + n_t n^\mu) T^{tt} n_t - (\delta_r^\mu + n_r n^\mu) T^{rt} n_t.
 \end{aligned} \tag{2.2.79}$$

The only non-vanishing component of the momentum density will then be

$$j^r = T^{rt} n_t = S^r \tag{2.2.80}$$

with covariant form

$$j_r = S_r. \tag{2.2.81}$$

Spacetime equations

We then write the evolution and constraint equations obtained by making the 3+1 split of the Einstein equations in polar-areal coordinates. From the Hamiltonian constraint, (2.2.34),

$$\mathcal{R} + K^2 - K^\mu_\nu K^\nu_\mu = 16\pi\varrho, \tag{2.2.82}$$

we obtain a constraint equation for the spacetime variable a (see appendix A.2.2 for this derivation),

$$\frac{\partial_r a}{a} = a^2 \left[4\pi r (\tau + D) - \frac{m}{r^2} \right]. \tag{2.2.83}$$

From the momentum constraint, (2.2.35),

$$D_\nu K^{\mu\nu} - D^\mu K = 8\pi j^\mu. \tag{2.2.84}$$

we can also derive an evolution equation for a ,

$$\partial_t a = -4\pi r \alpha a S_r. \tag{2.2.85}$$

as shown in appendix A.2.3. Finally, from the evolution equation for the extrinsic curvature, (A.1.112),

$$\begin{aligned} \mathcal{L}_t K_{\mu\nu} = & \mathcal{L}_\beta K_{\mu\nu} - D_\mu D_\nu \alpha + \alpha \left\{ \mathcal{R}_{\mu\nu} + K K_{\mu\nu} \right. \\ & \left. - 2K_{\mu\alpha} K^\alpha_\nu - 8\pi \left[\mathfrak{E}_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} (\mathfrak{E} - \varrho) \right] \right\}. \end{aligned} \quad (2.2.86)$$

we are able to derive a constraint equation for the lapse, α . In particular, this is obtained from the $K_{\theta\theta}$ component which, in appendix A.2.4 we find the result

$$\frac{\partial_r \alpha}{\alpha} = a^2 \left[4\pi r (S_r v^r + p) + \frac{m}{r^2} \right]. \quad (2.2.87)$$

This has then coupled the two spacetime variables a and α , and therefore the evolution of α can be described through this constraint equation and (2.2.83). The Hamiltonian constraint, (2.2.85), is therefore not required to solve for spacetime variable, which, as described in section 3.3.1, allows it to be used to check the accuracy of a NR simulation.

Fluid evolution equations

The evolution of the rest mass density in the Eulerian frame, D , is derived, as in SR, from the conservation of mass (2.2.14),

$$\nabla_\mu (\rho_0 u^\mu) = 0. \quad (2.2.88)$$

In appendix A.2.5 we show that this leads to the evolution equation

$$\partial_t (aD) + \frac{1}{r^2} \partial_r (r^2 \alpha a D v^r) = 0. \quad (2.2.89)$$

To obtain an equation for the total energy density E we consider the t component for the conservation of stress-energy (2.2.15),

$$\nabla_\mu T^{\mu t} = 0. \quad (2.2.90)$$

As in SR, we use the variable $\tau \equiv E - D$, and the evolution equation we

obtain, as derived in appendix A.2.6, is

$$\partial_t (a\tau) + \frac{1}{r^2} \partial_r [r^2 \alpha a (\tau + p) v^r] = -\frac{\alpha a m S_r}{r^2}. \quad (2.2.91)$$

Finally the evolution of the single non-vanishing momentum component, S^r , is obtained from the spatial components of (2.2.15),

$$\nabla_\mu T^{\mu r} = 0. \quad (2.2.92)$$

This results in

$$\partial_t (aS_r) + \frac{1}{r^2} \partial_r [r^2 \alpha a (S_r v^r + p)] = \alpha a \left[-\frac{a^2 m}{r^2} (S_r v^r + p + \tau + D) + \frac{2p}{r} \right] \quad (2.2.93)$$

which is derived in appendix A.2.7.

We now have balance-law equations for the conserved variables, (2.2.89), (2.2.93) and (2.2.91). As with the Newtonian and SR equations, we typically write these in vector form

$$\begin{aligned} \partial_t \left[a \begin{pmatrix} D \\ S_r \\ \tau \end{pmatrix} \right] + \frac{1}{r^2} \partial_r \left[r^2 \alpha a^2 \begin{pmatrix} Dv^r \\ S_r v^r + p \\ (\tau + p) v^r \end{pmatrix} \right] = \\ \alpha a \begin{bmatrix} 0 \\ -\frac{a^2 m}{r^2} (S_r v^r + p + \tau + D) + \frac{2p}{r} \\ -\frac{S_r m}{r^2} \end{bmatrix}. \end{aligned} \quad (2.2.94)$$

This equation can be rearranged to reduce problems with the coordinate singularity at the origin, in a similar manner to that used for the Newtonian fluid evolution equations in spherical symmetry in section 2.1.1. The pressure term in the momentum equation, (2.2.93), can be rearranged to give

$$\frac{1}{r^2} \partial_r (r^2 \alpha a p) = \partial_r (\alpha a p) + \frac{2\alpha a p}{r}. \quad (2.2.95)$$

Using this and by removing $1/r^2$ terms with (2.1.21),

$$\frac{1}{r^2} \frac{d}{dr} = 3 \frac{d}{d(r^3)}, \quad (2.2.96)$$

means we can write the system of equations, in vector form, as

$$\partial_t (a\mathbf{q}) + 3\partial_{r^3} (r^2 \alpha a \mathbf{f}^{(1)}) + \partial_r (\alpha a \mathbf{f}^{(2)}) = \alpha a \mathbf{s}. \quad (2.2.97)$$

Here we have

$$\mathbf{q} = \begin{pmatrix} D \\ S_r \\ \tau \end{pmatrix}, \quad (2.2.98)$$

$$\mathbf{f}^{(1)} = \begin{pmatrix} Dv^r \\ S_r v^r \\ (\tau + p) v^r \end{pmatrix}, \quad (2.2.99)$$

$$\mathbf{f}^{(2)} = \begin{pmatrix} 0 \\ p \\ 0 \end{pmatrix} \quad (2.2.100)$$

and

$$\mathbf{s} = \begin{pmatrix} 0 \\ -\frac{a^2 m}{r^2} (S_r v^r + p + \tau + D) \\ -\frac{S_r m}{r^2} \end{pmatrix}. \quad (2.2.101)$$

It is instantly apparent that these equations have a similar form to those for a Newtonian spherically symmetric fluid, (2.1.24). We can also see the coupling of the matter to the spacetime through the variables a and α , introduced through relativistic formulation of these equations. When solving these equations, we must therefore ensure that the equations governing the spacetime (e.g. (2.2.85) and (2.2.87)) are solved simultaneously.

2.3 Equation of state

In section 2.1, we demonstrated the necessity for an additional equation, the EOS, to close the system of evolution equations. This is an equation of the form

$$p = p(\rho_0, \varepsilon). \quad (2.3.1)$$

The EOS gives specific information about the type of fluid being modelled. As a result the choice of EOS will always be model dependent, and will be

based both on the desire to incorporate an accurate description of the fluid and the ease with which it can be implemented for the problem in question (be it analytical or a numerical). Ideally the EOS will be derived from the microphysical properties of the fluid to offer a global description of the fluid behaviour.

In this work one requirement of our choice of EOS is that it must provide a suitable description of matter within a NS. One fairly simple EOS that achieves this, whilst also offering a reasonable description of many common fluids, e.g. air or helium, is the ideal fluid EOS [9]. This can be derived from the standard relationship,

$$pV = nkT, \quad (2.3.2)$$

where V is the volume of the fluid, n the number of particles in this volume, k the Boltzmann constant and T the temperature of the fluid. The specific heat of the fluid is a quantity that can be defined as the amount of heat, Q , per unit mass required to raise T by one degree. This quantity does in fact have two possible values, which depend upon whether it is calculated whilst either V or p are kept constant. These two values give the specific heat at constant volume,

$$c_V = \frac{1}{M} \left(\frac{dQ}{dT} \right)_{V=\text{const.}} \quad (2.3.3)$$

and the specific heat at constant pressure

$$c_p = \frac{1}{M} \left(\frac{dQ}{dT} \right)_{p=\text{const.}} \quad (2.3.4)$$

where M is the total mass of the fluid within V . Assuming all particles are of equal mass, the total mass is given by $M = nm$, with m the mass of the individual particles. It is assumed that there is no mass transfer into or out of V , therefore M is a constant.

The first law of thermodynamics states

$$dQ = dU + p dV, \quad (2.3.5)$$

where U is the internal energy, related to the specific internal energy by $U = nm\varepsilon$. If we consider the case where V is constant, then (2.3.5) reduces to

$dQ = dU$, and we can therefore write (2.3.3) as

$$dU = Mc_V dT. \quad (2.3.6)$$

If it is assumed that c_V is constant, then this can be integrated to give

$$U = Mc_V T. \quad (2.3.7)$$

Alternatively, at constant pressure, we can use (2.3.2) to obtain the relationship

$$pdV = nk dT. \quad (2.3.8)$$

The two results, (2.3.6) and (2.3.8) allow the first law of thermodynamics, (2.3.5) to be written

$$dQ = Mc_V dT + nk dT. \quad (2.3.9)$$

Using this relationship in (2.3.4) we can now write c_p as

$$c_p = \frac{Mc_V dT + nk dT}{M dT} = c_V + \frac{k}{m} \quad (2.3.10)$$

which, when rearranged, gives

$$\frac{k}{m} = c_p - c_V. \quad (2.3.11)$$

We then use (2.3.7) to remove the explicit dependence on T from the EOS (2.3.2), which, when rearranged, gives

$$p = \frac{nkU}{Mc_V V} \quad (2.3.12)$$

which can then be written in terms of the individual particle mass m ,

$$p = \frac{kU}{mc_V V}. \quad (2.3.13)$$

We now define the adiabatic index, γ , as the ratio of the specific heats,

$$\gamma = \frac{c_p}{c_V} \quad (2.3.14)$$

which, using (2.3.11) allows this to be written

$$p = \frac{(\gamma - 1) U}{V}. \quad (2.3.15)$$

Finally, the definition $U = nm\varepsilon$, along with the relationship $\rho_0 = nm/V$ yields the EOS in the form,

$$p = (\gamma - 1) \rho_0 \varepsilon. \quad (2.3.16)$$

This provides a relationship that allows us to close the systems of evolution equations introduced so far in this chapter.

From the EOS we can define the specific entropy of the system, s , in terms of the fluid variables (see e.g. [116]). For an ideal fluid we have

$$s = c_V \ln p - c_p \ln \rho_0 + \text{const.} \quad (2.3.17)$$

It is useful to consider the case when entropy is constant (we shall see later some practical cases in which this occurs). From (2.3.17), we can see clearly that this implies

$$p = K \rho_0^\gamma \quad (2.3.18)$$

where K is a constant. This relationship is sometimes known as the polytropic EOS (a special case of the ideal fluid EOS), and K is referred to as the polytropic constant.

The ideal fluid EOS is described as a fairly simple example of an EOS since it uses a single parameter, γ , to differentiate between fluid types. To achieve this we must make many assumptions about the fluid being modelled such as taking constant values for the specific heats c_V and c_p . This has the advantage that despite this simplicity, it can offer a reasonable description of many common fluids (e.g. air or helium), and, of particular interest here, the $\gamma = 2$ case offers a surprisingly good description for a NS interior (based on e.g. mass-radius profiles and interior dynamics using this EOS) [177].

By introducing more parameters into the choice of EOS we hope to encompass a greater range of microphysical behaviour of the fluid. These parameters will often become specific to the fluid being modelled, for example

in a NS we may wish to include nuclear interactions between particles, which are negligible in air or helium.

For NSs, several more ‘realistic’ EOSs are available, the most commonly used are those by Shen et al. ([178]) and Lattimer and Swesty ([118]). These are specific EOSs for dense matter such as that in NSs, and parametrise the fluid through electron fraction, Y_e , and more realistic dependence on temperature, T , in addition to density. Using these EOSs can increase the accuracy of a NS model, however, in general, their more complex nature of means that there is not a simple analytic form for the EOS. Instead the fluid properties are prescribed through tabulated data. When making a numerical simulation of e.g. a NS, it is repeatedly necessary to convert between primitive variables (ρ_0 , v^i , ε and p) and conserved variables (D , S^i and τ for a relativistic fluid). Making this conversion with tabulated data is computationally expensive. Since the tabulated data cannot fully represent the continuum of the fluid, extrapolation of the physical quantities is required. As a result for many purposes, such as testing numerical techniques (such as those in this thesis), there is a clear advantage to using an EOS with a simple analytic form.

2.4 The Riemann problem

The motivation for considering the Riemann problem is twofold. Firstly, when developing numerical techniques for solving the fluid evolution equations, test cases with exact solutions are important for ensuring that these developments have been implemented correctly. For Newtonian and special relativistic fluids, the Riemann problem is one such test case. Secondly, and perhaps more importantly, many of the numerical techniques used for solving these equations rely on repeatedly solving the Riemann problem. The specific role of the problem in the numerical methods used in this thesis is described in detail in section 3.4. As a result, a detailed understanding of the Riemann problem is essential.

The Riemann problem is a class of 1+1 dimensional initial value problems for a system of evolution equations for any state vector \mathbf{u} . The initial conditions, if coordinates are chosen such that the problem is centred on $x = x_0$ at

$t = t_0$, are given by

$$\mathbf{u}(t_0, x) = \begin{cases} \mathbf{u}_L, & x < x_0 \\ \mathbf{u}_R, & x > x_0 \end{cases} \quad (2.4.1)$$

where \mathbf{u}_L and \mathbf{u}_R are constant vectors.

The solution of the Riemann problem for the fluid evolution equations is *self-similar*. A self-similar problem can be written in terms of a single independent variable ξ defined as a function of the coordinate variables, e.g. $\xi = f(t, x)$ for a 1+1 dimensional self-similar solution. For the Riemann problem with initial data (2.4.1), we have

$$\xi = \frac{x - x_0}{t - t_0}. \quad (2.4.2)$$

For any initial data \mathbf{u}_L and \mathbf{u}_R , then we have some information as to the structure of the solution to the Riemann problem. If our state vector \mathbf{u} has N variables, then the general solution will consist of N waves which separate $N + 1$ constant states. Therefore for the 1+1 dimensional fluid evolution equations, we would expect three waves which separate four constant states. In this solution, each wave will then have some associated velocity. There are three possible wave types that can occur. The waves with the largest (most positive) and smallest (most negative) velocities are determined by regions where characteristics either converge or diverge. These two waves are denoted \mathcal{W}_\rightarrow , \mathcal{W}_\leftarrow respectively. Diverging characteristics at either of these waves will lead to a rarefaction, whilst converging characteristics result in a shock wave. The characteristics at the remaining wave are parallel, and here we have a contact discontinuity, denoted \mathcal{C} . All three waves are self-similar: i.e. they follow a line $\xi = \text{const}$. The constant state to the left of \mathcal{W}_\leftarrow , and to the right of \mathcal{W}_\rightarrow , will be the initial states \mathbf{q}_L and \mathbf{q}_R respectively. The ‘intermediate’ constant states, between \mathcal{W}_\leftarrow and \mathcal{C} , and \mathcal{C} and \mathcal{W}_\rightarrow , are then denoted \mathbf{q}_{L*} and \mathbf{q}_{R*} respectively. An example of a solution for the Riemann problem in which all three wave types are formed is shown in figure 2.2.

The waveform of \mathcal{W}_\leftarrow is determined by

$$\mathcal{W}_\leftarrow = \begin{cases} \mathcal{R}_\leftarrow, & p_{L*} < p_L \\ \mathcal{S}_\leftarrow, & p_{L*} > p_L \end{cases} \quad (2.4.3)$$




Figure 2.2: One possible solution to the Riemann problem showing a rarefaction wave between \mathbf{q}_L and \mathbf{q}_{L*} , a contact discontinuity separating the resultant states \mathbf{q}_{L*} and \mathbf{q}_{R*} , and a shock wave separating \mathbf{q}_{R*} and \mathbf{q}_R . Since the solution to the Riemann problem is self-similar, all features propagate along lines $\xi = x/t = \text{constant}$.

where \mathcal{R}_\leftarrow is a rarefaction, and \mathcal{S}_\leftarrow a shock. Similarly, \mathcal{W}_\rightarrow can be determined by substituting R for L in the conditions given above.

Here we first consider the solution to the Riemann problem in Newtonian physics, as a reference solution. This allows for a greater understanding of the form of the solution, such that when we consider the SR problem, we have intuition for the more complex form of the solution in this case. The Riemann problem is set up in the absence of gravitational effects (a constant gravitational field is assumed). It is, however, still used in the numerical techniques for simulations of e.g. stars. In these techniques, the Riemann problem is solved many times, in each case over a small domain. Therefore in these regions, we can consider the spacetime to be locally flat, and hence the techniques remain valid [158].

In this section, many features will have a sign dependence depending on the direction of movement of the wave being considered, i.e. if they are associated with \mathcal{W}_\leftarrow or \mathcal{W}_\rightarrow . Unless otherwise stated, the positive case corresponds to a wave moving in the positive x -direction, and negative case to a wave moving in the negative x -direction.




Figure 2.3: Illustrating the characteristic form of a rarefaction fan. We see that the solution is continuous across the fan.

2.4.1 Wave types in the Newtonian Riemann problem

To calculate the solution to the Riemann problem, it is first useful to consider each type of wave individually. The results in this section are derived in e.g. [116, 201]. We consider each wave in terms of the two constant states either side of the wave. We denote these in terms of the state ahead of the wave \mathbf{q}_a and behind the wave \mathbf{q}_b . Therefore for a wave moving in the negative x -direction we have $a = L$, $b = L_*$, and if it is moving in the positive x -direction $a = R$, $b = R_*$. In general the waves resulting from the Newtonian Riemann problem are described in terms of the state ahead of the wave, \mathbf{q}_a , and pressure behind the wave, p_b .

Rarefactions

Here we consider a situation in which we have a rarefaction wave between the states \mathbf{q}_a and \mathbf{q}_b . At a rarefaction, characteristic lines diverge and there is more than one mathematically viable solution describing the characteristics within the wave. The wave illustrated in figure 2.2 is a *rarefaction fan*, across which the solution is continuous, the full characteristic structure around this wave is shown in figure 2.3.

Alternatively the wave could be described by a *rarefaction shock*, the characteristics of which are illustrated in figure 2.4. This solution, although mathematically allowed, is unstable (see e.g. [201]), with small perturbations leading to major changes in the solution. Therefore the rarefaction fan shown in figure 2.2 is indeed the physical solution when characteristics diverge.




Figure 2.4: Illustrating the characteristic form of a rarefaction shock. We see the characteristics remain parallel, and 'emerge' from the single self-similar line, the rarefaction shock.

A rarefaction is bounded by the characteristics

$$\xi = v_a \pm c_{s,a}, \quad \xi = v_b \pm c_{s,b}, \quad (2.4.4)$$

where c_s is the speed of sound, given by

$$c_s^2 = \frac{\gamma p}{\rho_0}. \quad (2.4.5)$$

To obtain these characteristics, we first write the evolution equations in the form

$$\mathbf{C} \partial_\xi \mathbf{w} = \mathbf{0} \quad (2.4.6)$$

where we have the matrix $\mathbf{C} = \mathbf{C}(\xi, \mathbf{w})$ and \mathbf{w} , as previously defined, is the vector of the primitive variables. The characteristics are then given by the requirements on ξ such that non-trivial solutions exist for this system of equations, i.e. $\det(\mathbf{C}) = 0$ ([94]).

Across a rarefaction, the Riemann Invariant

$$v \mp \frac{2c_s}{\gamma - 1} \quad (2.4.7)$$

is constant. By equating the invariant quantity at the boundaries of the rarefaction we get

$$v_b = v_a \pm \frac{2}{\gamma - 1} (c_{s,b} - c_{s,a}). \quad (2.4.8)$$

One feature of a rarefaction is that it is adiabatic - there is no entropy change

across it. In section 2.3 we saw that for an ideal fluid with constant entropy we have the relationship ((2.3.18))

$$p \propto \rho_0^\gamma. \quad (2.4.9)$$

Using the definition of the speed of sound can then yield another relationship,

$$c_s \propto \rho_0^{(\gamma-1)/2}. \quad (2.4.10)$$

Assuming knowledge of the speed of sound at the boundaries of the rarefaction we get

$$\rho_{0,b} = \rho_{0,a} \left(\frac{c_{s,b}}{c_{s,a}} \right)^{2/(\gamma-1)} \quad (2.4.11)$$

and

$$p_b = p_a \left(\frac{c_{s,b}}{c_{s,a}} \right)^{2\gamma/(\gamma-1)}. \quad (2.4.12)$$

By using the latter of these relationships in (2.4.8) to eliminate $c_{s,b}$ we get

$$v_b = v_a \pm \frac{2c_{s,a}}{\gamma - 1} \left[\left(\frac{p_b}{p_a} \right)^{(\gamma-1)/(2\gamma)} - 1 \right]. \quad (2.4.13)$$

We could obtain the obvious analogous result if we want to obtain v_a in terms of \mathbf{w}_b and p_a .

Shock waves

If a shock wave is present in the solution of the Riemann problem (as determined by the conditions given in (2.4.3)), then characteristic lines converge, shown in figure 2.5. This differs from the rarefaction wave case where characteristic lines diverged.

The discontinuity in the variables at a shock wave is governed by the Rankine-Hugoniot conditions,

$$\mathbf{f}_b - \mathbf{f}_a = v_s (\mathbf{q}_b - \mathbf{q}_a) \quad (2.4.14)$$

where v_s is the velocity of the shock and \mathbf{q} and \mathbf{f} are the state and flux




Figure 2.5: Illustrating the characteristic form of a shock wave. We see the characteristics merge along a self-similar line, forming the discontinuous shock wave.

vectors for the fluid evolution equations. It is this velocity that determines the characteristic line of the shock,

$$\xi = v_S. \quad (2.4.15)$$

The convergence of the characteristic lines results in a compression condition for v_S ,

$$v_a \pm c_{s,a} \geq v_S \geq v_b \pm c_{s,b}. \quad (2.4.16)$$

As with the rarefaction case, we assume \mathbf{w}_a and p_b are known, and want to calculate the remaining variables in \mathbf{w}_b . To achieve this we first consider the Rankine-Hugoniot conditions in the frame moving with v_S . The velocities of the two states in this frame are given by

$$\tilde{v}_{a(b)} = v_{a(b)} - v_S. \quad (2.4.17)$$

In this frame, the Rankine-Hugoniot conditions reduce to $\mathbf{f}_a = \mathbf{f}_b$ which, written explicitly, are

$$\rho_{0,a} \tilde{v}_a = \rho_{0,b} \tilde{v}_b, \quad (2.4.18a)$$

$$\rho_{0,a} \tilde{v}_a^2 + p_a = \rho_{0,b} \tilde{v}_b^2 + p_b, \quad (2.4.18b)$$

$$\tilde{v}_a (\tilde{E}_a + p_a) = \tilde{v}_b (\tilde{E}_b + p_b). \quad (2.4.18c)$$

We can then introduce the *mass flux* Q_a as the quantity given by (2.4.18a),

i.e.

$$Q_a = \mp \rho_{0,a} \tilde{v}_a = \mp \rho_{0,b} \tilde{v}_b. \quad (2.4.19)$$

Using (2.4.18b) and converting back to the Eulerian frame, the v_s component of the mass flux can be removed, giving

$$Q_a = \pm \frac{p_b - p_a}{v_b - v_a}. \quad (2.4.20)$$

Rearranging this gives a relationship for v_b ,

$$v_b = v_a \pm \frac{p_b - p_a}{Q_a}. \quad (2.4.21)$$

Alternatively we can use (2.4.18a) to obtain a relationship for Q_a which does not depend on v_b ,

$$Q_a^2 = \pm \frac{p_b - p_a}{\frac{1}{\rho_{0,b}} - \frac{1}{\rho_{0,a}}}. \quad (2.4.22)$$

Using the final Rankine-Hugoniot condition (2.4.18c) and the ideal fluid EOS we can obtain

$$\rho_{0,b} = \rho_{0,a} \left[\frac{\frac{\gamma-1}{\gamma+1} + \frac{p_b}{p_a}}{\left(\frac{\gamma-1}{\gamma+1} \right) \left(\frac{p_b}{p_a} \right) + 1} \right] \quad (2.4.23)$$

and eventually, using this in (2.4.22), get

$$Q_a = \left(\frac{p_b + B_a}{A_a} \right)^{1/2} \quad (2.4.24)$$

where

$$A_a = \frac{2}{(\gamma - 1) \rho_{0,a}} \quad (2.4.25)$$

and

$$B_a = \left(\frac{\gamma - 1}{\gamma + 1} \right) p_a. \quad (2.4.26)$$

This then allows us to obtain the post-shock velocity v_b by substituting this result for Q_a in (2.4.21),

$$v_b = v_a \pm \left(\frac{A_a}{p_a + B_a} \right)^{1/2} (p_b - p_a). \quad (2.4.27)$$

As with the rarefaction case, an analogous result could be obtained if we as-

sumed \mathbf{w}_b and p_a were known.

Contact discontinuities

Contact discontinuities are the final wave type that can result from the Riemann problem solution. At a contact discontinuity, characteristic lines are parallel. Here we find that pressure and (normal) velocity remain constant across the wave, whilst all other properties jump. Therefore for the 1+1 dimensional Riemann problem we have

$$v_{L*} = v_{R*} = v_*, \quad (2.4.28)$$

$$p_{L*} = p_{R*} = p_*, \quad (2.4.29)$$

and therefore the contact discontinuity occurs along the characteristic line

$$\xi = v_*. \quad (2.4.30)$$

This implies that a contact discontinuity moves with the velocity of the fluid, and hence there is no actual flow of fluid across the feature. As such, they can separate two distinct fluid regions, and the case in which the EOS changes across the contact discontinuity can be implemented with comparative ease.

2.4.2 Solving the Newtonian Riemann problem

We have stated that an exact solution to the Riemann problem can be found if we know the initial constant states \mathbf{q}_L and \mathbf{q}_R . By using the results for the wavetypes seen in the previous section, we can identify the variables in the intermediate states, \mathbf{q}_{L*} and \mathbf{q}_{R*} . There are many ways in which we can achieve this, but in general we must solve some implicit equation for one variable within the intermediate states, typically p_* . From this the remaining variables can be calculated. Here we follow the method given by [201], but make the extension to deal with a change in the EOS (a change in γ for an ideal fluid EOS) across the contact discontinuity.

In the previous section we derived results for the intermediate velocity (v_b , where we have either $b = L*$ or $b = R*$) for both rarefactions (2.4.13) and shock

waves (2.4.27). We also determined that $v_{L*} = v_{R*}$ (and $p_{L*} = p_{R*}$ (2.4.28)), hence these two results can be equated. We can therefore obtain the implicit equation for p_* ,

$$\mathcal{P}(p_*, \mathbf{w}_L, \mathbf{w}_R) \equiv \mathcal{P}_-(p_*, \mathbf{w}_L) + \mathcal{P}_-(p_*, \mathbf{w}_R) + \Delta v = 0 \quad (2.4.31)$$

where

$$\Delta v = v_R - v_L \quad (2.4.32)$$

and

$$\mathcal{P}_K(p_*, \mathbf{w}_K) = \begin{cases} (p_* - p_K) \left(\frac{A_K}{p_* + B_K} \right)^{1/2} & p_* > p_K \\ \frac{2c_{s,K}}{\gamma_K - 1} \left[\left(\frac{p_*}{p_K} \right)^{(\gamma_K - 1)/(2\gamma_K)} - 1 \right] & p_* < p_K \end{cases}. \quad (2.4.33)$$

where K represents either the left moving state (\mathcal{P}_- and \mathbf{w}_L) or right state (\mathcal{P}_+ and \mathbf{w}_R). The constants A_K and B_K are as given in (2.4.25) and (2.4.26), and for a multicomponent fluid are

$$A_K = \frac{2}{(\gamma_K - 1) \rho_{0,K}}, \quad B_K = \left(\frac{\gamma_K - 1}{\gamma_K + 1} \right) p_K. \quad (2.4.34)$$

The two cases of (2.4.33) are dependent on whether the wave in the direction K is a shock (i.e. $p_* > p_K$ or a rarefaction $p_* < p_K$, (2.4.3)).

We now have enough information to solve (2.4.31) iteratively, and obtain the intermediate pressure p_* . Furthermore, for the Newtonian Riemann problem, we do have information about what each wave type will be, and hence can ensure that our initial value for the iteration is reasonable. This is based on the pressures of the two initial constant states, and the velocity difference. We will assume, without loss of generality, that we have $p_L < p_R$. The wave type is determined by the sign of $\mathcal{P}(p_L, \mathbf{w}_L, \mathbf{w}_R)$ and $\mathcal{P}(p_R, \mathbf{w}_L, \mathbf{w}_R)$, which, for simplicity, are denoted \mathcal{P}_{p_L} and \mathcal{P}_{p_R} respectively. Note that in calculating e.g. \mathcal{P}_{p_L} , the $\mathcal{P}_-(p_L, \mathbf{w}_L)$ term will automatically vanish, but all the others can be non-zero.

The solution to the Riemann problem will be two rarefaction waves if

$$\mathcal{P}_{p_L} < 0, \quad \mathcal{P}_{p_R} < 0. \quad (2.4.35)$$

If we have

$$\mathcal{P}_{p_L} < 0, \quad \mathcal{P}_{p_R} > 0. \quad (2.4.36)$$

then we have a left-moving rarefaction and a right moving shock. Finally if we have

$$\mathcal{P}_{p_L} > 0, \quad \mathcal{P}_{p_R} > 0. \quad (2.4.37)$$

then we have two shocks.

Once we have calculated p_* , and in doing so determined the wavetypes within the solution, then obtaining the remaining variables is trivial. The intermediate velocity v_* can be calculated through either (2.4.13) or (2.4.27) and then the density in each intermediate state through (2.4.9) for a rarefaction or (2.4.23) for a shock. If a shock wave is present in the solution then we need to calculate the speed with which it moves. This is done through the Rankine-Hugoniot conditions (2.4.14), e.g. we can get

$$v_S = \frac{\rho_{0,b}v_b - \rho_{0,a}v_a}{\rho_{0,b} - \rho_{0,a}}. \quad (2.4.38)$$

If we have a rarefaction, then we can use the Riemann Invariant (2.4.7) to calculate the variables across the fan,

$$v_{\text{fan}} \mp \frac{2c_{s,\text{fan}}}{\gamma_a - 1} = v_a \mp \frac{2c_{s,a}}{\gamma_a - 1} = v_b \mp \frac{2c_{s,b}}{\gamma_a - 1} \quad (2.4.39)$$

where we note that γ takes the same value in both the initial and intermediate states. This gives us the velocity and speed of sound within the fan, and (2.4.11) and (2.4.12) then allow the density and pressure respectively to be calculated. The final information required for the full solution is the characteristic lines of the waves. These can be calculated using (2.4.4) for rarefactions, (2.4.15) for shocks and (2.4.30) for the contact discontinuity.

2.4.3 Wave types in the special relativistic Riemann problem

When we make the special relativistic extension of the Riemann problem, we find that qualitatively we obtain the same behaviours as experienced in the Newtonian case. That is the solution is still self-similar, with $\xi = x/t$ as in (2.4.2), and that there are three waves present in the solution, two of which can be either a rarefaction or a shock, and one a contact discontinuity. The solution to the Riemann problem in SR was found by Martí and Müller [133], and an extension to include non-zero tangential velocities in more than one dimension was made by Pons, Martí and Müller [159]. The case with non-zero tangential velocity is not a trivial extension in relativistic cases since the Lorentz factor $W^2 = 1/(1 - v^i v_i)$ couples all velocity components in addition to the coupling through the energy (which was also true in Newtonian physics). Due to this, we will first consider the 1+1 dimensional case, and then the multidimensional extension. As with the Newtonian case, we first consider each wave type individually, following [133].

Rarefactions

A rarefaction fan for the SR Riemann problem is bounded by the characteristics

$$\xi = \frac{v_a \pm c_{s,a}}{1 \pm v_a c_{s,a}}, \quad \xi = \frac{v_b \pm c_{s,b}}{1 \pm v_b c_{s,b}}, \quad (2.4.40)$$

As with the Newtonian case, we assume that we know the state ahead of the rarefaction wave, \mathbf{w}_a , and the pressure behind the wave, p_b . Again we wish to calculate the remaining primitive variables behind the rarefaction, \mathbf{w}_b . The velocity behind the rarefaction is given by

$$v_b = \frac{(1 - v_a) A_{\pm,b} - (1 - v_a)}{(1 - v_a) A_{\pm,b} + (1 - v_a)} \quad (2.4.41)$$

where $A_{\pm,b}$ is given by

$$A_{\pm,b} = \left\{ \left[\frac{(\gamma - 1)^{1/2} - c_{s,b}}{(\gamma - 1)^{1/2} + c_{s,b}} \right] \left[\frac{(\gamma - 1)^{1/2} + c_{s,a}}{(\gamma - 1)^{1/2} - c_{s,a}} \right] \right\}^{\pm 2(\gamma-1)^{-1/2}}. \quad (2.4.42)$$

The speed of sound of a relativistic fluid is given by

$$c_s^2 = \frac{\gamma p}{\rho_0 h}, \quad (2.4.43)$$

with h the specific enthalpy as defined by (2.2.3),

$$h = 1 + \varepsilon + \frac{p}{\rho_0}. \quad (2.4.44)$$

As with the Newtonian case, entropy is constant across a rarefaction, hence we can obtain the density through the relationship $p \propto \rho_0^\gamma$, hence again we have (2.4.9),

$$\rho_{0,b} = \rho_{0,a} \left(\frac{p_b}{p_a} \right)^{1/\gamma}. \quad (2.4.45)$$

This then gives us enough information to calculate v_b by expressing the speed of sound, and hence $A_{\pm,b}$, (2.4.42), in terms of p_b (and \mathbf{w}_a). A full solution will again require information about the fluid within the rarefaction fan, \mathbf{q}_{fan} . The characteristic lines within the rarefaction are given by

$$\xi = \frac{v_{\text{fan}} \pm c_{s,\text{fan}}}{1 \pm v_{\text{fan}} c_{s,\text{fan}}} \quad (2.4.46)$$

which can be rearranged to give an expression for $c_{s,\text{fan}}$,

$$c_{s,\text{fan}} = \pm \frac{v_{\text{fan}} - \xi}{1 - \xi v_{\text{fan}}}. \quad (2.4.47)$$

The velocity within the rarefaction is then given by (2.4.41),

$$v_{\text{fan}} = \frac{(1 - v_a) A_{\pm,\text{fan}} - (1 - v_a)}{(1 - v_a) A_{\pm,\text{fan}} + (1 - v_a)} \quad (2.4.48)$$

where $A_{\pm,\text{fan}}$ can be given by (2.4.42). Once we have v_{fan} , the remaining variables within the rarefaction fan can be found. From (2.4.47) we can determine $c_{s,\text{fan}}$, and then (2.4.43) and (2.4.45) allow p_{fan} and $\rho_{0,\text{fan}}$ to be obtained.

Shock waves

The jump in variables across a shock wave is governed by the relativistic Rankine-Hugoniot conditions ([195])

$$[\rho_0 u^\mu] n_\mu = 0, \quad [T^{\mu\nu}] n_\nu = 0, \quad (2.4.49)$$

where the notation

$$[F] = F_a - F_b \quad (2.4.50)$$

has been used, and n_μ is the unit normal to a spacetime slice Σ_τ as described in section 2.2. In the frame with which a shock has velocity v_S then this is given by

$$n_\mu = W_S (-v_S, 1, 0, 0), \quad (2.4.51)$$

where W_S is the Lorentz factor associated with the shock. As with the rarefaction case, we assume knowledge of the pre-shock region \mathbf{q}_a and the post shock pressure p_b . The post shock velocity, v_b , is given by

$$v_b = \frac{\frac{W_S}{j} (p_b - p_a) + h_a W_a v_a}{h_a W_a + (p_b - p_a) \left(\frac{W_S v_a}{j} + \frac{1}{W_a \rho_{0,a}} \right)}, \quad (2.4.52)$$

where j is the mass flux across the shock. The mass flux can again be obtained through the Rankine-Hugoniot conditions, and in relativistic situations is given by

$$j = \pm \sqrt{\frac{p_b - p_a}{\frac{h_a}{\rho_{0,a}} - \frac{h_b}{\rho_{0,b}}}}. \quad (2.4.53)$$

A relation for the post-shock enthalpy, h_b is given by the Taub adiabat ([196]),

$$[h^2] = \left(\frac{h_b}{\rho_{0,b}} + \frac{h_a}{\rho_{0,a}} \right) [p]. \quad (2.4.54)$$

Using the EOS to express $\rho_{0,a}$ in terms of p_a and rearranging gives a quadratic equation for h_b ,

$$\left[1 + \frac{\gamma - 1}{p_b \gamma} (p_a - p_b) \right] h_b^2 - \left[\frac{\gamma - 1}{p_b \gamma} (p_a - p_b) \right] h_b + \frac{h_a}{\rho_{0,a}} (p_a - p_b) - h_a^2 = 0. \quad (2.4.55)$$

Solving this will yield one positive (and hence physical) value for h_b , and from this, $\rho_{0,b}$ and then j can be calculated.

To obtain v_b we then need the velocity of the shock, which, as in the Newtonian case, is obtained from the Rankine-Hugoniot conditions,

$$v_s = \frac{1}{j^2 + (\rho_{0,a} W_a)^2} \left[(\rho_{0,a} W_a)^2 v_a \pm j^2 \sqrt{1 + \left(\frac{\rho_{0,a}}{j} \right)^2} \right]. \quad (2.4.56)$$

This allows the Lorentz factor associated with the shock to be computed in the obvious manner, hence all information to calculate v_b is available.

Contact discontinuities

There is no difference in how a contact discontinuity behaves in 1+1 dimensional SR when compared to the Newtonian case. Therefore we still have

$$p_a = p_b = p_*, \quad v_a = v_b = v_* \quad (2.4.57)$$

and characteristic line

$$\xi = v_* \quad (2.4.58)$$

2.4.4 Solving the special relativistic Riemann problem

As with the Newtonian Riemann problem, in SR we obtain an implicit equation for the intermediate pressure p_* by equating the expressions for the intermediate velocities for each wave, i.e. we have

$$\tilde{\mathcal{P}}_{\leftarrow}(p_*, \mathbf{w}_L) - \tilde{\mathcal{P}}_{\rightarrow}(p_*, \mathbf{w}_R) = 0. \quad (2.4.59)$$

Note that there is now no explicit Δv term for SR, the velocity terms are incorporated in the $\tilde{\mathcal{P}}$ terms. These functions are then given explicitly by

$$\tilde{\mathcal{P}}_K(p_*, \mathbf{w}_K) = \begin{cases} \frac{\frac{W_{S,K}}{j_K}(p_* - p_K) + h_K W_K v_K}{h_K W_K + (p_* - p_K) \left(\frac{W_{S,K} v_K}{j_K} + \frac{1}{W_K \rho_{0,K}} \right)} & p_* > p_K \\ \frac{(1 - v_K) A_{K*} - (1 - v_K)}{(1 - v_K) A_{K*} + (1 - v_K)} & p_* < p_K \end{cases} \quad (2.4.60)$$

where again K represents either the left moving state (\mathcal{P}_\leftarrow and \mathbf{w}_L) or right state (\mathcal{P}_\rightarrow and \mathbf{w}_R). Here we have defined A_{K*} to be either $A_{-,L*}$ or $A_{+,R*}$ for left or right moving rarefactions respectively, and that j_K has a sign dependence shown in (2.4.53). Also note that if the EOS changes across the contact discontinuity then $W_{S,K}$, j_K , h_K and A_{k*} are all dependent on the value of γ .

To solve this implicit equation, we can select an initial guess for the intermediate pressure, \tilde{p}_* , and iteratively solve (2.4.59). At each iteration we must determine whether the left and right moving waves need to be solved as shocks or rarefactions for the current guess \tilde{p}_* .

It is possible to determine the expected wave types *a priori* in a similar manner to the Newtonian problem. This is done by comparing the *special relativistic relative velocity*,

$$v_{LR} \equiv \frac{v_L - v_R}{1 - v_L v_R}, \quad (2.4.61)$$

to some relative limiting velocities for the different wavetypes. This procedure is detailed for single component fluids in [166], and it is also noted that this procedure can be extended to the situations considered here, where the EOS jumps across the contact discontinuity.

2.4.5 The special relativistic Riemann problem with non-zero tangential velocities

We will now consider the solution to the (1+1 dimensional) Riemann problem in which we have more than one spatial dimension, i.e. with the discontinuity in the initial data perpendicular to the x -axis. This case is non-trivial in SR if we allow for non-zero tangential velocities, since we have coupling of velocities through the Lorentz factor $W = 1/\sqrt{1 - v_i v^i}$. The results given here were

derived in [159]. When considering the different wave types, the relationships

$$hWv^y = \text{const.}, \quad hWv^z = \text{const.}, \quad (2.4.62)$$

always hold, which implies $v^y/v^z = \text{const}$, and hence means the direction of the tangential velocity does not change with time, only its magnitude does. Therefore the waves produced from the initial data will remain perpendicular to the x -axis, and hence we can still write the solution in terms of the similarity variable $\xi = x/t$.

Rarefactions

For a non-zero tangential velocity, the Riemann fan has characteristic lines

$$\xi = \frac{v^x (1 - c_s^2) \mp c_s \sqrt{(1 - v^2) [1 - v^2 c_s^2 - (v^x)^2 (1 - c_s^2)]}}{1 - v^2 c_s^2} \quad (2.4.63)$$

where the notation used throughout this section holds, i.e. minus or plus corresponds to \mathcal{R}_\rightarrow and \mathcal{R}_\leftarrow respectively.

As with the previous situations, velocity within the rarefaction fan can be given in terms of the pressure. In this case, however, we have an ordinary differential equation to solve,

$$\frac{dv^x}{dp} = \mp \frac{1}{\rho_0 h W^2 c_s} \frac{1}{\sqrt{1 + g(\xi_\mp, v^x, v^T)}} \quad (2.4.64)$$

where we have defined v^T to be the absolute value of the tangential velocity,

$$v^T = \sqrt{(v^y)^2 + (v^z)^2}, \quad (2.4.65)$$

and the function g is given by

$$g(\xi_\mp, v^x, v^T) = \frac{(v^T)^2 (\xi_\mp^2 - 1)}{(1 - \xi_\mp v^x)^2}. \quad (2.4.66)$$

When solving the Riemann problem for the non-zero tangential velocity case, we will follow the same general procedure as outlined for the 1+1 dimensional case given in 2.4.3. Therefore we assume knowledge of the fluid ahead

of the rarefaction, \mathbf{w}_a and the post-rarefaction pressure p_b . We then obtain an expression for the velocity behind the rarefaction, v_b^x in terms of \mathbf{w}_a and p_b . Using the condition $hWv^T = \text{const}$ we obtain

$$v_b^T = h_a W_a v_a^T \left[\frac{1 - (v_b^x)^2}{h_b^2 + (h_a W_a v_a^T)^2} \right]^{1/2}. \quad (2.4.67)$$

We can use the EOS and the fact that $p \propto \rho_0$ to express h in terms of p , and then use this to express c_s also in terms of p . This then allows the differential equation for v^x , (2.4.64), to be written in terms in terms of p only. The differential equation then needs to be solved across the rarefaction until we have v_b^x in terms of p_b .

Shock waves

The special relativistic Rankine-Hugoniot conditions given by (2.4.49) still hold for non-zero tangential velocity. The dependence of the tangential velocity when calculating the post-shock variables is not as obvious as it was in the rarefaction case. The velocity behind the shock, v_b^x is given by a similar equation to (2.4.52),

$$v_b^x = \frac{\frac{W_s}{j} (p_b - p_a) + h_a W_a v_a^x}{h_a W_a + (p_b - p_a) \left(\frac{W_s v_a^x}{j} + \frac{1}{W_a \rho_{0,a}} \right)}. \quad (2.4.68)$$

Here the tangential velocity components only enter through the Lorentz factors. We obtain the mass flux j and the post shock enthalpy h_b in the same way as the 1+1 dimensional case ((2.4.53) and (2.4.55) respectively), and the obvious extension for the shock velocity v_s applies,

$$v_s = \frac{1}{j^2 + (\rho_{0,a} W_a)^2} \left[(\rho_{0,a} W_a)^2 v_a^x \pm j^2 \sqrt{1 + \left(\frac{\rho_{0,a}}{j} \right)^2} \right]. \quad (2.4.69)$$

The density in the post shock region can then be calculated using the relativistic Rankine-Hugoniot conditions as shown previously. The tangential velocity

components are also given by the relativistic Rankine-Hugoniot conditions,

$$v_b^{y,z} = h_a W_a v_a^{y,z} \left[\frac{1 - (v_b^x)^2}{h_b^2 + (h_a W_a v_a^{y,z})^2} \right]^{1/2}. \quad (2.4.70)$$

Contact discontinuities

There is no difference in the treatment of contact discontinuities in this case compared to those previously considered. The pressure and normal velocity remain constant across the contact discontinuity, whilst density and tangential velocity components are not constrained.

2.5 Approximate Riemann solvers

In section 2.4 we motivated our interest in the Riemann problem by stating that many numerical techniques require repeatedly solving the Riemann problem. These techniques effectively require the solution at $x = 0$ at a time $t = \Delta t$ for initial data as given in (2.4.1) centred on $x = 0$, as described in section 3.4. The exact solutions to the Riemann problem (in both the Newtonian and special relativistic cases) all require the solution of an implicit equation, and in the case of SR with tangential velocities an ordinary differential equation (ODE) as well. Solving such equations is computationally expensive, hence repeatedly obtaining the exact solution when using such numerical techniques is undesirable. To address this issue, approximate Riemann solvers have been developed. These attempt to provide a close approximation to the intermediate states produced in the Riemann problem (\mathbf{q}_{L*} and \mathbf{q}_{R*}) using the information available analytically, i.e. \mathbf{q}_L and \mathbf{q}_R . The advantage these approximate solvers have over the analytic solutions is they are very much cheaper computationally, hence practical for repeatedly obtaining solutions. Many approximate Riemann solvers have been developed, here we shall describe those used in this work, see e.g. [201] for further examples. Typically these solvers have been developed for the Newtonian Riemann problem, and from these, extensions to the relativistic problem were then developed.

2.5.1 Newtonian HLLE solver

One of the simplest approximate solvers is the Harten-Lax-van-Leer-Einfeldt (HLLE) solver, developed by Einfeldt [61], based on the Harten-Lax-van-Leer (HLL) solver [91]. This solver greatly simplifies the expected solution from the Riemann problem, assuming there are just two waves formed. Therefore we have a single constant intermediate state $\mathbf{q}_{*,\text{HLLE}}$, separating the two waves. It is then assumed that each wave moves with a single velocity (i.e. behaves like a shock or contact discontinuity). These are the maximum or minimum wave velocities for the left or right moving wave moving respectively, which are denoted λ_{\pm} respectively. Despite this simplicity, the HLLE solver is, by construction, consistent with the conservation form of the evolution equations [91], the importance of which is discussed in section 3.3. The solution can then be expressed as

$$\hat{\mathbf{q}} = \begin{cases} \mathbf{q}_L, & \xi < \lambda_- \\ \hat{\mathbf{q}}_{*,\text{HLLE}}, & \lambda_- < \xi < \lambda_+ \\ \mathbf{q}_R, & \xi > \lambda_+ \end{cases} \quad (2.5.1)$$

where $\hat{\mathbf{q}}$ represents an approximate solution to the Riemann problem. Since the assumption is made that the waves have a single velocity, the approximate intermediate state can be calculated using the Rankine-Hugoniot conditions (2.4.14), which, for the two waves can be expressed as

$$\mathbf{f}_R - \hat{\mathbf{f}}_* = \lambda_+ (\mathbf{q}_R - \hat{\mathbf{q}}_*), \quad \hat{\mathbf{f}}_* - \mathbf{f}_L = \lambda_- (\hat{\mathbf{q}}_* - \mathbf{q}_L). \quad (2.5.2)$$

Rearranging these to remove \mathbf{f}_* gives the intermediate state as

$$\hat{\mathbf{q}}_{*,\text{HLLE}} = \frac{\lambda_+ \mathbf{q}_R - \lambda_- \mathbf{q}_L - \mathbf{f}(\mathbf{q}_R) + \mathbf{f}(\mathbf{q}_L)}{\lambda_+ - \lambda_-}. \quad (2.5.3)$$

The maximum and minimum wavespeeds then require some approximation since the explicit form of the two waves is not calculated. One suggestion, [51], is to use the characteristic speed of the leading front of a rarefaction, (2.4.4),

$$\lambda_- = v_L - c_{s,L}, \quad \lambda_+ = v_R + c_{s,R}. \quad (2.5.4)$$

Another suggestion for λ_{\pm} , [51, 201], which is more commonly used is to consider only the maximum (absolute) wavespeed, λ_{\max} ,

$$\lambda_{\max} = \max \{ |v_L| + c_{s,L}, |v_R| + c_{s,R} \} \quad (2.5.5)$$

and then set $\lambda_{\pm} = \pm \lambda_{\max}$.

When using an approximate Riemann solver in a numerical simulation, it is in fact the fluxes within the region $\hat{\mathbf{f}}_*$ that are required, not the conserved variables. These can be calculated from $\hat{\mathbf{q}}_*$, however it is simpler to calculate them directly from the Rankine-Hugoniot conditions, (2.5.2), giving

$$\hat{\mathbf{f}}_{*,\text{HLLE}} = \frac{\lambda_+ \mathbf{f}_L - \lambda_- \mathbf{f}_R + \lambda_+ \lambda_- (\mathbf{q}_R - \mathbf{q}_L)}{\lambda_+ - \lambda_-}. \quad (2.5.6)$$

It is clear that using the HLLE solver to approximate the entire solution to the Riemann problem will contain serious errors. When used in a numerical method, it is, however, only the solution at a single point that is required. In this case the accuracy to which this value is obtained is often sufficient for a numerical simulation to be able to reproduce the physical problem being modelled.

2.5.2 Relativistic HLLE solver

The extension of the HLLE method to relativistic situations (the relativistic HLLE, or RHLLE method) was first proposed by Schneider et al. [175]. Many of the techniques from the HLLE method carry over to the relativistic case. The intermediate state vector, $\hat{\mathbf{q}}_*$, and flux vector, $\hat{\mathbf{f}}_*$, are again obtained using (2.5.3) and (2.5.6) respectively. The maximum and minimum wavespeeds in relativity are given by those of a rarefaction in SR, (2.4.46), i.e.

$$\lambda_- = \frac{\tilde{v}_L - \tilde{c}_{s,L}}{1 - \tilde{v}_L \tilde{c}_{s,L}}, \quad \lambda_+ = \frac{\tilde{v}_R + \tilde{c}_{s,R}}{1 + \tilde{v}_R \tilde{c}_{s,R}}. \quad (2.5.7)$$

When solving the special relativistic fluid evolution equations using an approximate Riemann solver, two values are obtained for \mathbf{q}_L and \mathbf{q}_R . The tilde notation then represents the arithmetic mean of these two values. More details about how these left and right states are obtained is in section 3.4.

Working in relativity often allows a simplification for the RHLLE solver

however, since wave speeds are constrained by the speed of light. For SR and in many cases in GR this allows us to set $\lambda_{\pm} = \pm 1$. In some cases, gauge choices may result in $c \neq 1$, but in such cases, the value for the speed of light will still provide the appropriate wavespeeds.

2.5.3 Roe and Marquina solvers

The Roe solver, [168], and Marquina solver, [57], are two approximate Riemann solvers which are based on a linearisation of the fluid evolution equations. They are constructed by first writing the 1+1 dimensional fluid evolution equations as

$$\partial_t \mathbf{q} + \mathbf{A}(\mathbf{q}) \partial_x \mathbf{q} = 0. \quad (2.5.8)$$

where \mathbf{A} (by using the chain rule) is the Jacobian matrix,

$$\mathbf{A}(\mathbf{q}) = \frac{\partial \mathbf{f}}{\partial \mathbf{q}}. \quad (2.5.9)$$

The Roe solver then linearises this form of the evolution equations by replacing \mathbf{A} with a constant Jacobian matrix $\tilde{\mathbf{A}}(\tilde{\mathbf{q}}(\mathbf{q}_L, \mathbf{q}_R))$. The point at which the Jacobian is evaluated, $\tilde{\mathbf{q}}$, is known as the *Roe mean*, and must be chosen according to the problem being considered. In many cases is given simply by the arithmetic mean of the initial states,

$$\tilde{\mathbf{q}} = \frac{1}{2} (\mathbf{q}_L + \mathbf{q}_R). \quad (2.5.10)$$

To find the full solution for the Riemann problem, we then can solve the linear equation

$$\partial_t \mathbf{q} + \tilde{\mathbf{A}}(\tilde{\mathbf{q}}) \partial_x \mathbf{q} = 0. \quad (2.5.11)$$

To solve this, we must find the *wave strengths* $\tilde{\alpha}_i$ which are found through the projection of $\Delta \mathbf{q} \equiv \mathbf{q}_R - \mathbf{q}_L$ onto the right eigenvector of $\tilde{\mathbf{A}}$, \mathbf{r}^i . For a system of k variables, this is given by solving

$$\Delta \mathbf{q} = \sum_{i=1}^k \tilde{\alpha}_i \mathbf{r}^i. \quad (2.5.12)$$

For this linear system, the flux along the characteristic $\xi = 0$ can then be

given by

$$\hat{\mathbf{f}}_{*,\text{Roe}} = \frac{1}{2} \left(\mathbf{f}_L + \mathbf{f}_R - \sum_{i=1}^k |\lambda_i| \tilde{\alpha}_i \mathbf{r}^i \right) \quad (2.5.13)$$

where λ_i is the i^{th} eigenvalue.

The Roe solver will not always give the correct solution since (2.5.13) assumes that the characteristic $\xi = 0$ is within the intermediate state of the Riemann problem. The Marquina solver is an “extension” of the Roe solver which allows the solution along the characteristic $\xi = 0$ to be approximated when this characteristic is contained within a rarefaction. As with the Roe solver, the characteristic matrix and its eigenvalues and eigenvectors are required. We then define left and right characteristic variables and fluxes,

$$\omega_{L,R}^i = \mathbf{l}^i(\mathbf{q}_{L,R}) \cdot \mathbf{q}_{L,R}, \quad \varphi_{L,R}^i = \mathbf{l}^i(\mathbf{q}_{L,R}) \cdot \mathbf{f}_{L,R}, \quad (2.5.14)$$

where \mathbf{l}^i is the left eigenvector of $\tilde{\mathbf{A}}$. associated with λ_i . The Marquina flux formula then defines two fluxes, φ_{\pm}^i , by

$$\varphi_{+}^i = \begin{cases} \varphi_L^i, & \lambda_i(\mathbf{q}_L) \lambda_i(\mathbf{q}_R) > 0 \text{ and } \lambda_i(\mathbf{q}_L) > 0 \\ 0, & \lambda_i(\mathbf{q}_L) \lambda_i(\mathbf{q}_R) > 0 \text{ and } \lambda_i(\mathbf{q}_L) < 0 \\ \frac{1}{2} (\varphi_L^i + \alpha_i \omega_L^i), & \lambda_i(\mathbf{q}_L) \lambda_i(\mathbf{q}_R) < 0 \end{cases} \quad (2.5.15)$$

$$\varphi_{-}^i = \begin{cases} 0, & \lambda_i(\mathbf{q}_L) \lambda_i(\mathbf{q}_R) > 0 \text{ and } \lambda_i(\mathbf{q}_L) > 0 \\ \varphi_R^i, & \lambda_i(\mathbf{q}_L) \lambda_i(\mathbf{q}_R) > 0 \text{ and } \lambda_i(\mathbf{q}_L) < 0 \\ \frac{1}{2} (\varphi_R^i + \alpha_i \omega_R^i), & \lambda_i(\mathbf{q}_L) \lambda_i(\mathbf{q}_R) < 0 \end{cases} \quad (2.5.16)$$

where

$$\alpha_i = \max \{ |\lambda_i(\mathbf{q}_L)|, |\lambda_i(\mathbf{q}_R)| \}. \quad (2.5.17)$$

The flux along the characteristic line $\xi = 0$, which is what is required in a numerical simulation, is given by

$$\hat{\mathbf{f}}_{*,\text{Marq.}} = \sum_i [\varphi_{+}^i \mathbf{r}^i(\mathbf{q}_L) + \varphi_{-}^i \mathbf{r}^i(\mathbf{q}_R)]. \quad (2.5.18)$$

Note that the Marquina solvers does not provide any estimates for the intermediate states $\hat{\mathbf{q}}_*$, it simply provides the approximate flux $\hat{\mathbf{f}}_*$ along the

characteristic line $\xi = 0$. This solver therefore has been designed specifically for use in numerical methods as described in section 3.4.

2.6 The TOV equations

When undertaking a numerical simulation of the evolution of a star, or other self-gravitating body, it is essential to know whether the initial data for the system is stable. That is whether the body will, over some timescale, achieve hydrostatic equilibrium. The Tolman-Oppenheimer-Volkoff (TOV) equations describe the structure of a static spherically symmetric self-gravitating perfect fluid body in GR, and hence allow for initial data for a star in hydrostatic equilibrium to be produced for 1+1 dimensional GR. As a result they are an important resource for providing simple test cases for the numerical techniques we implement in this thesis, i.e. they are a case in GR in which an exact solution is known. They can also be used to provide static initial data upon which a (non-linear) perturbation can be added for dynamic tests of numerical techniques. The TOV equations are a coupled system of first-order ODEs, discovered by Tolman, Oppenheimer and Volkoff in 1939 [200, 147]. They can be derived by considering the Einstein equations and conservation equations for a static spherically symmetric metric, i.e. we have line element

$$ds^2 = -\alpha^2 dt^2 + a^2 dr^2 + r^2 d\Omega^2 \quad (2.6.1)$$

where α and a are now functions of r only. The metric associated with this line element is, however, just a special case of the spherically symmetric metric 1+1 dimensional we considered when deriving the evolution equations for GR in section 2.2.3. We can therefore obtain the TOV equations by considering the static case of the GR fluid evolution equations and constraint equations (i.e. all time derivatives and velocities vanish). Typically the TOV equations are written in terms of three variables, m , the mass aspect function, (2.2.75),

$$m(r) = \frac{r}{2} \left(1 - \frac{1}{a^2} \right), \quad (2.6.2)$$

the spacetime variable, $\Phi = \ln \alpha$, and the pressure, p . This choice of variables was made such that in the Newtonian limit, the TOV equations reduce to the

definitions of mass, the Poisson equation for gravitational potential and the equation for hydrostatic equilibrium (see e.g. [136]),

$$\frac{dm}{dr} = 4\pi r^2 \rho_0, \quad (2.6.3a)$$

$$\nabla^2 \Phi = 4\pi \rho_0, \quad (2.6.3b)$$

$$\frac{1}{r^2} \frac{d}{dr} \left(\frac{r^2}{\rho_0} \frac{dp}{dr} \right) = -4\pi \rho_0. \quad (2.6.3c)$$

By considering the Newtonian limits, we see why the mass aspect function can be interpreted as the gravitational mass within the radius r , and similarly Φ can be thought of as some ‘relativistic equivalent’ of gravitational potential.

We obtain the TOV equation for m from the Hamiltonian constraint (2.2.83). This gives

$$\frac{dm}{dr} = 4\pi r^2 \rho_0 (1 + \varepsilon). \quad (2.6.4)$$

The equation for Φ comes from the evolution equation from the slicing condition, (2.2.87), and results in

$$\frac{d\Phi}{dr} = \frac{m + 4\pi r^3 p}{r(r - 2m)}. \quad (2.6.5)$$

Finally, the ODE for the pressure uses the evolution equation for momentum, S_r , (2.2.97) (in particular from the pressure balance term in this equation).

This then yields

$$\frac{dp}{dr} = [\rho_0 (1 + \varepsilon) + p] \frac{d\Phi}{dr}. \quad (2.6.6)$$

The full derivation of these equations is given in appendix A.3.

In order to obtain initial data for a numerical simulation of a star, we must provide initial data for the three ODEs ((2.6.4), (2.6.5), (2.6.6)). At $r = 0$, we have $m = 0$ and $p = p_c$, some chosen central pressure. The initial condition for Φ is not so obvious. Since Φ can be thought of as being a relativistic equivalent of gravitational potential in Newtonian physics, it is standard to set Φ to zero as $r \rightarrow \infty$. To achieve this, an arbitrary initial value is picked for Φ , and then once (2.6.5) has been solved, the solution is scaled such that $\Phi \rightarrow 0$ as $r \rightarrow \infty$.

Once the pressure has been obtained, the density can then be calculated

using the EOS. For the ideal fluid EOS, the entropy within a static star is chosen to be constant hence we use the polytropic EOS,

$$p = K\rho_0^\gamma. \quad (2.6.7)$$

We are able to use the TOV equations to determine whether static initial data will be stable in an evolution. Once we have picked an EOS, the TOV equations are then governed by a single parameter, typically the central density $\rho_{0,c}$ is used. We can then investigate the relationship between e.g. $\rho_{0,c}$ and the total mass of the star, M . As motivated in section 1.1, GR provides a maximum mass for a NS. When considering the mass-central density plot, there is indeed a maximum value above which an increase in $\rho_{0,c}$ leads to a decrease in M . This is identified as the maximum mass of a TOV star, see e.g. [136, 177].

2.6.1 The TOV equations for multicomponent stars

In this thesis we will need initial data for static, stable multicomponent stars, and therefore we consider how the TOV equations extend to such cases. Across the interface it is obvious that mass will not jump, and since the interface is modelled as a contact discontinuity, pressure will also remain constant. From (2.6.5) we see that the derivative of Φ depends only on continuous variables across the interface, and hence Φ must also be continuous. The TOV equations can now be solved as in the single component star case, providing the jump in density at the interface is known.

Since pressure is continuous across the interface, we can use (2.6.7) to obtain a relationship for the density jump,

$$p = K_{\mathcal{E}}\rho_{0,\mathcal{E}}^{\gamma_{\mathcal{E}}}, \quad (2.6.8)$$

where subscript \mathcal{E} represents an exterior quantity at the interface. We therefore can obtain $\rho_{0,\mathcal{E}}$ by

$$\rho_{0,\mathcal{E}} = \left(\frac{p}{K_{\mathcal{E}}}\right)^{1/\gamma_{\mathcal{E}}}. \quad (2.6.9)$$

We therefore need to know $\gamma_{\mathcal{E}}$ and $K_{\mathcal{E}}$ for the exterior fluid. We can choose a

value for $\gamma_{\mathcal{E}}$ based on the type of fluid that we wish to model in the exterior. If we have freedom to choose a value for $K_{\mathcal{E}}$ too, then the magnitude of the density jump is unconstrained, and this can lead an unstable solution. That is, under evolution, we would either experience collapse or exponential growth. We found that by ensuring the derivative of the pressure is continuous over the interface, then stable solutions could be achieved. It is important to note that this may not necessarily generalise to multidimensions, where instabilities due to e.g. shearing effects may be important. Using (2.6.6) this gives

$$[\rho_{0,\mathcal{I}}(1 + \varepsilon_{\mathcal{I}}) + p] \frac{d\Phi}{dr} = [\rho_{0,\mathcal{E}}(1 + \varepsilon_{\mathcal{E}}) + p] \frac{d\Phi}{dr}. \quad (2.6.10)$$

where subscripts \mathcal{I} represents an interior quantity at the interface. We know that the derivative of Φ is continuous across the interface, therefore simplifying this equation and using the ideal fluid EOS, $p = (\gamma - 1) \rho_0 \varepsilon$, we get

$$\rho_{0,\mathcal{I}} + \frac{p}{\gamma_{\mathcal{I}} - 1} = \rho_{0,\mathcal{E}} + \frac{p}{\gamma_{\mathcal{E}} - 1}. \quad (2.6.11)$$

Using (2.6.9) and rearranging we get

$$\left(\frac{p}{K_{\mathcal{E}}} \right)^{1/\gamma_{\mathcal{E}}} = \rho_{0,\mathcal{I}} + \frac{(\gamma_{\mathcal{E}} - \gamma_{\mathcal{I}}) p}{(\gamma_{\mathcal{I}} - 1)(\gamma_{\mathcal{E}} - 1)} \quad (2.6.12)$$

hence we can obtain $K_{\mathcal{E}}$ by

$$K_{\mathcal{E}} = p \left[\rho_{0,\mathcal{I}} + \frac{(\gamma_{\mathcal{E}} - \gamma_{\mathcal{I}}) p}{(\gamma_{\mathcal{I}} - 1)(\gamma_{\mathcal{E}} - 1)} \right]^{-1/\gamma_{\mathcal{E}}}. \quad (2.6.13)$$

From (2.6.9) and this result we then have enough information to provide the exterior density profile. We now have all the information required for initial data for a stable multicomponent star in 1+1 spherically symmetric GR. As with the single-component TOV equations, the maximum stable central density for given fluid components and interface location can again be obtained through e.g. a mass-central density plot.

2.7 Recovering the primitive variables

When solving the fluid evolution equations, we need to be able to convert between the conserved variables, (ρ_0, S_i, E) in Newtonian or (D, S_i, τ) in relativity, to the primitive variables, (ρ_0, v_i, p) . For the ideal fluid EOS, obtaining the conserved variables from primitive variables is trivial, based on the definition of these quantities. Similarly the reverse conversion in Newtonian physics is equally trivial, with a simple closed form,

$$v_i = \frac{S_i}{\rho_0}, \quad p = (\gamma - 1) \left(E - \frac{1}{2} \rho_0 v_i v^i \right). \quad (2.7.1)$$

The simple nature of this conversion is, in part, due to the analytic form of the EOS, it is clear that for more complicated EOSs, there is no guarantee an explicit form of this conversion will exist.

In the relativistic case, the coupling of variables through the Lorentz factor means that even for the ideal fluid EOS there is no guarantee that an explicit conversion from conserved to primitive variables can be found. Even when such a conversion is available, computationally expensive methods will often be required to calculate the primitive variables. This is an issue specifically during a numerical simulation, where conversion between variables needs to happen repeatedly. In order for the simulation to be computationally achievable it is then desirable to use a rapidly converging iterative scheme for this conversion (see section 3.3.1 for a description of convergence). Here we use the method in [10] where a Newton-Raphson scheme is used to solve a root for finding the pressure.

In order to start the iteration, we first need to guess a value for the pressure, \tilde{p} , and using this, with the known values for the conserved variables, we obtain estimates for the remaining primitive variables. Once the density and specific internal energy have been estimated, the EOS can be used to provide an equation for the root finding method,

$$f(\tilde{p}) = (\gamma - 1) \rho_{0,*}(\tilde{p}) \varepsilon_*(\tilde{p}) - \tilde{p}, \quad (2.7.2)$$

where a subscript $*$ represents a quantity estimated using \tilde{p} . Note that this

determines the difference between the pressure as calculated using the EOS compared to that of the estimated value. It is therefore obvious to make an extension to this equation should a different EOS be required.

We first note that from (2.2.53), $S_i = (E + p)v_i$, we can obtain an estimate of the velocity components,

$$v_{i,*}(\tilde{p}) = \frac{S_i}{\tau + D + \tilde{p}}, \quad (2.7.3)$$

where the obvious equivalent relationship can be used to obtain v_*^i . This can then be used to obtain a Lorentz factor

$$W_*(\tilde{p}) = \frac{1}{\sqrt{1 - v_{i,*}(\tilde{p})v_*^i(\tilde{p})}}. \quad (2.7.4)$$

The density and specific internal energy can then be estimated by

$$\rho_{z,*}(\tilde{p}) = \frac{D}{W_*(\tilde{p})} \quad (2.7.5)$$

and

$$\varepsilon_*(\tilde{p}) = \frac{\tau + D(1 - W_*(\tilde{p})) + \tilde{p}[1 - (W_*(\tilde{p}))^2]}{DW_*(\tilde{p})}. \quad (2.7.6)$$

From these equations, we can now calculate $f(\tilde{p})$ for a given \tilde{p} .

The Newton-Raphson method (sometimes referred to simply as the Newton method) calculates the root of an equation through the iterative scheme

$$u_{n+1} = u_n - \frac{f(u_n)}{f'(u_n)} \quad (2.7.7)$$

where prime represents a derivative with respect to u . In this case we use $u = p$ with f given by (2.7.2). It is possible to calculate the derivative of f , but it is cheaper computationally to use the approximation given in [10],

$$f'(\tilde{p}) \simeq v_{i,*}(\tilde{p})v_*^i(\tilde{p})(c_{s,*}(\tilde{p}))^2 - 1 \quad (2.7.8)$$

where the speed of sound for the estimated variables is given by

$$(c_{s*}(\tilde{p}))^2 = \frac{\gamma \tilde{p}}{\rho_{0,*}(\tilde{p}) \left(1 + \varepsilon_*(\tilde{p}) + \frac{\tilde{p}}{\rho_{0,*}(\tilde{p})}\right)}. \quad (2.7.9)$$

The Newton-Raphson method is a useful method for such conversions since it converges rapidly, the error in the approximated root falls off proportionally to the square of the number of iterations.

The drawback of such an accurate method is that the correct root will be obtained only for ‘appropriate’ selection of the initial guess for \tilde{p} . This guess must be close enough to the actual root or the method may not converge. Effectively if the gradient of the root is not ‘pointing’ towards the actual root (i.e. if the distance from the actual root to the intersect of the tangent to $f(\tilde{p})$ is greater than that to \tilde{p} itself) then the scheme will not find a root. In most cases when evolving the fluid evolution equations this is not a major issue since the most recently known value of the pressure will provide a suitably close initial estimate for \tilde{p} .

Chapter 3

Numerical methods

This chapter concerns the numerical techniques required to solve the fluid evolution equations. The necessity for numerical solutions of the equations that govern the behaviour of a NS has previously been discussed; we now outline the techniques for solving such equations. In particular we will give careful consideration as to why the conservation form (or balance-law form) in which we expressed these equations is necessary.

The use of numerical techniques for solving a system of equations will always require some balance between accuracy of the solution and limitations based on the computational cost. The use of these techniques for the fluid evolution equations requires some discretisation of a continuum system and as a result, this will obviously introduce some level of error. For a sensible choice of numerical method we would expect this error to decrease as the resolution of the simulation increases. Increasing resolution, however, means an increase in computational cost of the simulation, which is obviously limited through practical requirements. Our requirements for a suitable numerical technique are that we will have close to have minimal error within the solution per unit of time spent on a simulation. Note that this could lead to a crude, but computationally cheap technique being more useful than a very accurate but expensive one. The computational requirements for implementing a scheme with high accuracy may be greater than running a less accurate scheme at high resolution.

Many techniques have been developed for the numerical solution of PDEs. Amongst the most commonly used are finite volume and finite difference ([116,

123, 201]), spectral methods ([80, 39]) and smoothed particle hydrodynamics (SPH) ([73, 131, 137]). These methods all have both advantages and disadvantages, and the choice of method depends strongly on the physical situation being modelled along with the expected behaviour. For example, spectral methods can offer high accuracy with modest computing power, but have severe issues dealing with non-linear behaviour. When using such methods it is standard to add some artificial dissipation which prevents true shocks from forming. In this work, the interaction of non-linear behaviour with interfaces is a key aspect that we intend to investigate. As a result, our choice of numerical technique is heavily dictated by this, and finite volume methods are a natural choice. For obtaining comparable levels of accuracy in smooth regions of a fluid simulations to e.g. spectral methods, these methods are computationally expensive. Despite this, the major advantage of finite volume methods is that techniques have been constructed in order to deal accurately with non-linear behaviour (without dissipating features), hence with care we can provide accurate solutions across the entire computational domain.

We will now consider the use of finite volume methods in solving a 1+1 dimensional system of hyperbolic PDEs in balance law form,

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{q})}{\partial x} = \mathbf{s}(\mathbf{q}), \quad (3.0.1)$$

with arbitrary source vector \mathbf{q} , flux vector \mathbf{f} and source vector \mathbf{s} . In chapter 2 we saw that the evolution equations in both Newtonian and relativistic hydrodynamics could be expressed in this form. We will describe how the numerical methods introduced in this chapter generalise to multidimensional situations in section 5.1. Detail on many of the standard techniques introduced can be found in detail in e.g. [116, 123, 201].

Obtaining the numerical solution to this system of equations means that we must first discretise our domain. Here we chose to split the spatial domain into N cells of width Δx , denoted X_i , $i \in [1, N]$ where x_i denotes the value of x at the centre of the cell. The time domain is then discretised into ‘levels’ separated by a timestep Δt , which are denoted t^n , $n \in [0, (t^{\text{end}} - t^0) / \Delta t]$ where t^0 and t^{end} represent the initial and final times of the numerical solution respectively. The magnitude of this discretisation determines the resolution

of the numerical simulation. Here we have made the choice to use equally spaced discretisation, however, in general there is no specific requirement that either the time or space discretisations are kept constant. Advanced numerical methods have been developed which use more cells in the regions where there is complex behaviour (e.g. across a shock) than in smooth regions. One such common method is adaptive mesh refinement (AMR) which can identify regions in which high resolution is required, and use this to obtain a more accurate solution. When using AMR in a simulation, the desired accuracy can be achieved with significant reductions in computational power, since smooth regions can be simulated at a lower resolution than the fine features. This is of particular importance in large 3+1 dimensional simulations where computational resources are a major limiting factor. Details of this technique can be found in e.g. [31]. In this work we are aiming to test a numerical technique, therefore it is desirable that we keep all other features of our model as simple as possible. This will allow for greater ease in both analysing the effects of this technique, and identifying and dealing with any problems that may arise. Therefore we will assume a fixed grid and timestep for each simulation (Δx and Δt are constant). The domain of a cell X_i is given by $X_i \in [x_{i-1/2}, x_{i+1/2}]$ where $x_{i\pm 1/2} = x_i \pm \frac{1}{2}\Delta x$.

The discretised values for the vector quantities in (3.0.1) are denoted by $\hat{\mathbf{q}}_i^n$. Finite volume methods assume that $\hat{\mathbf{q}}_i^n$ is the integral average of the cell X_i at time t^n , i.e.

$$\hat{\mathbf{q}}_i^n = \frac{1}{\Delta x} \int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{q}(t^n, x) \, dx. \quad (3.0.2)$$

When choosing a numerical technique to use, stability is a very important consideration. We must ensure that the technique does not introduce oscillations or other unphysical behaviour into the solution. To determine whether a technique is stable, we can consider its *total variation*, TV , given, for an arbitrary discretised vector $\hat{\mathbf{u}}_i^n$ by

$$TV(\hat{\mathbf{u}}^n) = \sup_{\forall x_i} \sum_{i=-\infty}^{\infty} |\hat{\mathbf{u}}_{i+1}^n - \hat{\mathbf{u}}_i^n| \quad (3.0.3)$$

where the supremum means we take all possible samples from the domain of the calculation. A method is said to be *strong stability preserving* (SSP), [81],

if the total variation does not increase after any timestep,

$$TV(\hat{\mathbf{u}}^{n+1}) \leq TV(\hat{\mathbf{u}}^n). \quad (3.0.4)$$

This property is also commonly referred to as *total variation diminishing* (TVD), and this is particularly the case when only the variation of the spatial derivative are being considered [123]. A SSP method ensures that the evolution itself (moving to the next timestep) also introduces no increase in variation. A SSP method can be implemented such that unphysical oscillations are always avoided. A weaker condition can be placed on a technique, that it is *essentially non-oscillatory* (ENO). Here we have increase in total variation proportional to some power r ,

$$TV(\hat{\mathbf{u}}^{n+1}) \leq TV(\hat{\mathbf{u}}^n) + \mathcal{O}[(\Delta x)^r, (\Delta t)^r]. \quad (3.0.5)$$

In the limit $\Delta x \rightarrow 0$, $\Delta t \rightarrow 0$ we will recover the SSP condition. See e.g. [116] for further details on stability.

The discretisation of the numerical domain cannot be undertaken arbitrarily, i.e. Δx and Δt must be chosen with some care. For many numerical methods the choice of Δt is constrained by the choice of Δx (or vice versa). This is related to the *numerical domain of dependence* of the method versus the *physical domain of dependence* of the physical system.

The numerical domain of dependence comprises the region from which information is required for obtaining a solution at a single point. For example, the value for $\hat{\mathbf{q}}_i^{n+1}$ may require information from $\hat{\mathbf{q}}_j^n$, $j \in [i - \delta_1, i + \delta_2]$, where δ_1 and δ_2 are integers dependent on the numerical method implemented and $\delta_1 \leq \delta_2$. It would be this range that then gives the numerical domain of dependence, since we cannot obtain $\hat{\mathbf{q}}_i^{n+1}$ without knowledge of these integral averages.

The physical domain of dependence is the range over which the behaviour of \mathbf{q} at a time e.g. $t = t^n$ can affect the solution at a point e.g. $\mathbf{q}(t^{n+1}, x_i)$. The size of this domain is based on the maximum (most positive) and minimum (most negative) wavespeeds of the fluid. The left boundary of the physical domain of dependence, ε_L is defined by the point where information from $\mathbf{q}(t^n, \varepsilon_L)$,

travelling at the maximum wavespeed, will reach x_i at time t^{n+1} . This therefore is the minimum value on the domain at time t^n at which $\mathbf{q}(t^n, x)$ can affect $\mathbf{q}(t^{n+1}, x_i)$. The right boundary of the physical domain of dependence can be obtained in an analogous manner using the minimum wavespeed. In order for our numerical solution to include all possible behaviour we would therefore expect to have $\varepsilon_1 \geq \delta_1$ and $\varepsilon_2 \leq \delta_2$.

This therefore implies that for a given grid size Δx there is a maximum timestep Δt that can be taken for the entire physical domain of dependence to be included. This maximum is given by some multiple of the grid size, λ , using information about the maximum wavespeed, u_{\max} ,

$$\frac{u_{\max} \Delta t}{\Delta x} = \lambda. \quad (3.0.6)$$

This constraint is known as the *Courant-Friedrichs-Lowy* (CFL) condition [48, 49]. The value for λ is typically given by the numerical method being used. When working in relativity and using geometric units, since all velocities are bounded by the speed of light, it is standard to take $u_{\max} = 1$. In almost all cases, if the CFL condition is not satisfied, then the solution will not be stable. There are, however, (very few) exceptions to this, such as the *weighted average flux* CFL2 method of Toro and Billett [202], which allows for a CFL factor $\lambda = 2$. There are also a series of methods, *semi-Lagrangian* methods, which can achieve a CFL factor $\lambda > 6$, though these methods are typically only appropriate for e.g. advection equation type problems [190].

3.1 Method of lines

In order to solve a PDE numerically it is common to first convert it into an ODE using the method of lines (MOL). Many techniques then exist to solve ODEs numerically [201], e.g. Runge-Kutta methods, the Euler and backward Euler methods or the trapezoidal method.

To demonstrate this we apply the MOL to a hyperbolic PDE as given in (3.0.1) for which $\mathbf{s} = \mathbf{0}$. Discretising the state vector only in space, using

integral averages, gives

$$\hat{\mathbf{q}}_i(t) = \frac{1}{\Delta x} \int_{x_{i-1/2}}^{x_{i+1/2}} \mathbf{q}(t, x) dx, \quad (3.1.1)$$

which can be inserted into the hyperbolic PDE (3.0.1) to give

$$\frac{d}{dt} \hat{\mathbf{q}}_i(t) = \frac{1}{\Delta x} [\mathbf{f}(\mathbf{q}(x_{i-1/2}, t)) - \mathbf{f}(\mathbf{q}(x_{i+1/2}, t))]. \quad (3.1.2)$$

This is then an ODE for $\hat{\mathbf{q}}_i$. In order to solve it, we need to know the value of $\mathbf{f}(\mathbf{q}(x_{i\pm 1/2}, t))$. Again we make some approximation of these quantities, known as the *intercell fluxes*, $\hat{\mathbf{f}}_{i\pm 1/2}(t)$. A variety of methods exist for determining these approximated intercell fluxes, some of which are discussed in detail in section 3.4, and methods for including the source terms in section 3.6.

3.2 Runge-Kutta methods

Once we have used the MOL to convert the fluid evolution equations into an ODE, we obtained using the MOL as in (3.1.2). Here we need to solve an ODE of the form

$$\frac{d\mathbf{u}}{dt} = -\mathbf{F}(\mathbf{u}) \quad (3.2.1)$$

for arbitrary \mathbf{u} and \mathbf{F} . Many factors must be considered when selecting an appropriate ODE solver to use. Most importantly, the technique must be compatible with the reconstruction techniques used for obtaining the intercell fluxes $\hat{\mathbf{f}}_{i\pm 1/2}$ such that the evolution is numerically stable. Additionally we must also consider computational accuracy, cost and data storage requirements (how many time levels t^n must be stored). Here we use Runge-Kutta (RK) methods, a family of fairly simple, but accurate ODE solvers.

The first developments towards the RK class of ODE solvers came in 1895 [171], and since then there have been many further improvements and developments to these methods [37]. RK methods are an ideal tool for solving the fluid evolution equations since they allow accurate numerical evaluation of an ODE for modest computational cost. They also require only one time level to be stored (only $\hat{\mathbf{u}}^n$ is required to calculate $\hat{\mathbf{u}}^{n+1}$) hence minimising the amount of

data storage required. The need for only one time level is often referred to as self-starting, since the initial data alone is sufficient to provide the solution at later times.

RK methods are often referred to as a type of predictor-corrector method, since they work by taking an initial estimated solution (the prediction) at some time $t^n < t \leq t^{n+1}$. A series of corrections are then made to this prediction, leading to a solution with the desired level of accuracy at time $t = t^{n+1}$. Different RK methods are characterised by their level of accuracy, their order of convergence, which will be discussed in detail section 3.3.1.

A general N^{th} order RK method can be written

$$\begin{aligned} \hat{\mathbf{u}}^{(0)} &= \hat{\mathbf{u}}^n, \\ \hat{\mathbf{u}}^{(i)} &= \sum_{k=0}^{i-1} \left(\alpha_{ik} \hat{\mathbf{u}}^{(k)} + \Delta t \beta_{ik} \hat{\mathbf{F}}(\mathbf{u}^{(k)}) \right), \quad i = 1, \dots, N, \\ \hat{\mathbf{u}}^{n+1} &= \hat{\mathbf{u}}^{(N)}. \end{aligned} \quad (3.2.2)$$

Here $\hat{\mathbf{u}}^{(i)}$ is the predicted value ($i = 1$) or the i^{th} order correction ($i > 1$) and α_{ik} and β_{ik} are constant coefficients. These coefficients are chosen such that the solver obtains the desired level of accuracy, and ensures numerical stability.

From this general formalism, a variety of Runge-Kutta methods have been developed. In this thesis work we consider a standard second-order method,

$$\begin{aligned} \hat{\mathbf{u}}^{(1)} &= \hat{\mathbf{u}}^n + \Delta t \hat{\mathbf{F}}(\mathbf{u}^n), \\ \hat{\mathbf{u}}^{n+1} &= \frac{1}{2} \left(\hat{\mathbf{q}}^n + \hat{\mathbf{q}}^{(1)} + \Delta t \hat{\mathbf{F}}(\hat{\mathbf{q}}^{(1)}) \right), \end{aligned} \quad (3.2.3)$$

third-order method,

$$\begin{aligned} \hat{\mathbf{u}}^{(1)} &= \hat{\mathbf{u}}^n + \Delta t \hat{\mathbf{F}}(\mathbf{u}^n), \\ \hat{\mathbf{u}}^{(2)} &= \frac{1}{4} \left(3\hat{\mathbf{q}}^n + \hat{\mathbf{q}}^{(1)} + \Delta t \hat{\mathbf{F}}(\hat{\mathbf{q}}^{(1)}) \right), \\ \hat{\mathbf{u}}^{n+1} &= \frac{1}{3} \left(\hat{\mathbf{q}}^n + 2\hat{\mathbf{q}}^{(2)} + 2\Delta t \hat{\mathbf{F}}(\hat{\mathbf{q}}^{(2)}) \right), \end{aligned} \quad (3.2.4)$$

and fourth-order method

$$\begin{aligned}\hat{\mathbf{u}}^{(1)} &= \hat{\mathbf{u}}^n + \frac{1}{2} \Delta t \hat{\mathbf{F}}(\hat{\mathbf{u}}^n), \\ \hat{\mathbf{u}}^{(2)} &= \hat{\mathbf{u}}^n + \frac{1}{2} \Delta t \hat{\mathbf{F}}(\hat{\mathbf{u}}^{(1)}), \\ \hat{\mathbf{u}}^{(3)} &= \hat{\mathbf{u}}^n + \Delta t \hat{\mathbf{F}} \hat{\mathbf{u}}^{(2)}, \\ \hat{\mathbf{u}}^{n+1} &= \frac{1}{6} \left(-2\hat{\mathbf{u}}^n + 2\hat{\mathbf{u}}^{(1)} + 4\hat{\mathbf{u}}^{(2)} + 2\hat{\mathbf{u}}^{(3)} + \Delta t \hat{\mathbf{F}} \hat{\mathbf{u}}^{(3)} \right).\end{aligned}\tag{3.2.5}$$

It has been shown that the second- and third-order methods are SSP, [184], and, furthermore, that fourth-order methods cannot be SSP, [82]. It has, however, been found practically that the fourth-order method given here does not introduce oscillatory behaviour for the majority of situations.

3.3 The Lax-Wendroff Theorem and convergence

In both Newtonian physics and GR, the fluid evolution equations can lead to the formation of non-linear behaviour, even from smooth initial data. In such cases, the differential equation form, such as in (3.0.1), is no longer valid, since spatial derivatives become infinite. However, the integral form of the equations, which can be written

$$\int_{t_0}^{t_1} \int_a^b [\mathbf{q}_t + \mathbf{f}(\mathbf{q})_x] dx dt = 0,\tag{3.3.1}$$

still holds.

When this form, known as the *weak form*, is being solved, non-linear behaviour can be described, with the solution then termed a *weak solution*. It is possible for a single differential equation to have more than one weak solution, the choice of a rarefaction wave or rarefaction shock, shown in section 2.4.1 is an example of this. Care must be taken in order to ensure that the correct behaviour is obtained when weak solutions are formed, (i.e. to ensure the correct equation is actually being solved), which can be achieved through additional information e.g. some entropy condition [123].

We now consider why we have expressed the evolution equations in chapter 2 in either conservation or balance law form, and why this helps us to find at least a weak solution to the system. These equations can then be solved using a conservative numerical method (one that is also written in conservation form). Such methods guarantee that if there is any discontinuous behaviour, either in the initial data, or formed during the evolution, then the location of the discontinuity will be correctly positioned throughout the evolution. The exact location of the shock may not be known, but a tightly constrained region containing the shock is identifiable, this behaviour is known as *shock capturing*.

This feature of conservative numerical methods can be seen by considering the MOL ODE, (3.1.2), on an arbitrary domain $[x_M, x_N]$. Once we have obtained the intercell fluxes $\hat{\mathbf{f}}_{i\pm 1/2}$, the MOL ODE is

$$\frac{d}{dt} \hat{\mathbf{q}}_i(t) = \frac{1}{\Delta x} \left(\hat{\mathbf{f}}_{i-1/2} - \hat{\mathbf{f}}_{i+1/2} \right). \quad (3.3.2)$$

If we consider the sum of this ODE across the entire domain noting that $\hat{\mathbf{f}}_{i+1/2} = \hat{\mathbf{f}}_{(i+1)-1/2}$, we get

$$\begin{aligned} \frac{d}{dt} \sum_{i=M}^N \hat{\mathbf{q}}_i &= \sum_{i=M}^N \frac{1}{\Delta x} \left(\hat{\mathbf{f}}_{i-1/2} - \hat{\mathbf{f}}_{i+1/2} \right) \\ &= \frac{1}{\Delta x} \left(\hat{\mathbf{f}}_{M-1/2} - \hat{\mathbf{f}}_{N+1/2} \right). \end{aligned} \quad (3.3.3)$$

This means that the overall change in $\hat{\mathbf{q}}$ can be described purely by the flow through the boundaries of the domain, i.e. $\hat{\mathbf{q}}$ has been conserved over the domain. The advantages this has for non-linear behaviour can be seen by assuming there is a discontinuous feature within the domain at x_S , i.e. $x_M < x_S < x_N$. If a conservative numerical method is used then the result given above means that if we let x_M or x_N approach x_S then $\hat{\mathbf{q}}$ is still conserved. This means that the shock must be correctly positioned within the simulation. Such behaviour is not guaranteed around a discontinuity if a non-conservative method is used. We illustrate this property in figure 3.1.

In general, it is not essential for a numerical method to be conservative. Non-conservative methods can produce the correct solution if the solution is smooth for the entire simulation. In fact in the absence of strong non-linear behaviour (discontinuities), many non-conservative methods can offer significant




Figure 3.1: Showing the potential differences between conservative and non-conservative numerical methods. In each case the solid line represents the exact solution to some non-linear PDE, and the dashed line a numerical solution. The left panel shows the conservative method, where (3.3.3) means that the location of the discontinuity has been captured correctly. The right panel shows what can happen if a non-conservative method is used in a situation involving a discontinuity, where the speed of this feature has not been obtained correctly.

improvements in the accuracy of solutions obtained, e.g. spectral methods [80]. These advantages in accuracy mean that even for some situations in which non-linear behaviour is expected (though is not a major feature in the simulation) then there may be advantages to using non-conservative methods. Care must then be taken such that large errors are not introduced. Typically these methods can be designed with some artificial dissipation that is introduced around non-linear features. This should be implemented in such a manner that these features are smoothed sufficiently to avoid the issues shown in figure 3.1. For such techniques to work there is often the need to introduce some additional technique to determine the location of the shock. In this thesis we are interested in modelling the effects of shock interactions with other features within a fluid (particularly sharp interfaces). Therefore being able to model shocks as accurately as possible, by using conservative finite volume numerical schemes, is essential.

Although conservative methods are guaranteed to obtain the correct shock location, this is not sufficient to ensure that the correct solution to the system of equations is also obtained. We have made no statements as to the accuracy with which the shock is captured, or whether the correct behaviour local to the shock is observed, simply that it is in the correct position. Many conservative methods introduce spurious oscillations around shocks, hence simply knowing

the location of the shock offers no advantages. We therefore need to place additional conditions on our choice of numerical method.

It is essential to ensure that the numerical method we use is stable, e.g. we ensure the method is TVD or ENO. Assuming we do have such a method, there are two other essential properties we require the method to possess. Firstly the solution must converge to some function, i.e. as we increase the amount of cells used in a simulation, the solution obtained should increasingly resemble some configuration. As with other properties of numerical methods, by itself a solution which is convergent is not sufficient for accuracy since the convergence may be towards an incorrect solution. The second property is that the method must be consistent. This states that in the limit $\Delta x \rightarrow 0$ and $\Delta t \rightarrow 0$, the numerical method is equivalent to the system of equations we are solving.

The Lax-Wendroff theorem then states that for a conservative numerical method with a solution that converges to some function in the limit $\Delta x \rightarrow 0$ and $\Delta t \rightarrow 0$ and is consistent, then it is guaranteed that this is a weak solution to the system of equations [121]. In fact it can also be shown that non-conservative schemes will not converge to even the weak solution if a shock wave is present [98]. As stated previously, care then needs to be taken to ensure that this weak solution is the correct one. We must ensure that this solution satisfies the physical problem that we are considering. To ensure that the correct weak solution is obtained then we must ensure that the numerical method obeys some entropy condition (i.e. that entropy is non-decreasing with time), for more information see e.g. [123].

In order to know whether our choice of conservative numerical method satisfies the Lax-Wendroff theorem, we need to be able to test these properties. Consistency can be tested relatively easily and usually analytically by simply taking the appropriate limits $\Delta x \rightarrow 0$ and $\Delta t \rightarrow 0$. Whether the method is convergent is not as trivial to calculate. In all but a few cases, which generally lack physical relevance, convergence of a scheme must be tested practically. By using the numerical method to solve scenarios for which an exact solution exists, and then comparing the simulated results at different resolutions to this exact solution we can obtain some measure of the convergence of the method. Alternatively, we may have constraint equations (e.g. the Hamiltonian constraint in GR), which, as we detail below, can also be used to determine convergence. We

now detail how the convergence of a method can be quantified.

3.3.1 Measuring convergence

The rate of convergence of a numerical method is a description of how fast it approaches a solution. When developing a numerical method it is important to know this rate, both for ensuring the Lax-Wendroff theorem is obeyed, and to give an indication to some of the advantages (and disadvantages) of the method. For example, a computationally expensive method would not be practical to implement in many cases unless it approaches a solution very rapidly, i.e. has a high rate of convergence. In situations where a new numerical method is not being developed, the rate of convergence of the established methods used will be known. In these cases convergence testing is still important since it allows the specific implementation of a numerical method to be tested by ensuring it displays the correct rate of convergence.

Convergence of a numerical method can be quantified by determining the rate at which a numerically approximated solution to some system of equations approaches the exact solution as a function of the resolution of the grid. In general, the order of a numerical method for a PDE depends on both the methods used to obtain both the spatial and time derivatives. It will usually be the method with the slowest rate of convergence that dominates the overall convergence of the method. For the numerical techniques considered in this thesis, the timestep and grid spacing are not independent, being related through the CFL factor λ given in (3.0.6). Therefore increasing the resolution of one of these discretisations automatically leads to a similar increase in the other. The resolution, and hence rate of convergence of the entire solution, can then be given in terms of just one of the discretised quantities, typically the grid spacing Δx .

When describing the accuracy of a numerical method, we use the *order of convergence*. This details the rate at which the errors in a solution decrease as resolution is increased. For example, the solutions obtained with a first-order method will approach the exact solution linearly as resolution increases, whilst for a second-order method, the error decreases in proportion to the

square of the resolution. In general, a k^{th} -order method means that we have the relationship between cell width and error

$$\text{error} \propto (\Delta x)^k. \quad (3.3.4)$$

It is this measure of error that characterises the Runge-Kutta methods in section 3.2. It is possible to select the constant coefficients in (3.2.2) to achieve the desired level of convergence, with the explicit equations shown being some examples of this.

Two measures of the error in a solution are available, the *local error*, and the *global error*. The local error takes the error of a solution over a single timestep as a function of space. By comparing how two resolutions scale, we can easily obtain the order of convergence of this local error. We do not, however, have any information as to whether this level of convergence is maintained at other times. This requires some measure of how the error evolves as a function of time, hence we can obtain the global error. Various techniques exist which allow the global error to be quantified. One such technique is to use the p -norms of some averaged error at each timestep [123]. We can then compare the scaling of p -norms at different resolutions to obtain the order of convergence of this global error.

The p -norms

The name p -norm refers to a family of such norms, $\|\Delta\|_p$, which have the general formula

$$\|\Delta\|_p = \left(\Delta x \sum_{i=-\infty}^{\infty} |\Delta_i|^p \right)^{1/p} \quad (3.3.5)$$

for any integer p where Δ_i is the error at point x_i given by $\hat{q}_i - q(x_i)$. Note that the p -norm is calculated for each variable in \mathbf{q} independently. Although there are in theory an infinite number of available norms, in practice there are three standard choices of norm, the 1-norm

$$\|\Delta\|_1 = \Delta x \sum_{i=M}^N |\Delta_i| \quad (3.3.6)$$

the 2-norm

$$\|\Delta\|_2 = \sqrt{\Delta x \sum_{i=M}^N |\Delta_i|^2} \quad (3.3.7)$$

and the ∞ -norm

$$\|\Delta\|_\infty = \max_{M < i < N} |\Delta_i|. \quad (3.3.8)$$

This final case is the obtained from the limit $p \rightarrow \infty$, and is simply a measure of the largest error across the domain. Note that for these three definitions we have assumed a computational domain $i \in [M, N]$.

The choice of norm can greatly affect the order of convergence obtained when comparing resolutions. As the value of p increases, the dominance of the largest error in the calculation of the corresponding p -norm also increases (up until the ∞ -norm where it is the single measurement required). This then means that in these cases the overall convergence measured by the p -norm is going to be dominated by the regions in which convergence is slowest (which are therefore likely to have the highest errors). As a result, we will often denote the 1-norm as a *weaker* measure of convergence, it will not necessarily highlight if there are small regions where the method is failing to converge at the correct order.

Relying solely on the higher-order norms, however, is also not guaranteed to give a useful measure of convergence. Where non-linear features such as shocks are present, there is often a known local reduction in the order of convergence of a numerical method (see section 3.4), and therefore we will also want to know how the technique is converging in the smoother regions of the simulation. This is particularly an issue if we attempt to use the ∞ -norm for such cases, since we will often have some constant (but not increasing) error introduced around a shock. The amplitude of these errors are often non-convergent, and as a result, the ∞ -norm would show low- to zeroth-order convergence.

Knowledge of the exact solution is not always required to ensure a numerical technique is convergent for given initial data. It is possible to evaluate the errors for results run at two different resolutions compared to a run at a third resolution. In such cases some care must be taken since although we may know that the numerical simulations are converging to a solution, we do not necessarily know that this is the correct solution. If the Lax-Wendroff

theorem is satisfied by the numerical technique, however, we can at least be certain that convergence is towards a weak solution of the system of equations.

Convergence testing through constraint equations

When working in GR, we saw in section 2.2.3 that constraint equations are introduced. In polar-areal coordinates we had the Hamiltonian constraint, (2.2.83),

$$\frac{\partial_r a}{a} = a^2 \left[4\pi r(\tau + D) - \frac{m}{r^2} \right] \quad (3.3.9)$$

and slicing condition, (2.2.87),

$$\frac{\partial_r \alpha}{\alpha} = a^2 \left[4\pi r (S_r v^r + p) + \frac{m}{r^2} \right]. \quad (3.3.10)$$

When solving the evolution equations, the slicing condition is used to calculate the lapse, α , once a is known. However, a itself is obtained from an evolution equation arising from the momentum constraint, (2.2.85)

$$\partial_t a = -4\pi r \alpha a S_r. \quad (3.3.11)$$

The Hamiltonian constraint is then not required for describing how the space-time changes through the evolution. This allows it to be used as a simple method of checking the numerical error of the solution.

Once evolved data has been obtained at a timestep we can compare the derivative of a , which we must calculate from the numerical solution, to the quantity on the right hand side of the Hamiltonian constraint, (3.3.9). As stated by the constraint, these two quantities should be equal, however, since this equation is not being solved during the evolution, numerical errors will be highlighted. The errors used here can again give the local error at a single timestep, or we can take p -norms to give a measure of global error. Note that in this case, if the order of the method is greater than the order to which the derivative of a is calculated then it is again this lower order which will dominate. Providing a is smooth, a Taylor series expansion allows the derivative to be taken to arbitrarily high-order with relative ease and in a computationally cheap manner. Therefore as long as the order to which the derivative is taken

is greater than or equal to that of the numerical method, then an accurate measure of convergence can be obtained.

We also note here that since the two equations used to obtain α and a , (3.3.9) and (3.3.10), are ODEs, then the RK methods in section 3.2 can be used to solve them during the evolution of the fluid equations.

3.4 Reconstruction

In section 3.1 we reduced our hyperbolic system of PDEs to a system of ODEs, (3.1.2). To solve these ODEs at between times t^0 and $t^1 = t^0 + \Delta t$ we need to approximate the intercell fluxes, $\hat{\mathbf{f}}_{i\pm 1/2}$, at this time using the information from t^0 . We obtain this information by setting up a Riemann problem at the cell boundary for the state vector $\hat{\mathbf{q}}(t^0, x)$. Using any Riemann solver of choice, such as those detailed in section 2.5, we can obtain an approximate solution for $\hat{\mathbf{q}}(t^1, x)$, and hence $\hat{\mathbf{f}}(t^1, x)$. It is only actually the intercell fluxes $\hat{\mathbf{f}}(t^1, x)$ that are required to solve the ODE, therefore it is computationally cheaper for all Riemann solvers that we implement to calculate this directly. In fact, some Riemann solvers, such as the Marquina solver, do not give the variables at time $t = t^1$, and instead calculate only the intercell flux.

In order to set up a Riemann problem at a cell boundary, we need to obtain a constant left and right state, obtained, or *reconstructed*, from the integral averages, $\hat{\mathbf{q}}_i^n$, of the cells local to the cell boundary. This technique was pioneered by Godunov [76], and therefore methods to make this reconstruction are often known as Godunov-type methods. The original Godunov method simply uses the integral average either side of the cell boundary to provide these constant states, i.e. the reconstructed variables at $x_{i+1/2}$, typically denoted $\bar{\mathbf{q}}_{i+1/2}^\pm$, can be given by

$$\bar{\mathbf{q}}_{i+1/2}^- = \hat{\mathbf{q}}_L = \hat{\mathbf{q}}_i, \quad \bar{\mathbf{q}}_{i+1/2}^+ = \hat{\mathbf{q}}_R = \hat{\mathbf{q}}_{i+1}. \quad (3.4.1)$$

From this initial data we can then obtain a single value for the intercell flux, $\bar{\mathbf{f}}_{i+1/2}$ by solving the Riemann problem (exactly or approximately).

Whilst Godunov's method provides an obvious method for obtaining a Riemann problem to give the intercell fluxes, the simple approximations to the constant states (the zeroth-order extrapolation of the left and right state)

result in it being overall a first-order convergent method. In practice, first-order methods are too computationally expensive if reasonable levels of accuracy are to be obtained from large multidimensional simulations. It is therefore essential that we can use better than first-order convergence in the numerical methods used. We therefore consider the techniques available for higher order reconstruction of $\bar{\mathbf{q}}_{i+1/2}^\pm$, which are known collectively as *high resolution shock capturing* (HRSC) methods. Here high resolution refers to a global order of convergence greater than one, hence greater resolution of the physical features of a simulation.

It is obvious that to achieve greater than first-order convergence, we will move beyond zeroth-order reconstruction of the left and right states. Although this means that the reconstruction will no longer be constant for these states (and hence the problem is not a true Riemann problem), as long as this reconstruction is no greater than first-order, unphysical behaviour is not introduced by making this assumption. Unfortunately it was shown by Godunov [76] that greater than first-order reconstruction leads to oscillations being introduced at non-linear features. The aim of HRSC methods then has been to provide a reconstruction which is better than first-order where the solution is smooth, but reduces to a first-order method around non-linear features. This way, we can achieve a global order of convergence which is better than first-order. Various techniques through which this reconstruction can be obtained are available. These limit the behaviour of the reconstructed function around extrema, and ensure derivatives near non-linear features are kept as small as possible. One class of such techniques, which achieve this in a transparent manner, are slope limiters.

3.4.1 Minmod slope limiter

One of the simplest available slope limiting techniques is the minmod limiter [169]. For this method, the reconstructions at the cell boundaries are made linearly, based on a first-order approximation for the gradient (slope) of $\hat{\mathbf{q}}_i$, which we denote \mathbf{s}_i . There are two obvious linear approximations that can be used to give the slope of $\hat{\mathbf{q}}_i$ across the cell x_i ,




Figure 3.2: Detailing how the minmod method selects which slope to use. In this illustration there is some discontinuity between x_{i-1} and x_i . Therefore we have $|s_{R,i}| < |s_{L,i}|$ and as a result this is selected for the reconstruction.

$$s_{L,i} = \frac{\hat{q}_i - \hat{q}_{i-1}}{\Delta x}, \quad s_{R,i} = \frac{\hat{q}_{i+1} - \hat{q}_i}{\Delta x}. \quad (3.4.2)$$

These slopes are known as the downwind and upwind slopes respectively. The minmod method then selects one of these slopes to make the reconstruction with. This is done by taking the smallest slope in absolute value. An exception to this occurs if the direction of the slope changes across the cell (i.e. the signs of the upwind and downwind slopes are different) in which case we assume that the slope can be approximated as being zero. We can express this choice of slope as

$$s_i = \begin{cases} \operatorname{sgn}(s_{L,i}) \min(|s_{L,i}|, |s_{R,i}|), & s_{L,i}s_{R,i} > 0 \\ 0 & \text{otherwise} \end{cases}. \quad (3.4.3)$$

This is illustrated in figure 3.2.

Note that we can arbitrarily choose which slope to take the sign of, both slopes must be in the same direction for a non-trivial result. From the slope s_i , the reconstructed variables for use at $x_{i+1/2}$ are then given by

$$\bar{q}_{i+1/2}^- = \hat{q}_i + \frac{1}{2}\Delta x s_i, \quad \bar{q}_{i+1/2}^+ = \hat{q}_{i+1} - \frac{1}{2}\Delta x s_{i+1}. \quad (3.4.4)$$

As resolution increases, the minmod method tends to second-order conver-




Figure 3.3: Showing how the MC limiter chooses a reconstruction based on the different combinations of $|\mathfrak{s}_{L,i}|$ and $|\mathfrak{s}_{R,i}|$. The situation is set up as in figure 3.2. In this illustration the interpolated value $2|\mathfrak{s}_{R,i}|$ is selected.

gence. However, it is a fairly diffusive limiter, and as a result, the overall accuracy of a simulation is lower when compared with other techniques, even those of comparable order of convergence.

3.4.2 MC limiter

The monotonized central (MC) limiter is a similar in many ways to minmod, but offers greater accuracy [122]. The upwind and downwind slopes, as given by (3.4.2), are again required here. The choice of slope for the reconstruction is, however, now calculated in a different manner,

$$\mathfrak{s}_i = \begin{cases} \operatorname{sgn}(\mathfrak{s}_{L,i}) \min [2|\mathfrak{s}_{L,i}|, 2|\mathfrak{s}_{R,i}|, \frac{1}{2}(|\mathfrak{s}_{L,i}| + |\mathfrak{s}_{R,i}|)], & \mathfrak{s}_{L,i}\mathfrak{s}_{R,i} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3.4.5)$$

This is illustrated in figure 3.3

Again we see that if the slope changes sign across a cell, then the MC

limiter returns a slope of zero. We can see that otherwise, in general, the arithmetic mean of the two slopes will be used, except in cases in which one of the two slopes is sufficiently large. Once this slope has been obtained, the reconstruction at cell $x_{i+1/2}$ is again given by (3.4.4).

The method by which minmod and the MC limiter ‘limit’ the slope is very apparent: if any slope gets too large, then it is effectively ignored. The MC limiter also clearly identifies how the order of the method is reduced at non-linear features. The arithmetic mean is the centred difference approximation of the slope, a second-order accurate approximation. As previously stated, this will lead to oscillation if used for reconstruction close to non-linear features. Therefore, when the slope is large, suggesting we are near such a feature, a first-order approximation is then used.

3.4.3 PPM

The two slope limiting techniques presented so far are fairly simple. Higher order convergence is achievable with more complex methods. In this work we consider one such method, the *piecewise parabolic method* (PPM) [46]. This makes use of higher order reconstruction to achieve greater accuracy. As a result of this, some additional dissipation is necessary to avoid the issues with high-order reconstruction around non-linear features. Here we outline the steps taken in the method, a more complete description as to how it is implemented is given in section B.1.

PPM first makes a fourth-order polynomial reconstruction at the intercell boundaries. This is undertaken in a manner that steepens this reconstruction at discontinuities. Once this polynomial has been obtained, it is steepened further at discontinuities. To achieve this, it must be implemented with some procedure which determines where such discontinuities are. In certain cases this may attempt to steepen regions which are actually continuous, but even in this case, the reconstruction remains second-order accurate. In order to prevent oscillatory behaviour, PPM then ‘flattens’ the reconstruction around shocks, using some artificial viscosity. Finally it is ensured that the reconstructions are monotonic (that the intercell value is not a maximum or minimum). An extension of PPM to relativistic situations was made in [134].

3.4.4 Alternative reconstruction methods

The techniques described so far to make the reconstruction can be described as a MUSCL scheme (*Monotone Upstream-centered Scheme for Conservation Laws*) [122]. These then make use of e.g. slope limiters or PPM to avoid oscillatory behaviour at sharp features. Other techniques within the MUSCL scheme framework exist, such as further slope limiters, including the superbee limiter [169] or the Sweby limiter [193]. It is possible to make improvements to convergence, above those attainable through MUSCL schemes, by using ENO methods to make the reconstruction. Such methods were introduced by Harten et al. [92, 90], with improvements made by Shu and Osher [184, 185]. These take several polynomial reconstructions of $\hat{\mathbf{q}}_i$ making use of neighbouring cells, and then pick the ‘smoothest’ for the final reconstruction. The standard ENO scheme achieved third-order convergence. In practice ENO methods can be constructed with arbitrary levels of convergence.

The standard ENO method will only make use of one of the polynomial reconstructions. It is possible for greater accuracy (and hence higher order convergence) to be achieved if a linear combination of all possible polynomial reconstructions (up to some desired order) can be used. This is the idea behind the *weighted ENO* (WENO) method [127, 128]. The contribution of individual reconstructed polynomials to this linear combination is determined by some weighting. This is carried out in a manner that over discontinuities, the contribution of the polynomials drops to very near zero, and as a result leads to non-oscillatory behaviour in these regions. As with the ENO method, the WENO method can theoretically achieve an arbitrary order of convergence. Practically, the initial implementation of the WENO method was fourth-order convergent, and, through improvements made in [104], a fifth-order convergent method was formulated.

3.5 Boundary conditions

At the beginning of this chapter we described the numerical domain of dependence of a method. Here we consider what issues arise if the domain of $\hat{\mathbf{q}}_i^{n+1}$ requires information from time $t = t^n$ which lie outside the computational domain. As an example of this we consider the numerical domain of dependence




Figure 3.4: This figure demonstrates the issues that arise at the boundaries of the computational domain. Here, the left boundary of this domain (at $x_{M-1/2}$) is indicated by the dashed line. The numerical domain of dependence required for obtaining $\hat{\mathbf{q}}_M^{n+1}$ is given by the shaded region. It is clear that this includes cells which are not defined (those to the left of x_M).

of $\hat{\mathbf{q}}_M^n$ where this is the integral average of the first cell on the computational domain. The solution for $\hat{\mathbf{q}}_M^{n+1}$ will require information from the cells x_j , $j \in [M - \delta_1, M + \delta_2]$, where δ_1 and δ_2 give the range of the numerical domain of dependence as defined by the numerical method being used. If $\delta_1 \neq 0$ then $\hat{\mathbf{q}}_M^{n+1}$ requires information from outside the computational domain. This is illustrated in figure 3.4

There are several possibilities for dealing with this issue. It may be possible to use a different numerical method near the boundaries such that the numerical domain of dependence lies entirely within the computational domain (see e.g. [116] for descriptions of such methods). In many cases, including when solving the fluid evolution equations, this will not, however, cover the physical domain of dependence, and as a result unphysical behaviour may be introduced into the solution. A more versatile alternative is to actually include this extra region required by the numerical domain of dependence into the computational domain. These cells are then artificially populated in such that the fluid behaves correctly at the edges of the domain. Since these regions are not actually part of the physical solution of a simulation, they are known as *ghost zones*. Typically we would introduce $n_g = \max(\delta_1, \delta_2)$ *ghost cells* into

this region. Once they have been populated, the same numerical method can be used across the entire domain.

In order to populate the ghost zones, we need to know how the fluid behaves at the edge of the computational domain, i.e. we need to know the boundary conditions of the physical problem. These are then incorporated into the ghost zones through some *boundary treatment*. Imposing a suitable boundary treatment is not necessarily a straightforward task. It must be implemented in such a manner that the numerical methods being used remain stable and they must not introduce (or limit the introduction of) unphysical behaviour. Here we follow [116] and consider techniques for two types of boundary, *solid* and *far-field* boundaries. The former of these describes the cases where the edge of the domain corresponds to some physical feature, such as a wall, which reflects the matter in the simulation. The latter is a purely numerical treatment where constraints on the computational cost of a simulation demand that the grid must be truncated before any physical boundary is reached (e.g. the only physical boundary in 3+1 dimensional NS simulations is at infinity, and therefore unfeasible to implement in a numerical simulation). In addition to this we will also briefly describe exact and periodic boundaries. In this section, all boundary conditions are populated for $j \in [1, n_g]$ cells beyond the computational domain in all directions.

3.5.1 Exact boundary conditions

Knowledge of the exact solution of a problem (if such a solution exists) means that this can be used to populate the ghost cells. For a numerical method that requires n_g ghost zones and using the boundary in figure 3.4 this treatment can then be expressed as

$$\hat{\mathbf{q}}_{M-j} = \mathbf{q}(x_j). \quad (3.5.1)$$

The obvious analogy applies for the condition on other boundaries within the system.

Although this should result in the correct treatment in all such cases, issues arise due to the numerical error inherent in any simulation. This error will lead to a discrepancy between the simulated value and the exact value at the boundary. This may then lead to unphysical behaviour which can propa-

gate back into the domain. Additionally, the practical uses of this treatment are limited since if an exact solution exists, a numerical simulation is largely unnecessary.

3.5.2 Periodic boundary conditions

When fluid leaves one boundary of the computational domain, if it then enters the domain again through a different boundary we have periodic boundaries. The boundary treatment for this case is then obvious, if there is a periodicity in the boundaries at x_M and x_N then the ghost zones will be populated by

$$\hat{\mathbf{q}}_{M-j} = \hat{\mathbf{q}}_{N-j}. \quad (3.5.2)$$

Here we are effectively populating the ghost zones with cells within the computational domain. As with the exact boundaries, physical uses of this condition are limited. If this treatment can be used, however, it will not introduce any additional behaviour or errors into the simulation. The most obvious application of such a treatment would then be for testing numerical methods without interference from the boundaries.

3.5.3 Solid boundary conditions

If we have a solid boundary then for a perfect fluid we would expect the boundary conditions here to be reflective, no matter should pass through the boundary. In terms of the ghost zone boundary treatment, the behaviour at the solid boundary is exactly that of there being an identical fluid (in terms of density, pressure and EOS) with the opposite normal velocity. This technique is known as *reflection* or *imaging*. The boundary treatment for the boundary illustrated in figure 3.4 would be given by

$$\hat{\rho}_{0M-j} = \hat{\rho}_{0M+j-1} \quad (3.5.3)$$

$$\hat{\mathbf{v}}_{n,M-j} = -\hat{\mathbf{v}}_{n,M+j-1} \quad (3.5.4)$$

$$\hat{\mathbf{v}}_{t,M-j} = \hat{\mathbf{v}}_{t,M+j-1} \quad (3.5.5)$$

$$\hat{p}_{M-j} = \hat{p}_{M+j-1} \quad (3.5.6)$$

where \mathbf{v}_n and \mathbf{v}_t are the normal and tangential components of the velocity, and an analogous treatment is applied to all other boundaries.

In many cases this boundary treatment is very successful. Issues can arise, however, if there non-zero derivatives of variables as they approach boundary. The reflection fluid in the ghost zones will then mean the derivative at the boundary does not exist. This can then lead to a reduction in the order of convergence of the numerical techniques. There can also be issues with this treatment when a shock hits a boundary. It was found that an unphysical spike can be introduced in a numerical solution. In this spike density reduces and specific internal energy and temperature increase. This behaviour is therefore sometimes known as *shock heating*, and is documented in [58].

3.5.4 Far-field boundary conditions

Far-field boundaries are imposed when a physical boundary lies outside of the computational domain. The treatment at these boundaries should allow waves from the computational domain to propagate outwards without reflection, an *outflow* condition. The treatment should also specify any ingoing waves that are required from outside of the computational domain, an *inflow* condition. In this work, we do not consider situations where any information should enter from outside the computational domain. Therefore we will restrict ourselves to considering purely outflow boundaries.

Determining how the ghost zones are populated in order to achieve outflow boundaries is far from straightforward. If the fluid variables are continuous at the boundary, then some polynomial extrapolation of the variables into the ghost zones could provide a suitable treatment. This should ensure that the fluid is continuous across the boundary, and hence we would not expect any reflected waves. Care must be taken if non-linear behaviour approaches the boundary since this treatment may no longer give appropriate conditions, since a polynomial extrapolation will no longer be appropriate.

In many cases, such as all those considered in this thesis, the interaction of the fluid with the boundary is not the focus of the numerical simulations. As a result, we can pick appropriate choices of computational domain such that all important features, particularly non-linear ones, remain within this domain. In these cases then we can make a simple zeroth-order extrapolation

to populate the ghost zones, i.e. for the situation in figure 3.4 we have

$$\hat{\mathbf{q}}_{M-i} = \hat{\mathbf{q}}_M. \quad (3.5.7)$$

In multidimensional situations this boundary condition is clearly inappropriate, the zeroth-order extrapolation (if it is to be used) should occur in the direction of the velocity at the boundary. In our preliminary work in multidimensional situations in section 5, it was beyond the scope of the work to consider a more appropriate boundary treatment. Instead the computational domain has been chosen such that we consider only the region suitably far from the boundaries that the effects of the inadequate treatment have not yet reached.

3.5.5 Vacuum boundary conditions

In astrophysical simulations, we will often be dealing with cases where the fluid boundary treatments are not of great importance (even in multidimensional situations) since the boundaries lie within the vacuum region. In GR, however, both the fluid and the spacetime can evolve, and therefore we must also consider the boundary conditions of the spacetime variables. At the boundaries there is no guarantee that the constraint equations are obeyed. This can then lead to errors in the spacetime propagating inwards from these boundaries (which appear as GWs). There are two commonly used methods to avoid these issues.

It is possible to pick gauge conditions within the boundaries such that they are *maximally dissipative*. This then prevents constraint violating errors from entering the numerical domain (these conditions are also known as *constraint preserving boundary conditions*) [71, 191].

Alternatively, the outer boundaries can be set at a sufficient distance that the effects of constraint violation do not propagate at a sufficient speed to effect the region in which results are obtained from a numerical simulation. This approach is only viable when using AMR, in which the outer boundaries can be set where only a low resolution grid is used, reducing the computational cost of a simulation.

In this thesis, since we restrict ourselves to 1+1 dimensions when working in GR, we do not have to worry about constraint violation at the outer boundaries. Since there is no GW formation in spherical symmetry, errors at the boundaries will not propagate inwards.

3.6 Source terms

When we outlined the MOL in section 3.1, we assumed that we were solving the evolution equations in the absence of source terms. Fortunately, from a numerical perspective, including source terms is relatively simple. Assuming we have some discretised source term, \mathbf{s}_i , then the balance law PDE (3.0.1)

$$\frac{\partial \mathbf{q}}{\partial t} + \frac{\partial \mathbf{f}(\mathbf{q})}{\partial x} = \mathbf{s}(\mathbf{q}), \quad (3.6.1)$$

can then be written in ODE form using the MOL, which gives

$$\frac{d}{dt} \hat{\mathbf{q}}_i(t) = \frac{1}{\Delta x} [\mathbf{f}(\mathbf{q}(x_{i-1/2}, t)) - \mathbf{f}(\mathbf{q}(x_{i+1/2}, t))] + \mathbf{s}_i. \quad (3.6.2)$$

Since the source vector depends on the state vector, $\mathbf{s} = \mathbf{s}(\mathbf{q})$, there are no problems presented by including these terms in the RK methods in section 3.2.

The only remaining issue is then how we discretise the source terms such that \mathbf{s}_i is the appropriate integral average. Fortunately the intuitive method by which we might attempt to do this,

$$\mathbf{s}_i = \mathbf{s}(\hat{\mathbf{q}}_i), \quad (3.6.3)$$

is accurate to second-order. By making some appropriate reconstruction of $\hat{\mathbf{q}}_i$ higher order approximations can be obtained, but we shall not consider them here.

3.7 Multicomponent fluids

When modelling multicomponent fluids all techniques described so far in this chapter apply to the fluid evolution, but further conditions must be taken. In section 1.4.2 we motivated the use of sharp interfaces as a method of dealing

with the different regions within a NS. It is expected that there will be a different EOS describing the fluid properties in each region. It is worth considering a quantitative example which details why it is impractical to explicitly model the transition layer. This model assumes we use AMR and we consider the lengthscales and resolution of the model that including a transition layer would require.

For this calculation, we consider a region for which there is a quantitative estimate of size, the Ekman layer in hot proto NS cores. This calculation is detailed in [14] and, to follow these results and allow for the result to be interpreted with ease, we use cgs units. The width of this region has been estimated as

$$\delta_{\text{Ekman}} \sim \left(\frac{\eta_\nu}{\rho_0 \Omega} \right)^{1/2} \quad (3.7.1)$$

where η_ν is the viscosity coefficient and Ω the rotation rate. There is a large temperature dependence in η_ν , but for a typical temperature of a proto NS $> 10^9$ K, we expect $\eta_\nu \sim 10^{13} - 10^{19}$ g cm $^{-1}$ s $^{-1}$. If we assume that at a density $\rho_0 \sim 10^{14}$ g cm $^{-3}$ and $\eta_\nu \sim 10^{16}$ g cm $^{-1}$ s $^{-1}$, then we have a lengthscale for the Ekman layer of

$$\delta_{\text{Ekman}} \sim 10\Omega^{-1/2} \text{cm.} \quad (3.7.2)$$

We now consider the numerical requirements for incorporating features of this size into a full 3+1 simulation of a NSNS binary inspiral for which a GW signal is to be extracted. We will assume HRSC methods (and hence the finite volume methods) will be used with computational power based on that achievable with near-future developments. Thus the calculation will therefore be based on a fully parallelised code using AMR to allow for the best possible spatial resolution. We assume that $\sim 10^{10}$ timesteps can be simulated, and that, in order to cover several orbits and the merger, we need this simulation to run for 10 ms. We therefore have the finest timestep at $\sim 10^{-12}$ s, which, using the CFL criteria (3.0.6),

$$\frac{u_{\max} \Delta t}{\Delta x} = \lambda, \quad (3.7.3)$$

gives us a constraint on the spatial resolution. The maximum wavespeed in a relativistic situation is c , and we assume that $\lambda = 1$, therefore we have $u_{\max} = 3 \times 10^{10}$ cm s $^{-1}$. This gives the smallest grid spacing of $\Delta x \sim 10^{-1}$ cm.

In standard numerical methods, there is a smearing of features which, if a feature is initially a contact discontinuity, will be spread over $T^{1/(p+1)}$ cells after T timesteps where p is the order of convergence of the method [89]. By making the (optimistic) assumption that fourth-order convergence can be achieved then by the final time this contact discontinuity will be spread over 10^2 cells. This will give this initial feature a width of ~ 10 cm. In order for the physics of a transition region to be modelled we require that the width of the feature is at least an order of magnitude greater than the smearing lengthscale. Therefore we would require $\delta_{\text{Ekman}} \sim 100$ cm for it to be resolved. It is clear that for standard NS rotation rates ~ 100 Hz this lengthscale will not be achieved, hence motivating the treatment of these regions as genuine interfaces. It is also worth noting that for a typical simulation of a NS, for which GW information is to be obtained at some suitable distance from the source, we will require the total size of the computational domain to be at least 100 km. Therefore achieving resolution of 10^{-1} cm will require ~ 20 levels of AMR, assuming resolution is doubled at each resolution. Although large, this would be feasible for near-future simulations if necessary. It must also be considered whether the computational cost of such a simulation justifies the use of this resolution. If the behaviour in the transition layer is largely stable, then the accuracy gained from ~ 20 levels of AMR would not provide physically interesting results. This suggests that regions, such as the Ekman layer, are too small to be included within a model of an entire NS and hence provides motivation for the use interface modelling techniques. In this case, the actual behaviour of the layer would never be explicitly calculated. Instead it would either be apparent through boundary conditions placed on the interface, or, if the effects did not effect the overall behaviour of the NS, not included in the model.

Further motivation for sharp interfaces comes from considering how a jump in the EOS can be implemented. Naively it may seem trivial to include an EOS which is a function of space. This means that the conversion between primitive and conserved variables (and calculation of quantities such as the speed of sound) would be dependent on this function. By introducing an evolution equation for this function we might then expect that this would continue to apply the correct EOS in the various fluid regions. At the sharp interfaces we would therefore have a jump in the EOS function. If we are using

HRSC methods (which are used for dealing with non-linear behaviour in the simulations) then there is going to be some smearing of the EOS function due to this discontinuity. This is obviously undesirable since it was the smearing of features that motivated the introduction of using sharp interfaces. There are further detrimental effects of using the discontinuous EOS as a function of space however.

The smearing of a discontinuous EOS function can be interpreted as the mixing of two fluid components. In some cases, this may be desirable (if such mixing is genuinely expected) and in these cases it must be ensured that this happens in a thermodynamically consistent manner. The smearing induced by HRSC methods is a purely numerical artifact, and therefore there is no guarantee that this is thermodynamically consistent. This then can lead to large errors being introduced at the interfaces [6]. In section 4.1.1, we use numerical simulations to explicitly show the problems associated with this technique.

In order to successfully model a system in which there is a change in EOS in a stable manner and a sharp interface, we must therefore consider other techniques. Having suggested that a single component model is unfeasible for such situations, we now consider a multicomponent method for modelling sharp interfaces. The idea behind such a model is to treat each fluid component separately, hence a system with n fluid types will have n components, each evolved using a single EOS. The final solution can then be ‘pieced’ back together from these components. In order for this technique to be viable, we must introduce additional features into our model. We must know the location of any interfaces, how they evolve with time and also how the fluid behaves at these interfaces. This final consideration must ensure that as waves in one fluid type reach an interface they interact with the other fluid type. Effectively we must place some boundary conditions on the fluid at the interfaces. This may appear to be significantly complicating the model, however successful and comparatively simple techniques have been developed for capturing the location of interfaces, and for providing these boundary conditions. We consider how the interface is modelled and how the boundary conditions are applied separately.

3.7.1 Level set methods

Level set methods present a valuable tool for tracking sharp features, such as interfaces [148, 176]. In order to locate such features, these methods introduce an external scalar function, typically denoted ϕ , known as the *level set function*. Here we shall explicitly assume that all features being tracked are interfaces, however the techniques described can generalise to any sharp feature. Assuming there are two fluid components, the level set function is initialised such that the contour $\phi(\mathbf{x}) = 0$ gives the position of the interfaces, i.e. we use the zeroes of ϕ . In this work we shall entirely focus on two-component fluids although it is possible to extend this to situations with additional fluid components. The sign of the level set function gives us the required information to ‘piece’ the solution together from the separate components. If we have $\phi > 0$ then one component will be present whilst for $\phi < 0$ we have the other component. In order to track the interfaces in time, we then need to evolve ϕ along with the rest of the fluid, and do so in a manner that is consistent with how the interface is expected to move. This effectively requires the evolution equation for ϕ is determined using the physics of the system being modelled.

Level set methods are not the only possibility which allow for the modelling of sharp interfaces. Other possibilities exist, such as the particle level set technique ([63, 148]), volume-of-fluid method ([141, 167]) or phase-field method ([102]). To provide an argument for why we will use level set methods, we describe some of the advantages to this technique, and how they apply to this work.

One key feature of level set methods is their ability to deal with changes in topology of an interface. If the interface becomes twisted (e.g. due to rotation) or two initially unconnected regions merge, (e.g. in binary NS merger) this is dealt with naturally in the level set formulation. In some cases, a physical NS situation may result in the case in which a new fluid model is ‘spontaneously’ required (e.g. if there is a transition to exotic model in the core). Here it is not obvious how to deal with this in the context of level set techniques without the additional information as to how, and when, this transition occurs. It may be possible to connect the function ϕ to e.g. some temperature variable, determining when a transition to exotic matter occurs, though we shall not consider this scenario.




Figure 3.5: Showing how a level set function, ϕ , models the location of interfaces. In the left panel we see example data for an early time in a 1+1 dimensional multicomponent fluid simulation in which there are two separate regions of fluid. The right panel then shows a possible solution for the late time evolution of this system. The level set function has naturally allowed the ‘fluid 1’ regions to merge.

In figure 3.5 we show how a level set method identifies interfaces, in a simple 1+1 dimensional multicomponent situation. The left panel shows an initial set up in which there are two regions of ‘fluid 1’, separated by ‘fluid 2’ i.e. the level set function has four zeroes, and hence we have four interfaces. At some future time, as shown in the right panel, these two regions of ‘fluid 1’ have merged, we now have only two interfaces. The evolution equations for the level set allow this to occur naturally. It is obvious that this is useful for NS simulations, in which, if we consider a two-component star, may at some point during merger require these components to merge.

Although figure 3.5 considers only the 1+1 dimensional level set function, the methods extend naturally for arbitrary dimensions, which is clearly necessary for NS simulations. Additionally, since the evolution equation for ϕ is determined from the behaviour of the fluid variables, many of the numerical techniques for evolving ϕ can similarly be ‘borrowed’ from those already being used if necessary. Further techniques which we use only for level set techniques are described in section 3.7.3. For a comprehensive description of the advantages of level set techniques, see e.g. [176].

Some final motivating features for the use of level set techniques are firstly that they are widely employed in determining the boundary conditions at an interface. Therefore for the purposes of this work, it is sensible to follow these examples. Secondly the use of level set methods in GR is well established, particularly in capturing the position of event and apparent horizons in black hole simulations [53, 199]. We therefore know that there should be no issues with their implementation to model a multicomponent NS.

3.7.2 The level set function

The first consideration when using level set methods is how the function ϕ is initialised. If we are considering a physical situation, then we must ensure that ϕ accurately gives the location of the interfaces present. In many cases this may not be a trivial task, particularly for a fluid modelled in 3+1 dimensions, with a complex initial structure to the interfaces. In this work, however, we predominantly consider simple physical situations or test cases, for which we can choose the initial shape of the interfaces in a much more arbitrary manner. As an example, the Riemann problem described in section 2.4 requires a single interface at $x = x_0$. Therefore we have initial data for ϕ given by

$$\phi(x) = x - x_0. \quad (3.7.4)$$

Determining the evolution equation for ϕ requires the physical behaviour of the fluid at the interface to be taken into account. For the interfaces we consider in this work there is an obvious form for the evolution of ϕ . The regions of a NS described in section 1.2 are those that are expected to be present in a stable isolated NS. For a sharp interface to be stable in this manner, there must then be pressure balance across it. We then make the assumption that there is *no matter transfer across the interface*. This effectively constrains the interface to be moving with the fluid. In general we therefore expect velocity to be continuous across the interface. It is then apparent that the interface can be described as a contact discontinuity (see section 2.4.1), hence we can

treat it as a being advected with the fluid. In the general case, the evolution of ϕ can be described as being Lie-dragged with the 4-velocity of the fluid

$$\mathcal{L}_{\mathbf{u}}\phi = 0. \quad (3.7.5)$$

In both the Newtonian Cartesian coordinate and SR cases this, when written explicitly, gives the standard advection equation form

$$\frac{\partial \phi}{\partial t} + v^i \frac{\partial \phi}{\partial x^i} = 0. \quad (3.7.6)$$

In GR, the spacetime effects enter the evolution equation through the Lie-derivative. When written in polar-areal coordinates (as described in section 2.2.3), the evolution of ϕ can be written explicitly as

$$\frac{\partial \phi}{\partial t} + \alpha v^r \frac{\partial \phi}{\partial r} = 0. \quad (3.7.7)$$

Note that these equations are not in conservation form. Whilst they can easily be converted to conservation form, e.g. (3.7.6) becomes

$$\frac{\partial (\rho_0 \phi)}{\partial t} + \frac{\partial (\rho_0 v^i \phi)}{\partial x^i} = 0, \quad (3.7.8)$$

and hence make use of the methods described so far in this section, this formalism is not ideal in this case. The key motivation for writing the fluid evolution equations in conservation form is to avoid the issues that arise due to non-linear behaviour. Since the only information that is required from ϕ is the location of the interface, its structure is therefore arbitrary. We are therefore free to enforce that ϕ is free from such behaviour, i.e. contains no discontinuities.

In fact, for level set functions, conservation form can actually be detrimental to the accuracy of the simulation. The conserved variable $(\rho_0 \phi)$ is discontinuous at the interface. Although HRSC methods mean that (3.7.8) can be evolved without introducing oscillations, there is still some unavoidable error due to the reduction in accuracy of these methods about discontinuities. This will manifest as an error in the location of the interface. In fact the initial work on level set methods was in fact done with ϕ discontinuous at the interfaces [149], and suffered from these issues. If a function does not contain discontinuous features however, an evolution of its non-conservative form,

e.g. (3.7.6), can be achieved without introducing errors [138].

If we ensure that the initial data for ϕ is continuous at the interface, there is no guarantee that it will remain as such. If the velocity profile of the fluid is not constant then the profile of ϕ can steepen. If this occurs then the errors associated with non-conservative methods will arise. In these cases, however, we can again use the fact that the only information required from ϕ is the interface locations. At any stage in the evolution we are free to redefine ϕ and, as long as the interface locations are not changed, the physical model will not be affected. Therefore *reinitialisation* techniques, which will ‘smooth’ out a level set function which gets too steep, have been developed [43]. In practice, rather than waiting for ϕ to steepen, a reinitialisation technique will be implemented after a certain number of timesteps. The number of timesteps will be chosen on a case-by-case basis, depending on how rapidly ϕ steepens.

In many cases we consider in this work, there we have a single interface, and we can use this to avoid the need for reinitialisation whilst retaining the ability to undertake long term evolutions. A single interface means the evolution of ϕ can be described everywhere using only the velocity of the interface. This then means we have a constant coefficient advection equation for ϕ , and are able to evolve it across the entire domain using simply the 4-velocity at the interface, $u_{\text{int}} = u|_{\phi(x)=0}$. We therefore have a general evolution equation given by

$$\mathcal{L}_{u_{\text{int}}} \phi = 0. \quad (3.7.9)$$

In this case, since ϕ is advected with constant velocity across the entire domain, it retains its initial profile, i.e. sharp features will never form.

Level set techniques capture the location of the interface, rather than explicitly tracking it (i.e. we know the location to within one cell on the numerical grid). In order to increase the accuracy of the evolution of (3.7.5) it is useful to make some interpolations to obtain the velocity at the interface itself. We must then first obtain the position of the interface. If we use a linear profile for ϕ , then this position can be calculated exactly. We first locate the cell x_i for which $\hat{\phi}_i \hat{\phi}_{i-1} < 0$, where $\hat{\phi}_i$ is the integral average of ϕ across cell x_i . This indicates that the interface lies in the range $x_{\text{int}} \in [x_{i-1}, x_i]$. The interface

location is then given by

$$x_{\text{int}} = \frac{x_{i-1}\hat{\phi}_i - x_i\hat{\phi}_{i-1}}{\hat{\phi}_i - \hat{\phi}_{i-1}}. \quad (3.7.10)$$

We can then make a linear extrapolation of the 3-velocity at the interface,

$$v_{\text{int}} = \frac{(v_i - v_{i-1})x_{\text{int}} + (v_{i-1}x_i - v_ix_{i-1})}{\Delta x}. \quad (3.7.11)$$

This then provides us with a second-order accurate approximation of v_{int} . Higher order extrapolation can be used if necessary, but since ϕ will usually be close to linear, these do not offer a significant advantage.

It is not always the case that the advection equation is sufficient for describing the evolution of ϕ . If additional physics is added to the model (i.e. we move away from an ideal fluid EOS) then this may impact the behaviour of the interface. For example, in some cases there will be non-negligible surface tension between the two fluid components. For further information about how such situations can be modelled see [148].

3.7.3 Numerical methods for the level set function

So far in this section we have stated that the level set function ϕ will be evolved using the non-conservative form of the advection equation ((3.7.6) or (3.7.7)). These will differ from the HRSC methods that have been the focus of this chapter, since these require conservation form of the evolution equations. Many of the features introduced so far, such as discretisation, still apply however.

The advection equations, given in (3.7.6) and (3.7.7) are specific examples of Hamilton-Jacobi equations [148]. The general form of such equations is

$$\frac{\partial \phi}{\partial t} + H(\nabla \phi) = 0. \quad (3.7.12)$$

In the cases we consider, we therefore have

$$H(\nabla\phi) = v^i \frac{\partial\phi}{\partial x^i} \quad (3.7.13)$$

for the Newtonian and SR advection equations, (3.7.6), and

$$H(\nabla\phi) = \alpha v^r \frac{\partial\phi}{\partial r} \quad (3.7.14)$$

for the GR equations in polar-areal coordinates, (3.7.7).

Many techniques have been developed to solve Hamilton-Jacobi equations and here, we shall employ one such technique, a Lax-Friedrichs scheme as described in [50]. As with the fluid evolution equations, we begin by making a discretisation of H such that (3.7.12) reduces to an ODE. At this point, any ODE solver, such as the RK methods outlined in section 3.2, can be used to solve this equation. One possible discretisation we can make for a Hamilton-Jacobi equation is

$$H(\nabla\phi) = \hat{H}(\phi_x^-, \phi_x^+, \phi_y^-, \phi_y^+, \phi_z^-, \phi_z^+), \quad (3.7.15)$$

where e.g. ϕ_x^\pm represents some choice of upwind and downwind approximations to the derivative of ϕ in the x -direction. For the Lax-Friedrichs scheme, \hat{H} is defined by

$$\begin{aligned} \hat{H} = & H \left[\frac{1}{2} (\phi_x^- + \phi_x^+), \frac{1}{2} (\phi_y^- + \phi_y^+), \frac{1}{2} (\phi_z^- + \phi_z^+) \right] \\ & - \hat{\alpha}^x \left[\frac{1}{2} (\phi_x^+ - \phi_x^-) \right] - \hat{\alpha}^y \left[\frac{1}{2} (\phi_y^+ - \phi_y^-) \right] - \hat{\alpha}^z \left[\frac{1}{2} (\phi_z^+ - \phi_z^-) \right] \end{aligned} \quad (3.7.16)$$

where $\hat{\alpha}^i$ are dissipation coefficients (note these are not related to the lapse function α). In order to determine these coefficients, we first define coefficients H_1 , H_2 and H_3 through

$$H_{1,(2,3)} = \frac{dH}{d(\partial_{x,(y,z)}\phi)}. \quad (3.7.17)$$

We then choose

$$\hat{\alpha}^x = \max(H_1) \quad (3.7.18)$$

where the maximum value over the entire grid is taken, and the obvious result for $\hat{\alpha}^y$ and $\hat{\alpha}^z$. It is not necessarily straightforward to calculate the $\hat{\alpha}^i$ coefficients, since taking derivatives of $H(\nabla\phi)$ with respect to $\partial_i\phi$ may not be obvious. Fortunately, in the case of the advection equation, we have a simple form for H , hence these coefficients are obvious. For Newtonian and SR, with the evolution of ϕ given by (3.7.13), we have

$$\hat{\alpha}^i = |v^i| \quad (3.7.19)$$

and for polar areal coordinates in GR, (3.7.14), we have

$$\hat{\alpha}^r = |\alpha v^r|. \quad (3.7.20)$$

In the 1+1 dimensional cases considered in section 4, ϕ should retain an approximately linear shape (or be reinitialised to such a shape) around the interface, it is sufficient for first-order approximations to be used for ϕ^\pm ,

$$\phi^+ = \frac{\hat{\phi}_{i+1} - \hat{\phi}_i}{\Delta x}, \quad \phi^- = \frac{\hat{\phi}_i - \hat{\phi}_{i-1}}{\Delta x}. \quad (3.7.21)$$

Higher order methods for obtaining these approximations are required for reasonable accuracy in multidimensional situations however, see section 5.3. This scheme is far from the only such method available for solving Hamiltonian-Jacobi equations, see e.g. [148] for further options.

3.7.4 Ghost Fluid Method

Level set methods, outlined in the previous section, are a robust and widely used set of methods for capturing a sharp feature such as an interface between two fluid components. Unfortunately, by themselves levels set methods do not give enough information to accurately model the physical behaviour of these fluid components. Level set methods simply give a location of the interface and do not in any way affect the evolution of the fluid. The actual behaviour of the fluid components at the interface are then modelled through the implementation of some dynamical boundary conditions. Several techniques have been developed in Newtonian CFD to implement suitable boundary conditions, e.g. [67, 129]. Whether these methods extend to relativistic situations has not

previously been investigated. We will make this extension, and therefore, it is advantageous for this work if we start from a simple, robust method. As a result, we will work with the Ghost Fluid Method of Fedkiw et al. [67].

The Ghost Fluid Method is the simplest successful method for providing boundary conditions at an interface using level set techniques. The techniques employed by the Ghost Fluid Method are reminiscent of those used to provide the boundary conditions at the edge of the computational domain in section 3.5. For the boundary between two fluid components, the conditions implemented must, however, be dynamic: they must move with the interface and must also capture the effects of the change in EOS at the interface.

When using level set methods to capture the location of the interface, we are able to model the fluid components separately. The final solution of an evolution using such methods is then pieced back together, with the sign of the level set function determining the physical extent of each component. This means that there are no actual constraints on the computational domain over which each fluid component is modelled.

In figure 3.6 we show the computational domain for one component of a multicomponent model. This component only enters the physical domain when $\phi < 0$, where it is shown by the solid line. When $\phi > 0$, this component will have no effect on the overall evolution of the fluid, even if it is included in the computational domain of these regions, as illustrated by the dashed line. This allows us to implement boundary conditions in a similar manner to the way ghost zones are implemented. That is we populate cells ‘behind’ and interface with artificial values such that the behaviour at the interface is appropriately modelled. This artificially altered fluid is the *ghost fluid*. Since the ghost fluid is created in the region not present in the physical model, this alteration will not manifest itself at any point other than through the provision of boundary conditions.

It should be noted that the Ghost Fluid Method is not actually a conservative method. When we reconstruct fluxes at a cell boundary, as described in section 3.4, we would normally obtain a single flux e.g. $\bar{\mathbf{f}}_{i+1/2}$. This then describes both the flux through the right boundary of the cell x_i and the left boundary of x_{i+1} . When using the Ghost Fluid Method, however, these two




Figure 3.6: Illustrating the region across which a single fluid component can be modelled. Although the component depicted only enters the physical model when $\phi < 0$, it can be modelled across the entire domain. When $\phi > 0$, this fluid component does not effect the physical solution, hence allowing it to be modelled across the entire domain does not present an issue. These regions then form the ghost fluid region described in this section.

fluxes are not equal. If an interface exists between x_i and x_{i+1} then the flux through the right boundary of x_i is computed using the ghost fluid for one fluid component. Similarly the flux through the left boundary of x_{i+1} is computed using the ghost fluid for the other component. Obviously there is no guarantee that these two fluxes will then be equal. Despite this reservation, the method has been used successfully in many cases. This does however suggest that care should be taken in some situations, particularly in some extreme scenarios (e.g. those involving high velocity shocks).

There are many advantages, beyond simplicity, for using the Ghost Fluid Method as the basis for this extension. It is an extremely versatile method, and many extensions we may wish to consider in NS modelling are investigated in the context of Newtonian CFD. The compatibility of the Ghost Fluid Method with AMR has been investigated, [143, 206], which is essential for use in e.g. NS merger simulations. It has also been shown that the method can be used in cases in which an Eulerian description of a fluid can be coupled to a Lagrangian description of a solid, [66], which would allow for e.g. a solid core to be simulated within a NS. We also note that the Ghost Fluid Method is not restricted solely to use with level set methods, but can be used with other sharp interface tracking techniques [198].

We now detail how the Ghost Fluid Method provides these boundary conditions. The conditions are based on the expected behaviour of the primitive variables (ρ_0 , v and p) across the interface. As we detailed in section 3.7.1, we know that normal velocity and pressure should remain continuous. To reflect this, the ghost fluid simply takes the corresponding values for these variables directly from the physical fluid. There is no such condition, however, on the density and tangential velocity at the interface as there are no constraints on the size of the jumps in these variables. The Ghost Fluid Method makes the *assumption* that entropy and tangential velocity are constant across the interface. As such, they are extrapolated into the ghost fluid region at zeroth-order in the normal direction to the an interface. From the entropy extrapolation, the density can then be recovered using this extrapolated value, the EOS and pressure. This assumption clearly will not always be correct; for example, when a shock hits an interface there is a jump in entropy governed by the Rankine-Hugoniot conditions. However, since this simple assumption has been used

successfully in a wide range of cases in Newtonian CFD, we expect that it will extend naturally to SR.

The manner in which the Ghost Fluid Method is implemented is considered explicitly for a 1+1 dimensional two-component fluid system. Note that in this case the single velocity component is always normal to the interface, so we do not have a tangential component to extrapolate. Here we assume a total computational domain of $x_i \in [x_{M+1}, x_N]$, in order to allow for consistent notation.

1. Beginning with an initial cell, x_{M+1} , which represents either the start of the domain, or the first cell after a previous interface. Identify the location of the next interface based on the point at which the sign of ϕ changes, i.e. $\hat{\phi}_{i-1}\hat{\phi}_i < 0$. Alternatively, the end of the domain could be reached. We assign $x_{\mathfrak{J}} = x_i$, the final cell before this next interface. This introduces a region, x_j , $j \in [M+1, \mathfrak{J}]$, for which the Ghost Fluid Method can be applied at the boundaries. In this region, we have a single fluid component with a single EOS.
2. The Ghost Fluid Method must be applied in the region x_l , $l \in [M, M - n_g]$ and x_r , $r \in [\mathfrak{J} + 1, \mathfrak{J} + n_g + 1]$, where n_g is the number of ghost zones required by the numerical method being used. If either x_{M+1} or $x_{\mathfrak{J}}$ are a cell at the edge of the physical domain, we ignore this region, since there is no interface present. The additional cell on top of n_g is required, since the interface may move by one cell during the evolution. If this happens before the final corrector step for solving the ODE provided by the MOL (see section 3.2) then the information provided by this additional cell will be required.
3. Copy the physical values of the velocity and pressure into the ghost fluid regions x_l and x_r . This comes from the other fluid component at these cells.
4. Make the zeroth-order extrapolation of entropy from x_{M+1} into x_l and $x_{\mathfrak{J}}$ into x_r . In section 2.3 we gave the entropy for an ideal fluid, (2.3.17),

$$s = c_v \ln p - c_p \ln \rho_0 + \text{const.} \quad (3.7.22)$$




Figure 3.7: An illustration of how the Ghost Fluid Method is applied. We show how the ghost fluid region of fluid 1 is populated. The crosses and squares represent the fluid that would be seen in the physical model, and the circles the artificial ghost fluid. Analogous modification for the boundaries of fluid 2 can easily be inferred from this figure.

If this is constant we have the relation $p \propto \rho_0^\gamma$. We can then implement this extrapolation directly, with the density within the ghost fluid region given by

$$\rho_{0,l} = \rho_{0,M+1} \left(\frac{p_l}{p_{M+1}} \right)^{1/\gamma}, \quad \rho_{0,r} = \rho_{0,3} \left(\frac{p_r}{p_3} \right)^{1/\gamma}. \quad (3.7.23)$$

5. Update the fluid within x_j using the desired HRSC method.

This procedure must be carried out for every region between two interfaces, or between an interface and the boundary of the computational domain, at every substep within the numerical method. To offer greater insight as to how the Ghost Fluid Method is applied, figure 3.7 illustrates this 1+1 dimensional implementation of one fluid component at one interface.

The description given so far was initially devised for a fluid in Newtonian physics. In relativistic situations we would expect that the 4-velocity normal to the interface should be copied from the real fluid to the ghost fluid. When working in 1+1 dimensions however, the continuity of the velocity (and space-time factors in GR) across the interface mean that the 3-velocity can be used

(as in the Newtonian case) without issue. When moving to multidimensions, the coupling of the velocity components through the Lorentz factor means that the full 4-velocity may need to be considered. This is discussed in further detail in section 5.

3.7.5 Alternative methods

The Ghost Fluid Method has been widely used in Newtonian CFD, with many applications shown in [67, 99]. Perhaps not surprisingly, the simplicity of the technique, in particular the zeroth-order entropy extrapolation, does result in cases in which the Ghost Fluid Method suffers inaccuracies, or fails altogether. These have been studied in e.g. [125, 126], where the interaction of strong shocks with the interface was seen to cause issues. The cause of these issues lies in the fact that the Ghost Fluid boundary conditions make the assumption that a shock impacting an interface is reflected as a rarefaction wave. If a shock is sufficiently strong, however, then a reflected shock wave can form instead (or, under certain conditions, no reflected wave is seen). When these cases were investigated using the original Ghost Fluid Method, there were spurious oscillations introduced at the interface, or in some cases failure to capture its location accurately. In [125], these issues were addressed with a modified Ghost Fluid Method. This effectively ensures that the boundary conditions are appropriate for the form of the reflected wave. This allows for the entropy to be extrapolated into the ghost fluid region from the exact location of the interface, instead of the first cell before the interface as shown in figure 3.7.

In addition to modifying the Ghost Fluid Method, further issues have been addressed through alternative treatments of the boundary conditions at the interface. One such method is the *Explicit Simplified Interface Method* (ESIM), introduced for acoustic problems in [157] and extended to multicomponent fluid problems in [129]. This method is concerned with correcting the problem the Ghost Fluid Method experiences due to the zeroth-order entropy extrapolation at the interface. In the original method, this lead to issues with the accuracy with which behaviour around the interface is captured if entropy gradients exist across it. The ESIM method makes a higher order reconstruction at the interface and hence addresses these issues. However, ESIM is still not a conservative method, hence care must still be taken in the presence of non-

linear features.

Further alternatives exist which do not use level set techniques. If a single component model can be devised in a thermodynamically consistent manner (through an appropriate choice of colour function), then a conservative method for dealing with fluid systems with a change in EOS will be available. Such a method was presented in [207], the *thermodynamically consistent and fully conservative* (TCFC) method. Despite having many of the advantages of a conservative method, the TCFC method has other issues, such as dealing with situations involving weak shocks. Additionally, by using a single component method, there will be unavoidable smearing of the interface.

3.8 Atmosphere treatment

One major problem with numerical evolutions of compact objects is how the surface is treated, and how we deal with the vacuum region beyond. Ideally we would want a NS model to include the outer regions, the crust, envelope and atmosphere, as described in section 1.2. If these could be included in a simulation, then the correct physics might be modelled at the surface. This is not necessarily straightforward for a model which includes the entire NS, however, especially for the thin regions (considering the argument given in section 3.7). In practice, current NS simulations assume that the region near the surface of the NS is composed of the same matter as the rest of the star. Therefore the surface will be determined purely by the point at which the density and pressure drop to zero. Here the assumption has been made that outside the star there is vacuum which, for most scenarios involving a cold NS, is reasonable.

Vanishing density and pressure can present several issues computationally, e.g when converting from conserved to primitive variables as described in section 2.7. Close to the surface, the magnitude of these variables means that small errors can lead to unphysical values for these variables, e.g. negative density or pressure. This can then lead to further unphysical behaviour which propagates into the star.

Currently no satisfactory techniques have been developed to tackle the

errors resulting from the surface in a physical manner. Here we use a standard treatment which, in many cases, does not have a significant effect on the behaviour of the system. An artificial atmosphere is imposed in the vacuum region of the star, as used in e.g. [54, 140, 142], and hence we remove the issues of vanishing quantities. This treatment typically involves setting a (constant) low density, static and polytropic atmosphere. In many cases this has been found to produce good results for reasonably long simulations. The artificial nature of this atmosphere can, however, lead to many undesirable effects. It is obvious that by adding additional matter to the exterior of the star, it is being moved away from true hydrostatic equilibrium. This can lead to matter accreting onto the star, and if this accretion is large enough then this will lead to unphysical behaviour. It has been found that if the atmosphere is a sufficiently small proportion of the central pressure, p_c , then the effects are negligible over many dynamical timescales of the star. Typically in this work we select an atmosphere pressure $p_{\text{atm}} \sim 10^{-12} p_c$. At the start of every timestep within an evolution, the atmosphere is then reset to this constant level. This means propagation of behaviour from within the atmosphere into the star is limited.

In addition to the issues with accretion onto the star, the behaviour of a perturbed star at the surface is not going to be accurately reproduced using the atmosphere technique. In the models used in this thesis, the surface of the star can be thought of as a boundary between fluid and vacuum, therefore whatever treatment is used should model these conditions. In the absence of satisfactory boundary conditions, we therefore cannot guarantee that there will be no unphysical reflections from the surface, or information lost into the atmosphere (which is indeed major issues in many simulations). It is possible that through use of level set techniques and appropriate Ghost Fluid-like boundary conditions that some suitable genuine vacuum surface may be able to be implemented. Some initial attempts to produce such a boundary condition were attempted, however these were not successful. Further considerations as to how this can be achieved in terms of the future work that may result from this thesis will be discussed in section 6.1.

Chapter 4

Results

In this chapter we present the results for modelling interfaces in relativistic hydrodynamics using the techniques outlined in the previous section. To enable better interpretation of the results from these relativistic scenarios, we first consider some reference solutions in Newtonian hydrodynamics. We then present the initial tests of the Ghost Fluid Method in SR, and then move on to a simple 2-component star in GR. Many of the results presented in this section are published in [135]. All plots in this chapter show tests in which a MUSCL scheme has been used with MC limiter to make the reconstruction, HLLE or the Marquina solver to obtain the fluxes and a RK2 time evolution, unless otherwise stated. Furthermore, the presence of a solid line in all plots in Newtonian physics and SR represents an exact solution, unless otherwise stated.

4.1 Preliminary results for Newtonian hydrodynamics

In order to test the Ghost Fluid Method in relativistic situations, it is imperative that we have some techniques available to gain both qualitative and quantitative analysis for the accuracy of the results. In section 2.4.3 we detailed how the Riemann shock tube problem was one of the few cases for which there was an exact solution for the ideal fluid evolution equations in SR. To fully test the Ghost Fluid Method in relativistic situations, however, further, more demanding tests must be considered, and these will not have exact so-

lutions. A qualitative analysis of the success of the relativistic Ghost Fluid Method can still be achieved in these solutions through the consideration of analogous cases in Newtonian hydrodynamics. This section details these results, obtained through an implementation of the original Ghost Fluid Method. All tests in this section are set up on a domain $x \in [0, 1]$.

Many results we will show in this chapter are based on the Riemann shock tube problem. It is therefore instructive as a first result to simply display the solution to a classical one component Riemann problem. In addition to this, it is useful to ensure that the fluid evolution equations are correctly implemented. The chosen test is often referred to as Sod's test problem [188], and is a test designed to produce a clear example of all three wave types that result from a Riemann problem, the rarefaction, shock and contact discontinuity.

The initial data for this test gives left and right states, \mathbf{q}_L and \mathbf{q}_R , by

$$\begin{cases} \rho_0 = 1, & v = 0, & p = 1, & x < 0.5 \\ \rho_0 = 0.1, & v = 0, & p = 0.125, & x \geq 0.5 \end{cases} . \quad (4.1.1)$$

The test is run to a final time of 0.2 for an ideal fluid with $\gamma = 1.4$.

Results from this test are shown in figure 4.1. We see from this test that all features, the rarefaction, shock and contact discontinuity have been correctly positioned, and the latter of which appears only in the density profile. This test shows very clearly the smearing issues that the numerical techniques for a single component fluid suffer from at a contact discontinuity. We also note that the location of the shock has been captured correctly without oscillatory behaviour for the two resolutions shown due to the use of HRSC methods. As expected, the results converge towards the exact solution with increasing resolution.

4.1.1 Motivation for use of the Ghost Fluid Method

The test shown in the previous section shows the smearing inherent within single component fluid models. The major problem, however, with modelling a fluid system as a single component comes when we desire different regions to be governed by individual EOSs. In such a system, we will need to introduce an additional equation which governs the movement of these different regions. The




Figure 4.1: Results for the Sod test. Two resolutions are shown, 200 cells and 800 cells, though for each resolution, only 100 cells are shown. All three expected features, the rarefaction, shock and contact discontinuity, are clearly identifiable, and correctly located. This test shows the smearing that single component simulations have around contact discontinuities. We do, however, see that there is convergence with resolution for all features, as desired.

obvious choice of such an equation is a “colour function”, as detailed in [123], which, for an ideal fluid, would simply give the value of γ as a function of time and space. In the situations we consider, the interfaces between regions behave as contact discontinuities, therefore the colour function is simply governed by an advection equation,

$$\partial_t (\rho_0(x, t) \gamma(x, t)) + \partial_i (\rho_0(x, t) \gamma(x, t) v^i(x, t)) = 0. \quad (4.1.2)$$

Here we investigate a simple situation where the only physical feature is this interface, visible through the corresponding jump in density. A non-trivial velocity is introduced, hence the interface should simply be advected with this fluid. This test is known to be physically stable, hence any failure within the solution must be numerical. The initial data for this test is

$$\begin{cases} \rho_0 = 1, & v = 0.1, & p = \frac{2}{3}, & x < 0.5 \\ \rho_0 = \frac{1}{2}, & v = 0.1, & p = \frac{2}{3}, & x \geq 0.5 \end{cases} \quad (4.1.3)$$

with the colour function given by

$$\gamma(x) = \begin{cases} \frac{5}{3} & x < 0.5 \\ \frac{4}{3} & x \geq 0.5 \end{cases} \quad (4.1.4)$$

Figure 4.2 shows the results for the colour function at a time $t = 0.2$. It is clear that even after this short time, severe problems have occurred with the solution. Around the interface in all three primitive variables large oscillations have been introduced. Furthermore, there is no clean convergence with resolution, hence this technique can not be used to get accurate results.

The reasons for this failure have been documented in e.g. [6]. As mentioned previously HRSC methods introduce some level of smearing, particularly around a contact discontinuity. This smearing will be apparent in the colour function, and hence effectively “mixes” the two fluid types. How this mixing occurs will in some way depend on the resolution of the model, but is unavoidable. This numerical (and hence unphysical) mixing leads to spurious generation of entropy at the interface, hence the method is not thermodynamically consistent. It is this inconsistency that leads to the formation of large




Figure 4.2: The failure of using the colour function (4.1.4) to model a stable moving contact discontinuity. This plot shows the results for two resolutions, 200 and 800 cells, after a time $t = 0.2$. It is clear that spurious non-convergent oscillations have been introduced after this short time, hence the colour function has failed to model the correct physical behaviour of the interface.

errors as shown in figure 4.2. Some of the alternative techniques for modelling multicomponent fluids discussed earlier, e.g. [207], lead to thermodynamically consistent methods to deal with the change in EOS. The issue with the loss of accuracy due to the smearing of the interface remains in this case however.

4.1.2 The Newtonian Ghost Fluid Method

The first reference solution we will show using the original Newtonian Ghost Fluid Method uses the same initial data as in (4.1.1), but we introduce a “trivial” interface, i.e. separating two fluid components with the same EOS. Therefore in order to include the initial discontinuity at $x = 0.5$, we include a level set function

$$\phi = x - 0.5. \quad (4.1.5)$$

The results of this test, shown in figure 4.3, can be compared directly to figure 4.1 and show the same behaviour for the shock and rarefaction. In this case, use of the Ghost Fluid Method means the contact discontinuity no longer smeared. We see convergence towards the exact solution across almost the entire domain. We do, however, see an undershoot in the density profile at the contact discontinuity. This is an artifact of combining the Ghost Fluid Method with initial data which is not in some equilibrium state across the interface, which results in the technique introducing an artificial ‘overheating’ error [67, 68]. This can be overcome through the application of the ‘isobaric fix’ discussed in these references. This technique, for an interface located within the cell x_i , uses the entropy from cell x_{i-2} to make the constant extrapolation, but otherwise is the same method discussed in section 3.7.4. The situations we consider within this thesis, beyond the initial tests, do not consider such initial data, hence we do not implement the isobaric fix. We can show that this is not necessary through a further test.

The test we shall consider here is an isolated shock hitting an interface. This was used to test the Ghost Fluid Method in the original paper [67], as “Test B”. In this test, we introduce a “genuine” interface, in which the EOS changes.




Figure 4.3: Results for the Sod test with an “artificial” interface initially positioned at $x = 0.5$ such that the initial data is equivalent to that in (4.1.1). These results should therefore be compared to figure 4.1. We see the same behaviour for the shock and rarefaction, as expected. The contact discontinuity is now been captured perfectly. There is, however, a clear undershoot either side of the contact discontinuity. This is not a physical effect, and is in fact due to the overheating errors experienced by methods such as the Ghost Fluid Method. Resolutions of 200 and 800 cells are shown, and only 100 cells of each resolution are shown.

The initial data is given by

$$\begin{cases} \rho_0 = 1.3333, & v = 0.3535\sqrt{10^5}, & p = 1.5 \times 10^5, & x < 0.05 \\ \rho_0 = 1, & v = 0, & p = 1 \times 10^5, & 0.05 \leq x < 0.5 \\ \rho_0 = 0.1379, & v = 0, & p = 1 \times 10^5, & x \geq 0.5 \end{cases} \quad (4.1.6)$$

with the jump in EOS given by

$$\begin{cases} \gamma = 1.4, & x < 0.5 \\ \gamma = 1.67, & x \geq 0.5 \end{cases} \quad (4.1.7)$$

The level set function is therefore initialised to (4.1.5). This test is run to a time $t = 0.0012$.

Results for an isolated shock hitting an interface are shown in figure 4.4. For such a test it is expected that once the shock has hit the interface then some of it will be transmitted through the interface, and some will be reflected as a rarefaction wave, along with the interface moving as a result of this interaction. This is clearly apparent in the results, which can be compared to an exact solution. This comparison is possible since the instant the shock hits the interface we have a Riemann problem with two constant left and right states. As expected, in this test, we do not suffer the overheating issues at the interface, hence we do not see the undershoots. As a result, we see convergence towards the exact solution everywhere. It is, however, clear that at lower resolution, the shock has not been captured very tightly.

A final reference test we shall consider is one for which an exact solution does not exist. The test we consider is, however, also considered in [207], hence we can ensure that our implementation of the Ghost Fluid Method gives comparable results. This test is set up such that an isolated shock hits a slab of lower density material. The behaviour in this test is more complex since the presence of two interfaces means the transmitted and reflected waves can interact with other interfaces and with each other. In addition to this, since the slab is of lower density, we would expect the impact to both move




Figure 4.4: Results for the “Test B” from [67] in which an isolated shock hits an genuine interface. We see the shock transmitted through the interface as well as the reflected rarefaction. The interface has been captured sharply, and the undershoots seen in figure 4.3 are no longer present due to the equilibrium across the interface in the initial data. We also see that at the lower resolution, the capturing of the shock is not particularly tight. Resolutions of 200 and 800 cells are shown, and only 100 cells of each resolution are shown.

and compress the slab. The initial data for this test is

$$\begin{cases} \rho_0 = 1.3795, & v = 0.3948, & p = 1.57, & x < 0.25 \\ \rho_0 = 1, & v = 0, & p = 1, & 0.25 \leq x < 0.4 \\ \rho_0 = 0.138, & v = 0, & p = 1, & 0.4 \leq x < 0.6 \\ \rho_0 = 1, & v = 0, & p = 1, & x \geq 0.6 \end{cases} \quad (4.1.8)$$

with the EOS parametrised by

$$\begin{cases} \gamma = 1.4, & x < 0.4 \\ \gamma = 1.67, & 0.4 \leq x < 0.6 \\ \gamma = 1.4, & x \geq 0.6 \end{cases} . \quad (4.1.9)$$

The level set function is initially given by

$$\phi = 0.1 - |x - 0.5| \quad (4.1.10)$$

and the test is run to time $t = 0.3$.

The results for an isolated shock hitting a low-density slab of material are presented in figure 4.5. The movement and compression of the slab are instantly obvious. The leading shock transmitted through both interfaces, and rarefaction from impact with the first interface are apparent. In addition, complex behaviour due to interactions between the further reflected and transmitted features is clearly apparent. Comparison to the original presentation of this test in [207] shows that this behaviour has been correctly modelled. It can also clearly be seen that the various features are more pronounced and accurately captured at higher resolution, with convergence towards some solution. In figure 4.6 results for a low-resolution simulation (100 cells) of this test are shown. This allows every point to be seen, hence we can ascertain that there are no features being missed by a ‘selection effect’ of the reduced sample of points shown in previous figures. This test was chosen to show this since due to its more complicated nature, we see that there are no issues with unphysical behaviour being generated for the Ghost Fluid Method in Newtonian physics. In this figure, we have also included the specific internal energy, ε , since errors are often prominent here.




Figure 4.5: Results for the “Test B” (as documented in [207]) in which an isolated shock hits a slab of low-density material. The leading transmitted shock and reflected rarefaction as a result of the initial shock hitting the two interfaces are apparent. There are many more details that can also be seen as a result of further interactions of shocks and rarefactions with the interfaces. The slab has moved from its initial location and contracted as a result of these effects. Resolutions of 200 and 800 cells are shown, and only 100 cells of each resolution are shown.




Figure 4.6: The results for a shock hitting a slab of material at a resolution of 100 cells. This allows us to see every cell, to ensure that there is no hidden behaviour missed in the higher resolution tests. We see the expected reduction in accuracy with which the features are captured. However, we also see there are no features that are hidden due to a selection effect in previous plots for the original Ghost Fluid Method.

4.2 The Ghost Fluid Method in special relativity

In this section we present the results from extending the Ghost Fluid Method to SR. We first consider some simple tests for which exact solutions exist, hence we can ensure that the correct results are being obtained. Since exact solutions typically only exist for comparatively simple situations, in order to fully test the robustness of the method, we will also consider some more complicated tests (such as an analogy of the “slab test” shown in figure 4.5). We will use these tests to both investigate how the Ghost Fluid Method copes in these non-trivial situations, and also to investigate some of the known sources of error that were identified for the Newtonian method, such as those suggested in [125, 126]. In this section, unless otherwise stated, tests are carried out on a domain $x \in [0, 1]$.

We will begin by ensuring that the Ghost Fluid Method can be used to model a Sod-type test in the same manner as given in figure 4.3. This test again has initial data

$$\begin{cases} \rho_0 = 1, & v = 0, & p = 1, & x < 0.5 \\ \rho_0 = 0.1, & v = 0, & p = 0.125, & x \geq 0.5 \end{cases} \quad (4.2.1)$$

and the initial level set function by

$$\phi = x - 0.5. \quad (4.2.2)$$

We use a fluid with $\gamma = 1.4$ and the test is now run to a time $t = 0.5$ to reflect the differences in the scales between Newtonian and relativistic situations.

Figure 4.7 shows the results of this Sod-type test in SR. For this test, we can clearly see that the features have been correctly positioned, and that these are converging towards the exact solution with resolution. As with the Newtonian Ghost Fluid Method, the overheating errors at the interface are apparent.

True tests of the Ghost Fluid Method will require a genuine interface over which the EOS changes. For the remainder of this section we will therefore




Figure 4.7: Results for a Sod-type test with an “artificial” interface in SR, as shown for an analogous Newtonian case in figure 4.3. Comparison with the exact solution reveals that the features have been correctly captured with similar undershoots observed due to the discontinuous initial data. Resolutions of 200 and 800 cells are shown, and only 100 cells of each resolution are shown.

focus solely on such cases. The first test to consider is still comparatively simple, and will investigate whether the Ghost Fluid Method really does have advantages over a single component model. To achieve this we recreate the scenario shown in figure 4.2, the advected contact discontinuity. For the Ghost Fluid Method to be viable in SR, it must be able to model this situation without introducing oscillations and maintaining a sharp contact discontinuity. The initial data will be the same as in (4.1.3),

$$\begin{cases} \rho_0 = 1, & v = 0.1, & p = \frac{2}{3}, & x < 0.5 \\ \rho_0 = \frac{1}{2}, & v = 0.1, & p = \frac{2}{3}, & x \geq 0.5 \end{cases} \quad (4.2.3)$$

with γ initially given by

$$\begin{cases} \gamma = \frac{5}{3}, & x < 0.5 \\ \gamma = \frac{4}{3}, & x \geq 0.5 \end{cases} \quad (4.2.4)$$

and the level set, as in (4.1.5), by

$$\phi = x - 0.5. \quad (4.2.5)$$

This test was run to a final time $t = 2$.

The results for the moving contact discontinuity in SR are shown in figure 4.8. The results in the left panels are plotted at $t = 0.2$, allowing for a direct comparison can be made to figure 4.2, where we saw the failure of this test in single component Newtonian model. The simple nature of this test, in which the pressure balance throughout the system means that the advection of the contact discontinuity is the only behaviour expected, hence this direct comparison between the two cases can be made. Here we see that by using the Ghost Fluid Method, no oscillations have been introduced. We also see that at later times, $t = 2$, the interface has been advected with the correct velocity, $v = 0.1$, hence now lies at $x = 0.7$.

In order to ensure that the interface has indeed been correctly located by the level set function, figure 4.9 shows a zoomed region of the density profile of this moving contact discontinuity test at $t = 2$. For all resolutions shown, it




Figure 4.8: A moving contact discontinuity for a multicomponent fluid in SR for resolutions of 200 and 800 cells. The left hand set of plots for this test can be compared to figure 4.2, where the results are plotted at the same time $t = 0.2$. It is clear that the oscillation experienced by the single component model are no longer present. The right hand set of plots shows the same set of variables plotted at the later time $t = 2$. We can see that even at this late time, no oscillations have been introduced.




Figure 4.9: Focusing on how accurately the contact discontinuity is captured for a moving contact discontinuity in SR. This is a zoom in on the $t = 2$ density profile shown in figure 4.8. We can see that there have been no issues in capturing the position of the interface correctly and without oscillation.

is clear that this position has been correctly located, with no oscillatory behaviour introduced.

The success of the Ghost Fluid Method in these preliminary tests means we now want to consider tests in which the fluid does not have constant velocity and pressure everywhere hence some evolution of the fluid variables will occur. In particular we focus on some examples of the Riemann problem in SR. In these cases, we will see that the Ghost Fluid Method suffers from the overheating issues at the interfaces, hence the undershoots seen in e.g. figure 4.7, will again be present in many of these tests. Despite this, these will still be valuable for testing the SR Ghost Fluid Method since, as detailed in section 2.4.4, exact solutions exist for these Riemann problem cases, we can still determine if the resultant waves are correctly captured.

The first test we consider is formulated such that it results in two rarefaction waves moving in opposite directions. This test has initial data

$$\begin{cases} \rho_0 = 1, & v = -0.3, & p = 1, & x < 0.5 \\ \rho_0 = 1, & v = 0.3, & p = 1, & x \geq 0.5 \end{cases} . \quad (4.2.6)$$

The EOS is parametrised by (4.2.4), and level set function given by (4.2.5), as was used for the moving contact discontinuity case. This test is run to time $t = 0.4$.

The results for the Riemann problem with two rarefactions moving away from each other are shown in figure 4.10. It is clear that the positions of both rarefactions have been correctly captured (note that the speed which the rarefactions themselves move is not equal due to the differences in EOS in each component), and the accuracy of the solution increases with resolution. As with the previous test, we will focus on the region around the contact discontinuity to ensure that its position has been located correctly without introducing oscillatory behaviour.

A zoomed view of the contact discontinuity in this test is shown in figure 4.11. The undershoots caused by the overheating are clearly apparent, but no other




Figure 4.10: A Riemann problem in SR in which two rarefactions are formed by two fluid components initially moving away from each other. We see that both rarefactions and the contact discontinuity have been correctly captured, and that the accuracy of the solution increases with resolution. For clarity only 100 cells of each resolution have been used.




Figure 4.11: The density profile around the contact discontinuity for the SR Riemann problem in which two rarefactions move away from each other, as shown in full in figure 4.10. We see that the jump in EOS means that the contact discontinuity moves slightly from its initial location.

unphysical behaviour.

A similar Riemann problem can be considered in which two shock waves move away from each other. This will test whether the Ghost Fluid Method can be used with the HRSC methods employed to capture shocks without oscillation. Initial data is similar to that given in the previous test,

$$\begin{cases} \rho_0 = 1, & v = 0.3, & p = 1, & x < 0.5 \\ \rho_0 = 1, & v = -0.3, & p = 1, & x \geq 0.5 \end{cases} \quad (4.2.7)$$

and again the value of γ is given by (4.2.4), the level set function by (4.2.5) and the test is run to $t = 0.4$.

The results for this test are shown in figure 4.12. We see that the position of the shocks is correct, and that the accuracy with which the shocks are captured increases with resolution. Similarly the contact discontinuity has been correctly located, and shows the expected overheating error. It is again worth focusing on the region around the contact discontinuity to ensure that the ac-




Figure 4.12: Results for a Riemann problem in SR in which two shock waves are moving away from each other (reflected off of each other). The shock position has been captured accurately without introducing oscillations, and the accuracy of the results increases with resolution. For clarity, only 100 cells of each resolution are shown.




Figure 4.13: The density profile around the interface for the test in which two shock waves move apart from each other, shown in full in figure 4.12. We see that aside from the expected overheating errors, there have been no issues with capturing the position of this feature.

curacy with which it is captured genuinely does improve with resolution.

In figure 4.13 we see the contact discontinuity has been correctly captured for the test in which two shock waves move apart. It is also apparent that there are no oscillations introduced at the contact discontinuity, with the only non-convergent error from the overheating discussed previously.

It is also worth ensuring that the SR Ghost Fluid Method experiences no issues in dealing with the two shock waves in this test. This will ensure that the HRSC methods are working correctly.

The behaviour around the two shock waves is shown in figure 4.14. We see that in all cases the shocks are captured over about four grid cells due to the use of HRSC methods. There are also no unexpected oscillations introduced at the shocks, showing the slope-limiting techniques employed are working successfully.

We now consider a test for which will induce more significant movement of




Figure 4.14: The density profile for the two shocks produced for the SR Riemann problem with initial data given by (4.2.7). The ability of the HRSC to capture the shocks over only a few grid cells is clearly visible. As a result, the accuracy with which these shocks are captured increases with resolution. We can also see that there is no severe oscillation at either of the shocks.

the interface between the two fluid components. This test is an analogy in SR to figure 4.4 (“test B” from [67]), where this scenario was presented for the Newtonian Ghost Fluid Method. As with this previous test, we set up an isolated shock moving towards an interface. Again the exact solution can be obtained by considering the Riemann problem that arises the instant the shock hits the interface. The initial data for this test is given by

$$\begin{cases} \rho_0 = 1.3346, & v = 0.1837, & p = 1.5, & x < 0.05 \\ \rho_0 = 1, & v = 0, & p = 1, & 0.05 \leq x < 0.5 \\ \rho_0 = 0.1379, & v = 0, & p = 1, & x \geq 0.5 \end{cases} \quad (4.2.8)$$

and γ by

$$\begin{cases} \gamma = 1.4, & x < 0.5 \\ \gamma = 1.67, & x \geq 0.5 \end{cases} \quad (4.2.9)$$

Since the interface is again positioned at $x = 0.5$, the level set function is again given by (4.2.5). The initial data for the isolated shock is given by solving the equations given in section 2.4.3 for given pre-shock variables and post-shock pressure. The test is run to time $t = 1$.

In figure 4.15 we show the results for an isolated shock hitting an interface. As with the Newtonian case shown in figure 4.4, we see that some of the shock wave has been transmitted through the interface, and some has been reflected as a rarefaction wave. It is also clear that the effects of the shock hitting the interface have moved it from its initial position of $x = 0.5$.

In order to ensure that the movement of the interface has been accurately captured when a shock wave hits an interface, we show a zoomed region about this feature in figure 4.16. Aside from showing that the interface movement has been correctly captured, we see that there are no undershoots for this test. This is a result of the interface in the initial data being in some equilibrium configuration, hence the overheating errors do not arise.

In figure 4.15 the magnitude of the rarefaction is too small to ascertain that




Figure 4.15: Results for an isolated shock hitting an interface. This test is an SR analogy of the results shown in figure 4.4. The transmitted shock is very clear in the velocity and pressure profiles, and hence we can see that there is convergence with resolution. The rarefaction is significantly smaller in all plots, but still clearly visible. The position of the contact discontinuity is shown by the dashed line. For clarity, only 100 cells of each resolution are shown.




Figure 4.16: A density plot for the region around the interface in a test in which an isolated shock has hit an interface in SR, as shown in figure 4.15. It is clear that the position has been accurately captured. We can also see the effects of beginning the simulation with the interface in an equilibrium configuration has resulted in there being no overheating issues, hence there are no undershoots present.




Figure 4.17: The pressure plots zoomed on the shock and rarefaction for the test shown in figure 4.15 in which an isolated shock hits an interface. In the left panel we see that there is convergence towards the exact solution for the rarefaction (which, in figure 4.15, is too small to allow any definitive statements about convergence to be made). We also confirm that there is convergence for the shock too, along with no additional oscillatory behaviour.

the simulation is converging to the exact solution. In figure 4.17 we zoom in on this region, as well as that around the shock wave (for the pressure only, since the rarefaction is most prominent for this variable, though any variable could have been used). We see the convergence towards the exact solution for both waves, and also that there is no oscillatory behaviour around the shock.

The tests shown so far suggest that the Ghost Fluid Method successfully extends to SR. The velocities in these plots have all been only mildly relativistic. We therefore now present a test very similar to that shown in figure 4.15, with an isolated shock hitting an interface, but with highly relativistic velocities. This will test the Ghost Fluid Method in the regime in which standard Newtonian physics is no longer suitable. The initial data for this test is

$$\begin{cases} \rho_0 = 10.2384, & v = 0.9411, & p = 50, & x < 0.05 \\ \rho_0 = 1, & v = 0, & p = 1, & 0.05 \leq x < 0.5 \\ \rho_0 = 0.1379, & v = 0, & p = 1, & x \geq 0.5 \end{cases} \quad (4.2.10)$$

with the parametrisation of the EOS given by (4.2.9), and level set function by (4.2.5). The high velocities in this test mean that we will use a domain of $x \in [0, 3]$ and run to a time $t = 2$ to ensure that all features are identifiable.

In figure 4.18 we show the results for a highly relativistic shock hitting a contact discontinuity. The same qualitative features are seen as in figure 4.15, where we have a similar test with lower a velocity shock, we see there is still a reflected rarefaction wave and transmitted shock. The high velocity in this test does however lead to some complications. In setting up an isolated shock, there is, due to discontinuous initial data, start-up error present (as happens at the contact discontinuity for previous tests), which follows with the shock velocity, v_S . This error is documented in [123], and, for the shocks so far considered in figures 4.4 and 4.15, has been to small to be an issue. Unfortunately, for this highly-relativistic shock, this initial error is non-negligible, and follows the shock very closely.

If there is sufficient separation between shock and start-up error then it can be manually removed, since all variables in the post-shock region are constant. We can also see in figure 4.18 that the coordinate distance between the




Figure 4.18: The results for a highly relativistic shock ($v \sim 0.94$) hitting an isolated contact discontinuity. The effects of this high velocity show that when compared to figure 4.15, in which the velocity was much lower, the coordinate distance between the transmitted shock and the interface is much smaller. The necessity for the extended domain in this test is evident by the extent of the movement of the interface from its initial position at $x = 0.5$. For clarity, only 200 cells of each resolution are shown.

interface and the transmitted shock is small. Therefore if the start up error is not removed then this small distance means that it interferes with the interface, giving an unreliable result. In order to distinguish, and hence remove, the start-up error, we must therefore run these simulations at high resolution, hence we have shown plots for 2000 and 4000 cells. It is important to note that this is a purely numerical issue, and would not occur with realistic initial data. We therefore see that the Ghost Fluid Method has continued to work in highly relativistic situations.

All tests so far have been chosen specifically since exact solutions are available and as a result, the behaviour seen has been fairly simple (from a physical point of view, even the highly relativistic test shown in figure 4.18 is simple, though this is not the case when considering it from a numerical point of view). A further test will now consider a situation in which we will see more complex behaviour. Since an exact solution will not be available for such a test, we will use an analogous test in Newtonian physics to ensure that the correct behaviour is observed. We therefore base this next test on the isolated shock hitting a slab of low-density material, as proposed by [207] and shown in figure 4.5. The initial data for this test is

$$\begin{cases} \rho_0 = 1.37795, & v = 0.17933, & p = 1.57, & x < 0.25 \\ \rho_0 = 1, & v = 0, & p = 1, & 0.25 \leq x < 0.45 \\ \rho_0 = 0.138, & v = 0, & p = 1, & 0.45 \leq x < 0.55 \\ \rho_0 = 1, & v = 0, & p = 1, & x \geq 0.55 \end{cases} \quad (4.2.11)$$

where γ is given by

$$\begin{cases} \gamma = 1.4, & x < 0.45 \\ \gamma = 1.67, & 0.45 \leq x < 0.55 \\ \gamma = 1.4, & x \geq 0.55 \end{cases} \quad (4.2.12)$$

and the level set function by

$$\phi = 0.05 - |x - 0.5|. \quad (4.2.13)$$

This initial data is not directly comparable to that used in the original Newto-

nian test, given by (4.1.8). In order to see the full range of behaviour, so that comparison can be made to figure 4.5, without the low-density slab leaving the domain, we found it was necessary to decrease the width of this slab. The test is run to time $t = 0.8$.

The results for an isolated shock hitting a low-density slab of material are shown in figure 4.19. Comparison to the Newtonian version of this test, figure 4.5, shows that the visible behaviour in this test is as expected. There are two clear reflected rarefactions and a large transmitted shock, along with a smaller, but still visible, second transmitted shock. It is clear that the magnitude of the behaviour resulting from further interactions between waves and the interface is too small to be fully identified in this plot. The expected movement and contraction of the slab is also apparent, and we see that for this more complex test, we still have convergence as resolution increases.

The finer behaviour which cannot be seen in figure 4.19 is shown in figure 4.20 where, to allow for accurate identification of all features, we have used increased resolution of 800 and 1600 cells. This shows that the same range of behaviour seen in the Newtonian version of this test is still apparent. Due to the magnitude of these features, the scales we consider also show the slight oscillatory behaviour that we would expect within a numerical simulation. In all other tests, at this scale, we would see this feature (both in the Newtonian and relativistic results), and it is not a source of worry for the Ghost Fluid Method in SR. In figure 4.21 we show results for this test at a lower resolution of 100 cells. As in figure 4.6, this ensures that no behaviour is missed as a ‘selection effect’ due to showing only a sample of points in higher resolution figures. Again we show specific internal energy, ε , and can see in all cases, the features are not showing unexpected behaviour.




Figure 4.19: Results for an isolated shock hitting a slab of low-density material in SR. Two reflected rarefactions and the leading transmitted shock are clearly visible. A second transmitted shock, though much smaller in magnitude, is still visible. Comparing to figure 4.5, we see that these major features are expected. We can see that the slab has been both moved and compressed when compared to the initial data. Convergence as resolution increases is apparent. For clarity, only 100 cells of each resolution are shown, and here, the solid line represents a reference solution with 6400 cells.




Figure 4.20: Zoomed results for the isolated shock hitting a slab of low density material, shown in full in figure 4.19. Higher resolutions of 800 and 1600 cells have been used to ensure the small magnitude behaviour is now clearly visible. The oscillatory behaviour we observe is expected due to the numerical simulation, and, due to its magnitude, not a severe issue.




Figure 4.21: The results for a shock hitting a slab of material at a resolution of 100 cells in SR. As with the Newtonian analogy, figure 4.6, we are able to show every cell for this test, and ensure that there is no behaviour missed by not showing every cell in previous plots. This shows that the Ghost Fluid Method is not introducing any spurious error in SR. At this resolution, the smearing of the level set function leads to a narrower slab than expected at this time.

4.2.1 Comparison of reconstruction methods for the Ghost Fluid Method in special relativity

The tests shown so far have used only MC Limiter as a reconstruction technique. In order to fully test the versatility of the Ghost Fluid Method, it is useful to consider how it performs under other techniques. We therefore will compare results using MC limiter, described in section 3.4.2, with those using PPM in section 3.4.3. We do not use minmod in any tests in this section, although practically it was found to give comparable results to MC limiter, but non-linear features suffered a greater degree of smearing. In order to highlight the differences between these two methods, we select tests for which the Newtonian Ghost Fluid Method is known to have issues. These issues arise from cases in which the entropy gradient is not constant across the interface. The zeroth order extrapolation of entropy implemented by the Ghost Fluid Method therefore leads to the introduction of errors.

To compare the two methods (MC limiter and PPM) we ideally need situations in which an exact solution exists. Unfortunately, cases in which the entropy gradient changes across the interface are not abundant. One of the few such cases we are able to consider is, in many ways, similar to the moving contact discontinuity test shown in figure 4.8. This test, suggested by [129], uses an advected sine wave, allowing entropy to remain continuous across the interfaces, but with non-zero gradient. The initial data for this test is

$$\begin{cases} \rho_0 = 1, & v = 0.5, \quad p = 1, \quad x < 0.16 \\ \rho_0 = 1 + 0.3 \sin(50(x - 0.16)), & v = 0.5, \quad p = 1, \quad 0.16 \leq x < 0.537 \\ \rho_0 = 1, & v = 0.5, \quad p = 1, \quad x \geq 0.537 \end{cases} \quad (4.2.14)$$

with EOS parametrised by

$$\begin{cases} \gamma = 1.4, & x < 0.16 \\ \gamma = 1.67, & 0.16 \leq x < 0.537 \\ \gamma = 1.4, & x \geq 0.537 \end{cases} \quad (4.2.15)$$




Figure 4.22: The density profile for the advected sine wave test for a 200 point resolution. The solid line shows the exact solution. It is clear that PPM (shown in the right panel) has more accurately captured the features of the sine wave. Both methods, however, suffer at the interfaces, with large errors introduced due to the entropy gradients. The dashed line indicates the location of the interfaces.

and the level set function by

$$\phi = 0.1885 - |x - 0.3485|. \quad (4.2.16)$$

This test is run to a time $t = 0.4$.

Figures 4.22 and 4.23 show the results for the advected sine wave test. Only the density is shown in these plots since the continuous velocity and pressure across the entire domain leads to the correct, and trivial, solution for these variables. The issue the Ghost Fluid Method suffers from the large entropy gradient across the interface is clear in these plots. There is a large error at the interfaces which does decrease with resolution, but not rapidly (the ∞ -norm of this test was calculated to be order 0.7 for both methods). Although the Ghost Fluid Method does suffer this source of error, no oscillatory behaviour is introduced. The differences between MC limiter and PPM are clearly apparent in the 200 point test, in which the accuracy with which the maxima




Figure 4.23: The density profile for the advected sine wave test for an 800 point resolution. Again the solid line shows the exact solution and the dashed line indicates the location of the interfaces. Both methods now capture all features of the sine wave well. The errors at the interfaces have clearly been reduced due to increase in resolution, but still exist.

and minima of the sine wave are captured is clearly better with PPM. We also note that the magnitude of the error at the interfaces is (slightly) smaller when using PPM, but the order of convergence is no higher.

We now present a test in which there is evolution of the fluid variables in a situation in which there is a change in entropy gradient across the interface. This is a shock-tube type test, in some ways similar to the Riemann problem though the density profile in the right state is a sine wave, rather than constant. This ‘perturbed’ shock tube test, proposed by [56] and used by e.g. [52], has initial data

$$\begin{cases} \rho_0 = 5, & v = 0, \quad p = 50, \quad x < 0.5 \\ \rho_0 = 2 + 0.3 \sin(50x), & v = 0, \quad p = 5, \quad x \geq 0.5 \end{cases} \quad (4.2.17)$$

with γ given by

$$\begin{cases} \gamma = 1.4, & x < 0.5 \\ \gamma = 1.67, & x \geq 0.5 \end{cases} \quad (4.2.18)$$

and the level set function by

$$\phi = x - 0.5. \quad (4.2.19)$$

This test was run to a time $t = 0.35$.

We show the results for the perturbed shock tube test for 200 cells in figure 4.24 and for 800 cells in figure 4.25. Since we cannot obtain an exact solution for this test, we instead include a very high resolution simulations, with 12800 cells, to check the solution converges as resolution increases. We see that, as in the previous test, PPM captures the features of the sine wave more accurately, though at high resolution, neither method struggles to resolve these features. Again the sharp entropy gradient at the interface has lead to errors here, which slowly converge, but are still present in the 800 point simulations. No spurious oscillations have been introduced however, it is be noted that the slight oscillatory behaviour which can be seen between the rarefaction and shock is indeed physical.

4.3 The Ghost Fluid Method in general relativity

The results in the previous section have shown that the Ghost Fluid Method can be successfully extended to SR. For it to be useful in modelling of neutron stars, we now need to show that it works in full GR. Here the tests are designed primarily to test the Ghost Fluid Method, rather than investigate physically relevant situations. Therefore the criterion for determining these test cases was for all effects of the interface to be obvious in the analysis of the results. As a result we will make no statements about the behaviour of multicomponent neutron stars based on these results.

In all tests in this section we can use the Hamiltonian constraint, as given in section 2.2.3,

$$\frac{\partial_r a}{a} = a^2 \left[4\pi r \tau + D - \frac{m}{r^2} \right] \quad (4.3.1)$$

to obtain a measure of the error introduced into the system. We can either look at a snapshot to gauge where errors are most prevalent within a star, or




Figure 4.24: Results for the perturbed shock tube test for MC limiter and PPM at a resolution of 200 cells. The solid line represents a reference solution of 12800 cells. We see that all features have been correctly located in this test. The greater accuracy of PPM in capturing the features of the sine wave is again apparent, as is the error at the interface due to entropy gradient.




Figure 4.25: Results for the perturbed shock tube test for MC limiter and PPM at a resolution of 800 cells, again the solid line represents a reference solution of 12800 cells. All features are now well captured, and a decrease in the error at the interface is visible.




Figure 4.26: The density and pressure profiles for a stable TOV star at time $t = 1000$. We see that the central density has remained at its initial value of $\rho_c = 1.28 \times 10^{-3}$. We have used a resolution of 1280 cells for this test.

use the norms of the error in the Hamiltonian constraint to test the convergence of the entire evolution.

In order to analyse the results obtained from these tests, we will first show a reference solution for an initial star configuration known to be stable. We choose a typical polytropic approximation to a NS (that is it has a mass of about $1.4M_\odot$ and radius about 10 km) obtained by solving the TOV equations given in section 2.6 for a given central density, $\rho_0(t = 0, r = 0) = \rho_c$, γ and K . The initial data required for such a star is

$$\rho_c = 1.28 \times 10^{-3}, \quad \gamma = 2, \quad K = 100. \quad (4.3.2)$$

This configuration provides a physically stable star, hence can also be used to ensure that the GR evolution code has been correctly implemented. In order to ensure that the evolution equations have been implemented correctly we must simulate it over several dynamical timescales (for this case we have $t_{\text{dyn}} \sim 150$). We therefore run this test until time $t = 1000$.

The density and pressure profiles for a stable single component star at time $t = 1000$ are shown in figure 4.26. We see that the central density retains

its initial value given by (4.3.2) and that there is no unexpected behaviour apparent in the star.

In order to ensure that this test is working correctly we must consider convergence results, shown in figure 4.27. The numerical methods used in this section should yield second order convergence of the results. The top panel of this figure shows the error in the Hamiltonian constraint as a function of radius at time $t = 1000$. This is shown for two resolutions, 640 and 1280 cells, and scaled to show that second order convergence is indeed achieved nearly everywhere. We do however see that at $r \sim 10$ there is a region of large error, and a loss of second order convergence. We can see from figure 4.26 that this is the location of the surface of the star. It is clear that this is an expected issue due that the atmosphere algorithm, as described in section 3.8 which causes this loss of convergence. The central panel shows the 1-norm of the Hamiltonian constraint as a function of time. This is scaled to show that the overall error in the star scales at an order of 1.55. Comparing this to the snapshot of convergence at $t = 1000$, it is clear that the errors associated with the surface and the interaction with the atmosphere are dominating in the 1-norm, hence we do not see second order convergence here. The lower panel shows the 1-norm of the error in the density. Since we are modelling a static star, the initial data therefore provides the exact solution, hence this error plot can be obtained. The low density of the atmosphere means that the errors that occur here do not dominate the 1-norm, and as a result we see that the convergence here is second order as desired. These plots show that the second order techniques used in the simulation are working correctly. This also provides a reference solution which will allow any errors introduced by the Ghost Fluid Method to be identified.

In order to test whether the introduction of an interface into a stable star causes any issues, we can, as in previous tests in Newtonian and SR cases, introduce a “trivial” interface. We therefore have the same initial data as (4.3.2), though we include a level set function

$$\phi = x - 3.015 \quad (4.3.3)$$




Figure 4.27: Showing measures of convergence for a single component reference solution. The upper panel shows second order convergence for the error in the Hamiltonian constraint as a function of radius at late times ($t = 1000$). The central panel shows the 1-norm of the error of the Hamiltonian constraint as a function of time. Here the order of convergence is scaled to 1.55, since the errors at $r \sim 10$, the surface of the star, dominate the 1-norm. The lower panel shows the 1-norm of the error in the density of the star as a function of time. Here the effects of the surface are less pronounced, hence this approaches second order convergence.




Figure 4.28: The density and pressure profiles for a stable TOV star with an artificial interface at $x \sim 3$ at time $t = 1000$. We see that the central density has remained at its initial value of $\rho_c = 1.28 \times 10^{-3}$. We also note that there is no visible behaviour introduced by this artificial interface. We have used a resolution of 1280 cells for this test.

which adds this trivial interface at $x = 3.015$. Since the EOS is the same both sides of the interface, there are no discontinuities in the initial data, hence we expect a solution as shown in figure 4.26.

In figure 4.28 we see the density and pressure profiles are indistinguishable from the true single component case shown in figure 4.26. There is no unphysical behaviour arising introduced through the interface, and the star has remained in its initial configuration.

To determine whether the presence of the artificial interface has had a detrimental effect on the rate of convergence, we present the error plots for this test in figure 4.29. These can be directly compared to the plots for a corresponding single component star in figure 4.27. The error in the Hamilton constraint as a function of time is virtually indistinguishable in the two cases. It is clear that the interface at $r = 3.015$ has not introduced unphysical behaviour. As a result of this, the plot of the 1-norm of the error in the Hamiltonian constraint, and the lower panel showing the 1-norm of the error in the density, are also




Figure 4.29: Convergence results for a static star with an artificial interface placed at $r \sim 3$. These plots should be compared to those for a genuinely single component star, shown in figure 4.27. Again the top panel shows second order convergence almost everywhere for the Hamiltonian constraint, when plotted as a function of radius. Importantly, there is no visible error introduced by the interface in this plots. As a result, the central plot, the 1-norm of the error in the Hamiltonian constraint, again has convergence of order 1.55. Similarly, there is no visible change in the 1-norm of the error in the density, shown in the lower panel. Again we see near perfect second order convergence.




Figure 4.30: Results for the velocity and entropy of a static star with an artificial interface at $r \sim 3$ at time $t \sim 270$. These plots show the (small) detectable effects of the interface. The left panel compares the velocity in the interior of the star for the two situations; with no interface and with an artificial interface. It is clear that the velocity remains constant across the interface. We can, however, see that there are some differences between the two. These arise from the surface, suggesting any minor differences in velocity in the two cases are magnified and reflected by the treatment of the surface and atmosphere. The right panel shows the error in the entropy of the star around the interface. There is a slight jump across the interface, due to the errors in the extrapolation of the Ghost Fluid boundary condition, at the 10^{-8} level.

indistinguishable from those in figure 4.27. That is these have convergence order of 1.55 and 2 respectively.

Further investigation as to whether there are any visible effects of the interface can be achieved by comparing velocity and entropy of the single component and artificial interface cases. These results are shown in figure 4.30. The left panel shows the velocity profiles in the interior of the star at time $t \sim 270$. Only the interior velocity is shown since the behaviour at the surface dominates the plots and other effects would be indistinguishable. It is clear that there is no discontinuity in the velocity profile at the interface. There is, however, a definite difference between the two cases, clearly originating from the surface and not the interface. The choice of $t \sim 270$ was chosen for this plot since at

at this point the source of this difference is so clear. At late times, this difference propagates inwards, hence the cause would be indeterminable. This velocity is still comparatively small, hence we still see the appropriate convergence results shown in figure 4.29. The differences emerge since any slight errors in the velocity introduced by the extrapolation at the interface are magnified and reflected by the surface and atmosphere treatment. The right panel in figure 4.30 shows the error in the entropy when compared to the exact solution. We see that there is a clear effect of the interface, but this is a very small effect, of order 10^{-8} . These suggest that there is no major source of error that occurs through the inclusion of a sharp interface in the context of GR.

The success of these preliminary results means we now consider cases for a static star with a genuine interface, across which the EOS changes. In these cases the interface is always positioned in a region in which the pressure and density are comparatively large. This allows for ease in analysing any behaviour that arises from the interface. Again we stress that these models are not based on physically relevant scenarios. The first test we consider has a small change in the value of γ across the interface. The initial data for this test is

$$\left\{ \begin{array}{llll} \rho_c = 1.28 \times 10^{-3}, & \gamma = 2, & K = 100, & r < 3.015 \\ & \gamma = 1.9, & K = 51.57, & r \geq 3.015 \end{array} \right. \quad (4.3.4)$$

with the level set equation given by (4.3.3). As with the previous static star cases, this is run to time $t = 1000$.

We show the results for the density and pressure for a star with an outer EOS parametrised by $\gamma = 1.9$ in figure 4.31. In the density we see the jump at the interface, again located at $r = 3.015$, with the pressure remaining continuous. The central density suggests that this star has indeed remained static.

In figure 4.32 we see how introducing a genuine interface has affected the convergence results for the star. The top panel shows that there is a noticeable




Figure 4.31: The density and pressure profiles for a stable TOV star with an interface at $r \sim 3$ over which there is a small change in γ . This is plotted at time $t = 1000$. We see that there is a jump in the density profile at the interface, but the pressure remains continuous. We have used a resolution of 1280 cells for this test.

effect from introducing a jump in the EOS. There are some errors introduced at the interface, although convergence is not greatly affected. This is not surprising due to the simple nature of the boundary conditions in the Ghost Fluid Method. The magnitude of these errors is small, however, and, as we see in the middle panel, there is still reasonable convergence in the 1-norm of the error for the Hamiltonian constraint. In fact we see that the order of convergence here, 1.65, is greater than for the single component star. The larger errors at the interface mean that these will have a greater effect on the 1-norm of the error. Since these are, to some degree convergent, this could raise the overall convergence indicated through the 1-norm of the Hamiltonian constraint. It is also possible that this increase in convergence arises from the softer exterior EOS leading to a slightly more extended star. This effectively reduces the influence of the surface on the convergence rate, thus leading to this improvement. The lower panel again shows second order convergence in the 1-norm of the density as we would expect.

We now consider a static star with a softer exterior EOS. This test has initial




Figure 4.32: Convergence plots of the errors in the Hamiltonian constraint and density for a static star with $\gamma = 1.9$ in the outer region. This plot can be compared to that of the reference solution in figure 4.27. We see that there is now a visible effect of the interface in the error plot for the Hamiltonian constraint at $t = 1000$. Although there are these errors, they are largely convergent, and hence do not present a major problem, with the 1-norm of the Hamiltonian constraint converging with order 1.65. We see that the 1-norm of the error in the density again converges at second order.




Figure 4.33: The density and pressure profiles for a stable TOV star with an interface at $r \sim 3$ and a softer exterior EOS, $\gamma = \frac{5}{3}$ plotted at time $t = 1000$. Again we see the jump in the density profile due to the EOS change, with a continuous pressure profile as expected. The softer exterior EOS results in this configuration having a larger radius. We have used a resolution of 1280 cells for this test.

data

$$\left\{ \begin{array}{llll} \rho_c = 6 \times 10^{-4}, & \gamma = 2, & K = 100, & r < 3.015 \\ & \gamma = \frac{5}{3}, & K = 11.17, & r \geq 3.015 \end{array} \right. . \quad (4.3.5)$$

Note that in order for this configuration to be stable we need a lower central density than used in previous tests. Again the interface is at $r = 3.015$, hence ϕ is given by (4.3.3), and we run to $t = 1000$.

We show the results for a static start with $\gamma = 5/3$ in the exterior in figure 4.33. The density jump at the interface is again apparent, with the pressure remaining continuous. The softer EOS means that the radius of the star with this configuration is greater, with the surface now at $r \sim 14$.

In figure 4.34 we show convergence plots for a star with a $\gamma = 5/3$ exterior component. Comparing these plots to the reference solution in figure 4.27 we




Figure 4.34: Convergence results for a static star with an interface at $r \sim 3$ and $\gamma = 5/3$ in the exterior region. Comparison to the reference solution, figure 4.27, shows that the effects of the interface are again noticeable. We can again see that they are not large enough to pose a serious problem, and do appear to converge with resolution. The softer exterior EOS leads to the errors at the surface of the star being smaller in magnitude. The 1-norm of the error in the density again shows second order convergence until late times, at which point the errors from the surface have increased sufficiently to start dominating.

see that, as in figure 4.32, the effects of a non-trivial interface are clearly visible. Again this error is predominantly convergent and not significantly large. The plot showing the error in the Hamiltonian constraint at $t = 1000$ also shows the effects of the softer external EOS. We see that the magnitude of the error is smaller, though again it does not converge at second order. The larger radius of the star means that this surface behaviour does not have such a large contribution to the 1-norm of the error in the Hamiltonian constraint. As a result we see a convergence order of 1.8, greater than that in the reference solution. Compared to the reference solution, the softer exterior EOS requires a lower central density to be stable. This, along with the more extended low-density region at the surface of the star, means that the errors in the density at the surface have a greater impact on the 1-norm convergence. As a result, we see the 1-norm of the error in the density drift away from second order convergence after $t \sim 800$.

The tests so far have shown that the Ghost Fluid Method does not introduce major unphysical behaviour into static stars in GR. In particular, the error in these stars is dominated by the surface algorithm. For the Ghost Fluid Method to be useful in more physically relevant situations, we also need to show that it can deal with interface movement and non-linear behaviour e.g. shock formation. In a 1+1 dimensional context, we can artificially perturb the star such that a shock wave forms in the interior and propagates radially outwards. We add this artificial perturbation onto a two-component star with initial data given by

$$\begin{cases} \rho_c = 1.28 \times 10^{-3}, & \gamma = 2, & K = 100, & r < 3.015 \\ & \gamma = 1.9, & K = 51.57, & r \geq 3.015 \end{cases} \quad (4.3.6)$$

as in (4.3.4). We then perturb the interior region using

$$\rho_0 = \rho_{0,\text{TOV}} [1 + h(r)], \quad p = p_{\text{TOV}} [1 + h(r)], \quad (4.3.7)$$

where

$$h(r) = \begin{cases} \frac{1}{20} \{1 - \tanh [50(r - 2)]\} & r < 2.5 \\ 0 & r \geq 2.5 \end{cases}. \quad (4.3.8)$$




Figure 4.35: Density and pressure profiles for a 2-component star with an initial perturbation. The left panel shows the initial data where the steep perturbation is clearly visible. The right panel shows the evolved data at time $t \sim 15$. At this point we can see that the initial perturbation has split into an outgoing shock, which has moved through the interface, and an ingoing feature which has been reflected at the origin and is now shocked, but has not yet reached the interface. Since the perturbation of the initial data made in an ad-hoc fashion, a ‘bulge’ in the density can be seen at $r \sim 2$, whilst the pressure has reverted to a smooth profile. The interface position has been marked by the vertical dashed line, and we can see that pressure remains continuous here.

This inserts a sharp (but not yet non-linear) feature which, as it propagates through the star, rapidly forms two shocks moving in opposite directions. This initial perturbation was chosen over a true discontinuity since this introduced problems, most likely from the evolution of the metric quantities at early times.

Figure 4.35 shows the density and pressure profiles for the initial and evolved data for this perturbed star case. The left panel shows the initial configuration with the large perturbation clearly visible. The right panel shows this star after evolution in which the perturbation has split into two components. We see an outward moving shock, which has passed through the interface. We also see a second shock, which started as an inward moving feature and has been reflected at the origin. We can see that the pressure profile remains continuous over the interface, even after the shock has passed through it. The initial data




Figure 4.36: Showing the errors in the Hamiltonian constraint at time $t \sim 15$. As with the previous tests, we can see the effects of the interface at $r \sim 3$ and the surface at $r \sim 10$. We also see errors present at the two shocks, as we would expect. Away from these features we see that there is the second order convergence we expect. Also plotted is the results for an 80 point run (appropriately scaled for second order convergence) showing that this convergence rate still holds for the lower resolutions possible in full 3+1 simulations.

has also left a “bulge” in the density profile due to the simple manner in which the perturbation was implemented. The results of this test have only been shown after a a comparatively short time (the right panel shows the results at $t \sim 15$) since the simple atmosphere algorithm applied here results in unreliable data once the shock reaches the surface. This is, however, sufficient to assess the effects the shock has on the interface. Note that in these tests the reflected features from the shocks hitting the interface, which were seen in the SR tests e.g. in figure 4.15, are too small to be identified in figure 4.35.

We show the errors in the Hamiltonian constraint for the perturbed star test in




Figure 4.37: The left panel shows the velocity profile of the star at $t \sim 15$. The locations of both shocks are clear, and we can see that there is no oscillatory behaviour generated at the interface as a result of the interaction with the outgoing shock. The right panel shows the movement of the interface as a function of time. Both the interpolated and discrete locations are plotted, as obtained from the level set function. We see that the movement of the interface is consistent with resolution.

figure 4.36. Since we have not used an equilibrium configuration in the initial data, we do not have an exact solution available and hence cannot determine the error in the density profile. This plot shows that in addition to the effects of the interface and surface, we see that there are also errors at the shock waves. These errors do not converge at second order since the HRSC methods reduce down to a first order technique at shocks. Away from these features, however, convergence is second order as we would expect.

The effects of the initial perturbation on the behaviour of the star and in particular the interface are highlighted in the results presented in figure 4.37. The left panel of this figure shows the velocity profile of the star at $t \sim 15$. We see the two shock waves very clearly, and also that the effect of the leading shock passing through the interface has left no visible shock, we can see no discontinuous behaviour or spurious oscillation. The comparative magnitude of the shock is much greater here than in figure 4.35, hence we can now confidently see that the shock locations are being captured consistently. We also see that

the pre-shocked outer regions of the star are falling inwards. This is a result of the initial perturbation of the star. By artificially increasing the density towards the centre of the star, we increase the gravitational mass of the star here. This then affects the spacetime and causes the outer layers to contract.

The right panel of figure 4.37 shows the interface location as a function of time. We include both the discretised location, which simply records when the zero of the level set function changes cell, and the interpolated position of the interface from the level set function. From the discretised plot we see that the interface genuinely does move change numerical cells within an evolution. By considering the other plots for the perturbed star, this shows that this movement does not introduce any unphysical behaviour. The interpolated position shows that initially, the additional gravitational mass in the interior causes the interface (along with the rest of the matter in the star) to fall inwards. The impact of the leading shock at $t \sim 4$ halts this infall and moves the interface outwards again. At $t \sim 13$ the effects of the interaction with the shock are no longer enough to prevent the infall of the interface again. The interaction of the trailing shock at $t \sim 18$ slows this infall, but the “bulge” in density left from the initial data visible in figure 4.35 means that there is still some infalling matter. This test has successfully demonstrated that the Ghost Fluid Method can deal with non-linear behaviour in GR.

4.4 Summary

The results in this section have demonstrated the extension of the Ghost Fluid Method in 1+1 dimensional relativistic situations. The results given in section 4.2 show the initial extension for SR. A variety of scenarios, based on Riemann shock-tube problems, were simulated to robustly test the capabilities of this extension to SR. We showed that the correct results could be obtained even in highly relativistic scenarios (see figure 4.18). We were also able to present results comparable to those in the Newtonian CFD literature (e.g. [207]) for a more complex test with multiple interfaces in figure 4.19. Some of the known shortcomings of the Ghost Fluid Method were investigated in section 4.2.1. Here we found that these issues, although still present, did

not pose a major issue to the relativistic extension.

Our extension of the Ghost Fluid Method was then tested in 1+1 dimensional GR in section 4.3. Here we ensured first that multicomponent stars in hydrostatic equilibrium could be successfully simulated. We then investigated the effects of non-linear behaviour within a star, see figure 4.35. In all these cases, we could analyse the convergence of our method. We showed that the second-order convergence expected from the numerical methods was obtained, and our method could successfully deal with non-linear features, such as shocks.

Chapter 5

The Ghost Fluid Method in multidimensions

A key test of the viability of the relativistic Ghost Fluid Method is whether it extends to multidimensional situations. This is essential for use with the 3+1 NS simulations required for obtaining GW templates. In this section we present preliminary results in which the Ghost Fluid Method is tested in 2+1 dimensional SR.

5.1 Numerical techniques for solving multidimensional equations

In chapter 3 all numerical techniques described pertained to the solution of 1+1 dimensional equations. We must now consider how these can be extended when dealing with multidimensional situations. When solving such equations numerically, one standard technique to use is *dimensional splitting* [192]. In order to use this technique, we must first be able to write our system of equations with state vector \mathbf{u} and spatial coordinates $\mathbf{x} = (x_1, x_2, x_3)^T$ in the form

$$\partial_t \mathbf{u} + \partial_{x_1} \mathbf{f}(\mathbf{u}, \mathbf{x}) + \partial_{x_2} \mathbf{g}(\mathbf{u}, \mathbf{x}) + \partial_{x_3} \mathbf{h}(\mathbf{u}, \mathbf{x}) = \mathbf{s}(\mathbf{u}, \mathbf{x}). \quad (5.1.1)$$

In chapter 2 we considered the fluid evolution equations in the general multi-dimensional case for both Newtonian and relativistic situations. In all cases, it was indeed possible to express the conservation form of the equations in this

manner. We can then make the reconstruction for each spatial derivative term, as described in section 3.4, individually.

In this thesis, our implementation of dimensional splitting differs from the original presentation in [192]. We will therefore contrast these two approaches assuming, without loss of generality, a 2+1 dimensional system with no source terms and using a first-order RK method for the time update, i.e. we use only a single step in the time evolution and $\mathbf{h} = \mathbf{s} = 0$. In 1+1 dimensions, the first order RK update is given by

$$\hat{\mathbf{u}}^{n+1} = \hat{\mathbf{u}}^n + \Delta t \left(\hat{\mathbf{f}}_{i-1/2} - \hat{\mathbf{f}}_{i+1/2} \right) \quad (5.1.2)$$

To make the update using dimensional splitting, we define the quantities

$$\mathcal{F}_{\Delta t} = \hat{\mathbf{f}}_{i-1/2}(\Delta t) - \hat{\mathbf{f}}_{i+1/2}(\Delta t), \quad \mathcal{G}_{\Delta t} = \hat{\mathbf{g}}_{j-1/2}(\Delta t) - \hat{\mathbf{g}}_{j+1/2}(\Delta t), \quad (5.1.3)$$

where we have assumed these quantities are evaluated for all (x_i, y_j) . In this work, we will then make the update at $t = t^{n+1}$ in a single step,

$$\hat{\mathbf{u}}^{n+1} = \hat{\mathbf{u}}^n + \Delta t (\mathcal{F}_{\Delta t} + \mathcal{G}_{\Delta t}). \quad (5.1.4)$$

By comparison, the update introduced in [192] is made through

$$\begin{aligned} \mathbf{u}^{*1} &= \hat{\mathbf{u}}^n + \frac{\Delta t}{2} \mathcal{F}_{\Delta t/2}, \\ \mathbf{u}^{*2} &= \hat{\mathbf{u}}^{*1} + \Delta t \mathcal{G}_{\Delta t}, \\ \mathbf{u}^{n+1} &= \hat{\mathbf{u}}^{*2} + \frac{\Delta t}{2} \mathcal{F}_{\Delta t/2}. \end{aligned} \quad (5.1.5)$$

The advantage of the form used here, (5.1.4), is that it preserves the symmetry of a solution if this exists and is not aligned with a coordinate axis.

When using dimensional splitting, the CFL factor described in section 3 must be chosen consistently for all spatial dimensions, i.e. it is determined by the overall maximum wavespeed.

5.2 Multicomponent fluids in more than one dimension

5.2.1 Level set methods

The level set methods introduced in section 3.7.1 were developed with the modelling of sharp features in multidimensional fluids in mind. In such situations the locations of these features are simply described by zeroes of some surface or hypersurface. Again the fluid components can simply be identified by the sign of the level set function, ϕ . Additionally, the methods for solving Hamilton-Jacobi equations presented in 3.7.3 also extend naturally to multidimensions. The need for the reinitialisation of ϕ due to steepening is, however, now likely to arise frequently due to non-trivial behaviour of the velocity components in multidimensional simulations. In this preliminary work, we shall not consider reinitialisation, and restrict our analysis of simulations to regions in which ϕ is suitably smooth. The techniques for reinitialisation are, however, described in e.g. [67, 148].

5.2.2 Ghost Fluid Method

In 1+1 dimensions, the Ghost Fluid boundary conditions are trivially implemented. In particular the extrapolation of the entropy and copying of normal velocity are greatly simplified by the single spatial dimension. In a multidimensional situation, we must be careful that any extrapolation occurs in the normal direction to the interface, and must also deal with the tangential velocity components in some way. In the original Ghost Fluid Method paper of Fedkiw et al. ([67]), the multidimensional Newtonian case was considered.

When there is more than one component to the velocity, then the Ghost Fluid Method assumes that across an interface the normal component of the velocity is copied from the actual fluid into the ghost fluid region. The tangential component is then extrapolated in the same manner as the entropy. Obviously there is no requirement that these components align with the coordinate axes. To obtain these components, we first determine the normal vector, \mathbf{N} , to the level set function. This is done at every grid point and is

given by

$$\mathbf{N} = \frac{\nabla\phi}{|\nabla\phi|}. \quad (5.2.1)$$

The normal is only required for providing information within the ghost fluid region, hence the structure of ϕ away from the interface will not lead to adverse effects. This does mean that if we were to reinitialise ϕ , we would have to ensure that this was done for the entire ghost fluid region. Care must also be taken in determining \mathbf{N} numerically if $\nabla\phi = 0$.

Once we have \mathbf{N} , we can define the normal and tangential velocity components. The magnitude of the normal component, v_N , is given by

$$v_N = \mathbf{v} \cdot \mathbf{N} \quad (5.2.2)$$

which, for the purpose of the dimensional splitting, is put into vector form,

$$\mathbf{v}_N = v_N \mathbf{N}. \quad (5.2.3)$$

The tangential velocity vector, \mathbf{v}_T , can then be defined by

$$\mathbf{v}_T = \mathbf{v} - v_N \mathbf{N}. \quad (5.2.4)$$

The normal velocity can now be copied into the ghost fluid region from the real fluid. The extrapolation of the entropy and tangential velocity should result in these variables being constant along \mathbf{N} . Therefore it is logical that we wish these variables to satisfy

$$\mathbf{N} \cdot \nabla I = 0 \quad (5.2.5)$$

where I is any variable to be extrapolated, i.e. the entropy and each component of \mathbf{v}_T . In order for this to be satisfied, we need some technique which allows us to make an initial guess for I , and then make some iteration that results in I satisfying (5.2.5). This can be achieved through the solution to the PDE

$$\partial_t I \pm \mathbf{N} \cdot \nabla I = 0 \quad (5.2.6)$$

where the sign of the PDE corresponds to the sign of ϕ in the ghost fluid region. The time derivative in this PDE is not necessarily the physical time in

the simulation. Instead we can solve for some *fictitious time* \tilde{t} , in which this PDE is evolved for several artificial timesteps until a static solution is reached, i.e. (5.2.5) is satisfied.

The velocity in the ghost fluid region can then be ‘built’ from the copied normal velocity \mathbf{v}_N and the extrapolated tangential velocity \mathbf{v}_T .

5.2.3 Relativistic effects in the Ghost Fluid Method

The relativistic extension of the Ghost Fluid Method in 1+1 dimensions was trivial. The coupling of velocity components through the Lorentz factor in multidimensional situations could complicate matters however. In relativity, it would be expected that the 4-velocity should be used in the Ghost Fluid boundary conditions. The coupling of the velocity components through the Lorentz factor means that making an extrapolation of the 3-velocity, as was done without issue in the 1+1 dimensional case, will not necessarily provide the correct result. This is particularly an issue for GR when spacetime factors enter the conversion between 3- and 4-velocities, as described in section A.1.5. In a Minkowski spacetime (i.e. SR), however, we would not expect such complications. In section 5.5.1 we investigate if there are any effects due to this choice.

5.3 Higher order methods for the level set evolution

In section 3.7.3 we detailed the numerical methods required for the evolution of the level set function ϕ . These used a Lax-Friedrichs scheme to evolve ϕ and, due to the simple nature of this function in 1+1 dimensional situations, it was sufficient to use first order approximations to the derivatives, ϕ^\pm . In multidimensional situations there is no such guarantee that this simple approximation will be suitable, the behaviour of the fluid, and hence the level set function, can be much more complex. It is therefore advantageous that the evolution of the level set function moves beyond this first-order approach. The extension of ENO methods to Hamilton-Jacobi equations was shown in [149], hence we now consider a third-order ENO approximation to ϕ_x^\pm as described in [148]. The

obvious result to obtain the derivatives in the y -direction, ϕ_y^\pm , holds. We begin by defining $D_{i+1/2,x}^1 \phi$ as the divided differences in the x -direction, equivalent to the first order approximations to the derivative of ϕ , i.e.

$$D_{i+1/2,x}^1 \phi = \frac{\phi_{i+1,j} - \phi_{i,j}}{\Delta x}. \quad (5.3.1)$$

In terms of the first order approximations we have $\phi_x^+ = D_{i+1/2,x}^1 \phi$ and $\phi_x^- = D_{i-1/2,x}^1 \phi$. We then have the second divided differences

$$D_{i,x}^2 \phi = \frac{D_{i+1/2,x}^1 \phi - D_{i-1/2,x}^1 \phi}{2\Delta x} \quad (5.3.2)$$

and third divided differences

$$D_{i+1/2,x}^3 \phi = \frac{D_{i,x}^2 \phi - D_{i-1,x}^2 \phi}{3\Delta x}. \quad (5.3.3)$$

The idea behind the ENO method is to then, if we consider a constant y value, make a polynomial reconstruction of ϕ ,

$$\hat{\phi}(x) = Q_0(x) + Q_1(x) + Q_2(x) + Q_3(x) \quad (5.3.4)$$

where Q_n is an n^{th} order polynomial. From this we obtain ϕ_x^\pm through the derivatives of this function,

$$\partial_x \hat{\phi}(x) = \partial_x Q_1(x) + \partial_x Q_2(x) + \partial_x Q_3(x) \quad (5.3.5)$$

Based on our definitions of the polynomials $Q_n(x)$ we can evaluate this approximation to at x_i to obtain ϕ_x^\pm . In the following definitions, we set $k = i - 1$ if we want to obtain ϕ_x^- and $k = i$ for ϕ_x^+ . We first define $Q_1(x)$ through

$$Q_1(x) = (D_{k+1/2,x}^1 \phi)(x - x_i) \quad (5.3.6)$$

hences we have

$$\partial_x Q_1(x_i) = D_{k+1/2,x}^1 \phi. \quad (5.3.7)$$

If this result is used solely in defining ϕ_x^\pm then we have the first-order method described in section 3.7.3. We can then make a second order correction, which could either use $D_{k,x}^2 \phi$ or $D_{k+1,x}^2 \phi$. In order to ensure that oscillatory behaviour

is kept to a minimum, we define

$$\mathcal{D} = \begin{cases} |D_{k,x}^2 \phi|, & |D_{k,x}^2 \phi| \leq |D_{k+1,x}^2 \phi| \\ |D_{k+1,x}^2 \phi|, & \text{otherwise} \end{cases} \quad (5.3.8)$$

and we also define

$$k^* = \begin{cases} k-1, & |D_{k,x}^2 \phi| \leq |D_{k+1,x}^2 \phi| \\ k, & \text{otherwise} \end{cases} \quad (5.3.9)$$

The second order polynomial correction is then given by

$$Q_2(x) = \mathcal{D}(x - x_k)(x - x_{k+1}) \quad (5.3.10)$$

with derivative

$$\partial_x Q_2(x_i) = \mathcal{D}[2(i-k) - 1] \Delta x. \quad (5.3.11)$$

Finally a third order correction is given in a similar manner. We first define

$$\mathcal{D}^* = \begin{cases} |D_{k^*+1/2,x}^3 \phi|, & |D_{k^*+1/2,x}^3 \phi| \leq |D_{k^*+3/2,x}^3 \phi| \\ |D_{k^*+3/2,x}^3 \phi|, & \text{otherwise} \end{cases} \quad (5.3.12)$$

where we note that k^* , which was defined, but not used, during the second order correction, is now used. The final polynomial correction is then given by

$$Q_3(x) = \mathcal{D}^*(x - x_{k^*})(x - x_{k^*+1})(x - x_{k^*+2}) \quad (5.3.13)$$

with derivative

$$\partial_x Q(x_i) = \mathcal{D}^* [3(i-k^*)^2 - 6(i-k^*) + 2] (\Delta x)^2. \quad (5.3.14)$$

Using these three polynomial results, we have a third-order accurate ENO scheme for evolving the level set function.

5.4 Results for the Newtonian multidimensional Ghost Fluid Method

The 2+1 dimensional Ghost Fluid Method was demonstrated in Newtonian physics [67], hence it is first worth using this to obtain a reference solution. This test considers the effect a shock wave of air has on a bubble of helium, previously considered both as a physical experiment (e.g. [85]) and numerically (e.g. [162]). This test is carried out on a domain $x \in [0, 325]$ and $y \in [-44.5, 44.5]$, with MC limiter for the fluid evolution and for accuracy we use the third order ENO scheme described in section 5.3 for the evolution of ϕ . The boundaries at $y = \pm 44.5$ are a solid wall, whilst those at $x = 0$ and $x = 325$ are constant extrapolation outflow boundaries. The computational domain effectively has three regions, the post-shocked air (\mathbf{q}_{post}), the pre-shocked air (\mathbf{q}_{pre}) and the helium (\mathbf{q}_{He}). The initial data for this test is given by

$$\left\{ \begin{array}{lll} \rho_{0,\text{post}} = 1.3764, & v_{x,\text{post}} = -0.394, & v_{y,\text{post}} = 0, \\ & p_{\text{post}} = 1.5698, & \gamma_{\text{post}} = 1.4, \\ \rho_{0,\text{pre}} = 1, & v_{x,\text{pre}} = 0, & v_{y,\text{pre}} = 0, \\ & p_{\text{pre}} = 1, & \gamma_{\text{pre}} = 1.4, \\ \rho_{0,\text{He}} = 0.138, & v_{x,\text{He}} = 0, & v_{y,\text{He}} = 0, \\ & p_{\text{He}} = 1, & \gamma_{\text{He}} = 1.67, \end{array} \right. \quad (5.4.1)$$

with the shocked condition initially in the region $x > 225$. The shock has been initialised such that it has $\mathfrak{M} = 1.22$ where \mathfrak{M} is the Mach number, given by

$$\mathfrak{M} = \frac{v_s}{c_{s,\text{pre}}} \quad (5.4.2)$$

with v_s the shock velocity, as defined in section 2.4.1. The level set function is initialised to give

$$\phi = 25 - \sqrt{(x - 175)^2 + y^2} \quad (5.4.3)$$

with $\phi > 0$ giving the helium bubble, and $\phi < 0$ the air. The symmetry of this problem about $y = 0$ means that the actual calculation need only use half of the domain (i.e. $y \in [0, 44.5]$). The test, shown in figure 5.1, was run with a resolution of 650×89 cells, and to a time $t \approx 185$.

The results for the shock wave hitting a cylindrical bubble of helium are shown for several stages of the evolution in figure 5.1. The low density region of the helium bubble (and hence the zeroes of the level set function) are clearly visible. The top panel shows the initial data, and the remaining panels show a succession of images as the shock wave passes through the bubble. We can clearly see the complex patterns of reflected features off the bubble, the outer walls and the symmetry axis. The bubble is also heavily distorted into a ‘C’ shape, as identified in previous similar simulations [22, 67]. Beyond the final time shown in this figure, $t \approx 185$, the physical behaviour of the system would cause the bubble to split. Numerically this presents a challenge in obtaining the correct behaviour for ϕ , which usually requires reinitialisation. For simplicity of implementation, it was therefore necessary to stop the simulation before this occurred.

5.5 Results for the special relativistic multidimensional Ghost Fluid Method

We now apply the Ghost Fluid Method to 2+1 dimensional situations in SR. One potential difference between Newtonian physics and relativity for the Ghost Fluid Method comes from the treatment of the velocities. In relativistic situations we may either use the 3-velocity or the 4-velocity to determine normal and tangential components. Differences may then arise when copying the normal and extrapolating the tangential components. The first test we consider in 2+1 dimensional is therefore designed to test this.

All tests in this section feature an isolated shock hitting some low-density region. We therefore will provide initial data for three constant states, the post-shock region, \mathbf{q}_{post} , the pre-shock region, \mathbf{q}_{post} , and the low density region, \mathbf{q}_{low} .

5.5.1 Comparing the velocity extrapolation techniques

In order to ensure that there is no obvious difference between extrapolating 3- and 4-velocities in SR we must consider a test with an exact solution, but




Figure 5.1: The 2+1 dimensional Ghost Fluid Method in Newtonian physics. This test shows a shock wave in air hitting a cylindrical bubble of helium. A succession of plots are shown from the evolution of this system, where we see a wide array of reflected and transmitted behaviour.

with both x and y components of the velocity non-zero. To achieve this, we use a ‘slab test’, as previously considered in section 4.2, propagating along a diagonal line within the domain. We set up a shock moving in the direction $x - y = \text{const.}$, and this hits a low density slab. The initial data for this test is

$$\left\{ \begin{array}{lll} \rho_{0,\text{post}} = 1.37795, & v_{x,\text{post}} = 0.12681, & v_{y,\text{post}} = 0.12681, \\ & p_{\text{post}} = 1.57, & \gamma_{\text{post}} = 1.4, \\ \rho_{0,\text{pre}} = 1, & v_{x,\text{pre}} = 0, & v_{y,\text{pre}} = 0, \\ & p_{\text{pre}} = 1, & \gamma_{\text{pre}} = 1.4, \\ \rho_{0,\text{He}} = 0.138, & v_{x,\text{He}} = 0, & v_{y,\text{He}} = 0, \\ & p_{\text{He}} = 1, & \gamma_{\text{He}} = 1.67, \end{array} \right. \quad (5.5.1)$$

with the post-shocked region initially given by

$$x + y < 0.75 \quad (5.5.2)$$

and the level set function by

$$\phi = 0.12 - |x + y - 1|. \quad (5.5.3)$$

The low density slab bounded is therefore by the contours $y = 1.22 - x$ and $y = 0.88 - x$.

The results for a shock hitting a slab in 2+1 dimensional SR are shown in figure 5.2. We show the results along the line $x = y$, orientated in the normal direction of the velocity. This test was implemented with outflow boundary conditions in the normal direction of the boundaries. As a result, severe errors propagate inwards from this boundary. We show these results before the errors have reached the line $x = y$. Qualitatively these results can be compared to the 1+1 dimensional analogy in figure 4.5. Quantitatively there are differences in scale (due to the increased maximum width of the domain along $x = y$). We see the expected reflected and transmitted shocks. The two panels in figure 5.2 compare the two cases using 3- and 4-velocity extrapolation, as explained in section 5.2.3, to determine normal and tangential components. It is clear that there is no difference between the 3-velocity extrapolation (upper panel) and 4-velocity extrapolation (lower panel) in the context of SR. As a result, all tests shown in this section will use only the 3-velocity extrapolation. We note




Figure 5.2: The density profile for a shock hitting a slab in 2+1 dimensional SR. The velocity is aligned with the line $x = y$, hence we have chosen this line for the plot. We see that qualitatively, all features expected in this test have been reproduced compared to the 1+1 dimensional test shown in figure 4.5. The upper panel shows the results when the 3-velocity is used in determining the normal and tangential components, and the lower panel uses the 4-velocity. It is clear that there is no qualitative difference between them. This test uses 200 by 200 cells and MC limiter.

that this result may not hold for GR.

5.5.2 Bubble tests for the Ghost Fluid Method in special relativity

In this section we present the results for various bubble tests (analogous to the Newtonian test shown in figure 5.1). To enable a more direct comparison between existing results in Newtonian physics, e.g. those in [22], we categorise these results though the relativistic Mach number of the shock, \mathfrak{M}_{SR} . This is given by ([109])

$$\mathfrak{M}_{\text{SR}} = \frac{\frac{v_s}{\sqrt{1-v_s^2}}}{\frac{c_{s,\text{pre}}}{\sqrt{1-c_{s,\text{pre}}^2}}}. \quad (5.5.4)$$

Due to the differences in scales between Newtonian physics and SR, all tests in this section have domain $x \in [-50, 325]$ and $y \in [-44.5, 44.5]$ where again the symmetry of this problem means we only simulate $y \in [0, 44.5]$. In all cases the level set function is given by (5.4.3),

$$\phi = 25 - \sqrt{(x - 175)^2 + y^2}. \quad (5.5.5)$$

We also do not use reinitialisation of the level set function in this preliminary work. As a result, in many cases the duration of the simulation is limited by the point at which the level set function can no longer be accurately modelled. Beyond this point, these errors can lead to unphysical behaviour.

The first test we consider has relativistic Mach number $\mathfrak{M}_{\text{SR}} = 1.22$, hence can be compared to the Newtonian test presented in figure 5.1. We can also use this test to compare different reconstruction methods in 2+1 dimensional SR, in particular focusing on the effects these have on the accuracy with which the interface is captured. The initial data for this test is given by

$$\left\{ \begin{array}{lll} \rho_{0,\text{post}} = 1.42182, & v_{x,\text{post}} = -0.19682, & v_{y,\text{post}} = 0, \\ & p_{\text{post}} = 1.64161, & \gamma_{\text{post}} = 1.4, \\ \rho_{0,\text{pre}} = 1, & v_{x,\text{pre}} = 0, & v_{y,\text{pre}} = 0, \\ & p_{\text{pre}} = 1, & \gamma_{\text{pre}} = 1.4, \\ \rho_{0,\text{low}} = 0.138, & v_{x,\text{low}} = 0, & v_{y,\text{low}} = 0, \\ & p_{\text{low}} = 1, & \gamma_{\text{low}} = 1.67, \end{array} \right. \quad (5.5.6)$$

with the shocked region when $x > 225$.

In figures 5.3 and 5.4 we show the results for a $\mathfrak{M}_{\text{SR}} = 1.22$ shock hitting a cylindrical bubble using MC limiter and PPM reconstruction respectively. The tests were run for resolution of 700 by 89 cells. For MC limiter, a final time of $t \approx 690$ was achieved, and for PPM, $t \approx 590$. Qualitatively many features in these two plots are similar. The reflected and transmitted features have been slightly more sharply captured using PPM, and in both cases are comparable to those seen in the analogous Newtonian case in figure 5.1. The key difference between the Newtonian and relativistic cases comes from the effects the shock has on the shape bubble. The differences between the two reconstruction methods is also apparent. There are finer details visible in figure 5.4 when PPM is used. This higher accuracy leads to stronger distortion of the bubble, and hence causes the shorter time for this simulation.

We would like to make a more direct comparison with the original results presented in [67]. In order to make this comparison we produce a Schlieren image of the density, obtained by plotting the quantity

$$\sqrt{(\partial_x \rho_0)^2 + (\partial_y \rho_0)^2} \quad (5.5.7)$$

on a logarithmic scale.

In figure 5.5 we compare the Schlieren image for the SR bubble test using MC limiter with that produced using the original Ghost Fluid Method. The differences in scale mean a direct comparison is not straightforward. We have shown the image for the SR test when the shock and bubble have moved by approximately the same amount as in the original image. The differences in bubble structure and speed of the reflected features is apparent. We can also clearly see a feature behind the shock only present in the SR test. This can be identified as start-up error in the shock (hence can be removed if necessary). The pattern of reflected and transmitted features is clearly consistent with the original results.

We now consider a test with $\mathfrak{M}_{\text{SR}} = 1.6$. The initial data for this test is given




Figure 5.3: Showing a $\mathfrak{M}_{SR}=1.22$ shock hitting a low density bubble in SR. MC limiter is used for this test, and we see that this leads to a very deformed ‘C’ shaped bubble. A series of snapshots are shown, and the pattern of reflected and transmitted features is clearly apparent.




Figure 5.4: Showing a $\mathfrak{M}_{\text{SR}} = 1.22$ shock hitting a low density bubble in SR with PPM reconstruction. This can be compared directly to figure 5.3. The greater accuracy of PPM leads to a more detailed bubble structure. The reflected and transmitted features are qualitatively similar.




Figure 5.5: Comparing the results for a $\mathfrak{M}_{\text{SR}} = 1.22$ shock wave hitting a bubble with analogous results in Newtonian physics presented in [67]. The differences in scale lead to very different behaviour of the bubble. The reflected and transmitted features are, however, comparable. In the SR case these features accurately follow the Newtonian results. Source for upper image: <http://physbam.stanford.edu/~fedkiw>

by

$$\left\{ \begin{array}{lll} \rho_{0,\text{post}} = 2.16527, & v_{x,\text{post}} = -0.42635, & v_{y,\text{post}} = 0, \\ & p_{\text{post}} = 3.04262, & \gamma_{\text{post}} = 1.4, \\ \rho_{0,\text{pre}} = 1, & v_{x,\text{pre}} = 0, & v_{y,\text{pre}} = 0, \\ & p_{\text{pre}} = 1, & \gamma_{\text{pre}} = 1.4, \\ \rho_{0,\text{low}} = 0.138, & v_{x,\text{low}} = 0, & v_{y,\text{low}} = 0, \\ & p_{\text{low}} = 1, & \gamma_{\text{low}} = 1.67, \end{array} \right. \quad (5.5.8)$$

with the post-shocked region initially located for $x > 300$ and the level set function given by (5.5.5). This test is shown for a resolution of 350 by 45 cells using PPM and run to a time $t \approx 550$.

In figure 5.6 we show the results for a $\mathfrak{M}_{\text{SR}} = 1.6$ shock hitting a low-density bubble. Many of the features are comparable to the milder shock case shown in figures 5.3 and 5.4. The bubble is again deformed into a curved shape, but is thinner in this case. It also follows the shock more closely than in this previous test. The interaction of waves off the bubble and boundaries




Figure 5.6: Results for a $\mathfrak{M}_{SR} = 1.6$ shock hitting a bubble in SR. We see that the larger shock has lead to greater deformation of the bubble, and that it follows the shock more closely. The general pattern of reflected and transmitted waves is similar to previous tests.

is similar to that seen in previous tests, as expected. There is also visible start up error present throughout the simulation.

The final test we consider has an even stronger shock, $\mathfrak{M}_{\text{SR}} = 2$. We now have initial data

$$\left\{ \begin{array}{lll} \rho_{0,\text{post}} = 2.93661, & v_{x,\text{post}} = -0.57968, & v_{y,\text{post}} = 0, \\ & p_{\text{post}} = 4.91233, & \gamma_{\text{post}} = 1.4, \\ \rho_{0,\text{pre}} = 1, & v_{x,\text{pre}} = 0, & v_{y,\text{pre}} = 0, \\ & p_{\text{pre}} = 1, & \gamma_{\text{pre}} = 1.4, \\ \rho_{0,\text{low}} = 0.138, & v_{x,\text{low}} = 0, & v_{y,\text{low}} = 0, \\ & p_{\text{low}} = 1, & \gamma_{\text{low}} = 1.67, \end{array} \right. \quad (5.5.9)$$

an initial post-shocked region $x > 225$ and a level set function again given by (5.5.5). This test uses 350 by 45 cells using PPM and is run to a time $t \approx 420$.

In figure 5.7 we show the results for a $\mathfrak{M}_{\text{SR}} = 2$ shock hitting a bubble. Here we see even greater contraction of the bubble than the previous $\mathfrak{M}_{\text{SR}} = 1.6$ test in figure 5.6. We see that this leads to the break up of the bubble. As expected, the level set methods have no problems in dealing with this change of topology. The interactions of the reflected and transmitted waves are again clearly apparent. It is clear that the x -velocity of the reflected features is very small, although there is still noticeable y -velocity, as features can be seen reflecting off the walls of the shock tube. As with the previous test there is clear start up error, visible in the post-shock region and trailing the bubble.

5.5.3 Summary

In this section we have presented some preliminary results for the extension of the Ghost Fluid Method to 2+1 dimensional SR. This work shows that in these (limited) cases of extensions from 1+1 dimensional tests and shock-tube tests with a low-density bubble, we obtain the correct qualitative behaviour. Before we can apply this to a full 3+1 dimensional NS in GR, a more rigorous investigation of this extension is required. In this thesis, time constraints limited both the quantity of tests that could be produced and the resolution




Figure 5.7: Results for a $\mathfrak{M}_{SR} = 2$ shock hitting a bubble in SR. The severe deformation of the bubble is now apparent, it has in fact split into four separate regions in the final panel. The qualitative behaviour of the reflected and transmitted features is comparable to previous tests, though we note that the x -velocity of the reflected features is now very small.




Figure 5.8: Labelling the features expected once a shock wave has hit a low-density bubble. We see the Mach stem, the side shock, the expansion fan and reflected interaction, as described by Bagabir and Drikakis in ‘Fig. 1. a’ in [22]

of those presented.

In addition to this, there are many cases which could not be simulated using the current methods implemented. If the behaviour of the bubble lead to complex features (as in figure 5.4) or if the bubble becomes compressed very rapidly (for shocks with $\mathfrak{M}_{SR} = 2$) then we could not reproduce this behaviour. This is liable to be partially due to the limited resolution, and also the lack of reinitialisation. Any future work would therefore make use of reinitialisation. Despite this room for improvement, these initial tests have been successful, and follow closely the qualitative behaviour shown from Newtonian shock-bubble interaction studies. We can compare our results to the expected features shown in Bagabir and Drikakis (‘Fig. 1. a’ in [22]).

In figure 5.8, we show the features visible as a result of the shock-bubble

interaction, as listed in [22]. We see the ‘side shock’ feature protruding from the ‘Mach stem’ and the expansion fan, which leads into the large reflected feature.

Chapter 6

Conclusions

The motivation for the work presented in this thesis was to improve the models used in numerical simulations of NSs, such that the internal behaviour could be better replicated. In order to achieve this, we wanted to develop techniques such that different regions within a NS could be included in a single model. Specifically we investigated techniques that allowed for these regions to be included in a multicomponent manner, based on relativistic extension of methods developed for Newtonian CFD.

Current models of the interior of NSs suggest that there are several regions (e.g. the crust and the inner and outer cores), and that these regions may be separated by thin transition layers. If we wish to simulate an entire NS (or binary system) in 3+1 dimensions, such that GW waveforms can be extracted, then the practical aspects of including these transition layers in the NS must be considered. In particular, for simulating a NS over a timescale long enough for e.g. several orbits and merger in a binary system, the inclusion of transition layers is prohibited through the computational requirements of such a simulation. In order to include different regions in a NS model, an alternative treatment is therefore necessary. We consider treating these transition layers as sharp (zero width) interfaces and implementing some appropriate boundary conditions to reproduce the behaviour across the layers.

The problem of developing numerical techniques to model multicomponent fluids separated by sharp interfaces has been tackled in Newtonian CFD. Whether these methods extend to relativistic situations has, however, not previously been considered. Therefore as a ‘proof-of-principle’ test we used the

simplest successful method from Newtonian CFD, the Ghost Fluid Method of Fedkiw et al. [67]. This method makes use of level set techniques to accurately capture the location of interfaces. The extension we present can effectively be split into three stages. We first ensure that the method extends to simple 1+1 dimensional situations in SR. Since a full NS simulation requires 3+1 dimensional GR, we then extend this to both 1+1 dimensional GR to ensure there are no issues introduced through the additional spacetime curvature introduced here. We finally must also consider the multidimensional extension of the method, which, due to limited computational resources, has currently been undertaken in 2+1 dimensional SR.

In chapter 2 we considered the theoretical background to modelling relativistic fluids. In this chapter we detail how initial data can be provided for multicomponent relativistic situations in which the EOS changes across the interface. Here we provide a novel presentation for the solution to multicomponent Riemann problems in SR (the feasibility of which was noted by [166]). We then detail the numerical methods required for HRSC methods and specific requirements for modelling stars in chapter 3. Section 3.7 details the specific techniques for modelling interfaces.

The results for tests using our relativistic extension of the Ghost Fluid Method are presented in chapter 4. In section 4.2 we apply the Ghost Fluid Method to multicomponent fluids in 1+1 dimensional SR. The initial tests for simple Riemann shock tube-type tests show that in these comparatively simple cases the exact solutions have been reproduced by the simulations correctly.

We then probe some more complicated scenarios to thoroughly test the versatility our extension to the Ghost Fluid Method in SR. We showed (in figure 4.18) that there were no issues in a highly relativistic test, with $v \sim 0.95c$. Similar success was achieved when cases which display more complex fluid behaviour are considered. This is demonstrated in figure 4.19, in which a shock has hit a slab of low density material. The expected pattern of resultant waves (transmitted shocks and reflected rarefactions) was correctly captured.

The simple nature of the Ghost Fluid Method can result in problems in dealing with some situations, notably those involving entropy gradients across an interface. This is well documented in Newtonian CFD [129]. To ensure

that these issues do not manifest in some more serious manner in relativistic situations, in section 4.2.1 we probed the behaviour of our extension to the Ghost Fluid Method in analogous relativistic situations to these ‘problem cases’. Again we find that the strong entropy gradients lead to small errors at the interface. These errors are, however, (slowly) convergent, and there is no additional unphysical behaviour introduced.

In section 4.3 we present tests for our extension of the Ghost Fluid Method in 1+1 GR. These results are based on two-component TOV-type stars. The density profiles have been set up such that they are similar to those of NSs, although the choice of interface location is selected with a purely numerical motivation. We ensure that it is in a region of sufficiently large density for the behaviour at the interface (and any undesirable affects that may also arise here) are visible.

Our results for static stars show that our method allows for a stable multicomponent star to be simulated without difficulty. We see that the convergence of these simulations is (in general) at second order, as expected from the numerical methods used. There is no significant loss of convergence near the interface, with the exception of the few grid cells adjacent to the interface location. We then investigate a perturbed star in which we investigate the effects of non-linear behaviour hitting the interface, shown in figure 4.35. We find that around the shocks, we see an expected reduction in convergence (see section 3.4 for details) and elsewhere the convergence remains at the desired level.

We have demonstrated that in 1+1 dimensions GR, our extension to the Ghost Fluid Method is successful in modelling for multicomponent situations. In order to be of use in current NS simulations, it must also be successful in 3+1 dimensional GR. In chapter 5, we present preliminary results to demonstrate the viability of our method in multidimensional relativistic situations. These results are presented for 2+1 dimensional SR, avoiding the computational cost of 3+1 GR and conceptual and numerical problems of 2+1 GR. We find that results based on extensions of problems investigated in Newtonian CFD provide the expected qualitative behaviour, as shown in figure 5.5.

The results presented in this thesis indicate that in general we can extend

interface modelling techniques from Newtonian CFD to relativity. This therefore suggests that such techniques will be applicable for full 3+1 dimensional NS simulations. We do note that some care may be needed when applying our method to such simulations. If strong entropy gradients form across an interface, then, as detailed in section 4.2.1, some error may be introduced into the results. This may lead to inaccuracies in the GW template obtained from these simulations. The success of our relativistic extension of the Ghost Fluid Method does, however, suggest that the techniques developed in Newtonian CFD to deal with such entropy gradients, e.g. [129], will also extend, allowing these issues to be tackled.

6.1 Future work

The work presented in this thesis provides a proof-of-principle test of multicomponent fluid modelling techniques in relativistic situations. We have demonstrated success in both 1+1 dimensional GR and 2+1 dimensional SR. The success of these early tests then leads to many obvious future applications. Most notable is that our extension to the Ghost Fluid Method should be applied to NS simulations in full 3+1 dimensional GR.

There is also obvious scope for improving the boundary conditions at the interface. As with the initial work this can follow from the Newtonian CFD literature, where improvements for certain aspects of the method have been considered, e.g. [125, 126].

In this work we have used level set methods to capture the location of interfaces between two fluid components. These techniques can, however, be applied to any such sharp boundary, e.g. the apparent horizon in black hole simulations. As a result, one natural extension of the techniques presented in this thesis would be to improve the treatment of the surface of a NS within numerical simulations. As we discuss in section 1.2, the atmosphere of a NS has a depth of ~ 1 cm, or may not exist at all. Therefore it is clear that the surface of a NS can be treated as a genuinely sharp matter/vacuum interface. In addition to using level set methods, we will need to apply boundary conditions. This case differs from the Ghost Fluid Method used in this thesis, since we do not

have a fluid interface. However, the appropriate boundary conditions could be applied using the framework of the method (i.e. by creating an artificial fluid which behaves as a surface). This could then provide a better boundary condition than the atmosphere treatment described in section 3.8. This is of obvious importance to current NS simulations since the lack of accuracy at the surface is a major source of error, as illustrated in the figures in section 4.3.

The use of level set methods to deal with the surface means we have genuine vacuum outside the NS and as a result, many of the issues associated with the atmosphere treatment, in particular the transfer of matter between the star and the atmosphere would no longer occur. The key challenge with implementing this matter/vacuum interface is the boundary conditions. Assuming the use of finite volume methods, we need to provide some correct value for the intercell fluxes in the vacuum region such that the matter behaves correctly at the surface. We would therefore expect these conditions to be dependent on both the values of the physical quantities at the surface (many of which drop trivially to zero) and the derivatives of these quantities.

Improving the physical model used in our simulations is another obvious extension of this work. Here we have restricted our model to the ideal fluid EOS, as described in section 2.3, however, there are a vast array of ‘realistic’ EOSs available for NS interiors. One of the advantages of the boundary conditions of the Ghost Fluid Method is that they are not constrained to a single choice of EOS, they depend on general properties of the fluid. In principle, therefore, these conditions can be applied to any choice of EOS. We may consider implementing a NS with a more realistic core and crust, or an inner and outer core. In these cases, an interface between two fluid components may not be sufficient to accurately reproduce the physical behaviour of these regions. We may wish to model the crust with an elastic lattice structure, or include a solid inner core. In these cases, as with the surface model described above, a Ghost Fluid Method-like method could still be used, with careful implementation of correct boundary conditions (see [66] for a description of a solid-fluid boundary Ghost Fluid Method in Newtonian CFD). The evolution of the level set function (and hence the interface) may also need to reflect the physics of these situations.

Appendix A

Calculations for relativistic fluid equations

A.1 Derivation of the 3+1 decomposition of the Einstein equations

In this section the 3+1 decomposition of the Einstein equations is derived in detail. Here we follow e.g. [44, 142].

A.1.1 Projection operators

In section 2.2.1 we describe the foliation of the spacetime \mathcal{M} into spacelike hypersurfaces Σ_τ separated by the normal vector n^μ . This dual-vector to this normal is given by (2.2.18),

$$n_\mu = -\alpha\Omega_\mu \quad (\text{A.1.1})$$

where α is the lapse function, and Ω_μ describes the hypersurfaces through (2.2.16),

$$\Omega_\mu = \nabla_\mu\tau. \quad (\text{A.1.2})$$

The torsion free nature of ∇_μ gives the result

$$\nabla_{[\mu}\Omega_{\nu]} = \nabla_{[\mu}\nabla_{\nu]}\tau = 0. \quad (\text{A.1.3})$$

The choice of sign for the normal dual-vector n_μ was chosen such that n^μ is future pointing, and as a result we have

$$n^\mu n_\mu = \alpha^2 g^{\mu\nu} \Omega_\mu \Omega_\nu = -1 \quad (\text{A.1.4})$$

where we have used the definition of α , (2.2.17),

$$g^{\mu\nu} \Omega_\mu \Omega_\nu = -\alpha^{-2}. \quad (\text{A.1.5})$$

The vector n^μ can be thought of as a 4-velocity field for observers moving orthogonally to Σ_τ and therefore move with 4-acceleration

$$a^\mu = n^\nu \nabla_\nu n^\mu. \quad (\text{A.1.6})$$

The spacelike projection tensor, \perp^μ_ν is defined in (2.2.19),

$$\perp^\mu_\nu \equiv \delta^\mu_\nu + n^\mu n_\nu. \quad (\text{A.1.7})$$

We can show that this tensor is purely spatial through the contraction of n^μ , a timelike vector,

$$\begin{aligned} \perp^\mu_\nu n^\nu &= \perp^\mu_\nu n^\nu \\ &= \delta^\mu_\nu n^\nu + n^\mu n_\nu n^\nu \\ &= n^\mu - n^\mu = 0. \end{aligned} \quad (\text{A.1.8})$$

We then define the spatial metric of the slices Σ_τ in (2.2.21),

$$\gamma_{\mu\nu} = g_{\mu\nu} + n_\mu n_\nu. \quad (\text{A.1.9})$$

Using this, it can be shown that the spatial derivative operator, $D_\mu = \perp \nabla_\mu$,

introduced in (2.2.23), has the expected property $D_\alpha \gamma_{\mu\nu} = 0$,

$$\begin{aligned} D_\alpha \gamma_{\mu\nu} &= \perp \nabla_\alpha (g_{\mu\nu} + n_\mu n_\nu) \\ &= \perp (n_\mu \nabla_\alpha n_\nu + n_\nu \nabla_\alpha n_\mu) \\ &= 0. \end{aligned} \tag{A.1.10}$$

In this derivation, we have used the fact that a spatial projection of the tensor product of n^μ or n_μ is zero, since this is the spatial projection of a timelike projection, as shown by (A.1.8).

Important results for the 3+1 split using the projection operators

It is instructive to consider some results that arise from the projection operators. We shall see that these results are important in making the 3+1 decomposition of the Einstein equations. The first such result we prove is

$$\perp \nabla_{[\mu} n_{\nu]} = 0. \tag{A.1.11}$$

This is achieved by considering the full form of $\perp \nabla_\mu n_\nu$,

$$\begin{aligned} \perp \nabla_\mu n_\nu &= \perp^\alpha_\mu \perp^\beta_\nu \nabla_\alpha n_\beta \\ &= (\delta^\alpha_\mu + n^\alpha n_\mu) (\delta^\beta_\nu + n^\beta n_\nu) \nabla_\alpha n_\beta \\ &= \nabla_\mu n_\nu + n_\nu n^\beta \nabla_\mu n_\beta + n_\mu n^\alpha \nabla_\alpha n_\nu + n^\alpha n_\mu n_\nu n^\beta \nabla_\alpha n_\beta. \end{aligned} \tag{A.1.12}$$

We then use the relationship

$$n^\beta \nabla_\mu n_\beta = n_\beta \nabla_\mu n^\beta = \frac{1}{2} \nabla_\mu (n_\beta n^\beta) = \frac{1}{2} \nabla_\mu (-1) = 0 \tag{A.1.13}$$

to simplify (A.1.12), giving

$$\perp \nabla_\mu n_\nu = \nabla_\mu n_\nu + n_\mu n^\alpha \nabla_\alpha n_\nu. \tag{A.1.14}$$

In order to show (A.1.11) holds, it is worth considering each term individually. The first term can be written

$$\begin{aligned}\nabla_\mu n_\nu &= -\nabla_\mu (\alpha \Omega_\nu) \\ &= -(\nabla_\mu \alpha) \Omega_\nu - \alpha (\nabla_\mu \Omega_\nu) \\ &= -(\nabla_\mu \alpha) \Omega_\nu - \alpha (\nabla_\mu \nabla_\nu \tau)\end{aligned}\quad (\text{A.1.15})$$

and from this, we can see that

$$\nabla_{[\mu} n_{\nu]} = -(\nabla_{[\mu} \alpha) \Omega_{\nu]}. \quad (\text{A.1.16})$$

Expanding the second term in a similar manner gives

$$\begin{aligned}n_\mu n^\alpha \nabla_\alpha n_\nu &= (-\alpha \Omega_\mu) (-\alpha \Omega^\alpha) \nabla_\alpha (-\alpha \Omega_\nu) \\ &= -\alpha^2 \Omega^\alpha (\nabla_\alpha \alpha) \Omega_\mu \Omega_\nu - \alpha^3 \Omega_\mu \Omega^\alpha (\nabla_\alpha \Omega_\nu) \\ &= -\alpha^2 \Omega^\alpha (\nabla_\alpha \alpha) \Omega_\mu \Omega_\nu - \alpha^3 \Omega_\mu \Omega^\alpha (\nabla_\nu \Omega_\alpha)\end{aligned}\quad (\text{A.1.17})$$

where the relation $\nabla_{[\mu} \Omega_{\nu]} = 0$ has been used to manipulate the indices in the last line of this expansion. The relation

$$\Omega^\nu \nabla_\mu \Omega_\nu = \Omega_\nu \nabla_\mu \Omega^\nu = \frac{1}{2} \nabla_\mu (\Omega_\nu \Omega^\nu) = \frac{1}{2} \nabla_\mu (-\alpha^{-2}) = \alpha^{-3} \nabla_\mu \alpha, \quad (\text{A.1.18})$$

can then be used to further rewrite (A.1.17),

$$n_\mu n^\alpha \nabla_\alpha n_\nu = -\alpha^2 \Omega^\alpha (\nabla_\alpha \alpha) \Omega_\mu \Omega_\nu - (\nabla_\nu \alpha) \Omega_\mu \quad (\text{A.1.19})$$

and from this, in a similar manner to that used for (A.1.16), we get

$$n_{[\mu} n^\alpha \nabla_\alpha n_{\nu]} = (\nabla_{[\mu} \alpha) \Omega_{\nu]}. \quad (\text{A.1.20})$$

Then finally, combining this and (A.1.16), we see that

$$\perp \nabla_{[\mu} n_{\nu]} = \nabla_{[\mu} n_{\nu]} + n_{[\mu} n^\alpha \nabla_\alpha n_{\nu]} = 0. \quad (\text{A.1.21})$$

Another useful result we will derive is

$$a_\mu = D_\mu (\ln \alpha). \quad (\text{A.1.22})$$

This derivation is best shown by considering an expansion of each side of this equation individually. Using (A.1.19), and the definition of a_μ , (A.1.6), we can see

$$a_\mu = n^\nu \nabla_\nu n_\mu = \alpha \Omega^\nu (\nabla_\nu \alpha) \Omega_\mu + \alpha^{-1} \nabla_\mu \alpha. \quad (\text{A.1.23})$$

The right hand side of (A.1.22) can be written

$$\begin{aligned} D_\mu (\ln \alpha) &= \perp_\mu^\nu \nabla_\nu (\ln \alpha) \\ &= (\delta_\mu^\nu + n^\nu n_\mu) \alpha^{-1} \nabla_\nu \alpha \\ &= \alpha^{-1} \nabla_\mu \alpha + (-\alpha \Omega^\nu) (-\alpha \Omega_\mu) \alpha^{-1} \nabla_\nu \alpha \\ &= \alpha \Omega^\nu (\nabla_\nu \alpha) \Omega_\mu + \alpha^{-1} \nabla_\mu \alpha. \end{aligned} \quad (\text{A.1.24})$$

This provides an identical result to (A.1.23), hence we find that (A.1.22), holds.

A.1.2 Extrinsic curvature tensor

Several forms exist for the extrinsic curvature tensor, the tensor which describes how the hypersurfaces Σ_τ are embedded in the spacetime manifold \mathcal{M} . In (2.2.26), we give the definition

$$K_{\mu\nu} = - \perp \nabla_\mu n_\nu \quad (\text{A.1.25})$$

which, using (A.1.14) gives

$$K_{\mu\nu} = -\nabla_\mu n_\nu - n_\mu a_\nu. \quad (\text{A.1.26})$$

To show that this definition is equivalent to those using the Lie-derivative of the spatial metric, (2.2.27),

$$K_{\mu\nu} = -\frac{1}{2} \mathcal{L}_n \gamma_{\mu\nu} = -\frac{1}{2} \perp \mathcal{L}_n g_{\mu\nu}. \quad (\text{A.1.27})$$

we first consider the Lie-derivative of $\gamma_{\mu\nu}$ along n^α ,

$$\begin{aligned}
 -\mathcal{L}_n \gamma_{\mu\nu} &= n^\alpha \nabla_\alpha (g_{\mu\nu} + n_\mu n_\nu) + (g_{\alpha\nu} + n_\alpha n_\nu) \nabla_\mu n^\alpha + \\
 &\quad (g_{\mu\alpha} + n_\mu n_\alpha) \nabla_\nu n^\alpha \\
 &= n^\alpha n_\mu \nabla_\alpha n_\nu + n^\alpha n_\nu \nabla_\alpha n_\mu + \nabla_\mu n_\nu + n_\nu n_\alpha \nabla_\mu n^\alpha \\
 &\quad + \nabla_\nu n_\mu + n_\mu n_\alpha \nabla_\nu n^\alpha.
 \end{aligned} \tag{A.1.28}$$

We can then use (A.1.13), and then (A.1.25), to reduce this to

$$\mathcal{L}_n \gamma_{\mu\nu} = 2 (\nabla_{(\mu} n_{\nu)} + n_{(\mu} a_{\nu)}) = 2 \perp \nabla_{(\mu} n_{\nu)} = -2K_{\mu\nu} \tag{A.1.29}$$

which gives the definition of extrinsic curvature (A.1.27). To express $K_{\mu\nu}$ using the Lie-derivative (again in the direction n^α) of $g_{\mu\nu}$, we begin with

$$\begin{aligned}
 \mathcal{L}_n g_{\mu\nu} &= n^\alpha \nabla_\alpha g_{\mu\nu} + g_{\alpha\nu} \nabla_\mu n^\alpha + g_{\mu\alpha} \nabla_\nu n^\alpha \\
 &= \nabla_\mu n_\nu + \nabla_\nu n_\mu \\
 &= 2\nabla_{(\mu} n_{\nu)}.
 \end{aligned} \tag{A.1.30}$$

It is then clear that we have

$$\perp \mathcal{L}_n g_{\mu\nu} = 2 \perp \nabla_{(\mu} n_{\nu)} = -2K_{\mu\nu}. \tag{A.1.31}$$

The further results for $K_{\mu\nu}$ presented in section 2.2.1 are found by first considering projections of the 4-dimensional Riemann tensor $R_{\mu\nu\alpha\beta}$. To obtain the spatial projection of $R_{\mu\nu\alpha\beta}$, we first note the following relationship, in which v_μ is an arbitrary spatial dual-vector,

$$\begin{aligned}
 v^\mu \perp R_{\mu\nu\alpha\beta} &= \perp (v^\mu R_{\mu\nu\alpha\beta}) \\
 &= \perp (R^\mu_{\nu\alpha\beta} v_\mu) \\
 &= \perp (R_{\beta\alpha\nu}^\mu v_\mu) \\
 &= \perp (\nabla_\beta \nabla_\alpha v_\nu - \nabla_\alpha \nabla_\beta v_\nu)
 \end{aligned} \tag{A.1.32}$$

where the symmetries of the Riemann tensor are taken into account to obtain

the third line. We can then consider

$$\begin{aligned}\perp \nabla_\mu v_\nu &= (\delta_\mu^\alpha + n^\alpha n_\mu) (\delta_\nu^\beta + n^\beta n_\nu) \nabla_\alpha v_\beta \\ &= \nabla_\mu v_\nu + n_\nu n^\beta \nabla_\mu v_\beta + n_\mu n^\alpha \nabla_\alpha v_\nu + n_\nu n_\mu n^\alpha n^\beta \nabla_\alpha v_\beta.\end{aligned}\quad (\text{A.1.33})$$

Using the result

$$\nabla_\nu (n^\mu v_\mu) = n^\mu \nabla_\nu v_\mu + v_\mu \nabla_\nu n^\mu = 0 \quad (\text{A.1.34})$$

and then (A.1.6), we can rearrange (A.1.33) to give

$$\perp \nabla_\mu v_\nu = \nabla_\mu v_\nu - n_\nu v_\alpha \nabla_\mu n^\alpha + n_\mu n^\beta \nabla_\beta v_\nu - n_\mu n_\nu v_\alpha a^\alpha. \quad (\text{A.1.35})$$

This can then be used to show that

$$\begin{aligned}D_\alpha D_\nu v_\mu &= \perp \nabla_\alpha (\perp \nabla_\nu v_\mu) \\ &= \perp \nabla_\alpha (\nabla_\nu v_\mu + n_\nu n^\gamma \nabla_\gamma v_\mu - n_\mu v_\gamma \nabla_\nu n^\gamma - n_\nu n_\mu v_\gamma a^\gamma) \\ &= \perp \nabla_\alpha \nabla_\nu v_\mu - \perp (\nabla_\alpha n_\mu) (\nabla_\nu n_\gamma) v^\gamma \\ &= \perp \nabla_\alpha \nabla_\nu v_\mu - K_{\alpha\mu} K_{\nu\gamma} v^\gamma\end{aligned}\quad (\text{A.1.36})$$

where this derivation has made use of the fact that $\perp n^\mu = 0$. Rearranging the above result, and using this, along with (2.2.24),

$$(D_\mu D_\nu - D_\nu D_\mu) w_\alpha = {}^{(3)}R_{\mu\nu\alpha}^\beta w_\beta, \quad (\text{A.1.37})$$

in (A.1.32), gives

$$\begin{aligned}v^\mu \perp R_{\mu\nu\alpha\beta} &= D_\beta D_\alpha v_\nu - D_\alpha D_\beta v_\nu + K_{\beta\nu} K_{\alpha\mu} v^\mu - K_{\alpha\nu} K_{\beta\mu} v^\mu \\ &= (\mathcal{R}_{\mu\nu\alpha\beta} + K_{\nu\beta} K_{\mu\alpha} - K_{\nu\alpha} K_{\mu\beta}) v^\mu.\end{aligned}\quad (\text{A.1.38})$$

Therefore we now have the spatial projection of the 4-dimensional Riemann tensor,

$$\perp R_{\mu\nu\alpha\beta} = \mathcal{R}_{\mu\nu\alpha\beta} + K_{\beta\nu} K_{\alpha\mu} - K_{\alpha\nu} K_{\beta\mu}. \quad (\text{A.1.39})$$

We now want to compute the spatial projection $\perp R_{\mu\nu\alpha\hat{n}}$ where we have

$$R_{\mu\nu\alpha\hat{n}} = R_{\mu\nu\alpha\beta} n^\beta. \quad (\text{A.1.40})$$

Using (A.1.32) and (A.1.11), we obtain

$$\begin{aligned}
\perp R_{\hat{n}\mu\nu\alpha} &= \perp (\nabla_\alpha \nabla_\nu n_\mu - \nabla_\nu \nabla_\alpha n_\mu) \\
&= - \perp [\nabla_\alpha (K_{\nu\mu} + n_\nu a_\mu) - \nabla_\nu (K_{\alpha\mu} + n_\alpha a_\mu)] \\
&= - \perp [\nabla_\alpha K_{\nu\mu} - \nabla_\nu K_{\alpha\mu} + (\nabla_\alpha n_\nu - \nabla_\nu n_\alpha) a_\mu - \\
&\quad n_\nu \nabla_\alpha a_\mu + n_\alpha \nabla_\nu a_\mu] \\
&= -(D_\alpha K_{\nu\mu} - D_\nu K_{\alpha\mu})
\end{aligned} \tag{A.1.41}$$

hence the desired spatial projection is

$$\perp R_{\mu\nu\alpha\hat{n}} = D_\nu K_{\mu\alpha} - D_\mu K_{\nu\alpha}. \tag{A.1.42}$$

These two spatial projections of the 4-dimensional Riemann tensor, (A.1.39) and (A.1.42), are collectively known as the Gauss-Codazzi equations.

A.1.3 Preliminary results for the 3+1 decomposition of the Einstein equations

When making the 3+1 split of an arbitrary symmetric tensor, we use (2.2.28),

$$\sigma_{\mu\nu} = \perp \sigma_{\mu\nu} - 2n_{(\mu} \perp \sigma_{\nu)\hat{n}} + n_\mu n_\nu \sigma_{\hat{n}\hat{n}}. \tag{A.1.43}$$

The Einstein tensor $G_{\mu\nu}$ and stress-energy tensor $T_{\mu\nu}$ are symmetric tensors, therefore this can be used to obtain the 3+1 split of the Einstein equations. However, before we do this, we first consider some important preliminary results.

The spatial contraction of the 4-dimensional Ricci tensor (which appears in $G_{\mu\nu}$), which, using (2.2.22), can be written

$$\begin{aligned}
\perp R_{\mu\nu} &= \perp (g^{\alpha\beta} R_{\mu\alpha\nu\beta}) \\
&= \perp (\gamma^{\alpha\beta} R_{\mu\alpha\nu\beta}) - \perp R_{\mu\hat{n}\nu\hat{n}}.
\end{aligned} \tag{A.1.44}$$

Since $\gamma^{\mu\nu}$ is a spatial tensor, the first term in the result above becomes

$$\perp (\gamma^{\alpha\beta} R_{\mu\alpha\nu\beta}) = \gamma^{\alpha\beta} \perp R_{\mu\alpha\nu\beta} = g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta} \tag{A.1.45}$$

which can then be used in (A.1.44) to give

$$\perp R_{\mu\nu} = g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta} - \perp R_{\mu\hat{n}\nu\hat{n}}. \quad (\text{A.1.46})$$

By using a 3+1 decomposition, as in (A.1.43), on $R_{\mu\hat{n}\nu\hat{n}}$ we get

$$R_{\mu\hat{n}\nu\hat{n}} = \perp R_{\mu\hat{n}\nu\hat{n}} - 2n_{(\mu} \perp R_{\nu)\hat{n}\hat{n}} + n_\mu n_\nu R_{\hat{n}\hat{n}\hat{n}\hat{n}} = \perp R_{\mu\hat{n}\nu\hat{n}} \quad (\text{A.1.47})$$

in which the final result is achieved since most terms vanish due to the anti-symmetry of the Riemann tensor over its first- and last-two indices. This allows us to write (A.1.46) as

$$\perp R_{\mu\nu} = g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta} - R_{\mu\hat{n}\nu\hat{n}}. \quad (\text{A.1.48})$$

It is then useful to consider the contraction of $\perp R_{\mu\nu}$ with $g^{\mu\nu}$, which, using the result above gives

$$g^{\mu\nu} \perp R_{\mu\nu} = -R_{\hat{n}\hat{n}} + g^{\mu\nu} g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta}. \quad (\text{A.1.49})$$

If we consider the 3+1 decomposition of $R_{\mu\nu}$, using (2.2.28), then we can obtain an alternative result for the contraction detailed above. A rearrangement of this decomposition gives

$$\perp R_{\mu\nu} = R_{\mu\nu} + 2n_{(\mu} \perp R_{\nu)\hat{n}} - n_\mu n_\nu R_{\hat{n}\hat{n}}. \quad (\text{A.1.50})$$

Contraction of this with $g^{\mu\nu}$ then gives

$$\begin{aligned} g^{\mu\nu} \perp R_{\mu\nu} &= R + \frac{1}{2} (n^\nu \perp R_{\nu\hat{n}} + n^\mu \perp R_{\mu\hat{n}}) - n^\mu n_\mu R_{\hat{n}\hat{n}} \\ &= R + R_{\hat{n}\hat{n}} \end{aligned} \quad (\text{A.1.51})$$

which, when combined with (A.1.49), gives a useful result for R ,

$$R = -2R_{\hat{n}\hat{n}} + g^{\mu\nu} g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta}. \quad (\text{A.1.52})$$

In (2.2.29), we introduce the physical interpretations of the contractions of the

stress energy tensor,

$$\varrho \equiv T_{\hat{n}\hat{n}} = T_{\mu\nu} n^\mu n^\nu, \quad (\text{A.1.53a})$$

$$j^\mu \equiv \perp T^{\mu\hat{n}} = - \perp (T^{\mu\nu} n_\nu), \quad (\text{A.1.53b})$$

$$\mathfrak{E}_{\mu\nu} \equiv \perp T_{\mu\nu}. \quad (\text{A.1.53c})$$

These quantities are the local energy density (ϱ) and momentum density (j^μ), and the local stress-energy tensor measured by observers moving along n^μ ($\mathfrak{E}_{\mu\nu}$). Note that when contracting a tensor with n_μ , a minus sign is introduced [217].

We now consider some useful results concerning an arbitrary spatial tensor of the form $\mathbf{S} = S_{\mu_1, \mu_2 \dots \mu_k}$. Firstly we show that for this spatial tensor, $\mathcal{L}_n \mathbf{S}$ is also a spatial tensor, i.e.

$$\perp \mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k} = \mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k}. \quad (\text{A.1.54})$$

We demonstrate this by first expanding the Lie-derivative,

$$\mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k} = n^\nu \nabla_\nu S_{\mu_1, \mu_2 \dots \mu_k} + \sum_{i=1}^k (\nabla_{\mu_i} n^\nu) S_{\mu_1, \mu_2 \dots \nu \dots \mu_k} \quad (\text{A.1.55})$$

and then contract the j^{th} index with n^μ ,

$$n^{\mu_j} \mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k} = n^{\mu_j} n^\nu \nabla_\nu S_{\mu_1, \mu_2 \dots \mu_k} + \sum_{i=1}^k (\nabla_{\mu_i} n^\nu) n^{\mu_j} S_{\mu_1, \mu_2 \dots \nu \dots \mu_k}. \quad (\text{A.1.56})$$

Because \mathbf{S} is a spatial tensor, all components within the sum are zero except when $i = j$. We can also use an analogous relationship to (A.1.34) to ‘move’ the derivative term onto n^{μ_j} , giving

$$n^{\mu_j} \mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k} = -n^\nu (\nabla_\nu n^{\mu_j}) S_{\mu_1, \mu_2 \dots \mu_k} + n^\nu (\nabla_\nu n^{\mu_j}) S_{\mu_1, \mu_2 \dots \mu_k} = 0, \quad (\text{A.1.57})$$

and since the projection of $\mathcal{L}_n \mathbf{S}$ with n^μ vanishes, it must be a spatial tensor, hence we obtain the desired result (A.1.54).

Secondly, we want to show that for any arbitrary function, f , and spatial

tensor \mathbf{S} , we have

$$\mathcal{L}_{fn} S_{\mu_1, \mu_2 \dots \mu_k} = f \mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k}. \quad (\text{A.1.58})$$

Again we begin this proof by expanding the Lie-derivative, and obtain

$$\begin{aligned} \mathcal{L}_{fn} S_{\mu_1, \mu_2 \dots \mu_k} &= fn^\nu \nabla_\nu S_{\mu_1, \mu_2 \dots \mu_k} + \sum_{i=1}^k \nabla_{\mu_i} (fn^\nu) S_{\mu_1, \mu_2 \dots \nu \dots \mu_k} \\ &= fn^\nu \nabla_\nu S_{\mu_1, \mu_2 \dots \mu_k} + \\ &\quad \sum_{i=1}^k [(\nabla_{\mu_i} f) n^\nu S_{\mu_1, \mu_2 \dots \nu \dots \mu_k} + f (\nabla_{\mu_i} n^\nu) S_{\mu_1, \mu_2 \dots \nu \dots \mu_k}] \\ &= f \left[n^\nu \nabla_\nu S_{\mu_1, \mu_2 \dots \mu_k} + \sum_{i=1}^k (\nabla_{\mu_i} n^\nu) S_{\mu_1, \mu_2 \dots \nu \dots \mu_k} \right] \\ &= f \mathcal{L}_n S_{\mu_1, \mu_2 \dots \mu_k}. \end{aligned} \quad (\text{A.1.59})$$

We shall now argue that in the decomposition of the Einstein equations, the vector

$$N^\mu = \alpha n^\mu \quad (\text{A.1.60})$$

is a natural choice for use when computing time derivatives. making use of the results given above. To aid this argument, we first will show that one property of N^μ is

$$\mathcal{L}_N \perp^\mu_\nu = 0 \quad (\text{A.1.61})$$

and that this implies

$$\perp \mathcal{L}_N S^{\mu_1, \mu_2 \dots \mu_k}_{\nu_1, \nu_2 \dots \nu_l} = \mathcal{L}_N S^{\mu_1, \mu_2 \dots \mu_k}_{\nu_1, \nu_2 \dots \nu_l}, \quad (\text{A.1.62})$$

where this is now a result for a general spatial tensor (rather than the covariant tensor used previously). The first result can be demonstrated, making use

of (A.1.6), (A.1.7), (A.1.13) and the expanded form of $D_\mu = \perp \nabla_\mu$, to give

$$\begin{aligned}
\mathcal{L}_N \perp^\mu_\nu &= N^\alpha \nabla_\alpha \perp^\mu_\nu - \perp^\alpha_\nu \nabla_\alpha N^\mu + \perp^\mu_\alpha \nabla_\nu N^\alpha \\
&= (\alpha n^\alpha) \nabla_\alpha (\delta^\mu_\nu + n^\mu n_\nu) - (\delta^\alpha_\nu + n^\alpha n_\nu) \nabla_\alpha (\alpha n^\mu) \\
&\quad + (\delta^\mu_\alpha + n^\mu n_\alpha) \nabla_\nu (\alpha n^\alpha) \\
&= \alpha n^\mu n^\alpha \nabla_\alpha n_\nu + \alpha n_\nu n^\alpha \nabla_\alpha n^\mu - \nabla_\nu (\alpha n^\mu) \\
&\quad - n_\nu n^\alpha \nabla_\alpha (\alpha n^\mu) + \nabla_\nu (\alpha n^\mu) + n^\mu n_\alpha \nabla_\nu (\alpha n^\alpha) \\
&= \alpha n^\mu a_\nu + \alpha n_\nu a^\mu - n_\nu n^\mu n^\alpha \nabla_\alpha \alpha - \alpha n_\nu a^\mu \\
&\quad + \alpha n^\mu n_\alpha \nabla_\nu n^\alpha - n^\mu \nabla_\nu \alpha \\
&= \alpha n^\mu [a_\nu - \alpha^{-1} (\nabla_\nu \alpha + n_\nu n^\alpha \nabla_\alpha \alpha)] \\
&= \alpha n^\mu (a_\nu - \alpha^{-1} D_\nu \alpha) \\
&= \alpha n^\mu [a_\nu - D_\nu (\ln \alpha)] \\
&= \alpha n^\mu (a_\nu - a_\nu) = 0.
\end{aligned} \tag{A.1.63}$$

To use this result to show (A.1.62), we first note that since \mathbf{S} is a spatial tensor,

$$S^{\mu_1, \mu_2, \dots, \mu_k}_{\nu_1, \nu_2, \dots, \nu_l} = \perp^{\mu_1}_{\alpha_1} \perp^{\mu_2}_{\alpha_2} \dots \perp^{\mu_k}_{\alpha_k} \perp^{\nu_1}_{\beta_1} \perp^{\nu_2}_{\beta_2} \dots \perp^{\nu_l}_{\beta_l} S^{\alpha_1, \alpha_2, \dots, \alpha_k}_{\beta_1, \beta_2, \dots, \beta_l}, \tag{A.1.64}$$

hence taking the Lie-derivative of this, and using (A.1.63), we see

$$\begin{aligned}
\mathcal{L}_N S^{\mu_1, \mu_2, \dots, \mu_k}_{\nu_1, \nu_2, \dots, \nu_l} &= \\
\perp^{\mu_1}_{\alpha_1} \perp^{\mu_2}_{\alpha_2} \dots \perp^{\mu_k}_{\alpha_k} \perp^{\nu_1}_{\beta_1} \perp^{\nu_2}_{\beta_2} \dots \perp^{\nu_l}_{\beta_l} \mathcal{L}_N S^{\alpha_1, \alpha_2, \dots, \alpha_k}_{\beta_1, \beta_2, \dots, \beta_l}, &
\end{aligned} \tag{A.1.65}$$

which is simply the expanded form of (A.1.62).

We now justify our choice of N^μ as the natural orthogonal vector. We introduced the dual-vector field, Ω_μ , that locally describes the hypersurface slices, and is related to n^μ through $n^\mu = -\alpha \Omega^\mu$ (the contravariant form of (A.1.1)). We then have

$$N^\mu \Omega_\mu = \alpha n^\mu \Omega_\mu = -\alpha^2 \Omega^\mu \Omega_\mu = 1. \tag{A.1.66}$$

since from (A.1.5) we have $\Omega^\mu \Omega_\mu = -\alpha^{-2}$. It is this normalisation which makes N^μ a natural choice for use in the Lie-derivatives (and hence evolution equations) of $\gamma_{\mu\nu}$ and $K_{\mu\nu}$. However, there is no reason to restrict the vector

for use in these Lie-derivatives to being normal to the slices. We can therefore generalise N^μ by the addition of an arbitrary spatial vector, β^μ , known as the shift vector,

$$t^\mu = N^\mu + \beta^\mu. \quad (\text{A.1.67})$$

Because β^μ is purely spatial, t^μ still has the same normalisation as N^μ ,

$$t^\mu \Omega_\mu = N^\mu \Omega_\mu + \beta^\mu \Omega_\mu = 1 + \alpha^{-1} \beta^\mu n_\mu = 1. \quad (\text{A.1.68})$$

With this generalised vector, the metric of the 3+1 decomposition can be written as ([9, 142]),

$$ds^2 = -(\alpha^2 - \beta_i \beta^i) dt^2 + 2\beta_i dx^i dt + \gamma_{ij} dx^i dx^j. \quad (\text{A.1.69})$$

for which we have normal vector components

$$n^\mu = \left(\frac{1}{\alpha}, -\frac{\beta^i}{\alpha} \right)^T \quad (\text{A.1.70})$$

and normal dual-vector components

$$n_\mu = (-\alpha, 0, 0, 0). \quad (\text{A.1.71})$$

In order to derive the evolution equations that result from the 3+1 split of the Einstein equations it is instructive to consider the spatial projection of $R_{\mu\hat{n}\nu\hat{n}}$. Based on the definition of the Riemann tensor, this is given by

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \perp [n^\alpha (\nabla_\nu \nabla_\alpha n_\mu - \nabla_\alpha \nabla_\nu n_\mu)]. \quad (\text{A.1.72})$$

Using the definition of the extrinsic curvature, (2.2.26), becomes

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \perp [n^\alpha \nabla_\alpha (K_{\nu\mu} + n_\nu a_\mu) - n^\alpha \nabla_\nu (K_{\alpha\mu} + n_\alpha a_\mu)], \quad (\text{A.1.73})$$

which, when expanded, and using (A.1.23), $a_\mu = \alpha \Omega^\nu (\nabla_\nu \alpha) \Omega_\mu + \alpha^{-1} \nabla_\mu \alpha$,

gives

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \perp (n^\alpha \nabla_\alpha K_{\nu\mu} + a_\nu a_\mu + n_\nu n^\alpha \nabla_\alpha a_\mu - n^\alpha \nabla_\nu K_{\alpha\mu} - a_\mu n^\alpha \nabla_\nu n_\alpha - n^\alpha n_\alpha \nabla_\nu a_\mu). \quad (\text{A.1.74})$$

Many terms then vanish under the spatial projection, and by using $n^\mu n_\mu = -1$, this then simplifies to

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \perp (n^\alpha \nabla_\alpha K_{\nu\mu} + a_\nu a_\mu - n^\alpha \nabla_\nu K_{\alpha\mu} + \nabla_\nu a_\mu). \quad (\text{A.1.75})$$

Since $K_{\mu\nu}$ is a spatial tensor, i.e. $n^\nu K_{\mu\nu} = 0$, the covariant derivative of this quantity also vanishes, hence using the product rule we have

$$-n^\alpha \nabla_\nu K_{\alpha\mu} = (\nabla_\nu n^\alpha) K_{\alpha\mu}. \quad (\text{A.1.76})$$

We then use this in (A.1.75), giving

$$\begin{aligned} \perp R_{\mu\hat{n}\nu\hat{n}} = & \perp [n^\alpha \nabla_\alpha K_{\mu\nu} + (\nabla_\nu n^\alpha) K_{\alpha\mu} \\ & + (\nabla_\mu n^\alpha) K_{\alpha\nu} - (\nabla_\mu n^\alpha) K_{\alpha\nu} + a_\nu a_\mu + \nabla_\nu a_\mu] \end{aligned} \quad (\text{A.1.77})$$

where we have added and then subtracted a $(\nabla_\mu n^\alpha) K_{\alpha\nu}$ term from the equation. This last step then allows this equation to be written as

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \perp [\mathcal{L}_n K_{\mu\nu} - (\nabla_\mu n^\alpha) K_{\alpha\nu} + a_\nu a_\mu + \nabla_\nu a_\mu]. \quad (\text{A.1.78})$$

From the definition of the extrinsic curvature, (2.2.26), we can write $-\perp (\nabla_\mu n^\alpha) K_{\alpha\nu} = K_\mu^\alpha K_{\nu\alpha}$, and since $K_{\mu\nu}$ is a spatial tensor, this can be brought back inside the spatial projection, giving

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \perp (\mathcal{L}_n K_{\mu\nu} + K_\mu^\alpha K_{\nu\alpha} + a_\mu a_\nu + \nabla_\nu a_\mu). \quad (\text{A.1.79})$$

The final two terms in this equation, when used with (A.1.22) can be written

$$\perp (a_\mu a_\nu + \nabla_\nu a_\mu) = \perp [D_\mu \ln \alpha D_\nu \ln \alpha + \nabla_\nu (D_\mu \ln \alpha)] \quad (\text{A.1.80})$$

which simplifies to

$$\perp (a_\mu a_\nu + \nabla_\nu a_\mu) = \perp [\alpha^{-2} D_\mu \alpha D_\nu \alpha + \nabla_\nu (\alpha^{-1} D_\mu \alpha)]. \quad (\text{A.1.81})$$

The first term is entirely spatial, so using this, and expanding the covariant derivative gives

$$\perp (a_\mu a_\nu + \nabla_\nu a_\mu) = \alpha^{-2} D_\mu \alpha D_\nu \alpha + \alpha^{-1} \perp \nabla_\nu D_\mu \alpha - \alpha^{-2} \perp \nabla_\nu \alpha D_\mu \alpha. \quad (\text{A.1.82})$$

The definition of the spatial derivative operator $D_\mu = \perp \nabla_\mu$, (2.2.23), means that the first and last terms here cancel and we obtain

$$\perp (a_\mu a_\nu + \nabla_\nu a_\mu) = \alpha^{-1} D_\mu D_\nu \alpha. \quad (\text{A.1.83})$$

We can also simplify the projection of the Lie-derivative in (A.1.79), using the results in (A.1.54) and (A.1.58), with $N^\mu = \alpha n^\mu$. From this, we see that since $K_{\mu\nu}$ is a spatial tensor, we have

$$\perp \mathcal{L}_n K_{\mu\nu} = \mathcal{L}_n K_{\mu\nu} = \alpha^{-1} \mathcal{L}_N K_{\mu\nu}. \quad (\text{A.1.84})$$

Using these two results, we obtain the spatial projection of $R_{\mu\hat{n}\nu\hat{n}}$,

$$\perp R_{\mu\hat{n}\nu\hat{n}} = \alpha^{-1} \mathcal{L}_N K_{\mu\nu} + K_{\mu\alpha} K_\nu^\alpha + \alpha^{-1} D_\mu D_\nu \alpha. \quad (\text{A.1.85})$$

A.1.4 The 3+1 split of the Einstein equations

Taking the 3+1 split of the Einstein equations results in two constraint equations and two evolution equations. Here we consider each of these in turn.

Hamiltonian constraint

If we consider the contraction of both indices of the Einstein equations with n^μ then we get

$$G_{\mu\nu}n^\mu n^\nu = R_{\mu\nu}n^\mu n^\nu - \frac{1}{2}g_{\mu\nu}n^\mu n^\nu R = 8\pi T_{\mu\nu}n^\mu n^\nu \quad (\text{A.1.86})$$

which, using (A.1.53a) and $n^\mu n_\nu = -1$, can be written

$$R_{\hat{n}\hat{n}} + \frac{1}{2}R = 8\pi\varrho. \quad (\text{A.1.87})$$

We then use (A.1.52) to replace the Ricci scalar, giving

$$R_{\hat{n}\hat{n}} + \frac{1}{2}(-2R_{\hat{n}\hat{n}} + g^{\mu\nu}g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta}) = 8\pi\varrho \quad (\text{A.1.88})$$

which simplifies to

$$\frac{1}{2}g^{\mu\nu}g^{\alpha\beta} \perp R_{\mu\alpha\nu\beta} = 8\pi\varrho. \quad (\text{A.1.89})$$

Using (A.1.39), one of the Gauss-Codazzi equations, this becomes

$$\frac{1}{2}g^{\mu\nu}g^{\alpha\beta}(\mathcal{R}_{\mu\alpha\nu\beta} + K_{\beta\nu}K_{\alpha\mu} - K_{\alpha\nu}K_{\beta\mu}) = 8\pi\varrho. \quad (\text{A.1.90})$$

The two contractions then bring this equation into the form

$$\mathcal{R} + K^2 - K^\mu_\nu K^\nu_\mu = 16\pi\varrho. \quad (\text{A.1.91})$$

where $K = K^\mu_\mu$. This final result is known as the Hamiltonian constraint. We refer to it as a constraint equations since it contains no time derivatives hence must always be satisfied.

Momentum constraint

Another contraction of the Einstein equations uses the contravariant form,

$$G^{\mu\nu} = R^{\mu\nu} - \frac{1}{2}g^{\mu\nu}R = 8\pi T^{\mu\nu}. \quad (\text{A.1.92})$$

We first contract with $-n_\mu$, giving

$$R^{\mu\hat{n}} + \frac{1}{2}n^\mu R = 8\pi T^{\mu\hat{n}}. \quad (\text{A.1.93})$$

and then we project this result onto the slice,

$$\perp R^{\mu\hat{n}} + \frac{1}{2}\perp n^\mu R = \perp R^{\mu\hat{n}} = 8\pi \perp T^{\mu\hat{n}} = 8\pi j^\mu \quad (\text{A.1.94})$$

where we have used (A.1.53b). Using (A.1.46), we see

$$\perp R_{\mu\hat{n}} = g^{\nu\alpha} \perp R_{\mu\nu\hat{n}\alpha} - \perp R_{\mu\hat{n}\hat{n}\hat{n}} = -g^{\nu\alpha} \perp R_{\mu\nu\alpha\hat{n}} \quad (\text{A.1.95})$$

which, with use of (A.1.42), one of the Gauss-Codazzi equations, then becomes

$$\perp R_{\mu\hat{n}} = -g^{\nu\alpha} (D_\nu K_{\mu\alpha} - D_\mu K_{\nu\alpha}) = D_\mu K - D^\nu K_{\mu\nu}. \quad (\text{A.1.96})$$

When indices are suitably raised (noting the minus sign introduced by the n_μ contraction), we can express (A.1.94) as

$$D_\nu K^{\mu\nu} - D^\mu K = 8\pi j^\mu. \quad (\text{A.1.97})$$

This relationship is known as the momentum constraint. As with the Hamiltonian constraint there are no time derivatives present, only spatial derivatives through D_μ .

Evolution of the spatial metric

The evolution equations for the spatial metric, $\gamma_{\mu\nu}$, follow from the definition of the extrinsic curvature, as given in (2.2.27),

$$K_{\mu\nu} = -\frac{1}{2}\mathcal{L}_n \gamma_{\mu\nu}. \quad (\text{A.1.98})$$

For full generality, we need to consider Lie-derivatives along t^μ rather than n^μ . We can use the property that for arbitrary vector fields v^μ and w^μ and tensor \mathbf{T} , we have

$$\mathcal{L}_{v+w} \mathbf{T} = \mathcal{L}_v \mathbf{T} + \mathcal{L}_w \mathbf{T}. \quad (\text{A.1.99})$$

For our general vector field t^μ , this implies

$$\mathcal{L}_t \gamma_{\mu\nu} = \alpha \mathcal{L}_n \gamma_{\mu\nu} + \mathcal{L}_\beta \gamma_{\mu\nu} \quad (\text{A.1.100})$$

which, using the definition in (A.1.98), gives the desired evolution equation

$$\mathcal{L}_t \gamma_{\mu\nu} = -2\alpha K_{\mu\nu} + \mathcal{L}_\beta \gamma_{\mu\nu}. \quad (\text{A.1.101})$$

Evolution of the extrinsic curvature tensor

To derive the evolution equations for the extrinsic curvature, we first note a result from the contraction of the Einstein equations,

$$g^{\mu\nu} G_{\mu\nu} = g^{\mu\nu} R_{\mu\nu} - \frac{1}{2} g^{\mu\nu} g_{\mu\nu} R = 8\pi g^{\mu\nu} T_{\mu\nu}, \quad (\text{A.1.102})$$

simplifying to give

$$G = R - \frac{1}{2} \delta_\mu^\mu R = -R = 8\pi T \quad (\text{A.1.103})$$

where, in 4-dimensions, we have $\delta_\mu^\mu = 4$. Using this, we can write the field equations as

$$R_{\mu\nu} = 8\pi T_{\mu\nu} + \frac{1}{2} g_{\mu\nu} R = 8\pi \left(T_{\mu\nu} - \frac{1}{2} g_{\mu\nu} T \right). \quad (\text{A.1.104})$$

A spatial projection of this then gives

$$\perp R_{\mu\nu} = 8\pi \left(\perp T_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} T \right). \quad (\text{A.1.105})$$

The 3+1 decomposition of the stress-energy tensor $T_{\mu\nu}$ can be rearranged as

$$\perp T_{\mu\nu} \equiv \mathfrak{E}_{\mu\nu} = T_{\mu\nu} + 2n_{(\mu} \perp T_{\nu)\hat{n}} - n_\mu n_\nu T_{\hat{n}\hat{n}}. \quad (\text{A.1.106})$$

in which the definition of $\mathfrak{E}_{\mu\nu}$, (A.1.53c), has been used. The contraction of this result then gives

$$\mathfrak{E}_{\mu}^{\mu} = \mathfrak{E} = T + 2n^{\mu} \perp T_{\mu\hat{n}} - n^{\mu} n_{\mu} T_{\hat{n}\hat{n}}. \quad (\text{A.1.107})$$

The middle term then vanishes since a spatial tensor is being projected with n^{μ} , and using $\varrho = T_{\hat{n}\hat{n}}$, (A.1.53a), and rearranging, we get

$$T = \mathfrak{E} - \varrho. \quad (\text{A.1.108})$$

We can therefore rewrite (A.1.105) as

$$\perp R_{\mu\nu} = 8\pi \left[\mathfrak{E}_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} (\mathfrak{E} - \varrho) \right]. \quad (\text{A.1.109})$$

The spatial projection of the Ricci scalar, $\perp R_{\mu\nu}$, can also be written in the form seen in (A.1.46), where the results for contractions of the Riemann tensor, (A.1.40) and (A.1.85), are used

$$\begin{aligned} \perp R_{\mu\nu} = & -(\alpha^{-1} \mathcal{L}_N K_{\mu\nu} + K_{\mu\alpha} K_{\nu}^{\alpha} + \alpha^{-1} D_{\mu} D_{\nu} \alpha) \\ & + g^{\alpha\beta} (\mathcal{R}_{\mu\nu\alpha\beta} + K_{\mu\nu} K_{\alpha\beta} - K_{\mu\alpha} K_{\beta\nu}). \end{aligned} \quad (\text{A.1.110})$$

Applying the contraction with $g^{\alpha\beta}$ and simplifying gives

$$\perp R_{\mu\nu} = -\alpha^{-1} \mathcal{L}_N K_{\mu\nu} - 2K_{\mu\alpha} K_{\nu}^{\alpha} - \alpha^{-1} D_{\mu} D_{\nu} \alpha + \mathcal{R}_{\mu\nu} + K K_{\mu\nu}. \quad (\text{A.1.111})$$

Since $N^{\mu} = t^{\mu} - \beta^{\mu}$, using the property of Lie-derivatives given in (A.1.99) and equating (A.1.109) and (A.1.111), gives, with some rearranging, an evolution equation for $K_{\mu\nu}$,

$$\begin{aligned} \mathcal{L}_t K_{\mu\nu} = & \mathcal{L}_{\beta} K_{\mu\nu} - D_{\mu} D_{\nu} \alpha + \alpha \left\{ \mathcal{R}_{\mu\nu} + K K_{\mu\nu} \right. \\ & \left. - 2K_{\mu\alpha} K_{\nu}^{\alpha} - 8\pi \left[\mathfrak{E}_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} (\mathfrak{E} - \varrho) \right] \right\}. \end{aligned} \quad (\text{A.1.112})$$

There is another useful form of the evolution equation for the extrinsic curvature which uses the ‘mixed’ form, K_{ν}^{μ} . This is derived from the relationship

given above, noting that $\mathcal{L}_N K_{\mu\nu} = \mathcal{L}_t K_{\mu\nu} - \mathcal{L}_\beta K_{\mu\nu}$, which can be written

$$\begin{aligned} -\alpha^{-1} \mathcal{L}_N K_{\mu\nu} - 2K_{\mu\alpha} K_{\nu}^{\alpha} - \alpha^{-1} D_{\mu} D_{\nu} \alpha \\ + \mathcal{R}_{\mu\nu} + K K_{\mu\nu} = 8\pi \left[\mathfrak{E}_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} (\mathfrak{E} - \varrho) \right]. \end{aligned} \quad (\text{A.1.113})$$

Since all terms in this equation are spatial, we can use $\gamma^{\mu\nu}$ to raise the indices giving

$$\begin{aligned} -\alpha^{-1} \gamma^{\mu\alpha} \mathcal{L}_N K_{\alpha\nu} - 2K^{\mu\alpha} K_{\alpha\nu} - \alpha^{-1} D^{\mu} D_{\nu} \alpha \\ + \mathcal{R}_{\nu}^{\mu} + K K_{\nu}^{\mu} = 8\pi \left[\mathfrak{E}_{\nu}^{\mu} - \frac{1}{2} \perp_{\nu}^{\mu} (\mathfrak{E} - \varrho) \right]. \end{aligned} \quad (\text{A.1.114})$$

We can then rewrite the Lie-derivative term as

$$\begin{aligned} \mathcal{L}_N K_{\nu}^{\mu} &= \mathcal{L}_N (\gamma^{\mu\alpha} K_{\nu\alpha}) \\ &= K_{\alpha\nu} \mathcal{L}_N \gamma^{\mu\alpha} + \gamma^{\mu\alpha} \mathcal{L}_N K_{\alpha\nu} \\ &= \alpha K_{\alpha\nu} \mathcal{L}_n \gamma^{\mu\alpha} + \gamma^{\mu\alpha} \mathcal{L}_N K_{\alpha\nu} \\ &= -2\alpha K_{\alpha\nu} K^{\mu\alpha} + \gamma^{\mu\alpha} \mathcal{L}_N K_{\alpha\nu} \end{aligned} \quad (\text{A.1.115})$$

where the definition of the extrinsic curvature, (A.1.98) has been used to obtain the final result. This result, once $N^{\mu} = t^{\mu} - \beta^{\mu}$, has been used, gives the ‘mixed’ evolution equation for the extrinsic curvature,

$$\mathcal{L}_t K_{\nu}^{\mu} = \mathcal{L}_\beta K_{\nu}^{\mu} - D^{\mu} D_{\nu} \alpha + \alpha \left\{ \mathcal{R}_{\nu}^{\mu} + K K_{\nu}^{\mu} + 8\pi \left[\frac{1}{2} \perp_{\nu}^{\mu} (\mathfrak{E} - \varrho) - \mathfrak{E}_{\nu}^{\mu} \right] \right\}. \quad (\text{A.1.116})$$

A.1.5 The Lorentz factor

As a final consideration we consider how some useful physical parameters are defined based on the 3+1 split of the Einstein equations. We can define a scalar parameter, W , as

$$W = -u^{\mu} n_{\mu} \quad (\text{A.1.117})$$

which, using (2.2.33), simplifies to

$$W = \alpha u^t \quad (\text{A.1.118})$$

In order to show the physical relevance of this factor, we use the 3-velocity, (A.1.119),

$$v^i = \frac{u^i}{\alpha u^t} + \frac{\beta^i}{\alpha}. \quad (\text{A.1.119})$$

A covariant form of the 3-velocity, (2.2.40) is then obtained by lowering the indices of u^i ,

$$u_i = g_{i\mu} u^\mu = g_{it} u^t + \gamma_{ij} u^j \quad (\text{A.1.120})$$

Using (2.2.31) we see $g_{it} = \beta_i$, and rearranging (2.2.39) to replace u^j gives

$$u_i = \beta_i u^t + \gamma_{ij} u^t (\alpha v^j - \beta^j) = W v_i \quad (\text{A.1.121})$$

where since v^i and β^i are spatial tensors, they can be contracted with γ_{ij} , and (A.1.118) has been used. This then yields (2.2.40),

$$v_i = \frac{u_i}{W}. \quad (\text{A.1.122})$$

The physical meaning of the scalar W arises from the fact that $u_\mu u^\mu = -1$. We first need to rewrite u_t ,

$$u_t = g_{t\mu} u^\mu = g_{tt} u^t + g_{ti} u^i \quad (\text{A.1.123})$$

which, using (2.2.31), (A.1.118) and (2.2.39) can be written

$$u_t = \frac{W}{\alpha} (-\alpha^2 + \beta_i \beta^i) + \beta_i \frac{W}{\alpha} (\alpha v^i - \beta^i) = \alpha W (-1 + \beta_i v^i). \quad (\text{A.1.124})$$

We then have

$$u_\mu u^\mu = u_t u^t + u_i u^i = \alpha W (-1 + \beta_i v^i) \frac{W}{\alpha} + \frac{W^2 v_i}{\alpha} (\alpha v^i - \beta^i) = -1. \quad (\text{A.1.125})$$

Simplifying gives

$$W^2 (-1 + v_i v^i) = -1 \quad (\text{A.1.126})$$

which, when rearranged becomes

$$W = \frac{1}{\sqrt{1 - v_i v^i}}. \quad (\text{A.1.127})$$

We can clearly see that this is the Lorentz factor as measured by Eulerian

observers.

A.2 Derivation of the spacetime and fluid evolution equations in spherical symmetry

Here we derive the fluid evolution equations from the conservation equations in polar-areal coordinates described in section 2.2. We also use the 3+1 split of the Einstein equations, shown in section 2.2.1, to obtain the constraint equations and describe the evolution of the spacetime.

A.2.1 Spacetime quantities

In this section, before undertaking any derivations, it is useful to consider the spacetime quantities in polar-areal form. We first note the non-vanishing Christoffel symbols in polar-areal coordinates,

$$\begin{aligned} \Gamma_{rr}^t &= a\partial_t a/\alpha^2, & \Gamma_{tt}^t &= \partial_t \alpha/\alpha, & \Gamma_{tr}^t &= \partial_r \alpha/\alpha, \\ \Gamma_{\theta r}^\theta &= \Gamma_{\phi r}^\phi = 1/r, & \Gamma_{\phi\phi}^\theta &= -\sin \theta \cos \theta, & \Gamma_{\phi\theta}^\phi &= \cot \theta, \\ \Gamma_{\phi\phi}^r &= -r \sin^2 \theta/a^2, & \Gamma_{\theta\theta}^r &= -r/a^2, & \Gamma_{tt}^r &= \alpha \partial_r \alpha/a^2, \\ \Gamma_{rr}^r &= \partial_r a/a, & \Gamma_{rt}^r &= \partial_t a/a. \end{aligned} \quad (\text{A.2.1})$$

The single non-vanishing term of the extrinsic curvature K_{rr} can then be calculated, starting with (A.1.31),

$$\begin{aligned} K_{\mu\nu} &= -\frac{1}{2} \mathcal{L}_n g_{\mu\nu} \\ &= -\frac{1}{2} (n^\alpha \nabla_\alpha g_{\mu\nu} + g_{\alpha\nu} \nabla_\mu n^\alpha + g_{\mu\alpha} \nabla_\nu n^\alpha) \\ &= -\frac{1}{2} (\nabla_\mu n_\nu + \nabla_\nu n_\mu) \\ &= -\frac{1}{2} (\partial_\mu n_\nu - \Gamma_{\mu\nu}^\alpha n_\alpha + \partial_\nu n_\mu - \Gamma_{\nu\mu}^\alpha n_\alpha) \\ &= -\frac{1}{2} (\partial_\mu n_\nu + \partial_\nu n_\mu - 2\Gamma_{\mu\nu}^\alpha n_\alpha). \end{aligned} \quad (\text{A.2.2})$$

Using $n_i = 0$, we then get

$$K_{rr} = \Gamma_{rr}^t n_t = -\frac{a\partial_t a}{\alpha}. \quad (\text{A.2.3})$$

From this, it follows that

$$K = K^i_i = g^{ij} K_{ij} = -\frac{\partial_t a}{\alpha a}. \quad (\text{A.2.4})$$

It is useful to note the non-zero components of \mathcal{R}^i_j ,

$$\mathcal{R}^r_r = 2\frac{\partial_r a}{ra^3} \quad (\text{A.2.5})$$

and

$$\mathcal{R}_\theta^\theta = \mathcal{R}_\phi^\phi = \frac{1}{r^2 a^3} (r\partial_r a + a^3 - a), \quad (\text{A.2.6})$$

and from these, we obtain

$$\mathcal{R} = \frac{2}{r^2 a^3} (2r\partial_r a + a^3 - a). \quad (\text{A.2.7})$$

A.2.2 Hamiltonian constraint

We can now consider the polar-areal form of the results from making the 3+1 decomposition of the Einstein equations in section 2.2. We begin with the Hamiltonian constraint

$$\mathcal{R} + K^2 - K^\mu_\nu K^\nu_\mu = 16\pi\varrho. \quad (\text{A.2.8})$$

In polar-areal coordinates, since the only non-vanishing component of the extrinsic curvature is K_{rr} , we have $K^\mu_\nu K^\nu_\mu = K^r_r K^r_r = K^2$, and using (A.2.7) gives the Hamiltonian constraint as

$$\frac{2}{r^2 a^3} (2r\partial_r a + a^3 - a) = 16\pi\varrho. \quad (\text{A.2.9})$$

This can then be rearranged to give

$$\frac{2r\partial_r a}{a^3} + 1 - \frac{1}{a^2} = 8\pi r^2 \varrho. \quad (\text{A.2.10})$$

By using the mass aspect function (2.2.75),

$$m(t, r) \equiv \frac{r}{2} \left(1 - \frac{1}{a(t, r)^2} \right), \quad (\text{A.2.11})$$

we can then arrange this to give the Hamiltonian constraint as

$$\frac{\partial_r a}{a} = a^2 \left(4\pi r \varrho - \frac{m}{r^2} \right). \quad (\text{A.2.12})$$

This is then a constraint equation for the function $a(r, t)$. We can write this in terms of the conserved variables using the definition of ϱ in polar-areal coordinates,

$$\varrho = T_{\hat{n}\hat{n}} = T_{\mu\nu} n^\mu n^\nu = T_{tt} n^t n^t = E. \quad (\text{A.2.13})$$

The polar-areal form of the Hamiltonian constraint can then be written

$$\frac{\partial_r a}{a} = a^2 \left(4\pi r (\tau + D) - \frac{m}{r^2} \right). \quad (\text{A.2.14})$$

A.2.3 Momentum constraint

Here we consider the polar-areal form of the momentum constraint, which uses (2.2.35), but the free index is lowered with $g_{\mu\nu}$,

$$D_\nu K^\nu_\mu - D_\mu K = 8\pi j_\mu. \quad (\text{A.2.15})$$

Using the definition of the spatial derivative, $D_\mu = \perp \nabla_\mu$, and expanding \perp for the free index, this can then be written

$$(\delta_\mu^\alpha + n_\mu n^\alpha) \nabla_\nu K^\nu_\mu - (\delta_\mu^\alpha + n_\mu n^\alpha) \nabla_\alpha K = 8\pi j_\mu. \quad (\text{A.2.16})$$

Expanding this and noting that the covariant derivative of a scalar (K) is equivalent to a partial derivative gives,

$$\nabla_\nu K^\nu_\mu + n_\mu n^\alpha \nabla_\nu K^\nu_\alpha - \partial_\mu K + n_\mu n^\alpha \partial_\alpha K = 8\pi j_\mu. \quad (\text{A.2.17})$$

Since (A.1.53b) defines $j^\mu = \perp T^{\mu\hat{n}}$, which is a spatial tensor, the above is non-zero when μ is a spatial index. All spatial components of n^μ and n_μ are zero in polar areal coordinates therefore upon expanding the covariant derivative, we get

$$\partial_\nu K^\nu_i + \Gamma_{\nu\alpha}^\nu K^\alpha_i - \Gamma_{\nu i}^\alpha K^\nu_\alpha - \partial_i K = 8\pi j_i \quad (\text{A.2.18})$$

For this to be non-vanishing we require $\mu = \nu = r$ and $i = r$ hence the equation above can be simplified to give

$$\Gamma_{\nu r}^\nu K^r_r - \Gamma_{rr}^r K^r_r = 8\pi j_r. \quad (\text{A.2.19})$$

Upon expanding the dummy index and simplifying we obtain

$$\Gamma_{\theta r}^\theta K^r_r + \Gamma_{\phi r}^\phi K^r_r = 8\pi j_r. \quad (\text{A.2.20})$$

Using (A.2.1) and (A.2.4), we can write this as

$$-\frac{2\partial_t a}{r\alpha a} = 8\pi j_r \quad (\text{A.2.21})$$

typically rearranged to give an evolution equation for a

$$\partial_t a = -4\pi r\alpha a j_r. \quad (\text{A.2.22})$$

Finally, this can be written in terms of the conserved variables using (2.2.81),

$$j_r = S_r, \quad (\text{A.2.23})$$

giving the polar-areal form of the momentum constraint,

$$\partial_t a = -4\pi r\alpha a S_r. \quad (\text{A.2.24})$$

A.2.4 Slicing condition

A constraint equation for α can be obtained from the evolution equation for the extrinsic curvature, in particular from the $K_{\theta\theta}$ component. We begin with (A.1.112),

$$\mathcal{L}_t K_{\mu\nu} = \mathcal{L}_\beta K_{\mu\nu} - D_\mu D_\nu \alpha + \alpha \left\{ \mathcal{R}_{\mu\nu} + K K_{\mu\nu} - 2 K_{\mu\alpha} K^\alpha_\nu - 8\pi \left[\mathfrak{E}_{\mu\nu} - \frac{1}{2} \gamma_{\mu\nu} (\mathfrak{E} - \varrho) \right] \right\}. \quad (\text{A.2.25})$$

For this derivation, it is useful to consider each component in this equation individually. We first note that for $\mu = \nu = \theta$, all terms involving the extrinsic curvature tensor vanish, as will any terms involving n_μ or n^μ . The spatial derivative term can be expanded to give

$$\begin{aligned} D_\mu D_\nu \alpha &= \perp \nabla_\mu (\perp \nabla_\nu \alpha) \\ &= \perp \nabla_\mu [(\delta_\nu^\alpha + n_\nu n^\alpha) \partial_\alpha \alpha] \\ &= \perp \nabla_\mu (\partial_\nu \alpha + n_\nu n^\alpha \partial_\alpha \alpha). \end{aligned} \quad (\text{A.2.26})$$

Expanding the remaining spatial projection operator gives

$$\begin{aligned} D_\mu D_\nu \alpha &= (\delta_\mu^\alpha + n_\mu n^\alpha) (\delta_\nu^\beta + n_\nu n^\beta) \nabla_\alpha \partial_\beta \alpha \\ &\quad + (\delta_\mu^\beta + n_\mu n^\beta) (\delta_\nu^\gamma + n_\nu n^\gamma) \nabla_\beta (n_\gamma n^\alpha \partial_\alpha \alpha) \end{aligned} \quad (\text{A.2.27})$$

and further manipulation yields

$$\begin{aligned} D_\mu D_\nu \alpha &= (\delta_\mu^\alpha + n_\mu n^\alpha) (\nabla_\alpha \partial_\nu \alpha + n_\nu n^\beta \nabla_\alpha \partial_\beta \alpha) \\ &\quad + (\delta_\mu^\beta + n_\mu n^\beta) [\nabla_\beta (n_\nu n^\alpha \partial_\alpha \alpha) + n_\nu n^\gamma \nabla_\beta (n_\gamma n^\alpha \partial_\alpha \alpha)]. \end{aligned} \quad (\text{A.2.28})$$

The only non-vanishing term in the above equation will be

$$\begin{aligned} \delta_\mu^\alpha \nabla_\alpha \partial_\nu \alpha &= \nabla_\mu \partial_\nu \alpha \\ &= \partial_\mu \partial_\nu \alpha - \Gamma_{\mu\nu}^\alpha \partial_\alpha \alpha \end{aligned} \quad (\text{A.2.29})$$

which, using (A.2.1) to obtain the non-vanishing Christoffel symbols, means we have

$$D_\theta D_\theta \alpha = \partial_\theta \partial_\theta \alpha - \Gamma_{\theta\theta}^r \partial_r \alpha = \frac{r}{a^2} \partial_r \alpha \quad (\text{A.2.30})$$

where we note that α is a function of r and t only.

We now consider the $\mathcal{R}_{\mu\nu}$ term, which we can write as

$$\mathcal{R}_{\mu\nu} = g_{\mu\alpha} \mathcal{R}^\alpha_\nu, \quad (\text{A.2.31})$$

the $\mu = \nu = \theta$ component of which is

$$g_{\theta\alpha} \mathcal{R}^\alpha_\theta = g_{\theta\theta} \mathcal{R}^\theta_\theta = \frac{1}{a^3} (r \partial_r a + a^3 - a). \quad (\text{A.2.32})$$

The $\mathfrak{E}_{\mu\nu}$ and $\mathfrak{E} = \mathfrak{E}^i_i$ terms use (2.2.29c), $\mathfrak{E}_{\mu\nu} = \pm T_{\mu\nu}$. This can then be expanded as

$$\mathfrak{E}_{\mu\nu} = (\delta_\mu^\alpha + n_\mu n^\alpha) (\delta_\nu^\beta + n_\nu n^\beta) T_{\alpha\beta}. \quad (\text{A.2.33})$$

In the $\mu = \nu = i$ case, we can see that the only non-vanishing term will be $\delta_i^\alpha \delta_i^\beta T_{\alpha\beta}$, hence we have

$$\mathfrak{E}_{ii} = T_{ii}. \quad (\text{A.2.34})$$

Using the results (A.2.30), (A.2.32) and $\gamma_{\theta\theta} = r^2$ in the $\mu = \nu = \theta$ component of (A.2.25), yields

$$-\frac{r}{a^2} \partial_r \alpha + \alpha \left\{ \frac{1}{a^3} (r \partial_r a + a^3 - a) - 8\pi \left[T_{\theta\theta} - \frac{r^2}{2} (T^i_i - \varrho) \right] \right\} = 0. \quad (\text{A.2.35})$$

This can then be rearranged to give the slicing condition for polar-areal coordinates as

$$\frac{\partial_r \alpha}{\alpha} = \frac{\partial_r a}{a} + \frac{1}{r} (a^2 - 1) + \frac{8\pi a^2}{r} \left[T_{\theta\theta} - \frac{r^2}{2} (T^i_i - \varrho) \right]. \quad (\text{A.2.36})$$

In order to write this in terms of the conserved variables, we first use the Hamiltonian constraint (A.2.14), to give

$$\frac{\partial_r \alpha}{\alpha} = a^2 \left(4\pi r E - \frac{m}{r^2} \right) + \frac{1}{r} (a^2 - 1) + \frac{8\pi a^2}{r} \left[T_{\theta\theta} - \frac{r^2}{2} (T^i_i - \varrho) \right]. \quad (\text{A.2.37})$$

We then need T^i_i , where, using the symmetry of the metric $g_{\mu\nu}$, we have

$$T^r_r = S_r v^r + p, \quad T^\theta_\theta = T^\phi_\phi = p \quad (\text{A.2.38})$$

hence we have

$$T^i_i = S_r v^r + 3p. \quad (\text{A.2.39})$$

Similarly, we also have

$$T_{\theta\theta} = r^2 p. \quad (\text{A.2.40})$$

Using these in (A.2.38), with (2.2.78) gives

$$\frac{\partial_r \alpha}{\alpha} = a^2 \left(4\pi r E - \frac{m}{r^2} \right) + \frac{1}{r} (a^2 - 1) + \frac{8\pi a^2}{r} \left[r^2 p - \frac{r^2}{2} (S_r v^r + 3p - E) \right]. \quad (\text{A.2.41})$$

Expanding terms and rewriting m term using the mass-aspect function definition, (2.2.75), gives

$$\frac{\partial_r \alpha}{\alpha} = 4\pi r a^2 E - \frac{a^2}{2r} \left(1 - \frac{1}{a^2} \right) + \frac{1}{r} (a^2 - 1) + 4\pi r a^2 (S_r v^r + p) - 4\pi r a^2 E. \quad (\text{A.2.42})$$

Simplifying this, we get

$$\frac{\partial_r \alpha}{\alpha} = 4\pi r a^2 (S_r v^r + p) + \frac{a^2}{r^2} \left[\frac{r}{2} \left(1 - \frac{1}{a^2} \right) \right] \quad (\text{A.2.43})$$

which can finally be written as the polar-areal constraint equation for α ,

$$\frac{\partial_r \alpha}{\alpha} = a^2 \left[4\pi r (S_r v^r + p) + \frac{m}{r^2} \right]. \quad (\text{A.2.44})$$

A.2.5 Conservation of rest mass density

Here we consider the polar-areal form of the conservation of mass, (2.2.14),

$$\nabla_\mu (\rho_0 u^\mu) = 0. \quad (\text{A.2.45})$$

Expanding the covariant derivative gives

$$\partial_\mu (\rho_0 u^\mu) + \Gamma_{\mu\nu}^\mu (\rho_0 u^\nu) = 0 \quad (\text{A.2.46})$$

which, since the only the t and r components of the 4-velocity are non-vanishing, (2.2.76), we have

$$\partial_t (\rho_0 u^t) + \partial_r (\rho_0 u^r) + \Gamma_{\mu t}^\mu (\rho_0 u^t) + \Gamma_{\mu r}^\mu (\rho_0 u^r) = 0. \quad (\text{A.2.47})$$

Writing the 4-velocity components in terms of the 3-velocity components ($u^t = W/\alpha$, $u^r = Wv^r$) and using the non-vanishing Christoffel symbols gives

$$\begin{aligned} \partial_t \left(\frac{\rho_0 W}{\alpha} \right) + \partial_r (\rho_0 W v^r) + (\Gamma_{tt}^t + \Gamma_{rt}^r) \frac{\rho_0 W}{\alpha} \\ + \left(\Gamma_{tr}^t + \Gamma_{rr}^r + \Gamma_{\theta r}^\theta + \Gamma_{\phi r}^\phi \right) \rho_0 W v^r = 0. \end{aligned} \quad (\text{A.2.48})$$

Using $D = \rho_0 W$, (2.2.63a), and the polar-areal forms of the Christoffel symbols, (A.2.1), this can be written

$$\partial_t \left(\frac{D}{\alpha} \right) + \partial_r (D v^r) + \left(\frac{\partial_t \alpha}{\alpha} + \frac{\partial_t a}{a} \right) \frac{D}{\alpha} + \left(\frac{\partial_r \alpha}{\alpha} + \frac{\partial_r a}{a} + \frac{2}{r} \right) D v^r = 0 \quad (\text{A.2.49})$$

which can then be expanded and rearranged to give

$$\begin{aligned} \frac{1}{\alpha} \partial_t D - \frac{D}{\alpha^2} \partial_t \alpha + \frac{D}{\alpha^2} \partial_t \alpha + \frac{D}{\alpha a} \partial_t a \\ + \partial_r (D v^r) + \frac{D v^r}{\alpha} \partial_r \alpha + \frac{D v^r}{a} \partial_r a + \frac{2 D v^r}{r} = 0. \end{aligned} \quad (\text{A.2.50})$$

Multiplying through by $r^2 \alpha a$ and rewriting some terms yields

$$\begin{aligned} r^2 a \partial_t D + r^2 D \partial_t a + r^2 \alpha a \partial_r (D v^r) + r^2 a D v^r \partial_r \alpha \\ + r^2 \alpha D v^r \partial_r a + \alpha a D v^r \partial_r (r^2) = 0. \end{aligned} \quad (\text{A.2.51})$$

This can finally be written as

$$\partial_t (a D) + \frac{1}{r^2} \partial_r (r^2 \alpha a D v^r) = 0. \quad (\text{A.2.52})$$

A.2.6 Conservation of total energy density

The conservation of total energy density, and hence the evolution of τ , is obtained from (2.2.15),

$$\nabla_\mu T^{\mu t} = 0. \quad (\text{A.2.53})$$

Expanding the covariant derivative gives

$$\partial_\mu T^{\mu t} + \Gamma_{\mu\nu}^\mu T^{\nu t} + \Gamma_{\mu\nu}^t T^{\mu\nu} = 0 \quad (\text{A.2.54})$$

which can be written for non-vanishing stress-energy tensor components as

$$\partial_t T^{tt} + \partial_r T^{rt} + \Gamma_{\mu t}^\mu T^{tt} + \Gamma_{\mu r}^\mu T^{rt} + \Gamma_{tt}^t T^{tt} + 2\Gamma_{rt}^t T^{rt} + \Gamma_{rr}^t T^{rr} = 0. \quad (\text{A.2.55})$$

Using (2.2.77) to write the components of the stress-energy tensor in terms of the conserved variables and considering only non-vanishing Christoffel symbols gives

$$\begin{aligned} \partial_t \left(\frac{E}{\alpha^2} \right) + \partial_r \left(\frac{S^r}{\alpha} \right) + (\Gamma_{tt}^t + \Gamma_{rt}^r) \frac{E}{\alpha^2} \\ + \left(\Gamma_{tr}^t + \Gamma_{rr}^r + \Gamma_{\theta r}^\theta + \Gamma_{\phi r}^\phi \right) \frac{S^r}{\alpha} + \Gamma_{tt}^t \frac{E}{\alpha^2} \\ + 2\Gamma_{rt}^t \frac{S^r}{\alpha} + \Gamma_{rr}^t \left(S^r v^r + \frac{p}{a^2} \right) = 0. \end{aligned} \quad (\text{A.2.56})$$

Expanding the derivative terms and giving the explicit form of the Christoffel symbols, (A.2.1), yields

$$\begin{aligned} \frac{1}{\alpha^2} \partial_t E - \frac{2E}{\alpha^3} \partial_t \alpha + \frac{1}{\alpha} \partial_r S^r - \frac{S^r}{\alpha^2} \partial_r \alpha + \frac{E}{\alpha^3} \partial_t \alpha \\ + \frac{E}{\alpha^2 a} \partial_t a + \frac{S^r}{\alpha^2} \partial_r \alpha + \frac{S^r}{\alpha a} \partial_r a + \frac{2S^r}{r\alpha} + \frac{E}{\alpha^3} \partial_t \alpha \\ + \frac{2S^r}{\alpha^2} \partial_r \alpha + \frac{a}{\alpha^2} \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a = 0. \end{aligned} \quad (\text{A.2.57})$$

We then rearrange this, and multiply through by $r^2 \alpha a$,

$$\begin{aligned} r^2 a \partial_t E + r^2 E \partial_t a + r^2 \alpha a \partial_r S^r + r^2 a S^r \partial_r \alpha + r^2 \alpha S^r \partial_r a \\ + \alpha a S^r \partial_r (r^2) = -r^2 a^2 \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a - r^2 a S^r \partial_r \alpha. \end{aligned} \quad (\text{A.2.58})$$

If we use (2.2.53), $S^i = (E + p) v^i$, to rewrite the S^r terms on the left hand side of this equation and simplify, we get

$$\partial_t (aE) + \frac{1}{r^2} \partial_r [r^2 \alpha a (E + p) v^r] = -a^2 \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a - a S^r \partial_r \alpha, \quad (\text{A.2.59})$$

a balance-law evolution equation for E . As in SR, we use the variable $\tau = E - D$, hence by subtracting (A.2.52) from the equation above we get

$$\partial_t (a\tau) + \frac{1}{r^2} \partial_r [r^2 \alpha a (\tau + p) v^r] = -a^2 \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a - a S^r \partial_r \alpha \quad (\text{A.2.60})$$

We can use the momentum constraint, (A.2.24), and the slicing condition, (A.2.44), to write the source term as

$$\begin{aligned} & -a^2 \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a - a S^r \partial_r \alpha \\ &= -a^2 \left(S^r v^r + \frac{p}{a^2} \right) (-4\pi r \alpha a S_r) \\ & \quad - \alpha a S^r \left\{ a^2 \left[4\pi r (S_r v^r + p) + \frac{m}{r^2} \right] \right\}. \end{aligned} \quad (\text{A.2.61})$$

By noting that the $4\pi r \alpha a^3 S^r S_r v^r$ terms cancel, this can be written

$$\begin{aligned} & -a^2 \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a - a S^r \partial_r \alpha \\ &= 4\pi r \alpha a S_r p - 4\pi r \alpha a^3 S^r p - \frac{\alpha a m S^r}{r^2}. \end{aligned} \quad (\text{A.2.62})$$

The right hand side can be simplified since $S_r = g_{\mu r} S^\mu = a^2 S^r$, and using this relationship to rewrite the only remaining term, we have

$$-a^2 \left(S^r v^r + \frac{p}{a^2} \right) \partial_t a - a S^r \partial_r \alpha = -\frac{\alpha a m S_r}{r^2}. \quad (\text{A.2.63})$$

Therefore the evolution equation for τ can be given by

$$\partial_t (a\tau) + \frac{1}{r^2} \partial_r [r^2 \alpha a (\tau + p) v^r] = -\frac{\alpha a m S_r}{r^2}. \quad (\text{A.2.64})$$

A.2.7 Conservation of momentum

The conservation of momentum can be expressed in polar-areal form using the spatial components of (2.2.15), $\nabla_\mu T^{\mu i}$, and noting that only the radial component is non-vanishing. Therefore this can be written

$$\nabla_\mu T^{\mu r} = \partial_\mu T^{\mu r} + \Gamma_{\mu\nu}^\mu T^{\nu r} + \Gamma_{\mu\nu}^r T^{\mu\nu} = 0. \quad (\text{A.2.65})$$

Considering only the non-vanishing Christoffel symbols and stress-energy tensor components, we can expand this to get

$$\begin{aligned} & \partial_t T^{tr} + \partial_r T^{rr} + (\Gamma_{tt}^t + \Gamma_{rt}^r) T^{tr} \\ & + (\Gamma_{tr}^t + \Gamma_{rr}^r + \Gamma_{\theta r}^{theta} + \Gamma_{\phi r}^\phi) T^{rr} + \Gamma_{tt}^r T^{tt} \\ & + 2\Gamma_{tr}^r T^{tr} + \Gamma_{rr}^r T^{rr} + \Gamma_{\theta\theta}^r T^{\theta\theta} + \Gamma_{\phi\phi}^r T^\phi = 0 \end{aligned} \quad (\text{A.2.66})$$

In order to obtain the evolution equation for S_r , we again use the result that $S_r = a^2 S^r$ to write the stress-energy components $T^{tr} = S_r / (\alpha a^2)$, $T^{rr} = (S_r v^r + p) / a^2$, hence the above equation can be written explicitly as

$$\begin{aligned} & \partial_t \left(\frac{S_r}{\alpha a^2} \right) + \partial_r \left(\frac{S_r v^r + p}{a^2} \right) + \left(\frac{\partial_t \alpha}{\alpha} + \frac{\partial_t a}{a} \right) \frac{S_r}{\alpha a^2} \\ & + \left(\frac{\partial_r \alpha}{\alpha} + \frac{\partial_r a}{a} + \frac{2}{r} \right) \frac{S_r v^r + p}{a^2} + \frac{\alpha \partial_r \alpha}{a^2} \frac{E}{\alpha^2} + \frac{2\partial_t a}{a} \frac{S_r}{\alpha a^2} \\ & + \frac{\partial_r a}{a} \frac{S_r v^r + p}{a^2} - \frac{r}{a^2} \frac{p}{r^2} - \frac{r \sin^2 \theta}{a^2} \frac{p}{r^2 \sin^2 \theta} = 0 \end{aligned} \quad (\text{A.2.67})$$

Expanding the derivative terms, and rearranging then gives

$$\begin{aligned} & \frac{1}{\alpha a^2} \partial_t S_r + \frac{S_r}{\alpha a^3} \partial_t a + \frac{1}{a^2} \partial_r (S_r v^r + p) + \frac{S_r v^r + p}{\alpha a^2} \partial_r \alpha \\ & + \frac{S_r v^r + p}{a^3} \partial_r a + \frac{S_r v^r + p}{r^2} \partial_r (r^2) + \frac{E}{\alpha a^2} \partial_r \alpha \\ & - \frac{(S_r v^r + p)}{a^3} \partial_r a - \frac{2p}{ra^2} = 0 \end{aligned} \quad (\text{A.2.68})$$

We then multiply through by $r^2\alpha a^3$ to get

$$\begin{aligned} r^2 a \partial_t S_r + r^2 S_r \partial_t a + r^2 \alpha a \partial_r (S_r v^r + p) + r^2 a (S_r v^r + p) \partial_r \alpha \\ + r^2 \alpha (S_r v^r + p) \partial_r a + \alpha a (S_r v^r + p) \partial_r (r^2) + r^2 a E \partial_r \alpha \\ - r^2 \alpha (S_r v^r + p) \partial_r a - 2r \alpha a p = 0 \end{aligned} \quad (\text{A.2.69})$$

which can then be written

$$\partial_t (a S_r) + \frac{1}{r^2} \partial_r [r^2 \alpha a (S_r v^r + p)] = \frac{2\alpha a p}{r} + \alpha (S_r v^r + p) \partial_r a - a E \partial_r \alpha. \quad (\text{A.2.70})$$

We can then use the Hamiltonian constraint, (A.2.14), and the slicing condition, (A.2.44) to rewrite the final two source terms,

$$\begin{aligned} \alpha (S_r v^r + p) \partial_r a - a E \partial_r \alpha = \alpha a^3 (S_r v^r + p) \left(4\pi r E - \frac{m}{r^2} \right) \\ - \alpha a^3 E \left[4\pi r (S_r v^r + p) + \frac{m}{r^2} \right]. \end{aligned} \quad (\text{A.2.71})$$

To simplify this we note that the $4\pi r \alpha a^3 E (S_r v^r + p)$ terms cancel, giving

$$\alpha (S_r v^r + p) \partial_r a - a E \partial_r \alpha = -\frac{\alpha a^3 m}{r^2} (S_r v^r + p + E). \quad (\text{A.2.72})$$

Using $\tau \equiv E - D$, the evolution equation for S_r can then be written

$$\partial_t (a S_r) + \frac{1}{r^2} \partial_r [r^2 \alpha a (S_r v^r + p)] = \alpha a \left[-\frac{a^2 m}{r^2} (S_r v^r + p + \tau + D) + \frac{2p}{r} \right]. \quad (\text{A.2.73})$$

A.3 Derivation of the TOV equations

In section 2.6 we introduced the TOV equations for deriving a hydrostatic equilibrium solution for a self gravitating spherically symmetric body in GR. Here we derive the TOV equations for m , Φ and p from the static case of the evolution equations in polar-areal coordinates, (2.2.97) and the respective constraint equations. The equation for m is obtained from the Hamiltonian

constraint, (2.2.83),

$$\frac{\partial_r a}{a} = a^2 \left(4\pi r E - \frac{m}{r^2} \right), \quad (\text{A.3.1})$$

where we note that for a static spacetime the partial derivative will become an ordinary derivative and we have $W = 1$, hence

$$E = \rho_0 h W^2 - p = \rho_0 h - p = \rho_0 (1 + \varepsilon) \quad (\text{A.3.2})$$

where the last step uses the definition of specific enthalpy, $h = 1 + \varepsilon + p/\rho_0$. Using the mass aspect function to express a in terms of m and r gives

$$\frac{d}{dr} \left[\left(1 - \frac{2m}{r} \right)^{-1/2} \right] = \left(1 - \frac{2m}{r} \right)^{-3/2} \left[4\pi r \rho_0 (1 + \varepsilon) - \frac{m}{r^2} \right]. \quad (\text{A.3.3})$$

Expanding the derivative term gives

$$-\frac{1}{2} \left(1 - \frac{2m}{r} \right)^{-3/2} \left(\frac{2m}{r^2} - \frac{2}{r} \frac{dm}{dr} \right) = \left(1 - \frac{2m}{r} \right)^{-3/2} \left[4\pi r \rho_0 (1 + \varepsilon) - \frac{m}{r^2} \right] \quad (\text{A.3.4})$$

which simplifies to

$$\frac{dm}{dr} = 4\pi r^2 \rho_0 (1 + \varepsilon). \quad (\text{A.3.5})$$

The equation for α (or Φ) arises naturally from the slicing condition, (2.2.87),

$$\frac{\partial_r \alpha}{\alpha} = a^2 \left[4\pi r (S_r v^r + p) + \frac{m}{r^2} \right], \quad (\text{A.3.6})$$

since

$$\frac{d\Phi}{dr} = \frac{d}{dr} \ln \alpha = \frac{1}{\alpha} \frac{d\alpha}{dr}. \quad (\text{A.3.7})$$

By considering the static case of (A.3.6), again writing a in terms of m and r , we obtain

$$\frac{d\Phi}{dr} = \left(1 - \frac{2m}{r} \right)^{-1} \left(\frac{m}{r^2} + 4\pi r p \right) \quad (\text{A.3.8})$$

which simplifies to

$$\frac{d\Phi}{dr} = \frac{m + 4\pi r^3 p}{r(r - 2m)}. \quad (\text{A.3.9})$$

Finally the equation for p comes from the evolution equation for the momentum, (2.2.97),

$$\partial_t a S_r + 3\partial_r^3 (r^2 \alpha a S_r v^r) + \partial_r \alpha a p = -\frac{\alpha a^3 m}{r^2} (S_r v^r + p + \tau + D), \quad (\text{A.3.10})$$

which, in the static case, becomes

$$\frac{d}{dr} (\alpha a p) = -\frac{\alpha a^3 m \rho_0 h}{r^2} \quad (\text{A.3.11})$$

where we have used $\tau = \rho_0 h W^2 - p - D$. This can be rearranged to give

$$\alpha a \frac{dp}{dr} + \alpha p \frac{da}{dr} + a p \frac{d\alpha}{dr} = -\frac{\alpha a^3 m \rho_0 h}{r^2} \quad (\text{A.3.12})$$

and further manipulation yields

$$\frac{dp}{dr} = -\frac{1}{\alpha} \frac{d\alpha}{dr} (\rho_0 h - E) - \frac{a^2 m (E + p)}{r^2} - \frac{1}{a} \frac{da}{dr} p \quad (\text{A.3.13})$$

where $E = \rho_0 h W^2 - p$ has been used. We can then use the Hamiltonian constraint and slicing condition to write

$$\frac{dp}{dr} = -\frac{1}{\alpha} \frac{d\alpha}{dr} \rho_0 h - \frac{a^2 m (E + p)}{r^2} + E a^2 \left(4\pi r p + \frac{m}{r^2} \right) - a^2 \left(4\pi r E - \frac{m}{r^2} \right) p \quad (\text{A.3.14})$$

and, using the definition of E , (A.3.2), and TOV equation for Φ , (A.3.7), we can simplify this, giving

$$\frac{dp}{dr} = [\rho_0 (1 + \varepsilon) + p] \frac{d\Phi}{dr}. \quad (\text{A.3.15})$$

Note all other evolution equations reduce to the trivial vanishing case for a static star.

Appendix B

Calculations for numerical techniques

B.1 PPM

Here we give a detailed description of PPM introduced in section 3.4.3. In [46] the method is given for arbitrary grid spacing. To keep our description consistent with the numerical techniques given throughout this thesis, we restrict ourselves to the fixed spacing Δx . We also apply the relativistic corrections given in [134].

Initially a central differenced approximation of the derivative across a cell, $\delta\hat{\mathbf{q}}_i$, is taken,

$$\delta\hat{\mathbf{q}}_i = \frac{\hat{\mathbf{q}}_{i+1} - \hat{\mathbf{q}}_{i-1}}{2\Delta x}. \quad (\text{B.1.1})$$

This approximation is then steepened if necessary, to a new estimate of the slope, $\delta_m\hat{\mathbf{q}}$. As with minmod and MC limiter, if the cell x_i contains a maximum or minimum of \mathbf{q} , then we set $\delta_m\hat{\mathbf{q}}_i = 0$. Otherwise, the slope may be steepened using

$$\delta_m\hat{\mathbf{q}}_i = \text{sgn}(\delta\hat{\mathbf{q}}_i) \min(|\delta\hat{\mathbf{q}}_i|, 2|\delta\hat{\mathbf{q}}_i - \delta\hat{\mathbf{q}}_{i-1}|, 2|\delta\hat{\mathbf{q}}_{i+1} - \delta\hat{\mathbf{q}}_i|). \quad (\text{B.1.2})$$

This ‘smooths’ the approximation to the derivatives near discontinuous features, in a similar manner to the method through which the slope was obtained for the MC limiter, (3.4.5). The region in which the discontinuity is captured is reduced by this, i.e. the feature has been steepened. From these steepened

derivatives, we can obtain first approximations for the reconstructed values at the cell boundaries,

$$\bar{\mathbf{q}}_{i+1/2}^- = \frac{\hat{\mathbf{q}}_i + \hat{\mathbf{q}}_{i+1}}{2} + \frac{\delta_m \hat{\mathbf{q}}_i - \delta_m \hat{\mathbf{q}}_{i+1}}{6}, \quad (\text{B.1.3})$$

$$\bar{\mathbf{q}}_{(i+1)-1/2}^+ = \bar{\mathbf{q}}_{i+1/2}^-. \quad (\text{B.1.4})$$

This is then a fourth-order polynomial reconstruction, and therefore leads to PPM having a third-order convergence as Δx is decreased (for uniform grid spacing).

The discontinuous regions can then be steepened further. We first define a parameter $\tilde{\eta}$ is used to determine how much steepening occurs. This parameter is non-zero if there is a sufficiently large discontinuity, if the density satisfies the condition

$$\delta_2 \rho_{0,i+1} \delta_2 \rho_{0,i-1} < 0, \quad |\rho_{0,i+1} - \rho_{0,i-1}| - \epsilon_S \min(|\rho_{0,i+1}|, |\rho_{0,i-1}|) > 0 \quad (\text{B.1.5})$$

where $\delta_2 \rho_{0,i}$ is an approximate second derivative of the density at x_i ,

$$\delta_2 \rho_{0,i} = \frac{\rho_{0,i+1} - 2\rho_{0,i} + \rho_{0,i-1}}{6(\Delta x)^2} \quad (\text{B.1.6})$$

and ϵ_S is a parameter that determines how large a gradient must be before it is considered a discontinuity. If this condition is satisfied then we have

$$\tilde{\eta}_i = \frac{\rho_{0,i-2} - \rho_{0,i+2} + 4\delta\rho_{0,i}}{12\delta\rho_{0,i}}. \quad (\text{B.1.7})$$

This parameter is then used to define a second parameter, η , which actually gives the extent of the steepening, and is non-zero if

$$\gamma k_0 |\delta\rho_{0,i}| \min(p_{i-1}, p_{i+1}) > |\delta p_i| \min(\rho_{0,i-1}, \rho_{0,i+1}) \quad (\text{B.1.8})$$

where k_0 is a problem dependent constant. Assuming this is satisfied, we then have

$$\eta_i = \max \left\{ 0, \min \left[1, \eta^{(1)} (\tilde{\eta}_i - \eta^{(2)}) \right] \right\} \quad (\text{B.1.9})$$

where $\eta^{(1)}$ and $\eta^{(2)}$ are two further constants and are selected such that the switch between the steepened and non-steepened reconstructions is smooth.

Typical values for these constants are $\epsilon_S = 0.1$, $k_0 = 1$, $\eta^{(1)} = 5$ and $\eta^{(2)} = 0.05$. We then use these parameters to steepen the reconstructed density profile in regions near a discontinuity, hence we redefine $\bar{\rho}_i^\pm$ through

$$\bar{\rho}_i^\pm = \bar{\rho}_i^\pm (1 - \eta_i) + \eta_i \left(\rho_{0,i\mp 1} \pm \frac{1}{2} \delta_m \rho_{0,i\mp 1} \right). \quad (\text{B.1.10})$$

By steepening the discontinuities, we have large derivatives over a small region of the discontinuity, and these can lead to oscillatory behaviour. The accuracy away from the discontinuity as a result of the steepening suggests that rather than altering these procedures, it is advantageous to introduce some dissipation across the region of the discontinuity. This is known as *zone flattening*, and applies only to cells which have a large gradient across them. By flattening each cell individually, reconstructed discontinuities take on a ‘stepped’ appearance (note this is for finding the intercell fluxes, this stepped appearance will not be present in the numerical solution). We therefore introduce a flattening parameter ζ , which provides the degree of flattening that occurs, based on a further parameter w , which determines whether we are actually within a discontinuity. Specifically, we have $w_i = 1$ if

$$\delta p_i > \epsilon \min(p_{i-1}, p_{i+1}), \quad \delta v_i < 0 \quad (\text{B.1.11})$$

for a shock moving in the positive x -direction and $w_i = 0$ otherwise, where ϵ determines the jump required for a feature to be considered a shock. We then obtain an initial value for the flattening parameter, $\tilde{\zeta}$,

$$\tilde{\zeta}_i = \max \left\{ 0, 1 - w_i \max \left[0, \omega^{(2)} \left(\frac{p_{i+1} - p_{i-1}}{p_{i+2} - p_{i-2}} - \omega^{(1)} \right) \right] \right\} \quad (\text{B.1.12})$$

where $\omega^{(1)}$ and $\omega^{(2)}$ are constants which determine the level of dissipation. Typical values for these constants are $\epsilon = 1$, $\omega^{(1)} = 0.5$, $\omega^{(2)} = 10$. Care must be taken if there is a sudden spike in the density in otherwise smooth data, i.e. $p_{i+2} = p_{i-2}$. In such cases we manually set $\tilde{\zeta} = 1$.

The flattening parameter ζ is then determined based on local values of $\tilde{\zeta}$.

We define σ_i such that

$$\sigma_i = \begin{cases} 1 & \delta p_i < 0 \\ -1 & \delta p_i > 0 \end{cases} \quad (\text{B.1.13})$$

from which we have

$$\zeta_i = \max \left(\tilde{\zeta}_i, \tilde{\zeta}_{i+\sigma_i} \right). \quad (\text{B.1.14})$$

The reconstructed variables are then flattened by

$$\bar{\mathbf{q}}_i^\pm = \zeta \bar{\mathbf{q}}_i^\pm + (1 - \zeta) \hat{\mathbf{q}}_i. \quad (\text{B.1.15})$$

In case of a strong discontinuity, we will have $\zeta = 0$, and hence in these regions the reconstruction is equivalent to the original Godunov scheme i.e. piecewise constant reconstruction.

The various stages of PPM can lead to the reconstruction not being monotonic. If this is the case then errors could enter the evolution. The final stage of PPM therefore ensures that the reconstruction is monotonic everywhere. This is ensured through

$$\begin{cases} \bar{\mathbf{q}}_i^\pm = \hat{\mathbf{q}}_i & (\bar{\mathbf{q}}_i^+ - \hat{\mathbf{q}}_i) (\hat{\mathbf{q}}_i - \bar{\mathbf{q}}_i^-) \leq 0 \\ \bar{\mathbf{q}}_i^- = 3\hat{\mathbf{q}}_i - 2\bar{\mathbf{q}}_i^+ & (\bar{\mathbf{q}}_i^+ - \bar{\mathbf{q}}_i^-) \left(\hat{\mathbf{q}}_i - \frac{\bar{\mathbf{q}}_i^+ + \bar{\mathbf{q}}_i^-}{2} \right) > \frac{(\bar{\mathbf{q}}_i^+ - \bar{\mathbf{q}}_i^-)^2}{6} \\ \bar{\mathbf{q}}_i^+ = 3\hat{\mathbf{q}}_i - 2\bar{\mathbf{q}}_i^- & (\bar{\mathbf{q}}_i^+ - \bar{\mathbf{q}}_i^-) \left(\hat{\mathbf{q}}_i - \frac{\bar{\mathbf{q}}_i^+ + \bar{\mathbf{q}}_i^-}{2} \right) < -\frac{(\bar{\mathbf{q}}_i^+ - \bar{\mathbf{q}}_i^-)^2}{6} \\ \bar{\mathbf{q}}_i^\pm = \bar{\mathbf{q}}_i^\pm & \text{otherwise} \end{cases} \quad (\text{B.1.16})$$

We then have SSP reconstructed states that can be used in the Riemann solver of choice.

Bibliography

- [1] Einstein telescope project. <https://www.et-gw.eu/et/>.
- [2] LIGO's Fifth Science Run. http://www.ligo.caltech.edu/~11_news/s5_news/s5article.htm.
- [3] Lorene code. <http://www.lorene.obspm.fr>.
- [4] VIRGO Sensitiviy curves. <http://wwwcascina.virgo.infn.it/DataAnalysis/Calibration/Sensitivity/>.
- [5] B. P. Abbott et al. Ligo: the laser interferometer gravitational-wave observatory. *Reports on Progress in Physics*, **72**(7):076901 (25pp), 2009.
- [6] R. Abgrall and S. Karni. Computations of compressible multifluids. *J. Comput. Phys.*, **169**:594–623, 2001.
- [7] A. Abramovici et al. LIGO - The Laser Interferometer Gravitational-Wave Observatory. *Science*, **256**:325–333, 1992.
- [8] C. Alcock and A. Illarionov. The surface chemistry of stars. I - Diffusion of heavy ions in white dwarf envelopes. II - Fractionated accretion of interstellar matter. *Astrophys. J.*, **235**:534–553, 1980.
- [9] M. Alcubierre. *Introduction to 3 + 1 Numerical Relativity*. Oxford University Press, 2008.
- [10] J. M. Aloy, M. H. and Ibáñez, J. M. Martí, and E. Müller. GENESIS: A High Resolution Code for Three-dimensional Relativistic Hydrodynamics. *Astrophys. J.*, **122**:151–166, 1999.

- [11] P. Amaro-Seoane et al. TOPICAL REVIEW: Intermediate and extreme mass-ratio inspirals - astrophysics, science applications and detection using LISA. *Classical and Quantum Gravity*, **24**:113, 2007.
- [12] M. Anderson et al. Simulating binary neutron stars: Dynamics and gravitational waves. *Phys. Rev. D*, **77**(2):024006, 2008.
- [13] N. Andersson. TOPICAL REVIEW: Gravitational waves from instabilities in relativistic stars. *Classical and Quantum Gravity*, **20**:105, 2003.
- [14] N. Andersson, G. L. Comer, and K. Glampedakis. How viscous is a superfluid neutron star core? *Nucl. Phys.*, **A763**:212–229, 2005.
- [15] N. Andersson et al. Gravitational waves from neutron stars: Promises and challenges. *ArXiv e-prints*, 2009.
- [16] N. Andersson and K. D. Kokkotas. Towards gravitational wave asteroseismology. *Mon. Not. R. Astron. Soc.*, **299**:1059–1068, 1998.
- [17] N. Andersson, T. Sidery, and G. L. Comer. Mutual friction in superfluid neutron stars. *Mon. Not. R. Astron. Soc.*, **368**:162–170, 2006.
- [18] S. Ando, J. F. Beacom, and H. Yüksel. Detection of Neutrinos from Supernovae in Nearby Galaxies. *Phys. Rev. Lett.*, **95**(17):171101, 2005.
- [19] R. Arnowitt, S. Deser, and C. W. Misner. *Gravitation: an Introduction to Current Research*, chapter The Dynamics of General Relativity. John Wiley and Sons, 1962.
- [20] B. Aylott et al. Testing gravitational-wave searches with numerical relativity waveforms: results from the first Numerical INjection Analysis (NINJA) project. *Classical and Quantum Gravity*, **26**(16):165008, 2009.
- [21] W. Baade and F. Zwicky. Cosmic Rays from Super-Novae. *Proc. Nat. Acad. Sci.*, **20**(5):259–263, 1934.
- [22] A. Bagabir and D. Drikakis. Mach number effects on shock-bubble interaction. *Shock Waves*, **11**:209–218, 2001.

- [23] L. Baiotti, B. Giacomazzo, and L. Rezzolla. Accurate evolutions of inspiralling neutron-star binaries: Prompt and delayed collapse to a black hole. *Phys. Rev. D*, **78**(8):084033, 2008.
- [24] L. Baiotti, B. Giacomazzo, and L. Rezzolla. Accurate evolutions of inspiralling neutron-star binaries: assessment of the truncation error. *Classical and Quantum Gravity*, **26**(11):114005, 2009.
- [25] L. Baiotti, I. Hawke, P. J. Montero, and L. Rezzolla. A new three-dimensional general-relativistic hydrodynamics code. *Memorie della Societa Astronomica Italiana Supplement*, **1**:210, 2003.
- [26] F. Banyuls, J. A. Font, J. M. Ibáñez, J. M. Martí, and J. A. Miralles. Numerical $\{3 + 1\}$ General Relativistic Hydrodynamics: A Local Characteristic Approach. *Astrophys. J.*, **462**:839–854, 1996.
- [27] L. Barack. TOPICAL REVIEW: Gravitational self-force in extreme mass-ratio inspirals. *Classical and Quantum Gravity*, **26**(21):213001, 2009.
- [28] J. M. Bardeen and T. Piran. General relativistic axisymmetric rotating systems: Coordinates and equations. *Phys. Rep.*, **96**(4):205 – 250, 1983.
- [29] G. Baym, C. Pethick, and P. Sutherland. The Ground State of Matter at High Densities: Equation of State and Stellar Models. *Astrophys. J.*, **170**:299, 1971.
- [30] O. Benhar. Neutron Star Matter Equation of State and Gravitational Wave Emission. *Modern Physics Letters A*, **20**:2335–2349, 2005.
- [31] M. J. Berger and J. Oliger. Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations. *J. Comput. Phys.*, **53**:484–512, 1984.
- [32] H. A. Bethe, G. Borner, and K. Sato. Nuclei in Neutron Matter. *Astron. Astrophys.*, **7**:279, 1970.
- [33] R. M. Bionta, G. Blewitt, C. B. Bratton, D. Caspere, and A. Ciocio. Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud. *Phys. Rev. Lett.*, **58**:1494–1496, 1987.

- [34] G. D. Birkhoff and R. E. Langer. *Relativity and modern physics*. Cambridge, Harvard University Press, 1923.
- [35] L. Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Rev. Rel.*, **9**:4, 2006.
- [36] A. Burrows, E. Livne, L. Dessart, C. D. Ott, and J. Murphy. An acoustic mechanism for core-collapse supernova explosions. *New Astronomy Reviews*, **50**(7-8):487 – 491, 2006.
- [37] J. C. Butcher. A history of Runge-Kutta methods. *Applied Numerical Mathematics*, **20**(3):247 – 260, 1996.
- [38] L. Cadonati et al. Status of NINJA: the Numerical INjection Analysis project. *Classical and Quantum Gravity*, **26**(11):114008, 2009.
- [39] C. Canuto, M. Y. Hussaini, A. Quarteroni, and T. A. Zang. *Spectral Methods in Fluid Mechanics*. Springer, Berlin, 1988.
- [40] B. Caron et al. The Virgo interferometer. *Classical and Quantum Gravity*, **14**:1461–1469, 1997.
- [41] B. Carter and H. Quintana. Stationary elastic rotational deformation of a relativistic neutron star model. *Astrophys. J.*, **202**:511–522, 1975.
- [42] J. Chadwick. Possible Existence of a Neutron. *Nature*, **129**:312, 1932.
- [43] D. Chopp. Computing Minimal Surfaces via Level Set Curvature Flow. *J. Comput. Phys.*, **106**:77–91, 1993.
- [44] M. W. Choptuik. The 3+1 Einstein Equations. *Unpublished Work*, <http://laplace.physics.ubc.ca/People/matt/Teaching/98Spring/Phy387N/Doc/3+1.ps>, (1998).
- [45] M. W. Choptuik. Consistency of finite-difference solutions of Einstein's equations. *Phys. Rev. D*, **44**(10):3124–3135, 1991.
- [46] P. Colella and P. R. Woodward. The Piecewise-Parabolic Method (PPM) for Gas-Dynamics. *J. Comput. Phys.*, **54**:174, 1984.

- [47] G. Cook. Initial Data for Numerical Relativity. *Living Rev. Rel.*, **3**:5, 2000.
- [48] R. Courant, K. O. Friedrichs, and H Lewy. Über die partiellen Differenzengleichungen der mathematischen Physik. *Math. Ann.*, **100**:32–74, 1928.
- [49] R. Courant, K. O. Friedrichs, and H Lewy. On the Partial Difference Equations of Mathematical Physics. *IBM J.*, **11**:215–234, 1967.
- [50] M. Crandall and P.-L. Lions. Two Approximations of Hamilton-Jacobi Equations. *Math. Comput.*, **43**:1–19, 1984.
- [51] S. F. Davis. Simplified Second-Order Godunov-Type Methods. *SIAM J. Sci. Stat. Comput.*, **9**:445–473, 1988.
- [52] L. Del Zanna and N. Bucciantini. An efficient shock-capturing central-type scheme for multidimensional relativistic flows. I. Hydrodynamics. *Astron. Astrophys.*, **390**:1177–1186, 2002.
- [53] P. Diener. A new general purpose event horizon finder for 3D numerical spacetimes. *Class. Quant. Grav.*, **20**:4901–4918, 2003.
- [54] H. Dimmelmeier, J. A. Font, and E. Müller. Relativistic simulations of rotational core collapse II. Collapse dynamics and gravitational radiation. *Astron. Astrophys.*, **393**:523–542, 2002.
- [55] H. Dimmelmeier, C. D. Ott, A. Marek, and H.-T. Janka. Gravitational wave burst signal from core collapse of rotating stars. *Phys. Rev. D*, **78**(6):064056, 2008.
- [56] A. Dolezal and S. S. M. Wong. Relativistic Hydrodynamics and Essentially Non-oscillatory Shock Capturing Schemes. *J. Comput. Phys.*, **120**(2):266 – 277, 1995.
- [57] R. Donat and A. Marquina. Capturing shock reflections: An improved flux formula. *J. Comput. Phys.*, **125**:42–58, 1996.
- [58] R. Donat and A. Marquina. Capturing shock reflections: An improved flux formula. *J. Comput. Phys.*, **125**(1):42 – 58, 1996.

- [59] M. D. Duez. Numerical relativity confronts compact neutron star binaries: a review and status report. *ArXiv e-prints*, 2009.
- [60] R. C. Duncan and C. Thompson. Formation of Very Strongly Magnetized Neutron Stars: Implications for Gamma-Ray Bursts. *Astrophys. J.*, **392**:9–13, 1992.
- [61] B. Einfeldt. On Godunov-Type Methods for Gas Dynamics. *SIAM J. Numer. Anal.*, **105**:92–107, 1993.
- [62] A. Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin)*, pages 844–847, 1915.
- [63] D. Enright, R. Fedkiw, J. Ferziger, and I. Mitchell. A hybrid particle level set method for improved interface capturing. *J. Comput. Phys.*, **183**(1):83 – 116, 2002.
- [64] Z. B. Etienne et al. Fully general relativistic simulations of black hole–neutron star mergers. *Phys. Rev. D*, **77**(8):084002, 2008.
- [65] J. Faber. Status of neutron star-black hole and binary neutron star simulations. *Classical and Quantum Gravity*, **26**(11):114004, 2009.
- [66] R. P. Fedkiw. Coupling an Eulerian Fluid Calculation to a Lagrangian Solid Calculation with the Ghost Fluid Method. *J. Comput. Phys.*, **175**:200–224, 2002.
- [67] R. P. Fedkiw, T. Aslam, B. Merriman, and S. A. Osher. A Non-oscillatory Eulerian Approach to Interfaces in Multimaterial Flows (the Ghost Fluid Method). *J. Comput. Phys.*, **152**:457–492, 1999.
- [68] R. P. Fedkiw, A. Marquina, and B. Merriman. An Isobaric Fix for the Overheating Problem in Multimaterial Compressible Flows. *J. Comput. Phys.*, **148**:545–578, 1999.
- [69] J. A. Font. Numerical Hydrodynamics and Magnetohydrodynamics in General Relativity. *Living Rev. Rel.*, **11**:7, 2008.

- [70] J. A. Font, M. Miller, W.-M. Suen, and M. Tobias. Three-dimensional numerical general relativistic hydrodynamics: Formulations, methods, and code tests. *Phys. Rev. D*, **61**(4):044011, 2000.
- [71] H. Friedrich and G. Nagy. The Initial Boundary Value Problem for Einstein's Vacuum Field Equation. *Commun. Math. Phys.*, **201**:619–655, 1999.
- [72] B. Giacomazzo, L. Rezzolla, and L. Baiotti. Can magnetic fields be detected during the inspiral of binary neutron stars? *Mon. Not. R. Astron. Soc.*, **399**:L164–L168, 2009.
- [73] R. A. Gingold and J. J. Monaghan. Smoothed particle hydrodynamics - Theory and application to non-spherical stars. *Mon. Not. R. Astron. Soc.*, **181**:375–389, 1977.
- [74] K. Glampedakis, N. Andersson, and D. I. Jones. Stability of Precessing Superfluid Neutron Stars. *Phys. Rev. Lett.*, **100**(8):081101, 2008.
- [75] K. Glampedakis, N. Andersson, and D. I. Jones. Do superfluid instabilities prevent neutron star precession? *Mon. Not. R. Astron. Soc.*, **394**:1908–1924, 2009.
- [76] S. Godunov. Finite Difference Method for Numerical Computation of Discontinuous Solutions of the Equations of Fluid Dynamics. *Mat. Sbornik*, **47**:3–271, 1959.
- [77] T. Gold. Rotating Neutron Stars as the Origin of the Pulsating Radio Sources. *Nature*, **221**:25, 1968.
- [78] P. Goldoni, C. Musso, P. A. Caraveo, and G. F. Bignami. Multiwave-length phenomenology of isolated neutron stars. *Astron. Astrophys.*, **298**:535, 1995.
- [79] J. A. González, U. Sperhake, and B. Brügmann. Black-hole binary simulations: The mass ratio 10:1. *Phys. Rev. D*, **79**(12):124006, 2009.
- [80] D. Gottlieb and S. A. Orszag. *Numerical analysis of spectral methods: theory and applications*. CBMS-NSF Regional Conference Series in Applied Math. SIAM, 1977.

- [81] S. Gottlieb, D. I. Ketcheson, and C.-W. Shu. High order strong stability preserving time discretizations. *J. Sci. Comput.*, **38**(3):251–289, 2009.
- [82] S. Gottlieb and C. W. Shu. Total variation diminishing Runge-Kutta schemes. *Math. Comp.*, **67**:73–85, 1998.
- [83] E. Gourgoulhon. Construction of initial data for 3+1 numerical relativity. *Journal of Physics Conference Series*, **91**(1):012001, 2007.
- [84] E. H. Gudmundsson, C. J. Pethick, and R. I. Epstein. Neutron star envelopes. *Astrophys. J. Lett.*, **259**:L19–L23, 1982.
- [85] J.-F. Haas and B. Sturtevant. Interaction of weak shock waves with cylindrical and spherical gas inhomogeneities. *J. Fluid Mech.*, **181**:41–76, 1987.
- [86] P. Haensel. Equation of State of Dense Matter and Maximum Mass of Neutron Stars. In C. Motch & J.-M. Hameury, editor, *EAS Publications Series*, volume **7** of *EAS Publications Series*, page 249, 2003.
- [87] P. Haensel, A. Y. Potekhin, and D. G. Yakovlev. *Neutron Stars 1. Equation of State and Structure*. Springer, 2007.
- [88] M. Hannam. Status of black-hole-binary simulations for gravitational-wave detection. *Classical and Quantum Gravity*, **26**(11):114001, 2009.
- [89] A. Harten. The Artificial Compression Method for Computation of Shocks and Contact Discontinuities: III. Self-Adjusting Hybrid Schemes. *Math. Comp.*, **32**(142):363–389, 1978.
- [90] A. Harten, B. Engquist, S. Osher, and S. Chakravarthy. Uniformly High-Order Accurate Essentially Non-Oscillatory Schemes III. *J. Comput. Phys.*, **71**:231–303, 1987.
- [91] A. Harten, P. D. Lax, and B. Van Leer. On upstream differencing and godunov-type schemes for hyperbolic conservation laws. *SIAM Rev.*, **25**(1):35–61, 1983.
- [92] A. Harten and S. Osher. Uniformly High-Order Accurate Nonoscillatory Schemes I. *SIAM J. Numer. Anal.*, **24**:279, 1987.

- [93] B. Haskell, D. I. Jones, and N. Andersson. Mountains on neutron stars: accreted versus non-accreted crusts. *Mon. Not. R. Astron. Soc.*, **373**:1423–1439, 2006.
- [94] I. Hawke. *Computational Ultrarelativistic Hydrodynamics*. PhD thesis, University of Cambridge, 2001.
- [95] C. W. Helstrom. *Statistical theory of signal detection*. International series of monographs in electronics and instrumentation, v. 9. Pergamon Press, Oxford, New York, 1968.
- [96] A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a Rapidly Pulsating Radio Source. *Nature*, **217**:709, 1968.
- [97] C. J. Horowitz and K. Kadau. Breaking Strain of Neutron Star Crust and Gravitational Waves. *Phys. Rev. Lett.*, **102**(19):191102, 2009.
- [98] T. Y. Hou and P. G. Le Floch. Why nonconservative schemes converge to wrong solutions: Error analysis. *Math. Comp.*, **62**(206):497–530, 1994.
- [99] X. Y. Hu and B. C. Khoo. An interface interaction method for compressible multifluids. *J. Comput. Phys.*, **198**:35–64, 2004.
- [100] J. Isern, E. García-Berro, J. Guerrero, P. Lorén-Aguilar, and J. A. Lobo. White dwarf merging and the emission of gravitational waves. In L. Burderi, L. A. Antonelli, F. D’Antona, T. di Salvo, G. L. Israel, L. Piersanti, A. Tornambè, & O. Straniero, editor, *Interacting Binaries: Accretion, Evolution, and Outcomes*, volume **797** of *American Institute of Physics Conference Series*, pages 463–470, 2005.
- [101] W. Israel and J. M. Stewart. Transient relativistic thermodynamics and kinetic theory. *Annals of Physics*, **118**(2):341–372, 1992.
- [102] Boettinger W. J., J. A. Warren, C. Beckermann, and A. Karma. Phase-field simulation of solidification. *Ann. Rev. Mater. Res.*, **32**:163 – 194, 2002.
- [103] H.-T. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller. Theory of core-collapse supernovae. *Phys. Rep.*, **442**:38–74, 2007.

- [104] G.-S. Jiang and C.-W. Shu. Efficient Implementation of Weighted ENO Schemes. *J. Comput. Phys.*, **126**:202–228, 1996.
- [105] D. I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. R. Astron. Soc.*, **331**:203–220, 2002.
- [106] M. Karlovini and L. Samuelsson. Elastic stars in general relativity: I. Foundations and equilibrium models. *Classical and Quantum Gravity*, **20**:3613–3648, 2003.
- [107] W. J. Kaufmann. *Universe*. W. H. Freeman, New York, 1994.
- [108] M. M. Komesaroff. Possible Mechanism for the Pulsar Radio Emission. *Nature*, **225**:612–614, 1970.
- [109] A. Konigl. Relativistic gasdynamics in two dimensions. *Physics of Fluids*, **23**:1083–1090, 1980.
- [110] R. K. Kopperapu et al. Host Galaxies Catalog Used in LIGO Searches for Compact Binary Coalescence Events. *Astrophys. J.*, **675**:1459–1467, 2008.
- [111] K. Kotake, K. Sato, and K. Takahashi. Explosion mechanism, neutrino burst and gravitational wave in core-collapse supernovae. *Reports on Progress in Physics*, **69**:971–1143, 2006.
- [112] A. Krawczyk, A. G. Lyne, J. A. Gil, and B. C. Joshi. Observations of 14 pulsar glitches. *Mon. Not. R. Astron. Soc.*, **340**:1087–1094, 2003.
- [113] A. G. Kuranov and K. A. Postnov. Neutron stars in globular clusters: Formation and observational manifestations. *Astronomy Letters*, **32**(6):393–405, 2006.
- [114] L. D. Landau and E. M. Lifshitz. Course of Theoretical Physics. In *Fluid Mechanics*, volume 6. Butterworth-Heinemann Ltd., 1987.
- [115] S. K. Lander, D. I. Jones, and A. Passamonti. Oscillations of rotating magnetized neutron stars with purely toroidal magnetic fields. *Mon. Not. R. Astron. Soc.*, page 363, 2010.

- [116] C. B. Laney. *Computational Gasdynamics*. Cambridge University Press, 1998.
- [117] M. I. Large, A. E. Vaughan, and B. Y. Mills. A Pulsar Supernova Association? *Nature*, **220**:340–341, 1968.
- [118] J. M. Lattimer and F. Douglas Swesty. A generalized equation of state for hot, dense matter. *Nuclear Physics A*, **535**:331–376, 1991.
- [119] J. M. Lattimer and M. Prakash. The physics of neutron stars. *Science*, **304**:536–542, 2004.
- [120] J. M. Lattimer and M. Prakash. Neutron star observations: Prognosis for equation of state constraints. *Phys. Rep.*, **442**:109–165, 2007.
- [121] P. D. Lax and B. Wendroff. Systems of conservation laws. *Commun. Pure Appl. Math.*, **13**:217–237, 1960.
- [122] B. Van Leer. Towards the ultimate conservative difference scheme. III. Upstream centered finite difference schemes for ideal compressible flows. *J. Comput. Phys.*, **23**:263–275, 1977.
- [123] R. J. Leveque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, 2002.
- [124] B. Link and R. I. Epstein. Thermally Driven Neutron Star Glitches. *Astrophys. J.*, **457**:844, 1996.
- [125] T. G. Liu, B. C. Khoo, and C. W. Wang. The ghost fluid method for compressible gas-water simulation. *J. Comput. Phys.*, **204**:193–221, 2005.
- [126] T. G. Liu, B. C. Khoo, and K. S. Yeo. Ghost fluid method for strong shock impacting on material interface. *J. Comput. Phys.*, **190**:651–681, 2003.
- [127] X.-D. Liu, S. Osher, and T. Chan. Weighted essentially non-oscillatory schemes. *J. Comput. Phys.*, **115**(1):200 – 212, 1994.
- [128] X.-D. Liu, S. Osher, and T. Chan. Weighted Essentially Non-Oscillatory Schemes. *J. Comput. Phys.*, **115**:200–212, 1994.

- [129] B. Lombard and R. Donat. The Explicit Simplified Method for Compressible Multicomponent Flows. *SIAM J. Sci. Comput.*, **27**:208–230, 2005.
- [130] H. Lück and the GEO600 Team. The GEO600 project. *Classical and Quantum Gravity*, **14**:1471–1476, 1997.
- [131] L. B. Lucy. A numerical approach to the testing of the fission hypothesis. *Astrophys. J.*, **82**:1013–1024, 1977.
- [132] A. G. Lyne, S. L. Shemar, and F. G. Smith. Statistical studies of pulsar glitches. *Mon. Not. R. Astron. Soc.*, **315**:534–542, 2000.
- [133] J. M. Martí and E. Müller. The analytical solution of the Riemann problem in relativistic hydrodynamics. *J. Fluid Mech.*, **258**:317–333, 1994.
- [134] J. M. Martí and E. Müller. Extension of the piecewise parabolic method to one-dimensional relativistic hydrodynamics. *J. Comput. Phys.*, **123**(1):1 – 14, 1996.
- [135] S. T. Millmore and I. Hawke. Numerical simulations of interfaces in relativistic hydrodynamics. *Classical and Quantum Gravity*, **27**(1):015007, 2010.
- [136] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. W. H. Freeman and Company, 1973.
- [137] J. J. Monaghan. Smoothed particle hydrodynamics. *ARA&A*, **30**:543–574, 1992.
- [138] W. Mulder, S. Osher, and J. Sethian. Computing Interface Motion in Compressible Gas Dynamics. *J. Comput. Phys.*, **100**:209–228, 1992.
- [139] E. Nakar. Short-hard gamma-ray bursts. *Phys. Rep.*, **442**(1-6):166 – 236, 2007. The Hans Bethe Centennial Volume 1906-2006.
- [140] D. W. Neilsen. *Extremely Relativistic Fluids in Strong-Field Gravity*. PhD thesis, University of Texas at Austin, 1999.

- [141] B. D. Nichols, C. W. Hirt, and R. S. Hotchkiss. SOLA-VOF: A solution algorithm for transient fluid flow with multiple free boundaries. *NASA STI/Recon Technical Report N*, **81**:14281, 1980.
- [142] S. C. Noble. *A Numerical Study of Relativistic Fluid Collapse*. PhD thesis, University of Texas at Austin, 2003.
- [143] R. R. Nourgaliev, T. N. Dinh, and T. G. Theofanous. Adaptive characteristics-based matching for compressible multifluid dynamics. *J. Comput. Phys.*, **213**:500–529, 2006.
- [144] M. Obergaulinger, P. Cerdá-Durán, E. Müller, and M. A. Aloy. Semi-global simulations of the magneto-rotational instability in core collapse supernovae. *Astron. Astrophys.*, **498**:241–271, 2009.
- [145] R. Oechslin, H.-T. Janka, and A. Marek. Relativistic neutron star merger simulations with non-zero temperature equations of state. I. Variation of binary parameters and equation of state. *Astron. Astrophys.*, **467**:395–409, 2007.
- [146] A. Ohnishi, D. Jido, T. Sekihara, and K. Tsubakihara. Possibility of an s-wave pion condensate in neutron stars reexamined. *Phys. Rev. C*, **80**(3):038202, 2009.
- [147] J. R. Oppenheimer and G. M. Volkoff. On Massive Neutron Cores. *Phys. Rev.*, **55**(4):374–381, 1939.
- [148] S. Osher and R. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*, volume 153 of *Applied Mathematical Sciences*. Springer-Verlag, 2003.
- [149] S. Osher and J. Sethian. Fronts Propagating with Curvature Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulations. *J. Comput. Phys.*, **79**:12–49, 1988.
- [150] C. D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Classical and Quantum Gravity*, **26**(20):204015, 2009.

- [151] C. D. Ott. TOPICAL REVIEW: The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity*, **26**(6):063001, 2009.
- [152] B. Paczynski. GB 790305 as a very strongly magnetized neutron star. *Acta Astronomica*, **42**:145–153, 1992.
- [153] D. Page and S. Reddy. Dense Matter in Compact Stars: Theoretical Developments and Observational Constraints. *Annual Review of Nuclear and Particle Science*, **56**:327–374, 2006.
- [154] A. Passamonti, B. Haskell, N. Andersson, D. I. Jones, and I. Hawke. Oscillations of rapidly rotating stratified neutron stars. *Mon. Not. R. Astron. Soc.*, **394**:730–741, 2009.
- [155] J. Peitz and S. Appl. Dissipative fluid dynamics in the 3 + 1 formalism. *Classical and Quantum Gravity*, **16**:979–989, 1999.
- [156] C. J. Pethick and D. G. Ravenhall. The Physics of Neutron Star Crusts. In Z. Arzoumanian, F. Van der Hooft, & E. P. J. van den Heuvel, editor, *Pulsar Timing, General Relativity and the Internal Structure of Neutron Stars*, page 177, 1999.
- [157] J. Piraux and B. Lombard. A new interface method for hyperbolic problems with discontinuous coefficients. one-dimensional acoustic example. *J. Comput. Phys.*, **168**:227–248, 2001.
- [158] J. A. Pons, J. A. Font, J. M. Ibáñez, J. M. Martí, and J. A. Miralles. General relativistic hydrodynamics with special relativistic Riemann solvers. *Astron. Astrophys.*, **339**:638–642, 1998.
- [159] J. A. Pons, J. Ma Martí, and E. Müller. The exact solution of the Riemann problem with non-zero tangential velocities in relativistic hydrodynamics. *J. Fluid Mech.*, **422**:125–139, 2000.
- [160] F. Pretorius. Evolution of Binary Black-Hole Spacetimes. *Phys. Rev. Lett.*, **95**(12):121101, 2005.
- [161] D. J. Price and S. Rosswog. Producing Ultrastrong Magnetic Fields in Neutron Star Mergers. *Science*, **312**:719–722, 2006.

- [162] J. J. Quirk and S. Karni. On the dynamics of a shock-bubble interaction. *J. Fluid Mech.*, **318**:129–163, 1996.
- [163] K. Rajagopal and R. Sharma. Crystallography of three-flavor quark matter. *Phys. Rev. D*, **74**(9):094019, 2006.
- [164] E. C. Reifenstein, W. D. Brundage, and D. H. Staelin. Crab Nebula Pulsar NPO527. *Phys. Rev. Lett.*, **22**:311, 1969.
- [165] E. C. Reifenstein and D. H. Staelin. Pulsating Radio Sources near the Crab Nebula. *Science*, **162**:1481–1483, 1968.
- [166] L. Rezzolla and O. Zanotti. An Improved Exact Riemann Solver for Relativistic Hydrodynamics. In *The Second National Conference on Astrophysics of Compact Objects*, page 54, 2001.
- [167] W. J. Rider and D. B. Kothe. Reconstructing volume tracking. *J. Comput. Phys.*, **141**(2):112 – 152, 1998.
- [168] P. L. Roe. Approximate Riemann solvers, parameter vectors, and finite difference schemes. *J. Comput. Phys.*, **43**:357–372, 1981.
- [169] P. L. Roe. Characteristic-based Schemes for the Euler Equations. *Ann. Rev. Fluid. Mech.*, **18**:337–365, 1986.
- [170] M. Ruderman. Pulsar spin-down induced phenomena: Heating; magnetic-field evolution; glitches; pulse-period modulations. *ASP Conference Series*, **308**:251–260, 2003.
- [171] C. Runge. Über die numerische Auflösing von Differentialgleichungen,. *Math. Ann.*, **46**:167–178, 1895.
- [172] B. S. Sathyaprakash and B. F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Rev. Rel.*, **12**:2, 2009.
- [173] J. Sauls. Superfluidity in the interiors of neutron stars. In H. Ögelman & E. P. J. van den Heuvel, editor, *Timing Neutron Stars*, page 457, 1989.
- [174] J. A. Sauls, D. L. Stein, and J. W. Serene. Magnetic vortices in a rotating $P23$ neutron superfluid. *Phys. Rev. D*, **25**(4):967–975, 1982.

- [175] V. Schneider et al. New algorithms for ultra-relativistic numerical hydrodynamics. *J. Comput. Phys.*, **105**:92–107, 1993.
- [176] J. A. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, 1996.
- [177] S. Shapiro and S. Teukolsky. *Black Holes, White Dwarfs and Neutron Stars*. John Wiley and Sons, 1983.
- [178] H. Shen, H. Toki, K. Oyamatsu, and K. Sumiyoshi. Relativistic Equation of State of Nuclear Matter for Supernova Explosion. *Progress of Theoretical Physics*, **100**:1013–1031, 1998.
- [179] M. Shibata and K. Taniguchi. Merger of binary neutron stars to a black hole: Disk mass, short gamma-ray bursts, and quasinormal mode ringing. *Phys. Rev. D*, **73**:064027, 2006.
- [180] M. Shibata, K. Taniguchi, and K. Uryū. Merger of binary neutron stars of unequal mass in full general relativity. *Phys. Rev. D*, **68**(8):084020, 2003.
- [181] M. Shibata, K. Taniguchi, and K. Uryū. Merger of binary neutron stars with realistic equations of state in full general relativity. *Phys. Rev. D*, **71**(8):084021, 2005.
- [182] M. Shibata and K. Uryū. Simulation of merging binary neutron stars in full general relativity: $\Gamma=2$ case. *Phys. Rev. D*, **61**(6):064001, 2000.
- [183] M. Shibata and K. Uryū. Merger of black hole neutron star binaries in full general relativity. *Classical and Quantum Gravity*, **24**:125, 2007.
- [184] C. W. Shu and S. Osher. Efficient Implementation of Essentially Non-Oscillatory Shock Capturing Schemes. *J. Comput. Phys.*, **77**:439–471, 1988.
- [185] C. W. Shu and S. Osher. Efficient Implementation of Essentially Non-Oscillatory Shock Capturing Schemes II. *J. Comput. Phys.*, **83**:32–78, 1989.

- [186] T. Sidery, N. Andersson, and G. L. Comer. Waves and instabilities in dissipative rotating superfluid neutron stars. *Mon. Not. R. Astron. Soc.*, **385**:335–348, 2008.
- [187] T. Sidery, A. Passamonti, and N. Andersson. The dynamics of pulsar glitches: Contrasting phenomenology with numerical evolutions. *ArXiv e-prints*, 2009.
- [188] G. A. Sod. A Survey of Several Finite-Difference Methods for Systems of Nonlinear Hyperbolic Conservation Laws. *J. Comput. Phys.*, **27**:1–31, 1978.
- [189] I. H. Stairs, A. G. Lyne, and S. L. Shemar. Evidence for free precession in a pulsar. *Nature*, **406**:484–486, 2000.
- [190] A. Staniforth and J. Côté. Semi-Lagrangian Integration Schemes for Atmospheric Models - A Review. *Monthly Weather Review*, **119**:2206, 1991.
- [191] J. M. Stewart. The Cauchy problem and the initial boundary value problem in numerical relativity. *Classical and Quantum Gravity*, **15**:2865–2889, 1998.
- [192] G. Strang. On the construction and comparison of difference schemes. *SIAM J. Numer. Anal.*, **5**(3):506–517, 1968.
- [193] P. K. Sweby. High Resolution Schemes Using Flux Limiters for Hyperbolic Conservation Laws. *SIAM J. Numer. Anal.*, **21**:995–1011, 1984.
- [194] K. Taniguchi and E. Gourgoulhon. Various features of quasiequilibrium sequences of binary neutron stars in general relativity. *Phys. Rev. D*, **68**(12):124025, 2003.
- [195] A. H. Taub. Relativistic rankine-hugoniot equations. *Phys. Rev.*, **74**(3):328–334, 1948.
- [196] A. H. Taub. Relativistic fluid mechanics. *Annual Review of Fluid Mechanics*, **10**:301–332, 1978.

- [197] J. H. Taylor and J. M. Weisberg. A New Test of General Relativity: Gravitational Radiation and the Binary Pulsar PSR 1913+16. *Astrophys. J.*, **253**:908–920, 1982.
- [198] H. Terashima and G. Tryggvason. A front-tracking/ghost-fluid method for fluid interfaces in compressible flows. *J. Comput. Phys.*, **228**:4012–4037, 2009.
- [199] J. Thornburg. Event and apparent horizon finders for 3+1 numerical relativity. *Living Rev. Rel.*, **10**(3), 2007.
- [200] R. C. Tolman. Static Solutions of Einstein’s Field Equations for Spheres of Fluid. *Phys. Rev.*, **55**(4):364–373, 1939.
- [201] E. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer, 1999.
- [202] E. F. Toro and S. J. Billett. A unified Riemann-problem-based extension of the Warming-Beam and Lax-Wendroff schemes. *IMA J. Numer. Anal.*, **17**(1):61–102, 1997.
- [203] Y. Totsuka et al. Observation of a neutrino burst from the supernova SN1987a. *Nuclear Physics A*, **478**:189–195, 1988.
- [204] K. Tsubono and The TAMA Collaboration. TAMA Project. In K. Tsubono, M.-K. Fujimoto, & K. Kuroda, editor, *Gravitational Wave Detection*, page 183, 1997.
- [205] R. M. Wald. *General Relativity*. University of Chicago Press, 1984.
- [206] C. Wang, H. Tang, and T. Liu. An adaptive ghost fluid finite volume method for compressible gas-water simulations. *J. Comput. Phys.*, **227**(12):6385 – 6409, 2008.
- [207] S.-P. Wang, M. Anderson, J. Oakley, M. Corradini, and R. Bonazza. A Thermodynamically Consistent and fully Conservative Treatment of Contact Discontinuities for Compressible Multicomponent Flows. *J. Comput. Phys.*, **195**:528–559, 2004.

- [208] L. Warszawski and A. Melatos. A cellular automaton model of pulsar glitches. *Mon. Not. R. Astron. Soc.*, **390**:175–191, 2008.
- [209] F. Weber. Strange quark matter and compact stars. *Prog. Part. Nucl. Phys.*, **54**:193–288, 2005.
- [210] F. Weber, R. Negreiros, and P. Rosenfield. Neutron Star Interiors and the Equation of State of Superdense Matter. *ArXiv e-prints*, 2007.
- [211] J. Weber. Gravitational Radiation. *Phys. Rev. Lett.*, **18**:498–501, 1967.
- [212] J. Weber. Evidence for Discovery of Gravitational Radiation. *Phys. Rev. Lett.*, **22**:1320–1324, 1969.
- [213] R. Wijnands and M. van der Klis. A millisecond pulsar in an X-ray binary system. *Nature*, **394**:344–346, 1998.
- [214] J. R. Wilson and G. J. Mathews. *Relativistic Numerical Hydrodynamics*. Cambridge University Press, 2003.
- [215] D. G. Yakovlev and C. J. Pethick. Neutron Star Cooling. *Annu. Rev. Astron. Astrophys.*, **42**:169–210, 2004.
- [216] T. Yamamoto, M. Shibata, and K. Taniguchi. Simulating coalescing compact binaries by a new code (SACRA). *Phys. Rev. D*, **78**(6):064054, 2008.
- [217] J. W. York. *Sources of Gravitational Radiation: Proceedings of the Batelle Seattle Workshop*, chapter Kinematics and Dynamics of General Relativity. Cambridge University Press, 1979.
- [218] V. E. Zavlin and G. G. Pavlov. Modeling Neutron Star Atmospheres. In W. Becker, H. Lesch, & J. Trümper, editor, *Neutron Stars, Pulsars, and Supernova Remnants*, page 263, 2002.
- [219] M. Zimmermann and E. Szedenits. Gravitational waves from rotating and precessing rigid bodies: Simple models and applications to pulsars. *Phys. Rev.*, **20**(2):351–355, 1979.