UNIVERSITY OF SOUTHAMPTON

Multiscale modelling of neutron star oceans

by

Alice Harpole

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

in the
Faculty of Social, Human and Mathematical Sciences
Department of Mathematical Sciences

June 2018

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL, HUMAN AND MATHEMATICAL SCIENCES
DEPARTMENT OF MATHEMATICAL SCIENCES

Doctor of Philosophy

by Alice Harpole

Type I X-ray bursts are thermonuclear burning events which occur on the surface of accreting neutron stars. Burning begins in a localised spot in the star's ocean layer before spreading across the entire surface. By gaining a better understanding of X-ray bursts, tighter limits can be determined for other neutron star properties such as the mass, radius, spin frequency and magnetic field. The ocean environment is very extreme, involving much higher pressure, temperature and magnetic field strength compared to the conditions typically found in terrestrial systems. We shall be looking at the effects of the strong gravitational field, modelling the ocean using general relativistic hydrodynamics.

The physics of X-ray bursts acts over a wide range of scales, which introduces a number of challenges when modelling them. In this work, we use the multiscale approach to couple together multiple physical models in order to best capture the physics across these various scales. On the smallest scales, the physics is dominated by turbulent burning. The speed of propagation of the burning front is much slower than the acoustic speed, making it difficult to model this with conventional numerical schemes. We therefore instead use the low Mach number approximation, which we have derived and implemented for the relativistic fluid equations based on the existing approach developed for the Newtonian case. On larger scales, the burning front can be thought of as a discontinuity. To model this, we investigate the reactive Riemann problem for relativistic deflagrations and detonations and develop a numerical solver. The large scale propagation of the burning front is believed to be dominated by the Coriolis force. To capture this behaviour, we have derived and implemented a model for the relativistic form of the shallow water equations. Finally, we construct a hybrid scheme to combine the best features of these approximations, extending existing adaptive mesh refinement techniques to include different physical models at different scales.

Contents

List of figures

List of tables

Nomenclature

Symbols used in this thesis are as below unless otherwise specified.

ρ	density
p	pressure
ε	specific internal energy
c_s	speed of sound
h	specific enthalpy / fluid height
T	temperature
γ	adiabatic index
Γ_1	first adiabatic coefficient, $\Gamma_1 = \mathrm{d}\log p/\mathrm{d}\log\rho\|_s$
t	time
ν	frequency
$\boldsymbol{u}$	velocity
v^i	components of the 3-velocity
X_k	species mass fraction of species k
$\dot{\omega}_k$	production rate of species k
q_k	energy release per unit mass of species k
Q	reaction energy per unit mass
M	neutron star mass
$M_\odot$	mass of the Sun, $M_\odot = 1.989 \times 10^{33}\,\mathrm{g}$
$\dot{X}$	time derivative of quantity X
L_{Edd}	Eddington luminosity
c	speed of light in a vacuum, $c = 2.998 \times 10^{10}\,\mathrm{cm\,s^{-1}}$
g	gravitational acceleration
G	Gravitational constant, $G = 6.674 \times 10^{-2}\,\mathrm{cm^3\,g^{-1}\,s^{-2}}$
Φ	geopotential
κ	opacity / thermal conductivity
R	neutron star radius
σ	Stefan-Boltzmann constant, $\sigma = 5.670 \times 10^{-5}\,\mathrm{g\,s^{-3}\,K^{-4}}$
T_{eff}	effective surface temperature
B	magnetic field density
z	redshift
z_s	redshift at neutron star surface

Ma	Mach number, $\mathrm{Ma} = v/c_s$
Da	Damköhler number, $\mathrm{Da} = \omega\ell/u$
Fr	Froude number, $\mathrm{Fr} = u/\sqrt{g\ell}$
π	dynamic pressure
E	total energy per unit volume
$g_{\mu\nu}$	metric coefficients
γ_{ij}	coefficients of spatial 3-metric
${}^{(3)}\gamma$	determinant of spatial 3-metric
α	lapse function
β^i	shift vector
$T^{\mu\nu}$	stress-energy tensor
W	Lorentz factor
$h_\mu{}^\nu$	projection operator, $h_\mu{}^\nu = u_\mu u^\nu + \delta_\mu{}^\nu$
$\Gamma^\mu{}_{\nu\rho}$	connection coefficients (also referred to as Christoffel symbols)
$R_{\mu\nu}$	Ricci tensor
$\boldsymbol{q}$	state vector
$\boldsymbol{f}^i$	flux vector in direction i
$\boldsymbol{s}$	source vector
$q^n_{i,j}$	data q at grid cell (i, j) at timestep n

Declaration of authorship

I, Alice Harpole, declare that this thesis entitled *Multiscale modelling of neutron star oceans* and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

- 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 Harpole and Hawke (2016).

Signed: .....

Date: .....

Acknowledgements

Firstly, I would like to thank my supervisor, Ian Hawke, for his help and guidance throughout the past 4 years. In particular, I am extremely grateful to him for allowing me the freedom to go off on tangents and for encouraging me to have the faith in my own abilities to apply for opportunities I would never have considered myself suitable for.

I would not be here today without the support from my family. Many thanks go to my sister for putting up with me spamming her with pointless messages throughout the day. I'd like to thank my parents for teaching my 6 year old self what an astrophysicist is and for supporting me ever since to realise my childhood dream. Special thanks go to my mother for providing me with adorable photos of my cat Arya whenever I needed them.

Finally, I'd like to thank all the friends I have made here in Southampton. Thank you to my friends in the gravity group for their support, thought-provoking discussions and always allowing to vent my programming frustrations. Thank you also to the members of the Southampton University Taekwondo Club. When I started this PhD, I would never have imagined myself part of a university sports team, let alone team captain of a martial arts club, but taekwondo has proved to be the best possible distraction from work-related worries, has shown me what I can achieve if I push myself and allowed me to meet some amazing people.

This work was made possible with the following Python packages and development tools: `numpy` (Oliphant, 2006), `scipy` (Jones et al.), `matplotlib` (Hunter, 2007), `yt` (Turk et al., 2011) and Jupyter (Kluyver et al., 2016).

1 Background & motivation

1.1 Introduction

Type I X-ray bursts are thermonuclear burning events which occur on the surface of accreting neutron stars. Burning is thought to begin in a localised spot in the ocean of the star before spreading across the entire surface (Joss, 1978; Shara, 1982). By gaining a better understanding of X-ray bursts, tighter limits can be determined for other neutron star properties such as the mass, radius, spin frequency and magnetic field strength. In this thesis, we aim to develop and build a numerical model of type I X-ray bursts in order to improve our understanding of the physics that produces them.

Neutron stars are compact objects, so the effects of strong gravity must be taken into account when investigating the ocean physics. We shall therefore include general relativistic effects in our model of this system, using the fluid equations on a curved background spacetime. The speed of propagation of the burning front is much slower than the acoustic speed in the ocean, which introduces challenges when modelling these systems numerically (Cavecchi, 2013). This is addressed here by applying the low Mach number approximation to the fluid equations, extending the approach developed for the Newtonian case by Nonaka et al. (2010) to relativistic fluids. There are additional problems when modelling the burning front itself, as the physics involved acts over a very large range of both spatial and temporal scales (Almgren et al., 2008). On the scale of the entire neutron star ocean, the burning front can be thought of as a discontinuity. This motivates taking a closer look at the physics of detonations and deflagrations in general relativistic fluids. The large scale propagation of the front is believed to be dominated by the Coriolis force (Spitkovsky et al., 2002; Cavecchi et al., 2016). To investigate this behaviour, we use the shallow water equations to model the large scale motion of the ocean. To accurately model the propagation of burning fronts, we need to represent both the small scale and large scale physics. To do this we use multiscale methods, building a hybrid scheme that incorporates multiple physical models to capture the physics across a range of scales.

In this thesis, we shall begin by discussing the physics and theoretical background of X-ray bursts and the motivation for studying them in Chapter 1. We shall then derive

the low Mach number approximation of the relativistic fluid equations used to describe the burning front propagation, following this with a description of their numerical implementation in Chapter 2. Next, we will present a discussion of relativistic detonations and deflagrations and the work done to model them in Chapter 3. In Chapter 4, we shall present the shallow water equations for modelling the large scale propagation of the burning front. We shall combine these with the compressible and low Mach number equations to produce a multiscale model, with the implementation of this described in Chapter 5. Finally, we conclude the thesis with a discussion of the work done and plans for the future in Chapter 6.

1.2 Neutron stars

Neutron stars are compact objects thought to form in one of two ways: either when a main sequence star of mass $8\ \mathrm{M}_\odot \lesssim M \lesssim 30\ \mathrm{M}_\odot$ undergoes gravitational collapse following a supernova, or from a collapsing white dwarf (Carroll and Ostlie, 1996). They are typically of mass $M \sim 1.4\ \mathrm{M}_\odot$ and radius $10\,\mathrm{km} \lesssim R \lesssim 15\,\mathrm{km}$ (Haensel et al., 2007), however there remains a large amount of uncertainty in these quantities as they are difficult to measure using existing techniques. Neutron stars' compactness, the ratio of mass and radius M/R, is $\sim 10^5$ times greater than that of the Sun: consequently, they have large enough surface gravity that general relativistic effects become important.

Accurate models of X-ray bursts will help provide a method to determine neutron star radii and masses more accurately (Miller, 2013; Watts et al., 2016; Özel and Freire, 2016). A number of different techniques for measuring mass and radii using bursts currently exist (see e.g. Özel and Freire (2016), Section 1.3.5), however they are limited by uncertainties in the underlying model. If we can accurately model the propagation across the neutron star surface of the burning fronts that produce these bursts, we will be able to put tighter limits on these quantities. This would help put tighter constraints on the interior's equation of state, giving a better idea of its composition and structure. As will be described in the next section, there are many uncertainties in existing models of the interior, so tighter constraints would help to strengthen these models. It may also be possible to put limits on other quantities such as the magnetic field strength, which is likewise difficult to measure but crucial to understanding the interior physics, as well as giving us information about the neutron star's binary companion.

1.2.1 Structure

The internal structure of neutron stars is poorly understood: the lack of certainty in the mass and radius means it is hard to put firm constraints on any particular equation of state (Lattimer and Prakash, 2016; Özel and Freire, 2016; Oertel et al., 2017). It

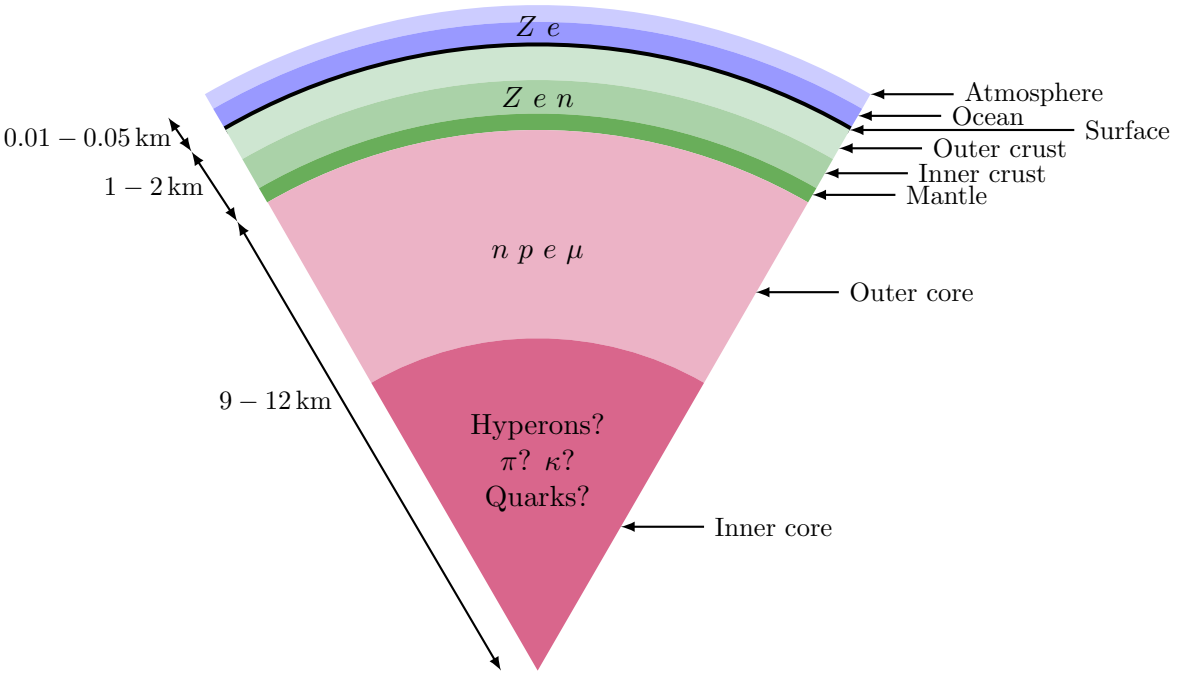

FIGURE 1.1: Neutron star structure, consisting of the exterior layers (blue), the crust (green) and the core (pink) (Haensel et al., 2007). Note that the relative thickness of the layers is not to scale.

is believed to consist of five main regions: the inner and outer cores, the crust (which can be further divided into the inner crust, outer crust and mantle), the ocean and the atmosphere (Lattimer and Prakash, 2004; Haensel et al., 2007). This is illustrated in Figure 1.1.

The outermost layer of a neutron star is the atmosphere. Its thickness varies with the effective surface temperature of the star, from $\sim 10\,\mathrm{cm}$ for a hot neutron star with effective surface temperature $T_s \sim 3 \times 10^6\,\mathrm{K}$, to $\sim 1\,\mathrm{mm}$ for a cold star with $T_s \sim 3 \times 10^5\,\mathrm{K}$. The atmosphere consists of a plasma of charged ions and molecules. Despite being so thin, the atmosphere is important in determining the star's electromagnetic spectrum.

Beneath the atmosphere is the ocean layer, which is about $10 - 50\,\mathrm{m}$ thick. This is not an ocean in the sense of liquid oceans on Earth, being composed of a plasma of electrons and nuclei. The Coulomb interaction energy between the ions is greater than the thermal energy, such the plasma behaves like a liquid (Medin and Cumming, 2011; Cavecchi, 2013). The ocean influences the transport and release of thermal energy from the neutron star's surface. It is in this ocean layer that the burning which produces X-ray bursts takes place. Due to the strong gravitational field in the ocean, the composition and density is highly stratified (Brown et al., 2002; Peng et al., 2007). This is illustrated in Figure 1.2. In the case of the accreting neutron stars that shall be studied in this thesis, the outermost layers are rich in the matter accreted from the donor star, so are

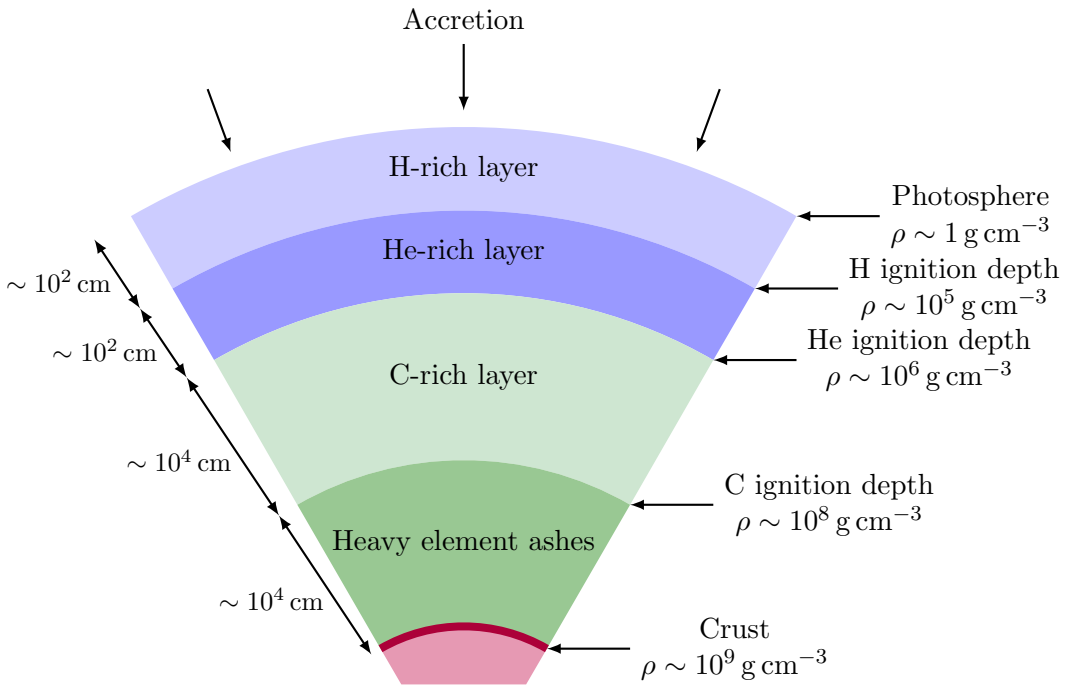

FIGURE 1.2: Structure of the atmosphere and ocean of an accreting neutron star, showing the depths and density at which thermonuclear ignition occurs. Note that the relative thickness of the layers is not to scale. Adapted from (Watts, 2012).

primarily made up of hydrogen and helium. Beneath them are layers made up of the ashes from previous bursts: carbon and heavier elements.

Beneath the ocean is the crust layer, which extends to $1 - 2\,\mathrm{km}$ below the surface and primarily consists of nuclei (Chamel and Haensel, 2008). This can be split into the outer crust (a few hundred metres thick) and the inner crust (up to a kilometre thick). The outer crust takes the form of a cubic iron lattice. As the density increases towards the centre of the star, the matter in the crust undergoes a number of transitions. At densities of $\rho \sim 10^4\,\mathrm{g\,cm^{-3}}$, the atoms become fully ionised, and at $\rho \sim 10^7\,\mathrm{g\,cm^{-3}}$ the composition of the nuclei becomes more neutron rich due to electron captures. The physics of the outer crust is important for X-ray bursts (especially for superbursts which occur deeper in the ocean, see Section 1.3.6), as heating and cooling of the crust will influence the thermal properties of the ocean (Deibel et al., 2016; Meisel and Deibel, 2017).

The boundary of the inner and outer crusts is defined by the neutron drip point at $\rho_{\mathrm{ND}} \simeq 4 \times 10^{11}\,\mathrm{g\,cm^{-3}}$. This corresponds to the point where the net neutron chemical potential of a neutron in a nucleus is zero, such that neutrons are no longer bound within nuclei and begin to 'drip out' (Bethe et al., 1970). The inner crust extends down to $\sim 0.1\rho_0 \simeq 10^{14}\,\mathrm{g\,cm^{-3}}$, where ρ_0 is the density at the centre of the star. In this region, the free neutrons may condense into a superfluid phase (Baym and Pethick, 1979). At the base of the crust there is a thin mantle layer, where it is thought that the matter transitions from the nuclear lattice to a neutron fluid in a series of 'pasta' phases between

$\sim 0.1\rho_0 - 0.3\rho_0$. This begins with a 3d lattice of voids, followed by 2d cylindrical voids, 1d slabs of nuclei interlaid with planar voids, 2d cylindrical nuclei and finally 3d nuclei (Hashimoto et al., 1984; Oyamatsu et al., 1984; Watanabe, 2007).

The core makes up to 99% of the total mass of the neutron star. Like the crust, it can also be split into outer and inner regions. The outer core consists primarily of a neutron fluid, with the remaining few percent made up of protons, electrons and muons. All components in the plasma are strongly degenerate. The electrons and muons produce an ideal Fermi gas, and the neutrons and protons form a strongly interacting Fermi liquid which may be in a superfluid state (Shternin et al., 2011). The central region of the star is known as the inner core and can be several kilometres in radius. Due to the extremely high densities in this region that go far beyond anything we can produce in the laboratory, its composition is largely unknown. It is possible that it may consist of exotic particles like hyperons, pions or kaons, or even that there may be a transition into a mixed phase of hadronic and deconfined quark matter (Potekhin, 2010).

The above description outlines our current understanding of the interior structure of neutron stars, however much uncertainty remains. The physics of the core is particularly challenging. It is not currently possible to recreate the extremely high densities of the core in the laboratory, and there does not yet exist a precise relativistic many body theory of strong interacting particles – we must instead rely on theoretical models. At present, the strongest constraint for the equation of state comes from the fact that neutron stars have been observed with masses $M > 2$ M$_\odot$ (Antoniadis et al., 2013).

Knowledge of the interior structure of neutron stars is relevant for our work as dynamics within the interior can have important consequences for physics on the surface. For example, the motion of superfluid neutrons in the outer core may alter the core's magnetic field, resulting in changes in the star's crust and surface magnetic field (Ruderman et al., 1998). The star's equation of state will determine its mass and radius, which in turn will determine the surface gravity. This interior physics will therefore directly affect processes on the surface such as X-ray bursts.

1.2.2 Low mass X-ray binaries

All known sources of type I X-ray bursts are believed to belong to low mass X-ray binaries (LMXBs). LMXBs are binary systems consisting of a degenerate star (a neutron star or a black hole) and a main sequence star of mass $M \lesssim$ M$_\odot$. Approximately 190 of the X-ray sources in the galaxy are believed to be LMXBs, with around 100 of these containing a neutron star (Liu et al., 2007; Lau et al., 2018). They are primarily visible in the X-ray part of the electromagnetic spectrum. The degenerate stars in these systems accrete hydrogen- and helium-rich matter from their companions at rates of $\dot{M} \sim 10^{-11} - 10^{-8}$ M$_\odot$ yr^{-1} (Cumming and Bildsten, 2000) via stable Roche lobe

overflow. This matter spirals inwards towards the degenerate star, forming an accretion disk. The accretion process releases gravitational potential energy, heating up the inner regions of the disk to temperatures of $T \sim 10^7\,\mathrm{K}$, which then emit the observed X-ray radiation (Cavecchi, 2013).

LMXBs are typically long lived, with lifetimes up to $\sim 10^{10}\,\mathrm{yr}$; by contrast, high mass X-ray binaries (HMXBs) consisting of a degenerate star and a main sequence star of mass $M \gtrsim 10\ \mathrm{M}_\odot$ have much shorter lifetimes of only $\sim 10^5\,\mathrm{yr}$ (Mineo et al., 2012; Walter et al., 2015). Most LMXBs are believed to be older than $10^6 - 10^7\,\mathrm{yr}$; consequently, they are old enough that the crusts of the neutron stars in these systems have been entirely replaced by accreted material and are therefore substantially different from the crusts of isolated neutron stars (Lau et al., 2018). The magnetic field of the degenerate stars in LMXBs is typically very weak, $B \lesssim 10^8 - 10^9\,\mathrm{G}$ (Mukherjee et al., 2015) (by contrast, magnetars are believed to have magnetic fields up to $B \sim 10^{15}\,\mathrm{G}$ (Esposito et al., 2018)), and is believed to be the result of accretion-induced field decay (Tauris and van den Heuvel, 2006). Consequently, very few of these systems contain X-ray pulsars (Casares et al., 2016).

1.3 X-ray bursts

About 110 (or 5%) of the neutron stars that have so far been observed in the Milky Way belong to LMXBs and have been seen to exhibit type I X-ray bursts (Galloway et al., 2008; Galloway and Keek, 2017) (see Zand (2018) for an up-to-date list of sources). Matter in the accretion disk eventually reaches the neutron star surface where it forms an ocean layer (see Figure 1.2). The large lateral pressure gradients caused by the strong gravitational field at the surface are expected to cause the matter to rapidly spread across the surface (Galloway and Keek, 2017). As matter is accumulated, the density and temperature in the ocean increases until it is high enough that nuclear fusion can commence. For the nuclear processes involved, the thermonuclear burning rate increases very sharply with temperature, so this leads to thermonuclear runaway. As a result, although the burning is thought to begin in a small region on the surface – a hot spot – it will very quickly spread throughout the entire ocean (Joss, 1978; Shara, 1982).

As shall be described in the subsequent sections, X-ray bursts display a wide range of features. Although there have been many models proposed to explain individual features, there remain many open questions and there exists no model which can consistently explain all observed features for all observed sources.

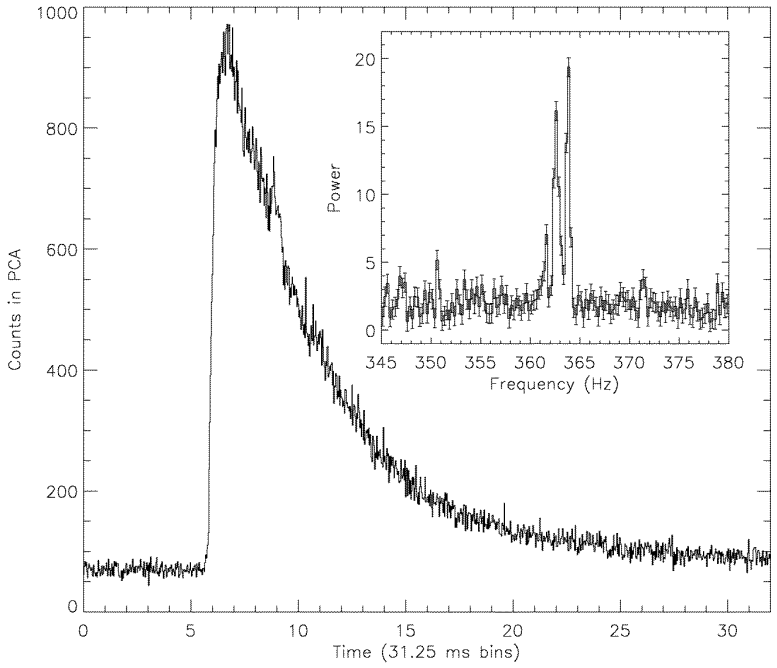

FIGURE 1.3: Observed light curve of an X-ray burst from 4U 1728-34 (Strohmayer et al., 1996). The rise, peak and decay phases of the burst can clearly be seen. The inset panel shows the power spectrum of the burst: there is a peak at around $\sim 364\,\mathrm{Hz}$, showing that the burst exhibits burst oscillations at a frequency which coincides with the neutron star's spin frequency.

1.3.1 Observations

The first type I X-ray bursts were detected in 1975 by the SAS-3 and ANS satellites (Grindlay et al., 1976; Clark et al., 1976). They have since been observed by almost all major X-ray satellites (Galloway et al., 2018), with previous and current missions including the Rossi X-ray Timing Explorer (RXTE) (Galloway et al., 2008), BeppoSAX (Cornelisse et al., 2003), Swift (Gehrels et al., 2004), INTEGRAL (Sguera et al., 2005), MAXI (Matsuoka et al., 2009), NuSTAR (Harrison et al., 2010), ASTROSAT (Singh et al., 2014) and NICER (Gendreau et al., 2016). 111 burst sources have been observed to date (Zand, 2018; Multi-INstrument Burst ARchive (MINBAR)).

The burning is observed as a very sharp increase in the X-ray radiation of the star: a type I X-ray burst. The bursts produce distinctive light curves (see Figure 1.3) which distinguish them from type II X-ray bursts (Hoffman et al., 1978). The light curve of a type I burst can be split into three distinct phases: the rise, lasting $\sim 1 - 10\,\mathrm{s}$, where the flux increases sharply as the flame spreads across the surface and the ocean ignites; the peak, where the flux reaches a maximum as the entire ocean undergoes burning; and the decay, lasting $\sim 20 - 200\,\mathrm{s}$, where the flux slowly decreases as the ocean cools again. In type II bursts, the peak is much flatter. Type I bursts occur every few hours to days and their spectra almost always soften (i.e. the energy in the spectra moves from harder, higher energy X-ray wavelengths to softer, lower energy X-ray wavelengths) during the

decay, with the average spectrum during the first few seconds of the burst generally harder than the persistent X-ray emission of the source. On the other hand, type II bursts occur every few seconds to minutes and their spectra do not soften during the decay. These differences between the two types of burst arise due to their different origins: type I bursts occur due to thermonuclear burning on the surface, whereas type II bursts result from instabilities in the accretion flow (Hoffman et al., 1978).

Less frequently, bursts with much longer durations have been observed. These intermediate duration bursts and superbursts are believed to occur deeper in the ocean, resulting from deep helium and carbon burning respectively. These bursts also have longer recurrence times as it takes longer to accumulate sufficient fuel. A comparison of the properties of these bursts can be seen in Table 2 of Galloway and Keek (2017).

1.3.2 Burst oscillations

Burst oscillations are coherent oscillations observed in the flux during bursts (see van der Klis (2000); Watts (2012); Chakraborty et al. (2017) for reviews). The power spectrum of a burst displaying burst oscillations can be seen in the inset panel of Figure 1.3. Burst oscillations generally have frequencies very close to the spin frequency of the star (Chakrabarty et al., 2003). They may appear in the rise part of the burst, the decay part of the burst or in both.

Burst oscillations have so far only been detected from 18 sources (see Watts (2015) for an up-to-date list), with possible claims for bursts from 9 more (Galloway and Keek, 2017). For some sources, oscillations have been detected during some bursts but not others. Why this is so is not currently clear, but there does seem to be a greater number detected at higher accretion rates, where oscillations can have larger amplitudes (Muno et al., 2004; Ootes et al., 2016).

There is some evidence that for a few sources the burst oscillation frequency is twice the spin frequency (van der Klis, 2000; Miller, 2000); the cause of this $m = 2$ azimuthal asymmetry is not known (Cumming and Bildsten, 2000), though it may be due to the excitation of some global mode in the ocean (Heyl, 2004; Lee and Strohmayer, 2005; Berkhout and Levin, 2008).

In the standard model of X-ray bursts, burning starts in a hot spot induced by the ignition of nuclear fusion. However, this does not explain why burst oscillations are not solely confined to the rise (Chamel and Haensel, 2008). In the rise, the burning front is still spreading across the surface; burst oscillations are likely to be a result of this temperature asymmetry on the star's surface, which is near stationary in the corotating frame and is modulated by its rotation (Watts et al., 2005). In the decay stage, the whole surface of the star is burning and so there is no longer this surface temperature

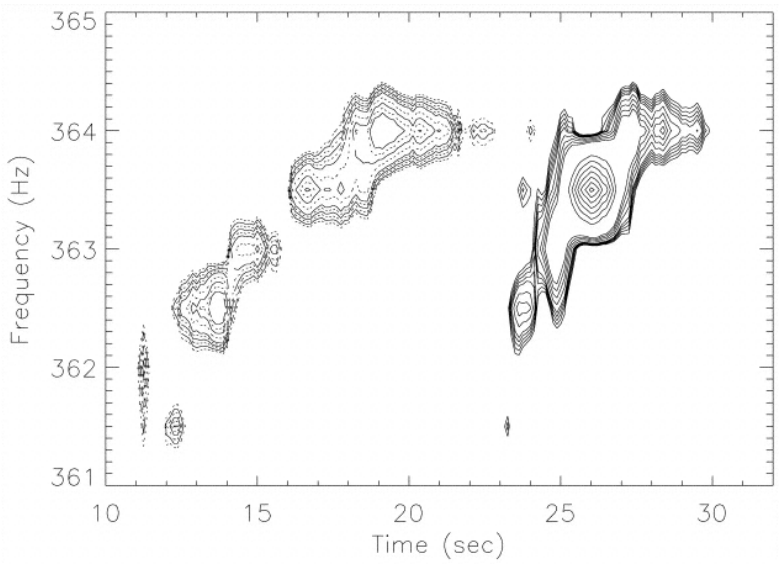

FIGURE 1.4: Dynamic power spectra computed from two different bursts from 4U 1728-34, separated by 1.6 yr (Strohmayer et al., 1998). The bursts both show frequency drifts of ~ 3 Hz, increasing to the same asymptotic frequency.

asymmetry and so no obvious mechanism to cause oscillations to continue (Chakraborty et al., 2017).

A number of models have been proposed to explain the observation of burst oscillations in the decay stage, including the ideas that they could be produced by the dynamic formation of vortices driven by the Coriolis force (Spitkovsky et al., 2002; Bhattacharyya and Strohmayer, 2007), by non-radial surface oscillations (Heyl, 2004; Lee and Strohmayer, 2005; Berkhout and Levin, 2008), or by latitudinal shear instabilities produced by differential rotation between the pole and equator (Cumming, 2005). (Cumming and Bildsten, 2000) suggested decay oscillations may be due to a 'cooling wake', a temperature asymmetry during the cooling of the neutron star that results from the fact that it takes a finite amount of time for the atmosphere to cool and for the burning to spread around the star. This was studied further by Mahmoodifar and Strohmayer (2016), who found that this model produced higher oscillation amplitudes than conventional cooling models and so is more consistent with observations.

Another model to explain oscillations in the decay stage is that the flame front does not spread across the entire surface of the star, and is instead stalled by some mechanism; the temperature asymmetry would then persist into the decay stage. Possible stalling mechanisms include confinement by the Coriolis force (Spitkovsky et al., 2002) or by the magnetic field (Cavecchi, 2013).

1.3.3 Frequency drifts

As described by Strohmayer and Bildsten (2003); Piro (2006), the frequency of burst oscillations can vary by $\sim \Delta 1 - 5\,\mathrm{Hz}$ over the course of the burst (see Figure 1.4). Typically, the frequency increases as the burst progresses, with the oscillations' frequency increasing to some limiting, asymptotic frequency, however there have been observations of decreases in the frequency as well. Bursts with detectable burst oscillations during the rise exhibit the largest frequency shifts, indicating that the process driving frequency evolution begins at the start of the burst (Muno et al., 2002; Watts, 2012). There is also evidence that frequency drifts occur more frequently in cases where the spin frequency is lower (Chakraborty et al., 2017).

There have been several models proposed to explain frequency drifts. Strohmayer et al. (1997b) suggested that they could be explained by the horizontal motion of the spreading flame front, however this would require the rotation to set a preferred direction in which the flame spreads. Alternatively, they could be caused by the conservation of angular momentum as the atmosphere expands then shrinks again as it cools (Strohmayer et al., 1997a; Cumming, 2000). Lee and Strohmayer (2005); Bhattacharyya et al. (2005) found that including general relativistic effects (gravitational redshift, light bending and frame dragging) gave corrections to frequency drift estimates of $5 - 10\%$.

1.3.4 Multiple peaks

There have been several observations of bursts with two or even three peaks (Fisker et al., 2003; Bhattacharyya and Strohmayer, 2006a,b; Cooper and Narayan, 2007; Zhang et al., 2009). It is believed that some kind of stalling mechanism halts the burning front for a few seconds before it becomes re-energised and continues propagating across the surface (Zhang et al., 2009; Cavecchi et al., 2015); the physical nature of this stalling mechanism is not currently known. It has been proposed by Bhattacharyya and Strohmayer (2006a) that multiple peaks may be produced by the flame stalling as it crosses the star's equator, where the Coriolis force decreases to zero. However, this effect does not stall flame propagation for long enough and does not explain observations of more than two peaks. Alternatively, Fisker et al. (2003) suggested that double-peaked bursts are caused by a waiting point in the thermonuclear reaction chain.

1.3.5 Photospheric radius expansion

As described by Strohmayer and Bildsten (2003), in bright bursts the local X-ray luminosity may reach the Eddington limit

$$L_{\mathrm{Edd}} = \frac{4\pi c G M}{\kappa}\left(1 - \frac{2GM}{c^2 R}\right)^{-1/2} = 4\pi R^2 \sigma T_{\mathrm{eff}}^4, \tag{1.1}$$

where M is the neutron star mass, R the radius, κ the atmospheric opacity, σ the Stefan-Boltzmann constant and T_{eff} the effective surface temperature. At this luminosity, the outwards force produced by the radiation pressure exceeds the inwards gravitational force, causing the photospheric layers to be lifted off the neutron star's surface and producing a photospheric radius expansion (PRE) burst. This causes the atmosphere to expand outwards by $\sim 10 - 50\,\text{m}$ ($\sim 5 - 25\,\text{m}$) during a mixed H/He (pure He) burst (Cumming and Bildsten, 2000). During such bursts, the blackbody temperature decreases as the inferred blackbody radius increases; the total X-ray flux stays approximately constant. The amount by which the photosphere can be lifted varies considerably between different bursts and different sources, and can sometimes be large enough to shift the effective surface temperature below the X-ray band entirely.

PRE bursts are of particular interest as they can act as (approximate) 'standard candles', allowing us to measure the mass and radius of the star. As we know that they are likely to reach the Eddington flux limit, this measurement can be combined with measurement of the apparent angular size and distance (Özel and Freire, 2016). The apparent angular size is found by measuring the bolometric thermal flux F_{bol}, and combining this with the effective surface temperature T_{eff} (calculated from the spectrum) and the distance measurement D to get

$$\frac{R_{\text{obs}}}{D} = \left(\frac{F_{\text{bol}}}{\sigma T_{\text{eff}}}\right)^{1/2}, \tag{1.2}$$

where σ is the Stephan-Boltzmann constant. Due to their extreme compactness, neutron stars will gravitationally lens their own emission. This means the observed radii found from the apparent angular size and distance must be modified to include this gravitational lensing effect,

$$R_{\text{obs}} = \left(1 - \frac{2GM}{c^2 R}\right)^{-1/2} R. \tag{1.3}$$

Combining this with (1.1), we can then solve for the mass and radius.

1.3.6 Superbursts

Superbursts are a type of X-ray burst which have larger luminosity (energy $\sim 10^{42}\,\text{erg}$), longer duration ($4 - 14$ hours) and occur less frequently (recurrence time ~ 1 year) (Cumming et al., 2006); see Zand (2017) for a review. All sources of observed superbursts are known to also be normal X-ray burst sources (Serino et al., 2016). They are believed to occur deeper in the ocean, where crustal heating triggers the unstable ignition of $^{12}\text{C} + {}^{12}\text{C}$ fusion (Deibel et al., 2016) (where ^{12}C has been produced as the ashes of burning during normal X-ray bursts). Because of this, superbursts are more sensitive to the thermal properties of the crust and core, so can be used to put tighter constraints

on the thermal conductivity of the neutron star crust and the neutrino cooling rate in the core (Brown, 2004; Schatz and Rehm, 2006).

After a superburst, normal X-ray bursts are quenched, with the next burst not occurring until several 10s-100s of days later (Cornelisse et al., 2000; Kuulkers et al., 2002; Serino et al., 2016). This occurs because the superburst deposits heat in the outer layers, causing freshly accreted hydrogen and helium to burn stably. Only once the envelope has cooled again does burning become unstable and the previous bursting behaviour resumes.

Current superburst models require a larger heat flux than can be supplied solely by crustal heating (Cumming et al., 2006); it is not known where this extra heating comes from. Burst oscillations have also been observed during superbursts at frequencies very close to those of the burst oscillations during regular bursts (Strohmayer and Markwardt, 2002). This indicates that despite burning occurring at a deeper ocean depth for superbursts, the physical processes producing burst oscillations are the same.

1.3.7 Ignition latitude

There is some uncertainty with regards to whereabouts on the surface of the star ignition first occurs. It has long been thought that it is very unlikely that all the accreted fuel over the entire surface will ignite simultaneously (Joss, 1978; Shara, 1982): it must instead begin at a localised 'hot spot' then spread across the surface of the star. The latitude of this hot spot is an open question. It was found by Spitkovsky et al. (2002) that the latitude of ignition is very important in determining the speed of the burning front's propagation and that ignition on a rapidly rotating neutron star is most likely to occur on the equator due to reduced effective gravitational acceleration. However, observations of burst rise oscillations and double-peaked bursts by Bhattacharyya and Strohmayer (2005, 2006a,b) suggest that the ignition must, at least in the systems studied, occur at or near the rotational pole. The case for ignition occurring at a latitude other than the equator was also made by Cooper and Narayan (2007) (where the ignition latitude was found to depend on the accretion rate) and Maurer and Watts (2008); Cavecchi et al. (2015).

1.3.8 Accretion rate

The accretion rate strongly influences whether thermonuclear burning will be stable or unstable (Fujimoto et al., 1981; Bildsten and Hall, 1998; Keek et al., 2009). A higher rate of accretion will lead to more heating due to compression, increasing the temperature. Assuming a neutron star of mass 1.4 $M_\odot$ with a radius of 10 km, a hydrogen mass fraction of $X = 0.7$ and a CNO mass fraction of $Z_{\rm CNO} = 0.01$, models predict that

hydrogen burning via the CNO cycle will be stable for accretion rates of $\dot{M} \gtrsim 10^{-2}\dot{M}_{\rm Edd}$. Theory predicts that helium burning via the triple-alpha process will be stable for $\dot{M} \gtrsim \dot{M}_{\rm Edd}$, but if the accreted matter is hydrogen-deficient this could be as high as $\dot{M} \gtrsim 10\dot{M}_{\rm Edd}$. However, observations indicate that stable helium burning in fact occurs at lower accretion rates of $\dot{M} \gtrsim 0.1\dot{M}_{\rm Edd}$. Models by Keek et al. (2009) suggested that this could be due to magnetic instabilities and rotation-induced mixing.

It has been found that there is a strong link between the accretion rate and the appearance of burst oscillations (Strohmayer and Bildsten, 2003; Watts, 2012; Zhang et al., 2015; Lampe et al., 2016). The accretion rate of a neutron star in an LMXB can vary with time, with burst oscillations only observed when a star is in a particular accretion state. Galloway et al. (2008) found that persistent pulsars (i.e. those that display accretion-powered pulsations throughout their accretion episodes, and not just in transient episodes), including those who have displayed burst oscillations, tend to remain in a low accretion state; sources that are not persistent pulsars tend to have detectable burst oscillations only when they are in a high accretion state. Zhang et al. (2013, 2015) concluded that bursts with oscillations ignite at high latitude when the mass accretion rate onto the NS surface is high, and those without oscillations ignite at low latitude when the mass accretion rate is low.

As described by Cavecchi et al. (2017), the relationship between rotation, accretion rates and burst properties such as burst rate is complex and unclear. Burst rates are generally expected to increase with increased accretion, however in certain 'delayed mixed burst' regimes the reverse is found to be true (Narayan and Heyl, 2003).

1.3.9 Rotation

Neutron stars in LMXBs tend to have very high rotation rates. Consequently, bursts have been detected for stars with spin frequencies in the range $11 - 620\,\mathrm{Hz}$, with the vast majority having frequencies $\gtrsim 300\,\mathrm{Hz}$. Such rapid rotation is likely to have an effect on the evolution of a burst.

The importance of rotation on ocean dynamics can be estimated from the Rossby number, R_0, the ratio of inertial to Coriolis forces in the fluid equations,

$$R_0 = \frac{v}{4\pi\nu_s \ell \cos\theta}, \tag{1.4}$$

where v is the velocity of motion, ℓ the characteristic lengthscale, ν_s the spin frequency of the star and θ the colatitude (Watts, 2012). When $R_0 \lesssim 1$, the Coriolis force becomes dominant. If we assume that the flame spreads due to pressure gradients, then we can take the velocity to be $v = \sqrt{gH}$, where g is the gravitational field strength and H the ocean scale height (Spitkovsky et al., 2002). We then find that the Coriolis force starts

to significantly slow the flame spread, balancing pressure gradients that develop as the hot fluid expands, for spin frequencies of

$$\nu_s > 25\,\mathrm{Hz}\left(\frac{g}{10^{14}\,\mathrm{cm\,s^{-2}}}\right)^{1/2}\left(\frac{H}{10\,\mathrm{m}}\right)^{1/2}\left(\frac{10\,\mathrm{km}}{R}\right), \tag{1.5}$$

implying that this will indeed be very important for almost all observed systems (Cavecchi et al., 2011).

The effect of the Coriolis force in flame spreading was first modelled by Spitkovsky et al. (2002). They found that initially the Coriolis force is unimportant due to the small size of the hot spot, and the flame spreads with the 'geostrophic speed'[1]. After the flame has spread to encompass a larger region, the Coriolis force will become important and the flame will spread at the much slower 'ageostrophic speed'[2], decreasing as the latitude of the burning front increases. This model was found by Bhattacharyya and Strohmayer (2007) to be consistent with the evolutionary structure of observed burst oscillation amplitudes.

Rapid rotation will also affect oscillation modes in the ocean and atmosphere (Watts, 2012), and could affect ignition conditions by lowering the effective gravity at the equator, leading to a higher fuel accumulation rate there and therefore making ignition more likely on the equator (Spitkovsky et al., 2002; Cooper and Narayan, 2007). Rotation induces mixing of newly accreted material with the ashes from previous bursts, which can change the temperature profile of the ocean (Piro and Bildsten, 2007; Keek et al., 2009; Cavecchi et al., 2017)

Another consequence of fast rotation is to deform the neutron star at the equator, making it oblate (Chandrasekhar and Fermi, 1953; Bocquet et al., 1995; Cadeau et al., 2007; Morsink et al., 2007; Haskell et al., 2008; Bauböck et al., 2013). The oblate shape means that the gravitational field near the poles is stronger than near the equator. This could have important consequences for e.g. bursts with multiple peaks, where it could help provide a stalling mechanism as the flame crosses the equator, enhancing the effect studied in spherical geometry by Cavecchi et al. (2015).

1.3.10 Magnetic fields

Neutron stars which exhibit X-ray bursts are believed to have relatively weak magnetic fields of $B \lesssim 10^8 - 10^9\,\mathrm{G}$ (Mukherjee et al., 2015); this contrasts with magnetars which can have magnetic fields up to $B \sim 10^{15}\,\mathrm{G}$. Although it is very difficult to measure

[1] The *geostrophic speed* is the speed of a flow in which the force of the pressure gradient is balanced by the Coriolis force.

[2] The *ageostrophic speed* is the speed of the component of the flow that deviates from the geostrophic flow and typically represents friction and other effects. For the burning front, it is likely to be dominated by the turbulent flame speed.

the magnetic field strength for these systems, an upper limit is inferred from the lack of persistent pulsations during accretion, implying that the spherical magnetospheric radius must be inside the neutron star (Strohmayer and Bildsten, 2003). Even if the magnetic field is much weaker than this, it could have significant consequences for the propagation of the burning front. Cumming and Bildsten (2000) argued that even a field as weak as $B \sim 10^6$ G could become dynamically important if an initially poloidal field becomes wound up by differential rotation. It is possible that the diffusive magnetorotational instability may play a role in vertical angular momentum transport within the burning layer (Menou, 2004; Cumming, 2005).

There are also several features of X-ray bursts which could be explained by the effects of magnetic fields, including frequency drifts (Cumming, 2000) and burst oscillations (Spitkovsky et al., 2002; Watts et al., 2005; Cavecchi, 2013). Cumming et al. (2001) suggested that while magnetic fields are screened by accretion in most LMXBs, this is not so in systems where the accretion rate is unusually low ($\dot{M} < 10^{-2}\dot{M}_{\rm Edd}$). The presence of dynamically important magnetic fields will also influence the accretion geometry (Kajava et al., 2014). Simulations by Cavecchi et al. (2016) found that even weaker magnetic fields of $B \sim 10^7$ G can have a non-negligible effect on the flame propagation, partially obstructing the Coriolis force such that the confinement decreases and the flame speed increases.

As described in Cavecchi (2013), if the Alvén speed is much less than the sound speed (i.e. it has low Mach number), then a hydrostatic scheme will filter out vertically-propagating sound waves but will not entirely filter out vertically-propagating magnetic waves. In the low Mach number regime to be investigated, it may therefore be the case that magnetic waves need to be taken into account when evolving the system. Although magnetic fields have not been included in the models developed in this thesis, this is something that we hope to do in future work.

1.3.11 General relativity

For compact objects such as neutron stars, general relativistic effects can significantly modify the physics. For example, a correction of up to 25% is needed in the surface gravity of a neutron star with a rotation frequency of 600 Hz when general relativity is taken into account (Cumming et al., 2002). However, no study has yet modelled the ocean dynamics during type I X-ray bursts using the fully general relativistic fluid equations.

In general relativity, a rotating body will experience frame dragging effects (Paschalidis et al., 2017). The effect of this on spectral lines was found by Bhattacharyya et al. (2006) to be small for spin frequencies of $300 - 700$ Hz, however it is likely that it will have important consequences for the propagation of the burning front itself.

There have previously been some studies of X-ray bursts which have included general relativistic effects, such as the conservation of angular momentum, as described in Abramowicz et al. (2001); Cumming et al. (2002). Including this effect, as well as gravitational redshift, light bending and frame dragging, were found to give corrections of $5-10\%$ in frequency drifts when calculating light curves (Lee and Strohmayer, 2005; Bhattacharyya et al., 2005). In Ayasli and Joss (1981), general relativistic corrections to the equations of stellar structure from Thorne (1977) were found to have a significant effect on the burst profile, increasing the accretion-driven luminosity by a factor of $2z_s(1+z_s)/[(1+z_s)^2-1]$ (where z_s is the surface redshift), while decreasing the peak burst luminosity by a factor of $\sim(1+z_s)^{-2}$. Wielgus et al. (2015b,a) showed that in general relativity, there exist radiation-supported atmospheres outside the surface of a radiating spherical body, close to the radius where the gravitational and radiative forces balance. These shells may be relevant for PRE bursts and could explain frequency drifts (Abarca and Kluźniak, 2016).

1.3.12 Burning

The burning front in a type I X-ray burst is a subsonic deflagration. Due to the high temperature sensitivity of the nuclear fusion reactions, the burning is confined to a narrow region in space (Klein, 1998; Hillebrandt et al., 2013). Consequently, this produces scale problems in both time and space when modelling the propagation of the front: if it is treated as a discontinuity (which it is on the scale of the whole star), then the internal structure of the flame cannot be resolved. On macroscopic scales, the burning front propagates due to turbulence (Reinecke et al., 1999), so it is not sufficient to merely model the microscopic flame structure. The timescales of the hydrodynamic flow and the nuclear reactions are also very different.

The reactions within the flame structure itself are very important in determining the propagation of the burning front (Klein, 1998; Schatz, 2011). They determine the conditions required for ignition to take place, the rate of burning and the products produced. These in turn determine the flame's behaviour, including the energy generation rate, the propagation speed and the turbulent flame brush thickness[3] (Smiljanovski et al., 1997; Peters, 2000).

The nuclear burning reactions depend on the thermodynamic properties of the ocean: the temperature, density and composition. These properties are strongly influenced by the accretion rate and composition of the accreted material. For lower accretion rates of $\dot{M} \lesssim 0.24\dot{M}_{\rm Edd}$, where $\dot{M}_{\rm Edd}$ is the Eddington limit[4], nuclear burning will occur at lower

[3] The *turbulent flame brush thickness* is defined to be the average distance between the mean front location and the instantaneous front location, and grows as turbulent eddies cause the flame front to become increasingly wrinkled.

[4] The *Eddington limit* is the maximum accretion rate beyond which the outwards radiation pressure produced by the infalling matter exceeds that of the inwards gravitational force.

temperatures $T \lesssim 7 \times 10^7\,\mathrm{K}$ where the CNO process dominates (Ayasli and Joss, 1981; Taam et al., 1996); consequently, the CNO abundance of accreted matter is important for such stars. For higher accretion rates, burning is instead dominated by the rp process, and so the CNO abundance is unimportant. For accretion rates $\dot{M} \gtrsim \dot{M}_{\mathrm{Edd}}$, burning is steady and so no bursts are seen (Fushiki and Lamb, 1987; Schatz et al., 1999).

In current state-of-the-art simulations, reactions in X-ray bursts are implemented using approximate reaction networks which track many different species. For example, the `MAESTRO` code (Nonaka et al., 2010) uses a network containing 10 species based on the description of Wallace and Woosley (1981), and approximates hot CNO burning, triple-α, rp-process breakout and burning up to ^{56}Ni (Malone et al., 2014).

1.3.13 Turbulence

The propagation of the burning front is driven by turbulent combustion (Reinecke et al., 1999). Turbulence operates differently in 2d and 3d systems. In 3d systems, there is a turbulent energy cascade where the energy in large scale features is transferred to smaller scales. In 2d systems this is reversed, with smaller scale features coalescing to form larger scale features, and transferring energy from smaller scales to larger scales (Kraichnan, 1967). In order to properly capture the physics of the turbulent burning front, it is therefore necessary to use a 3d model.

It has been suggested that relativistic turbulence may behave differently to Newtonian turbulence, which could influence the overall propagation of the front. In Radice and Rezzolla (2013), high order direct numerical simulations of driven relativistic turbulence of a hot plasma found that relativistic effects such as nonlinear couplings via the Lorentz factor need to be taken into account even when the Lorentz factor is small. The standard Kolmogorov model of the velocity power spectrum and intermittency[5] of the velocity field in the inertial range was found to carry over to the relativistic regime. However, the relativistic turbulent velocity was found to experience exponential decay rather than the classical power-law decay, an effect consistent with the findings of Zrake and MacFadyen (2011); Inoue et al. (2011). It was also found that relativistic turbulence is significantly more intermittent than its classical counterpart due to non-classical sources of intermittency such as the relativistic constraint that the velocity field cannot be Gaussian as the velocity field PDFs must have compact support in $(-c, c)$ (i.e. the speed cannot be greater than the speed of light). As the Lorentz factor increases, the PDFs flatten and the velocity field shows larger deviations from Gaussianity.

[5] *Intermittency* in a turbulent flow is characterised by the local appearance of anomalous, short-lived flow features. As first observed by Batchelor and Townsend (1949), fluctuations in velocity gradients become more extreme, with longer tails in their probability distrubution with increasing Reynolds number.

It should be noted that it remains uncertain whether these findings are genuine features of relativistic turbulence – it has been suggested that current simulations have not yet achieved a sufficiently long inertial range to yield reliable results (Eyink and Drivas, 2018).

1.4 Background physics

This section will give a brief overview of some of the physics that shall be used in subsequent chapters of this thesis. It shall begin by looking at the physics of fluids, then consider general relativity. It shall finish by looking at the physics of burning.

1.4.1 Hydrodynamics

In this work, we are mainly concerned with the physics of the neutron star ocean. Despite being primarily composed of H/He plasma, the extremely high pressure in the ocean causes interactions between the ions and degenerate electrons such that it behaves as a liquid. We can therefore describe it using hydrodynamics.

The time-dependent Euler equations of ideal Newtonian hydrodynamics are given by Toro (1999) as

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}) = 0, \tag{1.6}$$

$$\frac{\partial (\rho \boldsymbol{u})}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}\boldsymbol{u}) + \boldsymbol{\nabla} p = 0, \tag{1.7}$$

$$\frac{\partial E}{\partial t} + \boldsymbol{\nabla} \cdot (\boldsymbol{u}[E + p]) = 0, \tag{1.8}$$

where ρ, $\boldsymbol{u}$ and p are the density, velocity and pressure, and $E = \rho(e + \frac{1}{2}\boldsymbol{u} \cdot \boldsymbol{u})$ is the total energy per unit volume, where e is the specific internal energy. These equations describe compressible fluids at high pressures, where the effects of body forces, viscous stresses and heat flux are neglected. In order to close this system, an appropriate equation of state must be provided, for example the gamma law equation of state. This relates the pressure, density and specific internal energy by

$$p = (\gamma - 1)\rho e, \tag{1.9}$$

where γ is the adiabatic index, given by the ratio of the specific heat capacities at constant pressure and at constant volume.

In stellar environments, we obtain additional terms due to the presence of gravitational fields, heat transport and species evolution. These are given by Bell et al. (2004) as

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}) = 0, \tag{1.10}$$

$$\frac{\partial (\rho \boldsymbol{u})}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}\boldsymbol{u}) + \boldsymbol{\nabla} p = -\rho g \boldsymbol{e}_r, \tag{1.11}$$

$$\frac{\partial E}{\partial t} + \boldsymbol{\nabla} \cdot (\boldsymbol{u}[E + p]) = \boldsymbol{\nabla} \cdot (\kappa \boldsymbol{\nabla} T) - \rho g \left(\boldsymbol{u} \cdot \boldsymbol{e}_r\right) - \sum_k \rho q_k \dot{\omega}_k, \tag{1.12}$$

$$\frac{\partial (\rho X_k)}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \boldsymbol{u} X_k) = \rho \dot{\omega}_k, \tag{1.13}$$

where T is temperature, $g(r)$ is the radially dependent gravitational acceleration, $\boldsymbol{e}_r$ is the unit vector in the radial direction and κ is the thermal conductivity. X_k is the abundance of species k, which has the production rate $\dot{\omega}_k$ and reaction energy release per unit mass q_k. The system is closed with an equation of state $e = e(\rho, T, X_k)$. These equations assume that the effects of fluid viscosity and species diffusion are negligible compared to energy transport. Gravity is assumed to only act in the radial direction and be time independent (i.e. there is no back reaction from the fluid). This model also neglects the effects of magnetic fields (though as described in Section 1.3.10, it is important to be aware that this assumption may not be entirely valid for neutron star oceans).

1.4.2 Burning

As discussed in Section 1.3.13, the burning in type I X-ray bursts is extremely turbulent. The degree of turbulence in a system can be measured by the Reynolds number, the ratio of the inertial forces to viscous forces. Assuming equal diffusivities for all reactive scalars and that the Schmidt number $\mathrm{Sc} = \nu/D = 1$, where ν is the kinematic viscosity and D is the diffusivity, we can define the flame thickness ℓ_F and the flame time t_F to be

$$\ell_F = \frac{D}{s_L}, \quad t_F = \frac{D}{s_L^2}, \tag{1.14}$$

where s_L is the speed of the laminar flow. The (turbulent) Reynolds number can then be defined as

$$\mathrm{Re} = \frac{v\ell}{\nu} = \frac{v\ell}{s_L \ell_F}, \tag{1.15}$$

where ℓ is the characteristic length scale of the eddies (the *integral scale*[6]) and v is the turnover velocity of integral scale eddies (Peters, 2000). Taking $\rho \sim 10^5\,\mathrm{g\,cm^{-3}}$, $v \sim 10^5\,\mathrm{cm\,s^{-1}}$, $\ell \sim 10^2\,\mathrm{cm}$ (~ 0.1 times the ocean depth) and $\mu \sim 0.3$ (Niemeyer et al., 1999), we can estimate the Reynolds number in neutron star oceans to be $\mathrm{Re} \sim 10^9$. This is much higher than the Reynolds number of burning in terrestrial systems: for example,

[6] More specifically, the *integral scale* is defined as the length scale of the eddies in the turbulent cascade at which the power spectrum is at a maximum.

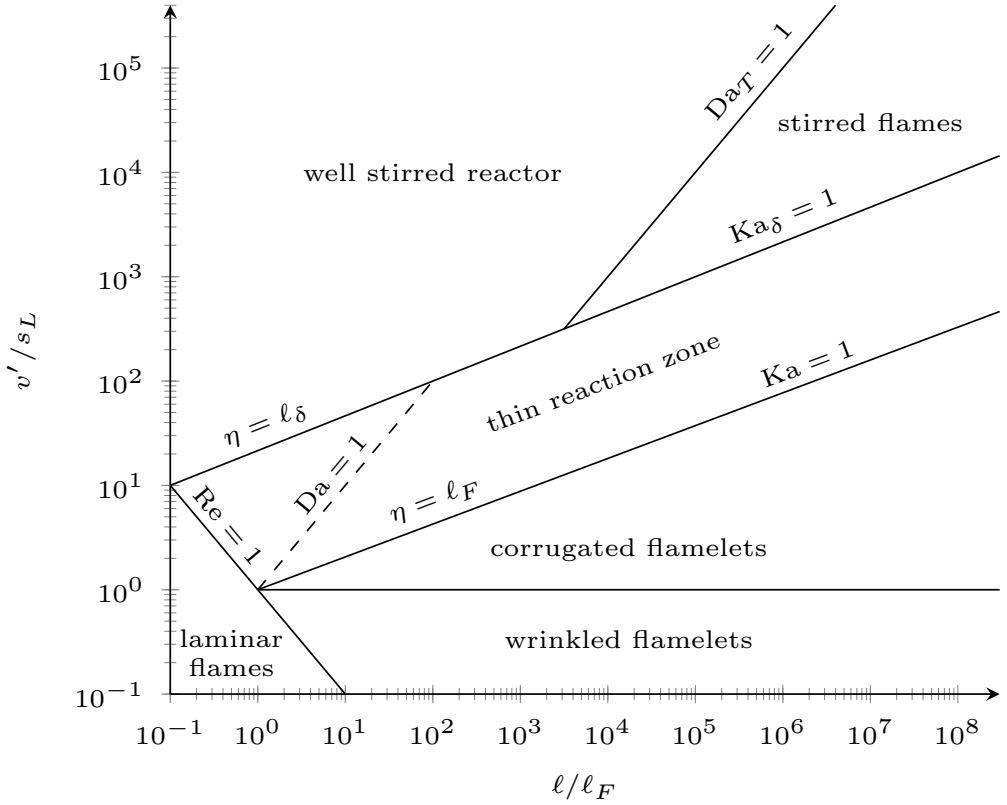

FIGURE 1.5: Regime diagram for premixed turbulent combustion, adapted from (Peters, 2000; Aspden et al., 2010). Labels indicate lines of constant Reynolds number, $\mathrm{Re} = v\ell/s_L\ell_F$, Damköhler number, $\mathrm{Da} = \omega\ell/v$ and Karlovitz number, $\mathrm{Ka} = \ell_F^2/\eta^2$ (where η is the Kolmogorov length scale, the size of the smallest eddies in the turbulent cascade). Burning in neutron star oceans is thought to have a very high Reynolds number and therefore take place in the stirred flames regime, where the reaction zone has become completely disrupted by the turbulent vortices.

the Reynolds number of hydrogen/air turbulent catalytic combustion is typically $\sim 10^4$ (Arani et al., 2018).

The relative size and speed of the turbulent eddies with respect to the burning length and timescales determines the properties of the burning, which can be split into several distinct burning regimes (see Figure 1.5). In the laminar flames regime, the Reynolds number is $\mathrm{Re} < 1$ and the effects of turbulence are negligible. In the flamelet regimes (wrinkled flamelets, corrugated flamelets and the thin reaction zone), the turbulent eddies are larger than the reaction zone thickness ℓ_δ, and cause the burning front to wrinkle. This increases the flame surface area and so increases the burning rate, but has no effect on the microphysics of the burning reactions.

In neutron star oceans, as in type Ia supernovae, it is thought that the burning takes place in the stirred flames regime (Aspden et al., 2008). Here, the turbulent eddies are smaller than the burning length scale, and so are able to penetrate the reaction zone (Niemeyer and Woosley, 1997). This causes the reaction zone to become completely disrupted, causing the burning front to broaden significantly (Klein, 1998; Peters, 2000),

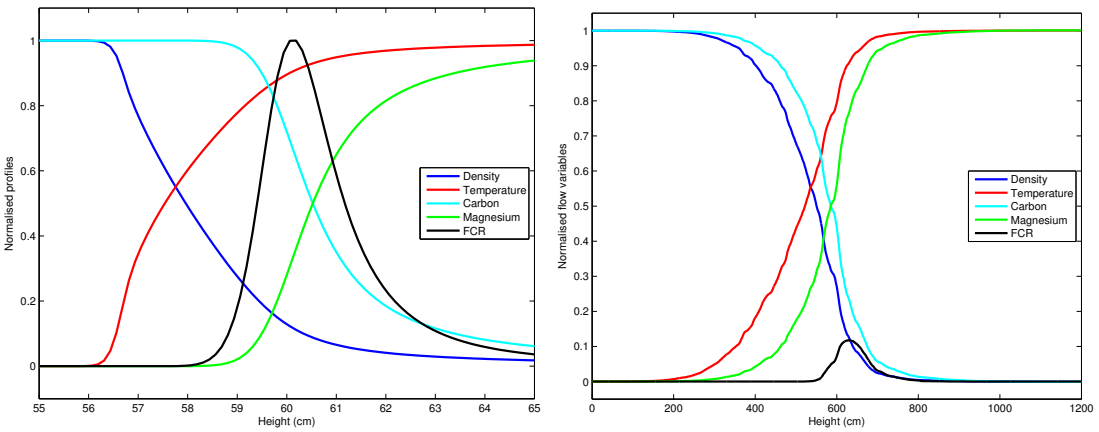

(A) Laminar profile (B) Turbulent profile

FIGURE 1.6: Comparison of laminar and turbulent burning profiles from Aspden et al. (2008). In these plots, the dark blue curves show the density, the red curves the temperature, the light blue curves show the baryon fraction of the fuel, carbon, the green curves the baryon fraction of the product, magnesium, and the black curves the fuel consumption rate (FCR). Note that the length scale on the x-axis spans a distance 120 times greater for the turbulent profile than it does for the laminar profile: the turbulent flame is ~ 60 times wider than the laminar flame.

as can be seen in Figure 1.6. Although the disruption of the reaction zone in the turbulent case greatly reduces the fuel consumption rate, the broadened flame width causes the overall fuel consumption rate across the entire flame to increase.

1.4.3 General relativity

The main aim of our work is to extend the current models of type I X-ray bursts to include general relativity. To do this, we shall use the fully general relativistic form of the fluid equations.

Before going any further, we shall briefly state the notation that shall be used. Throughout this work, the Greek indices μ, ν, ... will be used to indicate summation over all 4 spacetime dimensions, whereas the Latin indices i, j, ... will be used to indicate summation over the 3 spacelike dimensions. The signature $(-,+,+,+)$ will be used. Unless otherwise specified, for all relativistic calculations we shall use natural units, setting

$$G = \mathrm{M}_\odot = c = 1, \tag{1.16}$$

where G is Newton's gravitational constant, $\mathrm{M}_\odot$ is the solar mass and c is the speed of light.

Einstein's field equations are given by

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = 8\pi T_{\mu\nu}, \tag{1.17}$$

where $G_{\mu\nu}$ is the Einstein tensor, $R_{\mu\nu}$ the Ricci tensor, R the Ricci scalar, $g_{\mu\nu}$ the metric tensor and $T_{\mu\nu}$ the stress-energy tensor. For a perfect fluid, this can be written as

$$T^{\mu\nu} = \rho h u^\mu u^\nu + p g^{\mu\nu}, \tag{1.18}$$

where ρ is the proper rest-mass density, u^μ the 4-velocity and p the fluid pressure. The specific enthalpy is

$$h = 1 + \varepsilon + \frac{p}{\rho}, \tag{1.19}$$

where ε is the specific internal energy (Font, 2008; Rezzolla and Zanotti, 2013; Martí and Müller, 2015). As the system we are studying operates over short timescales with respect to any non-rotational time variation of the metric (caused by e.g. significant changes in the mass or structure of the neutron star), we will be solving the field equations with the approximation that the metric is static or stationary (the *Cowling approximation* (Cowling, 1941)).

The relativistic fluid equations are given by the five conservation laws

$$\nabla_\mu J^\mu = 0, \tag{1.20}$$

$$\nabla_\nu T^{\mu\nu} = 0, \tag{1.21}$$

where $J^\mu = \rho u^\mu$ is the density current. By projecting the second of these two conservation laws with the projection operator $h^\mu{}_\nu = u^\mu u_\nu + \delta^\mu{}_\nu$, we can obtain the fluid equations in a form closer to the Newtonian equations (Gourgoulhon, 2006):

$$\nabla_\mu (\rho u^\mu) = 0, \tag{1.22}$$

$$u^\mu \nabla_\mu (\rho h - p) + \rho h \nabla_\mu u^\mu = 0, \tag{1.23}$$

$$\rho h u^\nu \nabla_\nu u_\mu + \nabla_\mu p + u_\mu u^\nu \nabla_\nu p = 0. \tag{1.24}$$

In numerical relativity, it is useful to split spacetime into one timelike and three spacelike coordinates. We do this using the ADM 3+1 formalism of Arnowitt et al. (1962). This is done by foliating the spacetime with a set of non-intersecting spacelike hypersurfaces $\Sigma = \Sigma(t)$, with each parametrised by a constant value of the timelike coordinate t.

The general 3+1 metric is

$$\mathrm{d}s^2 = -(\alpha^2 - \beta_i \beta^i)\mathrm{d}t^2 + 2\beta_i \mathrm{d}x^i \mathrm{d}t + \gamma_{ij} \mathrm{d}x^i \mathrm{d}x^j, \tag{1.25}$$

where $\alpha(t, x^j)$ is the lapse function, which measures the proper time between adjacent hypersurfaces, $\beta^i(t, x^j)$ is the shift vector, which measures the change of coordinates between hypersurfaces, and γ_{ij} is the spacelike 3-metric on the hypersurfaces (Alcubierre, 2008). The contravariant components of the 3-velocity, $v^i = \gamma^{ij} v_j$, are defined in terms

of the 4-velocity u^μ as

$$v^i = \frac{u^i}{\alpha u^0} + \frac{\beta^i}{\alpha}, \tag{1.26}$$

and the Lorentz factor is given by $W = \alpha u^0 = (1 - v^i v_i)^{-1/2}$ (Font, 2008; Martí and Müller, 2015).

Our choice of foliation Σ is not unique. We can measure the curvature of a 3d hypersurface Σ_t with respect to the 4d spacetime in which it is embedded using the extrinsic curvature, which can be defined as

$$K_{ij} = -\frac{1}{2}\mathcal{L}_{\boldsymbol{n}}\gamma_{ij} = -\frac{1}{2\alpha}\left(\partial_t - \mathcal{L}_{\boldsymbol{\beta}}\right)\gamma_{ij}, \tag{1.27}$$

where $n^\mu \equiv -\alpha\nabla^\mu t$ is the future-pointing timelike unit normal vector to the hypersurface.

1.5 Numerical methods

In this section, we shall outline some of the numerical methods that shall be used later in this thesis to implement the various numerical models. We shall first describe the Riemann problem, then give an overview of its use in finite volume methods.

For many techniques in numerical modelling, we approximate a system defined in a continuous domain by discretising it, covering the domain in a grid and calculating the solution only at a finite number of points in this grid. For a time-dependent system, we also discretise in time, calculating the solution only at certain points in time, separated by the timestep Δt. When referring to data defined at a specific timestep and location on a numerical grid, we shall use the upper index to denote the timestep and the lower index to denote the spatial location. For example, the density ρ at timestep n and grid cell (i, j) would be denoted by $\rho^n_{i,j}$. The location x at the centre of a cell is given by $x_{i,j}$, where the cell's domain is $[x_{i-1/2,j}, x_{i+1/2,j}] \times [x_{i,j-1/2}, x_{i,j+1/2}]$. We define $x_{i\pm1/2,j} = x_{i,j} \pm \Delta x_i/2$, where Δx_i is the grid spacing in the i-direction.

1.5.1 The Riemann problem

The Riemann problem is a special type of initial value problem consisting of a system of m hyperbolic partial differential equations (PDEs) of the form

$$\partial_t \boldsymbol{q} + A(\boldsymbol{q})\partial_x \boldsymbol{q} = 0, \tag{1.28}$$

where $A(\boldsymbol{q})$ is some $m \times m$ matrix of coefficients (Toro, 1999). As the system is hyperbolic, $A(\boldsymbol{q})$ is diagonalisable with m real eigenvalues and m linearly independent

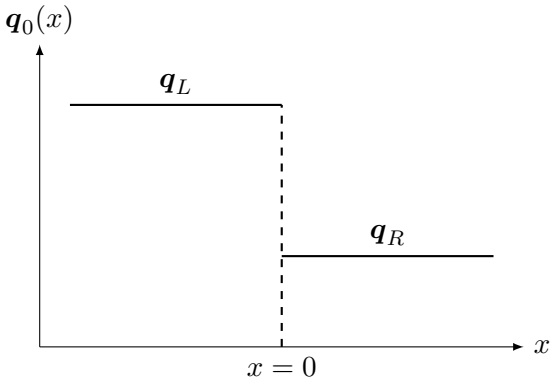

FIGURE 1.7: Initial data for the Riemann problem. At $t = 0$, the system consists of two states $\boldsymbol{q}_L$ and $\boldsymbol{q}_R$ separated by a discontinuity at $x = 0$.

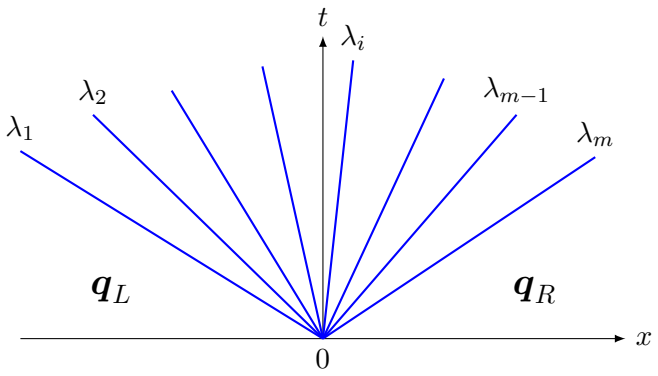

FIGURE 1.8: Structure of the solution of the Riemann problem for a system of m hyperbolic PDEs with constant coefficients. The solution consists of m characteristic curves propagating at constant speeds λ_i.

eigenvectors (LeVeque, 2002). This problem has the discontinuous initial data

$$\boldsymbol{q}(0, x) = \boldsymbol{q}_0(x) = \begin{cases} \boldsymbol{q}_L & \text{if } x \leq 0, \\ \boldsymbol{q}_R & \text{if } x > 0, \end{cases} \tag{1.29}$$

where $\boldsymbol{q}_L$ and $\boldsymbol{q}_R$ are some constant states separated by a discontinuity at $x = 0$. This is illustrated in Figure 1.7.

The solution of the Riemann problem consists of m characteristic curves[7] emanating from the origin, as shown in Figure 1.8. Across each characteristic is a discontinuity in the state propagating at constant speed. These constant speeds are the eigenvalues λ_i of A.

[7] *Characteristic curves* are curves $x = x(t)$, along which the PDE becomes an ODE such that the solution is constant along the curve.

1.5.1.1 Rankine-Hugoniot conditions

For a system of hyperbolic conservation laws,

$$\partial_t \boldsymbol{q} + \partial_x \boldsymbol{f}(\boldsymbol{q}) = 0, \tag{1.30}$$

where $\boldsymbol{q}$ is the conserved state vector and $\boldsymbol{f}(\boldsymbol{q})$ the flux, if the solution contains a discontinuity, then its speed of propagation v_S can be found using the *Rankine-Hugoniot conditions* (Toro, 1999; LeVeque, 1992). These are given by

$$[\![\boldsymbol{f}(\boldsymbol{q})]\!] = v_S [\![\boldsymbol{q}]\!], \tag{1.31}$$

where we define $[\![q]\!] \equiv q_R - q_L$ to be the difference in some quantity q to the right and left of the wave. In the case of a linear system where $\boldsymbol{f}(\boldsymbol{q}) = A\boldsymbol{q}$ and A is some matrix of constant coefficients, the Rankine-Hugoniot conditions become

$$A[\![\boldsymbol{q}]\!] = v_S [\![\boldsymbol{q}]\!]. \tag{1.32}$$

This is now an eigenvalue problem, and the speed of propagation of the discontinuity will therefore be an eigenvalue of the matrix of coefficients A, so $v_S = \lambda_i$.

1.5.1.2 Shocks, rarefactions & contact waves

The solution of the Riemann problem for an $m \times m$ hyperbolic system consists of $m+1$ constant states separated by m waves. For nonlinear systems such as the Euler equations, these waves may take the form of discontinuities such as shocks and contact waves, or smooth transitions such as rarefactions.

Shocks

A (right-going) shock wave occurs when the characteristic speeds in the left state are greater than those on the right, such that

$$\lambda(\boldsymbol{q}_L) > v_S > \lambda(\boldsymbol{q}_R), \tag{1.33}$$

where v_S is the speed of the shock. This is known as the *Lax entropy condition*. The converging shock characteristics can be visualised using Figure 1.9. Characteristics from either side of the initial discontinuity meet, forming a shock front propagating at speed v_S. The speed of this shock front can be found using the Rankine-Hugoniot conditions, as described in Section 1.5.1.1.

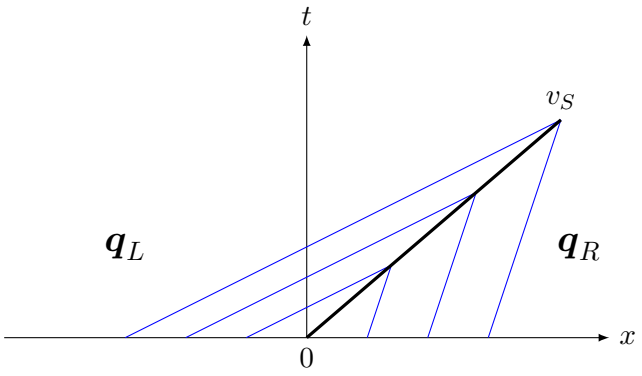

FIGURE 1.9: Converging characteristics of a shock. Characteristics (blue lines) from the initial left and right states meet, forming a shock front (thick black line) propagating at speed v_S.

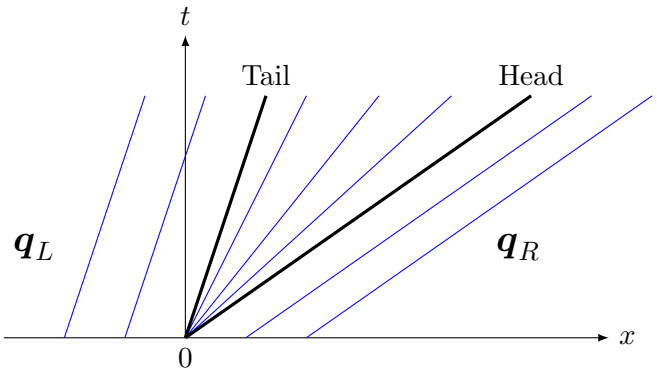

FIGURE 1.10: Diverging characteristics of a rarefaction. Characteristics (blue lines) from the initial left and right states move apart, producing a smooth transition region between them. The right edge of the rarefaction – the head – moves with speed $\lambda(\boldsymbol{q}_R)$, and the left edge – the tail – moves with speed $\lambda(\boldsymbol{q}_L)$.

Rarefactions

A (right-going) rarefaction occurs when the characteristic speeds in the left state are less than those on the right, such that

$$\lambda(\boldsymbol{q}_L) < \lambda(\boldsymbol{q}_R). \tag{1.34}$$

This results in the characteristics diverging, as illustrated in Figure 1.10. The initial constant states move apart, producing a smooth transition region (rather than a discontinuity) – the rarefaction wave. The right edge of this wave is known as the head of the rarefaction, travelling at $\lambda(\boldsymbol{q}_R)$. The left edge is the tail, travelling at $\lambda(\boldsymbol{q}_L)$.

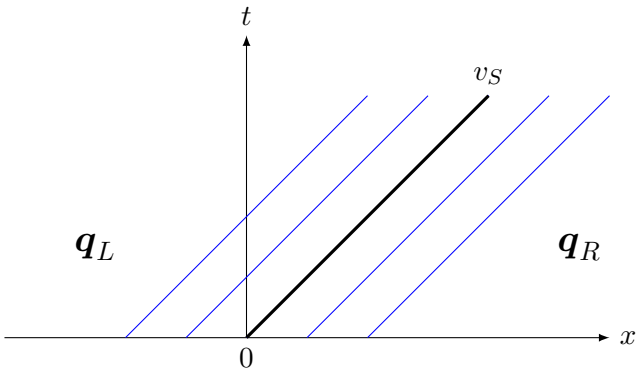

FIGURE 1.11: Parallel characteristics of a contact wave. Characteristics (blue lines) in both states run parallel to the wave (thick black line).

Contact waves

Contact waves are produced when there is a discontinuity in the initial data between the left and right states, but the characteristic speeds in the states are equal such that

$$\lambda(\boldsymbol{q}_L) = \lambda(\boldsymbol{q}_R) = v_S. \tag{1.35}$$

This is illustrated in Figure 1.11. The contact wave is parallel to the characteristics in the left and right states.

Detonations and deflagrations

Detonations and deflagrations are the counterparts of shocks and rarefactions respectively for the reactive Riemann problem. The initial left and right states now consist of the unburnt (or unreacted) material prior to the reaction and the burnt post-reaction material. Across detonations and deflagrations, a reaction takes place which may be endo- or exothermic. Consequently, there is an additional heat source across the wave. The left and right states may also have different equations of state. Detonations and deflagrations shall be described in more detail in Section 3.1.1.

1.5.2 Finite volume methods

In all systems modelled in this thesis, we shall be solving hyperbolic PDEs of the form given in (1.28) using finite volume methods.

Finite volume methods solve these systems by splitting the domain into a set of grid cells and approximating the total integral of the state over each cell. For a 1d domain split up into intervals, the average value $\bar{q}_i^n$ of some quantity $q(t^n, x)$ over the ith interval

at time t^n is approximated by q_i^n as

$$q_i^n \approx \bar{q}_i^n \equiv \frac{1}{\Delta x} \int_{x_{i-1/2}}^{x_{i+1/2}} \mathrm{d}x\, q(t^n, x), \tag{1.36}$$

where the interval length is given by $\Delta x = x_{i+1/2} - x_{i-1/2}$. If $q(t^n, x)$ is smooth, then this integral will always agree with q at the midpoint of the interval to $\mathcal{O}(\Delta x^2)$.

Note that in the rest of our discussion of finite volume methods in this section, for the sake of simplicity we shall continue to restrict ourselves to look at a 1d domain, however the schemes presented can also be extended to multidimensional domains.

Finite volume methods are among the most popular in computational fluid dynamics and in relativistic hydrodynamics. They offer many advantages over classical finite difference schemes, chief among them being that they are inherently conservative: for a system without source terms, the flux entering a given volume is always identical to that leaving the adjacent volume. This therefore makes them ideal for solving conservation laws, homogeneous hyperbolic equations that describe how a particular physical quantity remains constant in time over an entire physical domain. They also tend to be more robust for systems with discontinuities such as shocks.

However, they do have some disadvantages. High order finite volume schemes are much more complex to extend to the multidimensional case, and become particularly expensive in the general relativistic case where they require very high order schemes to calculate the metric and metric source terms (Radice, 2013). To obtain high order accuracy, they depend on reconstruction techniques with large stencils. They are difficult to implement for problems with complicated geometries and the treatment of boundary conditions can be non-trivial (Palaniappan et al., 2004).

In order to get a sufficiently accurate solution for our problem, we can use a MUSCL-type scheme (Monotonic Upstream-centred Scheme for Conservation Laws) (van Leer, 1979). This involves reconstructing the data to the cell boundaries and constraining this reconstruction so that spurious oscillations that will destabilise the solution are avoided.

1.5.2.1 Method of lines

The method of lines allows us to solve a system of PDEs by converting them into a set of ordinary differential equations (ODEs). This is done by discretising in terms of the spatial dimension(s) and leaving time continuous, such that for the PDE

$$\partial_t \boldsymbol{q}(t^n, x_i) + \partial_x \boldsymbol{f}(\boldsymbol{q}(t^n, x_i)) = 0, \tag{1.37}$$

we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t} \bar{\boldsymbol{q}}_i(t^n) + \frac{1}{\Delta x} \left[\boldsymbol{f}(\boldsymbol{q}(t^n, x_{i+1/2})) - \boldsymbol{f}(\boldsymbol{q}(t^n, x_{i-1/2})) \right] = 0, \tag{1.38}$$

where Δx is the grid spacing. To solve this we require the fluxes $\boldsymbol{f}(\boldsymbol{q})$, which can be approximated by the intercell fluxes $\boldsymbol{f}^n_{i\pm1/2} \approx \boldsymbol{f}(\boldsymbol{q}(t^n, x_{i\pm1/2}))$. This is done by first finding the approximation of the state vectors at the cell boundaries $\boldsymbol{q}^n_{i\pm1/2} \approx \boldsymbol{q}(t^n, x_{i\pm1/2})$, then solving the Riemann problem at the cell boundaries to get the intercell flux. The exact solution typically requires a computationally expensive root-finding operation, so approximate solvers are normally used instead. These will be described in more detail in Section 3.2.

1.5.2.2 Reconstruction

Given the approximate cell average state vector $\boldsymbol{q}_i$ given by (1.36), we need to find the state vectors at the cell boundaries, $\boldsymbol{q}_{i-1/2}$ and $\boldsymbol{q}_{i+1/2}$. We do this by reconstructing the state at the boundaries. We wish to do this in a way that minimises non-physical oscillations that will destabilise the code. This can be achieved by using *slope limiters*, as first introduced by van Leer (1979), such that the state at a point $x \in [x_i - \Delta x_i/2, x_i + \Delta x_i/2]$ is given by

$$\boldsymbol{q}(x) = \boldsymbol{q}_i + \boldsymbol{S}_i(x - x_i), \tag{1.39}$$

where $\boldsymbol{S}_i$ is the slope across cell i. The slope is calculated using a limiter function; this reduces the slope near a shock or discontinuity such that spurious oscillations are reduced.

The slope in cell i is given by

$$\boldsymbol{S}_i = \phi(r)\boldsymbol{S}^{\text{av}}_i, \tag{1.40}$$

where $\phi(r)$ is the limiter function, $\boldsymbol{S}^{\text{av}}_i$ is the average slope across the cell,

$$\boldsymbol{S}^{\text{av}}_i = \frac{1}{2}\left(\frac{\boldsymbol{q}_{i+1} - \boldsymbol{q}_i}{x_{i+1} - x_i} + \frac{\boldsymbol{q}_i - \boldsymbol{q}_{i-1}}{x_i - x_{i-1}}\right), \tag{1.41}$$

and r is the ratio of slopes either side of the cell,

$$r = \frac{\boldsymbol{q}_{i+1} - \boldsymbol{q}_i}{\boldsymbol{q}_i - \boldsymbol{q}_{i-1}}. \tag{1.42}$$

Slope limiters have the property that they are *total variation diminishing* (TVD) (LeVeque, 2002). If we define the total variation of the approximate cell average state vector $\boldsymbol{q}^n$ at time $t = t^n$ to be

$$\text{TV}(\boldsymbol{q}^n) = \sum_{i=-\infty}^{\infty} \left|\boldsymbol{q}^n_i - \boldsymbol{q}^n_{i-1}\right|, \tag{1.43}$$

then a method is total variation diminishing if for any set of data $\boldsymbol{q}^n$, the values $\boldsymbol{q}^{n+1}$ computed by the method satisfy

$$\text{TV}(\boldsymbol{q}^{n+1}) \leq \text{TV}(\boldsymbol{q}^n). \tag{1.44}$$

A consequence of this is that if data is initially monotonic, such that $\boldsymbol{q}_i^n \geq \boldsymbol{q}_{i+1}^n$ for all i, then it shall remain so for all time. This is known as the *monotonicity-preserving* property.

The choice of slope limiter $\phi(r)$ is important to the accuracy of the solution. If slopes are limited too much, then discontinuities will be smeared out as the simulation progresses. If slopes are not limited enough, non-physical oscillations will be able to develop and grow. Three possible functions for the slope limiter, going from most to least diffusive, are given by Toro (1999) as

- minmod
$$\phi = \max\left(0, \min\left(1, \frac{4}{1+r}\right)\right), \tag{1.45}$$
- monotised centred (MC)
$$\phi = \max\left(0, \min\left(\frac{2r}{1+r}, \frac{2}{1+r}\right)\right), \tag{1.46}$$
- and superbee
$$\phi = \begin{cases} \min\left(2, \min\left(r, \frac{2}{1+r}\right)\right) & r \geq 1, \\ 1 & 0.5 \leq r < 1, \\ 2r & 0 < r < 0.5, \\ 0 & \text{otherwise.} \end{cases} \tag{1.47}$$

An illustration of the slope limiting procedure using the minmod limiter can be seen in Figure 1.12. In the figure, it can be seen that the limited slopes are in general less steep than the gradient of the curve showing the 'real', continuous data. Where there are steep gradients in the original data, there are large jumps in the boundary data reconstructed from adjacent cells. As the simulation is evolved in time, this will lead to smearing of this feature.

1.5.2.3 Solution of the Riemann problem

Once the states have been reconstructed to the boundaries, the intercell flux is found by solving the Riemann problem with initial data given by

$$\boldsymbol{q}_L = \boldsymbol{q}^-_{+1/2}, \qquad \boldsymbol{q}_R = \boldsymbol{q}^+_{-1/2}, \tag{1.48}$$

where $\boldsymbol{q}^-$ is the reconstructed state of the neighbouring cell on the left and $\boldsymbol{q}^+$ is the reconstructed state of the neighbouring cell on the right. The intercell flux is then calculated using a numerical flux function. One such function is the Lax-Friedrichs flux,

$$\boldsymbol{f}_{\mathrm{LF}}(\boldsymbol{q}) = \frac{1}{2}\left(\boldsymbol{f}(\boldsymbol{q}_L) + \boldsymbol{f}(\boldsymbol{q}_R) + \alpha\left(\boldsymbol{q}_L - \boldsymbol{q}_R\right)\right), \tag{1.49}$$

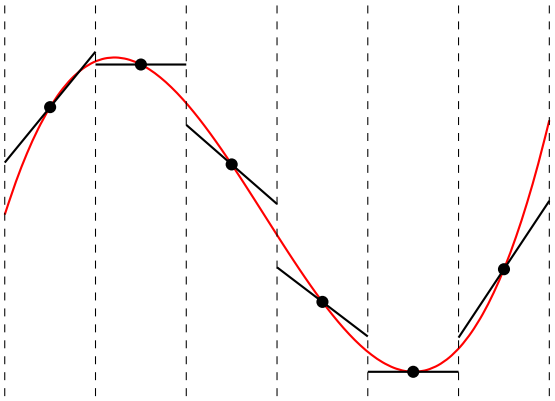

FIGURE 1.12: Illustration of slope limiters used to reconstruct boundary data. The red curve shows the 'real', continuous data, the black circles show the values of this data at the centres of the cells, and the black lines show the limited slopes calculated using the minmod slope limiter. The intersection of these lines with the cell boundaries (the vertical dashed lines) gives the reconstructed boundary data.

where α is the magnitude of the maximum wave speed across the whole domain (which in our relativistic simulations we shall set to $c = 1$ for simplicity). The Lax-Friedrichs flux is first order accurate and it is always guaranteed to converge to the correct solution as the grid is refined. However, it is also highly diffusive. This helps make the scheme more stable to spurious oscillations (LeVeque, 2002), however if higher accuracy is required, a higher order approximation for the intercell flux should be used.

1.5.2.4 Runge-Kutta methods

Once we have obtained the intercell fluxes, we will now be able to do the time integration. Using the method of lines, this can be done using Runge-Kutta methods, a family of implicit and explicit methods for the numerical solution of ODEs. For an ODE of the form

$$\frac{\mathrm{d}\boldsymbol{q}}{\mathrm{d}t} = \boldsymbol{f}, \tag{1.50}$$

the standard third-order SSP (strong stability preserving) Runge-Kutta method (RK3) is given by Gottlieb et al. (2009) as

$$\boldsymbol{q}^{(1)} = \boldsymbol{q}^n + \Delta t \boldsymbol{f}(\boldsymbol{q}^n), \tag{1.51}$$

$$\boldsymbol{q}^{(2)} = \frac{1}{4}\left(3\boldsymbol{q}^n + \boldsymbol{q}^{(1)} + \Delta t \boldsymbol{f}(\boldsymbol{q}^{(1)})\right), \tag{1.52}$$

$$\boldsymbol{q}^{n+1} = \frac{1}{3}\left(\boldsymbol{q}^n + 2\boldsymbol{q}^{(2)} + 2\Delta t \boldsymbol{f}(\boldsymbol{q}^{(2)})\right). \tag{1.53}$$

It is said that a method is strongly stable if there exists a value Δt_{FE} such that

$$\|\boldsymbol{q}^n + \Delta t \boldsymbol{f}(\boldsymbol{q}^n)\| \leq \|\boldsymbol{q}^n\| \quad \forall \boldsymbol{q}^n \text{ and for } 0 \leq \Delta t \leq \Delta t_{FE}. \tag{1.54}$$

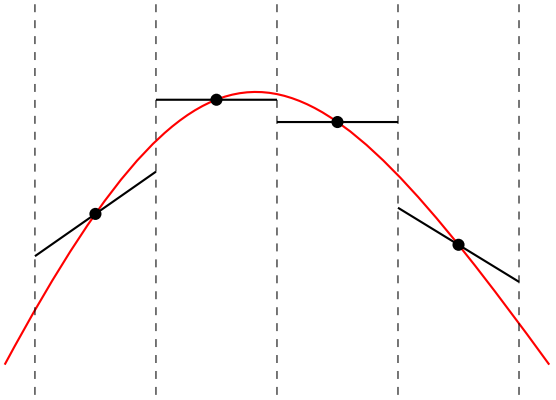

FIGURE 1.13: Accuracy of TVD methods at extrema (LeVeque, 2002). In order to remain TVD, slope limiters will set the slope to be zero in the region of extrema, which results in clipping of the peak value.

This is known as the *forward Euler condition*. Provided that this holds, a Runge-Kutta method is strong stability preserving if

$$\|\boldsymbol{q}^{n+1}\| \leq \|\boldsymbol{q}^{n}\| \quad \text{for } 0 \leq \Delta t \leq \mathcal{C}\Delta t_{FE} \tag{1.55}$$

for some coefficient $\mathcal{C} > 0$ (Hadjimichael et al., 2013).

SSP methods can be used to ensure the stability of hyperbolic PDEs with discontinuous solutions such as shocks.

As can be seen, for each timestep in the RK3 method it is necessary to calculate two intermediate states $\boldsymbol{q}^{(1)}$ and $\boldsymbol{q}^{(2)}$. This increases the computational cost of the method. Higher order methods will require the calculation of additional intermediate states.

1.5.2.5 Accuracy

In the previous sections, we have shown how we can use TVD methods in order to mitigate the development of non-physical spurious oscillations that may destabilise the simulation. However, by using these techniques, we will lose some accuracy in our solution. As can be seen in Figure 1.13, at extrema slope limiters will produce a slope of zero in order to ensure the method is TVD. Any other choice will produce reconstructed data $\tilde{\boldsymbol{q}}(t^n)$ with greater total variation than the real data, $\mathrm{TV}\left(\tilde{\boldsymbol{q}}(t^n)\right) > \mathrm{TV}(\boldsymbol{q}^n)$, which can then result in an increase in the total variation at the next timestep, $\mathrm{TV}(\boldsymbol{q}^{n+1}) > \mathrm{TV}(\boldsymbol{q}^n)$. However, setting the slope to zero leads to *clipping* of the solution, such that at each timestep the values at extrema will be reduced by $\mathcal{O}(\Delta x^2) = \mathcal{O}(\Delta t^2)$. This error grows as the simulation progresses, with the consequence that after $T/\Delta t$ timesteps it will produce a global error in the region of the extrema that is $\mathcal{O}(\Delta t)$. Even if the method is higher order, its accuracy will be reduced to first order near the extrema.

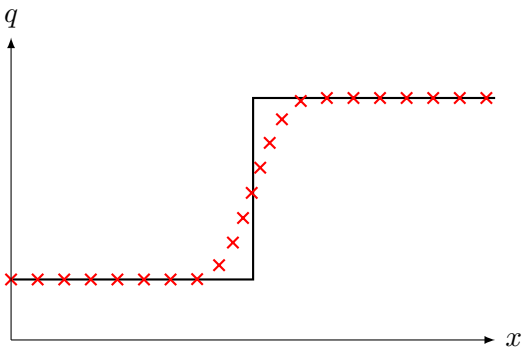

FIGURE 1.14: Accuracy of TVD methods at discontinuities. Slope limiters reduce the gradients in the reconstructed data, leading to smearing of discontinuities. In this example, the exact solution contains a discontinuity (black line). In the numerical solution calculated using slope limiters (red crosses), this has been smeared out across multiple cells.

It is possible to remedy this by relaxing the TVD property to allow for the total variation to slightly increase so that smooth peaks can be better represented. This must be done in a way that allows for small increases in the total variation around extrema while still preventing the growth of spurious oscillations. Essentially nonoscillatory (ENO) methods are one of such approach to this, and can obtain high accuracy even in the region of extrema.

As mentioned above, using slope limiters also leads to errors at discontinuities. As can be seen in Figure 1.12, the slope is limited the most where the gradient of the data becomes especially steep. This will lead to the reduction of the gradient of the reconstructed data around the discontinuity, so that over time the discontinuity will become smeared out across multiple cells, as shown in Figure 1.14. The choice of limiter is important here: for this example, we used the minmod limiter. If we had instead used a less diffusive limiter such as MC or superbee, the gradients of the reconstructed slopes would have been allowed to be greater, leading to smaller jumps in the reconstructed data at the boundary and therefore less smearing of the discontinuity in the solution.

1.5.2.6 Conservative schemes

The differential equation

$$\partial_t \boldsymbol{q}(t,x) + \partial_x \boldsymbol{f}\left(\boldsymbol{q}(t,x)\right) = 0 \tag{1.56}$$

is not valid for solutions containing discontinuities such as shocks, as it only holds if $\boldsymbol{q}(t,x)$ and $\boldsymbol{f}\left(\boldsymbol{q}(t,x)\right)$ are smooth functions. However, the original integral form of the conservation law from which this differential equation was derived does still hold when discontinuities are present. To recover this, we shall first integrate (1.56) in space

between x_1 and x_2 and in time between t_1 and t_2 to give

$$\int_{t_1}^{t_2} \mathrm{d}t \int_{x_1}^{x_2} \mathrm{d}x \left[\partial_t \boldsymbol{q} + \partial_x \boldsymbol{f}(\boldsymbol{q})\right] = 0. \tag{1.57}$$

We can make this more general by introducing the test function $\phi(t, x)$ such that the integral becomes

$$\int_{0}^{\infty} \mathrm{d}t \int_{-\infty}^{\infty} \mathrm{d}x \left[\partial_t \boldsymbol{q} + \partial_x \boldsymbol{f}(\boldsymbol{q})\right] \phi(t, x) = 0, \tag{1.58}$$

which reduces to (1.57) if

$$\phi(t, x) = \begin{cases} 1 & \text{if } x \in [x_1, x_2] \times [t_1, t_2], \\ 0 & \text{otherwise.} \end{cases} \tag{1.59}$$

If we assume $\phi(t, x)$ to be some smooth function with compact support[8], then we can integrate by parts to get

$$\int_{0}^{\infty} \mathrm{d}t \int_{-\infty}^{\infty} \mathrm{d}x \left[\boldsymbol{q}\partial_t \phi + \boldsymbol{f}(\boldsymbol{q})\partial_x \phi\right] = -\int_{-\infty}^{\infty} \mathrm{d}x \, \boldsymbol{q}(0, x)\phi(0, x). \tag{1.60}$$

This is known as the *weak form* of the equations. It no longer contains derivatives of $\boldsymbol{q}(t, x)$ and $\boldsymbol{f}(\boldsymbol{q}(t, x))$, so will still hold if one or both of them are discontinuous. The function $\boldsymbol{q}(t, x)$ with the initial data $\boldsymbol{q}(0, x)$ is considered a *weak solution* of the differential equation (1.56) if its weak form (1.60) holds for all test functions $\phi(t, x)$ that are continuously differentiable and have compact support.

Finite volume schemes are based on the integral form of the conservation law. As described above, using this form of the equations ensures that the scheme is conservative. If a conservative scheme is used, then the computed results are guaranteed to converge to a weak solution of the system; this is not true of nonconservative schemes for systems where discontinuities are present. This property of conservative schemes was first shown by Lax and Wendroff (1960), and is known as the *Lax-Wendroff theorem*. It states that if the solution computed with a consistent and conservative method converges to a function of $\boldsymbol{q}(t, x)$ as the resolution is increased, then $\boldsymbol{q}(t, x)$ is a weak solution of the conservation law.

This property is illustrated in Figure 1.15, where we show the location of a shock front computed using conservative and nonconservative schemes. For a conservative scheme, even if the computed solution is not particularly accurate due to e.g. clipping, the shock will still be centred about the correct location. A nonconservative scheme will not compute the correct shock speed, and so the computed location of the shock will not be correct even as the resolution is increased.

[8] The function $\phi(t, x)$ has *compact support* if it vanishes outside of some bounded region of the t-x plane.

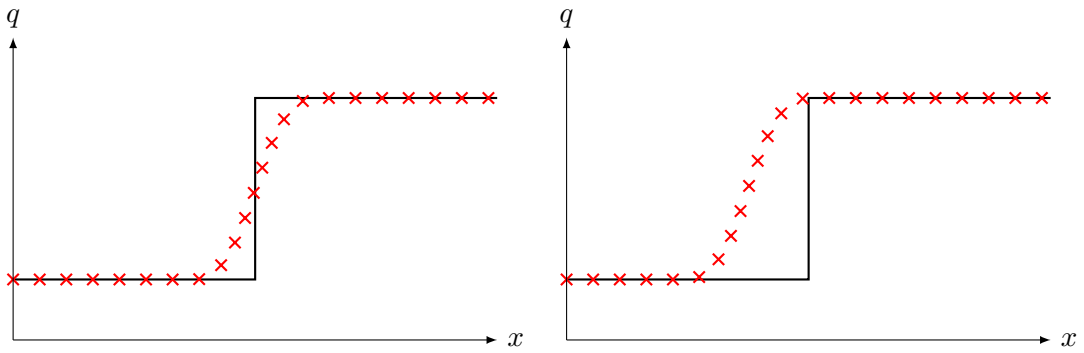

(A) Conservative (B) Nonconservative

FIGURE 1.15: Conservative schemes are important when modelling systems with discontinuities, as from the Lax-Wendroff theorem they are guaranteed to converge to a weak solution of the conservation law. This is not the case for nonconservative schemes. In the plots, the black line shows the real, discontinuous solution, and the red crosses show the computed solution. If a nonconservative scheme is used, the solution will not converge to a weak solution if there are discontinuities such as shocks present. Such a scheme will not calculate the correct shock speeds, such that the computed location of a shock will be incorrect even as the resolution of the simulation is increased.

1.5.2.7 Well-balanced schemes

For many of the systems we will be investigating in this thesis, the system will remain close to hydrostatic equilibrium (or the relativistic equivalent) throughout its evolution. The modelling of flows in systems such as these is challenging for standard finite volume schemes, as the discrete form of the equations used to evolve the system typically does not satisfy the discrete form of the equation for hydrostatic equilibrium (Käppeli and Mishra, 2016). Consequently, equilibrium states will not be preserved, and instead an error will be produced proportional to the truncation error of the scheme. In order to capture small perturbations of a steady state solution, the numerical resolution may have to be significantly increased (thereby making the simulation much more computationally expensive) such that the truncation errors are small enough not to obscure the features of interest.

One way that we can ameliorate this issue is by using a *well-balanced scheme*. First introduced by Greenberg and Leroux (1996), well-balanced schemes satisfy exactly the discrete form of the equations for hydrostatic equilibrium. A well-balanced scheme for the Euler equations with a gravitational source term can be implemented by modifying the reconstruction of the pressure at the cell interfaces (Käppeli and Mishra, 2016). Standard reconstruction of the pressure at cell interfaces is incompatible with the pressure stratification produced by a gravitational field. As the pressure profile produced by a gravitational field is typically not a polynomial, any reconstruction based on piecewise polynomials will not be able to accurately represent hydrostatic equilibrium. This causes pressure differences at the cell interfaces, which are interpreted by the Riemann solver as propagating waves. These waves may then go on to destroy the equilibrium state of the

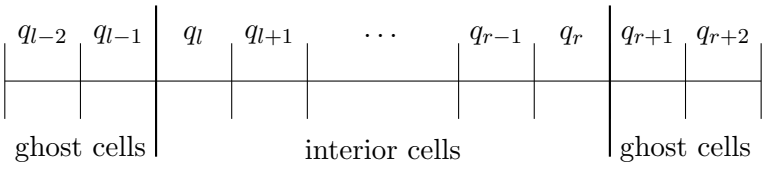

FIGURE 1.16: In order to calculate an update for the data at the edge of the domain, we can extend the grid to include ghost cells. Here, we have extended the 1d domain (with data at the left and right edges given by q_l and q_r) by adding two ghost cells outside of each of the boundaries.

system. The well-balanced scheme instead reconstructs the pressure using the discrete form of the equation for hydrostatic equilibrium, ensuring that this is always satisfied and so hydrostatic equilibrium is maintained.

1.5.2.8 Boundary conditions

As has been described in previous sections, in order to update a cell's data using finite volume methods, we require the data in neighbouring cells. This becomes non-trivial for cells at the edge of the domain, where the neighbouring cells would lie outside of the simulation's physical domain. The solution is to surround the physical grid by a region of *ghost cells*, whose values are calculated from data within the physical domain (the *interior cells*) by imposing boundary conditions. This is illustrated in Figure 1.16 for a 1d domain. The number of ghost cells required depends on the stencil of the numerical methods used in a cell's update, i.e. how many neighbouring cells are used in the calculation of a cell's data over a single timestep.

The boundary conditions can take a variety of forms – the choice of which to use depends on the nature of the problem to be solved. Some examples of commonly used boundary conditions in computational fluid dynamics are:

- *Reflective* – This is typically used for domains with a solid outer boundary, e.g. a wall. Incoming waves are reflected by the boundary, such that the data in the ghost cells is the mirror image of the data within the interior domain. The component of the velocity normal to the boundary is reversed. To illustrate this, consider the left boundary of the 1d domain illustrated in Figure 1.16. Let the state in the interior cell just inside the boundary be q_l (except the velocity, which we denote v_l), then the data in the ghost cell $i \geq 1$ cells away from the boundary is given by

$$q_{l-i} = q_{l+i-1}, \qquad v_{l-i} = -v_{l+i-1}. \tag{1.61}$$

- *Outflow* – Often used for domains with a free outer boundary, incoming waves simply flow straight through the boundary and leave the domain without reflection. The simplest form of outflow boundary conditions is zero order extrapolation: data

in the ghost cells is copied from the outermost cell in the interior region. Again considering the left edge of the 1d domain, the data in the ghost cells would be

$$q_{l-i} = q_l. \tag{1.62}$$

Higher order extrapolation may be used, however this can lead to stability problems (LeVeque, 2002).

- *Inflow* – Used to model a flow into the domain. This could be e.g. a flow of material or a flow of heat from some source outside the domain. Inflow boundary conditions are typically calculated by applying some function to the data at the edge of the interior domain.
- *Periodic* – Used to model (effectively) infinite domains. The data in the ghost cells on one edge of the domain is copied over from interior data on the opposite edge. Defining the interior cell just within the right boundary to be q_r, the ghost cells outside the left boundary would be

$$q_{l-i} = q_{r-i+1}. \tag{1.63}$$

More complex boundary conditions can be used, for example data can be reconstructed using the reconstruction techniques described in Section 1.5.2.2. For certain systems, it makes sense to impose physical boundary conditions, e.g. for fluid flowing through a pipe, 'non-slip' boundary conditions can be used, where the velocity tangential to the boundary is set to zero.

For most of the tests below, we shall be using outflow or periodic boundary conditions. However, for systems with a gravitational field, the case becomes more complex for the boundaries in the vertical direction. As discussed in the previous section, gravitational fields typically produce a non-linear pressure profile. If standard outflow boundary conditions or piecewise linear reconstruction are used, then the pressure at the bottom boundary will be underestimated and the pressure at the top boundary will be overestimated. This can lead to the development of unphysical oscillations at the boundaries and can cause the material in the simulation to gradually sink over time. This can be prevented by explicitly enforcing hydrostatic equilibrium in the ghost cells.

1.5.3 Codes in this thesis

Over the course of this thesis, a number of different codes shall be presented and discussed in detail. These codes have been designed to work at different scales and in different physical regimes, e.g. for low Mach number flows, for relativistic fluids, for shallow water environments. These regimes are illustrated in Figure 1.17. It is worth noting that the regimes illustrated are very approximate and are meant to demonstrate

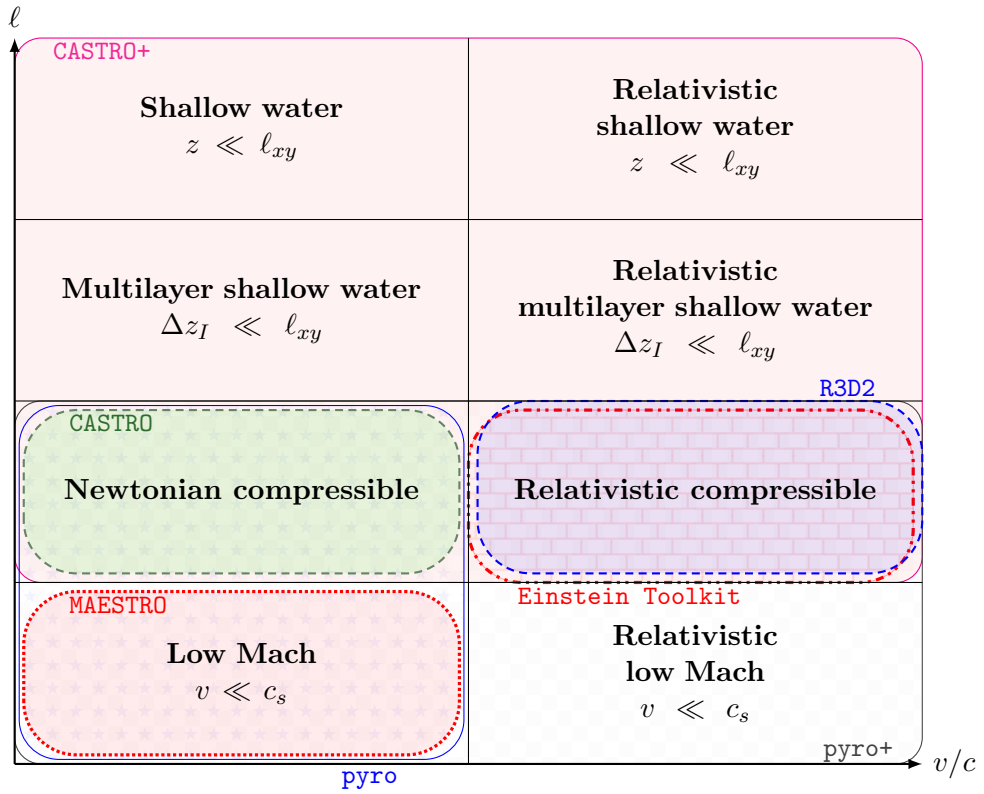

FIGURE 1.17: Regime diagram for the codes used in this thesis, demonstrating their regions of validity. The axes are the characteristic speed v as a fraction of the speed of light c, and the characteristic length scale ℓ. Note that the lengths here are not to scale, but are merely meant to roughly illustrate the various regions.

Indicated on the plot are some extra constraints for the models to be valid: the low Mach number models require that the characteristic speed is much less than the speed of sound c_s; the multilayer shallow water models require that each layer thickness Δz_I is much less than the characteristic horizontal length scale ℓ_{xy}; and the single layer shallow water models require that the fluid height z is much less than the characteristic horizontal length scale ℓ_{xy}.

In the case of `CASTRO` and `pyro`, two versions of the codes are shown: the original Newtonian codes, and the extended relativistic codes (labelled `CASTRO+` and `pyro+`).

where the codes work most effectively. For example, many of the relativistic codes will work for Newtonian systems, however it would be inefficient to use them there.

The codes that shall be discussed are as follows:

- `pyro` (Zingale, 2014) – A Python-based code for computational astrophysical hydrodynamics. It is used in Section 2.3 for modelling the low Mach number equations, where it is extended to model relativistic fluids. In the diagram above, we label this version as `pyro+`.
- The `Einstein Toolkit` (Löffler et al., 2012) – A computational infrastructure for relativistic astrophysics. It is used in Section 2.3 to compare low Mach number solutions against compressible solutions.

- `AMReX` (Zhang et al., 2016; Zingale et al., 2017) – A software framework for massively parallel, block-structured adaptive mesh refinement (AMR) applications, used by `MAESTRO` and `CASTRO`.
- `MAESTRO` (Almgren et al., 2007) – A low Mach number stellar hydrodynamics code. An extension of the Newtonian low Mach number solver in `pyro` to model relativistic fluids incorporated much of the algorithm used in this code.
- `R3D2` (Harpole and Hawke, 2016) – Reactive relativistic Riemann problem solver for detonations and deflagrations. Its development and application is described in Chapter 3.
- `CASTRO` (Almgren et al., 2010) – An adaptive mesh, astrophysical radiation hydrodynamics simulation code. The shallow water model in Chapter 4 and the multiscale model developed in Chapter 5 were implemented by extending `CASTRO`, labelled in the diagram above as `CASTRO+`.

1.6 Summary

As described in Section 1.3, X-ray bursts display a wide range of features, many of which cannot be fully explained using current models. Part of the challenge in modelling them is due to the extreme nature of the environment in which they occur: neutron star oceans have very high pressure, density and temperature and experience very strong gravitational and magnetic fields. They are therefore very different to any fluid we encounter here on Earth, and so it is challenging to determine which physical processes dominate.

Despite the fact that magnetic fields are likely to be important in the system we are looking at (as described in Section 1.3.10), we have chosen not to include them in the models presented here as this would make our calculations significantly more complex. The magnetic field of neutron stars is still poorly understood (Belvedere et al., 2015; Glampedakis and Lasky, 2016), so care would need to be taken in selecting a physically reasonable model for the field. A simple approximation would be to take the approach used by Cavecchi et al. (2016) in their study of X-ray bursts and model the initial magnetic field as purely vertical with some homogeneous value, though it would also be of interest to consider more complex configurations. Another difficulty is that there is no existing derivation of the low Mach number equations in Newtonian magnetohydrodynamics (MHD), so there would be no independent work against which we could test the Newtonian limit of our relativistic version. It is likely that the presence of Alvén waves in MHD may introduce some extra complexity for low Mach number MHD if these waves travel at velocities comparable to the slower flame propagation speed. In this case, care would need to be taken to ensure that these waves are not filtered out along with the sound waves.

In this thesis, we shall be concentrating on the strong gravitational field to see if general relativistic effects can have a significant effect on the propagation of the burning front. To do this, we shall develop a series of models valid on a range of different scales, with the aim to combine these models to produce a multiscale model capable of capturing both the large scale and small scale physics in a computationally efficient way.

We shall begin in Chapter 2 by extending the low Mach number approximation, used to model X-ray bursts in Newtonian physics by (Nonaka et al., 2010), to general relativistic fluids. We shall then implement these new equations numerically and test against existing Newtonian results to measure what effects our general relativistic model may have. In Chapter 3, we shall investigate relativistic detonations and deflagrations, implementing a solver for the reactive relativistic Riemann problem, `R3D2`. As described in Section 1.3.9, the large scale motion of the burning front is believed to be strongly affected by the Coriolis force. To capture this, in Chapter 4 we shall consider the shallow water equations, developing a relativistic formulation of these and implementing them numerically. Finally, in Chapter 5 we shall discuss mesh refinement methods and their extension to incorporate multiple numerical schemes and physical models. We shall then use this approach to combine the shallow water and compressible models to produce a multiscale model.

2 Relativistic low Mach number equations

When evolving hyperbolic systems of differential equations numerically, the size of the timestep, Δt, is limited by the Courant-Friedrichs-Lewy (CFL) limit (Courant et al., 1967):

$$\left|\frac{v\Delta t}{\Delta x}\right| \leq C_{\max}, \tag{2.1}$$

where Δx is the spatial resolution, v is the characteristic velocity of the system and $C_{\max}$ is the Courant number. This is typically less than one, but it can be greater than 100 for some implicit schemes. This limit is simply a result of causality: a scheme cannot be evolved such that information travels between timesteps at a speed greater than v. If this limit is violated, then spurious oscillations may develop in the solution, leading to instabilities. A consequence of the CFL limit is that if the characteristic velocity of a system is very high, it is very computationally expensive to get good spatial resolution (small grid spacing Δx) as this would require the timestep to be very small to compensate.

In neutron star oceans, the characteristic velocity is the sound speed. This is much higher than typical terrestrial sound speeds due to the high pressure and density, and so good spatial resolution is not possible using conventional methods without significant computational effort. We can illustrate this difficulty by calculating the Mach number $\mathrm{Ma} = v/c_s$, the ratio of the speed of propagation of the burning front v and the sound speed c_s. Simulations by Spitkovsky et al. (2002); Cavecchi (2013) found propagation speeds of $v \sim 10^5 - 10^6\,\mathrm{cm\,s^{-1}}$. Taking the pressure on the surface to be $p \sim 10^{22}\,\mathrm{erg\,cm^{-3}}$ and the density to be $\rho \sim 10^5\,\mathrm{g\,cm^{-3}}$ (Strohmayer and Bildsten, 2003), this gives a sound speed of

$$c_s \sim \sqrt{\frac{p}{\rho}} \sim \sqrt{\frac{10^{22}\ \mathrm{erg\,cm^{-3}}}{10^5\,\mathrm{g\,cm^{-3}}}} \sim 10^{8.5}\,\mathrm{cm\,s^{-1}}. \tag{2.2}$$

Consequently, the Mach number of the burning fronts is

$$\mathrm{Ma} = \frac{v}{c_s} \sim 10^{-3} \ll 1. \tag{2.3}$$

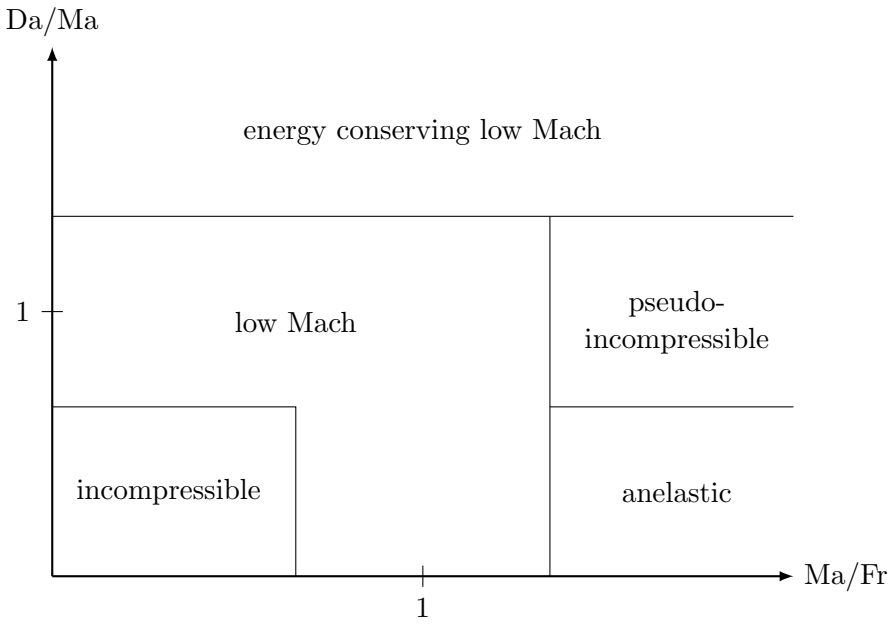

FIGURE 2.1: Regime diagram for sound-proof models demonstrating their approximate regions of validity. Along the horizontal axis, we plot the ratio of the Mach number to the Froude number. Along the vertical axis, we plot the ratio of the Damköhler number to the Mach number.
The Froude number $\mathrm{Fr} = \frac{u}{\sqrt{g\ell}}$ is the ratio of the flow inertia and the gravitational field. The Damköhler number is given by $\mathrm{Da} = \frac{\omega\ell}{u}$, where ω is the net production rate per unit mass, u the characteristic velocity and ℓ the characteristic length scales, and is the ratio of the characteristic flow time and the characteristic reaction time.

This motivates the use of sound-proof models, where we expand the hydrodynamic equations in order of the Mach number and take the limit as $\mathrm{Ma} \to 0$. By doing this, it is possible to effectively filter out the sound waves, allowing us to treat the much slower propagation speed of the burning front as the characteristic velocity of the system. This allows the use of much larger time steps in simulations, and so significantly better resolution and considerable computational speed up over conventional compressible schemes can be achieved. Several methods for doing this shall be described in the next section.

2.1 Sound-proof models

Over the years, there have been a number of different approaches to deriving formulations of the Newtonian hydrodynamics equations which filter out the sound waves (Nonaka et al., 2014; Kupka and Muthsam, 2017). A diagram illustrating the relationships between these models can be seen in Figure 2.1.

Incompressible

The simplest low Mach number approximation is incompressible hydrodynamics: the limit of the hydrodynamics equations as $\mathrm{Ma} \to 0$. This was first derived by Boussinesq (1903) and subsequently used by Rayleigh (1916) to study the onset of convective instability. It is valid provided that the vertical dimension of the fluid is much less than the scale height and that the relative horizontal variations of physical quantities are $\ll 1$ (Spiegel and Veronis, 1960). A result of this is that velocities in the problem must be small compared to the sound speed. In this limit, the velocity satisfies the constraint

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0, \tag{2.4}$$

such that the flow is instantaneously equilibrated and all sound waves are filtered out. This constraint also leads to the density being constant along particle paths, such that the system is incompressible:

$$\frac{\mathrm{D}p}{\mathrm{D}t} = 0. \tag{2.5}$$

Anelastic

In the anelastic approximation, as first derived by Batchelor (1953); Charney and Ogura (1960); Ogura and Phillips (1962); Gough (1969), the density is allowed to vary as small perturbations to a static hydrostatic background density ρ_0 which satisfies the constraint

$$\boldsymbol{\nabla} \cdot (\rho_0 \boldsymbol{u}) = 0. \tag{2.6}$$

As in the incompressible case, relative variations of physical quantities and the Mach number must be small. It is also required that the flow does not exhibit features that occur on timescales shorter than the crossing time of the domain. The anelastic approximation is not a good model for rapidly rotating systems, where the large Coriolis force can cause the time derivative of the density perturbation $\partial_t \rho'$ to become large (Calkins et al., 2015).

Low Mach number

The low Mach number approximation can be found by taking a low Mach number expansion of the fluid equations. It was first introduced by Rehm et al. (1978) and later rigorously derived from low Mach number asymptotic analysis by Klainerman and Majda (1982); Andrew and Sethian (1985). It was subsequently extended to stellar environments by Day and Bell (2000). In the approximation, the pressure is decomposed into a thermodynamic base state pressure, p_0, and a dynamic perturbational pressure,

π:

$$p(t,x) = p_0 + \pi(t,x), \tag{2.7}$$

where $\pi/p_0 = \mathcal{O}(\mathrm{Ma}^2)$. This decouples the pressure and density and filters out sound waves. It is required that the total pressure remains close to the background pressure, which is assumed to be constant. Unlike in the previous models, large variations in the density and temperature are permitted. The velocity constraint resembles that of the incompressible approximation (2.4), but now includes a source term which represents local compressibility due to energy generation and thermal diffusion:

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = S. \tag{2.8}$$

Unfortunately, as was pointed out by Guillard et al. (1999); Guillard and Murrone (2004), whilst this approximation holds for the continuous equations, for the discrete equations the pressure fluctuations scale with the Mach number rather than its square. Consequently, as the Mach number decreases results become worse and solutions of the low Mach number equations do not converge to the incompressible solution. This can be rectified by using a suitable preconditioning matrix, as shown by Guillard et al. (1999), or by using predictor corrector methods, as used by Day and Bell (2000).

Pseudo-incompressible

In the low Mach number approximation, the background pressure is assumed constant. This assumption will break down when modelling fluids in large enough domains that there is significant stratification. This can be remedied by using pseudo-incompressible (PI) methods, where the background pressure is allowed to vary (Durran, 1989). It is required that the Mach number and the normalised horizontal pressure perturbations are small, however no such restriction is made of the normalised horizontal temperature and density perturbations. The PI approximation gives the velocity constraint

$$\boldsymbol{\nabla} \cdot \left(p_0^{1/\gamma} \boldsymbol{u} \right) = H, \tag{2.9}$$

where γ is the adiabatic index and H is a source term representing heating effects.

2.1.1 Newtonian low Mach number equations

The pseudo-incompressible equations can be developed further to allow to for a non-trivial equation of state (which was assumed in Durran (1989)'s original derivation). These are (somewhat confusingly) also known as the low Mach number equations, though they are not the same as those described in the previous section. Following the approach used in the papers that describe the Newtonian low Mach number `MAESTRO` code

(Almgren et al., 2006b,a; Nonaka et al., 2010; Jacobs et al., 2016), the following set of equations can be derived:

$$\frac{\partial \rho}{\partial t} = -\boldsymbol{\nabla} \cdot (\rho \boldsymbol{u}) \,, \tag{2.10}$$

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\boldsymbol{u} \cdot \boldsymbol{\nabla} \boldsymbol{u} - \frac{\beta_0}{\rho} \boldsymbol{\nabla} \left(\frac{\pi}{\beta_0} \right) - \frac{(\rho - \rho_0)}{\rho} g \boldsymbol{e}_r, \tag{2.11}$$

$$\frac{\partial (\rho h)}{\partial t} = -\boldsymbol{\nabla} \cdot (\rho h \boldsymbol{u}) + \frac{\mathrm{D} p_0}{\mathrm{D} t} + \rho H_{\mathrm{nuc}} + \rho H_{\mathrm{ext}}, \tag{2.12}$$

where h is the specific enthalpy and the source terms H_{nuc} and H_{ext} describe the external heating rate and the nuclear energy generation rate per unit mass. The pressure has been decomposed into a hydrostatic base state pressure $p_0 = p_0(t, r)$ and a dynamic pressure perturbation $\pi = \pi(t, x^i)$, such that $|\pi|/p_0 = \mathcal{O}(\mathrm{Ma}^2)$. The base state density $\rho_0 = \rho_0(t, r)$ is defined by the equation for hydrostatic equilibrium of the base state pressure p_0, $\boldsymbol{\nabla} p_0 = -\rho_0 g \boldsymbol{e}_r$, where $\boldsymbol{e}_r$ is the unit vector in the outwards radial direction.

The buoyancy term in the momentum equation, $\frac{\beta_0}{\rho} \boldsymbol{\nabla} \left(\frac{\pi}{\beta_0} \right)$, is derived from a constraint that explicitly conserves energy, as proposed by Klein and Pauluis (2012) and by Vasil et al. (2013) in their 'generalised pseudo-incompressible' (GPI) model. The equation of state is cast into an elliptic constraint on the velocity field,

$$\boldsymbol{\nabla} \cdot (\beta_0 \boldsymbol{u}) = \beta_0 \left(S - \frac{1}{\bar{\Gamma}_1 p_0} \frac{\partial p_0}{\partial t} \right), \tag{2.13}$$

where β_0 is a density-like variable that captures the expansion of a fluid parcel as it rises due to the background stratification, defined as

$$\beta_0(t, r) = \rho_0(t, 0) \exp \left(\int_0^r \mathrm{d}r' \, \frac{1}{\bar{\Gamma}_1 p_0} \frac{\partial p_0}{\partial r'} \right). \tag{2.14}$$

$\bar{\Gamma}_1$ is the lateral average of $\Gamma_1 = \mathrm{d} \log p / \mathrm{d} \log \rho|_s$, and the expansion term, S, incorporates local compressibility effects due to heat release from reactions and external sources,

$$S = \sigma H_{\mathrm{nuc}} + \sigma H_{\mathrm{ext}}, \tag{2.15}$$

where $\sigma \equiv p_T / (\rho c_p p_\rho)$, $p_T \equiv \partial p / \partial T|_\rho$, $p_\rho \equiv \partial p / \partial \rho|_T$ and $c_p \equiv \partial h / \partial T|_p$.

The equations themselves do not enforce that the Mach number remains small. If the Mach number does become $\mathcal{O}(1)$ over the course of a numerical simulation, the solution to the equations will no longer be physically meaningful (Nonaka et al., 2010).

In the `MAESTRO` papers, further equations are derived for the base state and perturbational variables by splitting them in a similar way to the pressure (Almgren et al., 2006a), e.g. for the velocity

$$\boldsymbol{u} = w_0(t, r) \boldsymbol{e}_r + \tilde{\boldsymbol{u}}(t, x^i), \tag{2.16}$$

where the average of the radial component of the velocity perturbation $\overline{(\tilde{\boldsymbol{u}} \cdot \boldsymbol{e}_r)} = 0$ and the base state velocity is defined to be the average of the radial component of the velocity, $w_0 = \overline{(\boldsymbol{u} \cdot \boldsymbol{e}_r)}$.

In the rest of this thesis, we shall refer to (2.10)–(2.12) with the constraint (2.13) as the (Newtonian) low Mach number equations.

2.1.2 Alternative sound-proof methods

In this section, we have described how sound-proof models can be implemented by changing the equations of hydrodynamics to filter out sound waves. This is the technique we are going to use in the rest of this work. However, as described in the review of Kupka and Muthsam (2017), it is also possible to take an alternative approach when modelling low Mach number systems. Rather than developing new approximations of the hydrodynamics equations, these alternative methods solve the original equations, which means they have the advantage that solvers can be built using the vast quantity of existing numerical methods developed for the Euler and Navier-Stokes equations.

One such method is to change the model parameters so that the system is no longer in the low Mach number regime, while preserving the quantities of interest. This has been used by Wang et al. (2015) in their compressible `CHORUS` code to model thermal convection in stellar and planetary interiors, where they scale up the luminosity to achieve higher Mach numbers whilst leaving the Rossby number unchanged. Doing this made it computationally feasible to model low Mach number flows of Ma $\sim 10^{-2}$, however the simulation was still much more computationally expensive than an equivalent run using the anelastic `ASH` code.

Another approach is to instead adapt the numerical methods used to solve the hydrodynamics equations so that the timestep is no longer restricted by the strict CFL condition for low Mach number flows. Such methods include preconditioning (Miczek et al., 2014) and implicit methods (Viallet et al., 2011). These methods can be very powerful, however are typically more difficult to implement than standard schemes.

2.2 Relativistic low Mach number equations

Our aim is to construct the relativistic equivalent of the low Mach number equations described in the previous section, and put them in a form which allows for a numerical evolution scheme to be deduced. We have chosen the low Mach number approach as it is well suited to the system we wish to study and has previously been used in the Newtonian case to study type I X-ray bursts. Alternative relativistic sound-proof models exist, however these neglect physics which we believe to be important for the

system. For example, the traditional approximation of Maniopoulou and Andersson (2004) neglects the horizontal component of the Coriolis vector, so would not be able to account for the effects of the Coriolis force on the propagation of the burning front. The relativistic anelastic approximation of Bonazzola et al. (2007), like the Newtonian anelastic approximation, does not evolve the density perturbation.

Initially, the derivation was attempted using the Valencia formulation (see Section 2.2.1 below), however the form of the low Mach number equations derived was quite different from the Newtonian equations found in the `MAESTRO` papers. The equations were also derived using the 3+1 form of Alcubierre (2008) (see Section 2.2.2): this produced equations which were clearly of a similar form to the Newtonian ones, but which are not suitable for numerical implementation (at least if we are to base our scheme on the existing Newtonian one). Using the Wilson formulation in Section 2.2.3, the equations derived are much more similar to the Newtonian low Mach equations, making them both easier to compare and easier to implement numerically, as a scheme could be built on the existing Newtonian code.

2.2.1 Conservative derivation

The Valencia formalism of Martí et al. (1991); Banyuls et al. (1997) expresses the relativistic fluid equations in a balance law form. As given in Font (2008), the conserved quantities of the general relativistic hydrodynamic system can be written in terms of the primitive variables $\boldsymbol{w} = (\rho, v_i, \varepsilon)$, the proper rest-mass density, fluid 3-velocity and specific energy density of the fluid in its rest frame, as

$$D = \rho W, \qquad S_j = \rho h W^2 v_j, \qquad E = \rho h W^2 - p, \tag{2.17}$$

which can be interpreted physically as the rest-mass density, momentum density in the j-th direction and the total energy density respectively. The enthalpy is given by

$$h = 1 + \varepsilon + \frac{p}{\rho}. \tag{2.18}$$

We can write the relativistic hydrodynamic equations in 3+1 conservative Eulerian form as

$$\frac{1}{\sqrt{-g}}\left(\frac{\partial\sqrt{\gamma}\boldsymbol{q}(\boldsymbol{w})}{\partial x^0} + \frac{\partial\sqrt{-g}\boldsymbol{f}^i(\boldsymbol{w})}{\partial x^i}\right) = \boldsymbol{S}(\boldsymbol{w}), \tag{2.19}$$

with the metric determinant $g \equiv \det(g_{\mu\nu})$ and $\sqrt{-g} = \alpha\sqrt{\gamma}$, where $\gamma = \det(\gamma_{ij})$ is the determinant of the spatial 3-metric. We define conserved state vector to be

$$\boldsymbol{q}(\boldsymbol{w}) = (D, S_j, \tau), \tag{2.20}$$

where $\tau \equiv E - D$. Finally, the flux vector is defined as

$$\boldsymbol{f}^i(\boldsymbol{w}) = \left(D\left[v^i - \frac{\beta^i}{\alpha}\right], S_j\left[v^i - \frac{\beta^i}{\alpha}\right] + p\delta^i_j, \tau\left[v^i - \frac{\beta^i}{\alpha}\right] + pv^i \right), \tag{2.21}$$

and the sources are given by

$$\boldsymbol{S}(\boldsymbol{w}) = \left(0, T^{\mu\nu}\left[\frac{\partial g_{\nu j}}{\partial x^\mu} - \Gamma^\delta{}_{\nu\mu} g_{\delta j}\right], \alpha\left[T^{\mu 0}\frac{\partial \ln\alpha}{\partial x^\mu} - T^{\mu\nu}\Gamma^0{}_{\nu\mu}\right] \right), \tag{2.22}$$

where for an ideal fluid in thermodynamic equilibrium, the energy-momentum tensor takes the form

$$T^{\mu\nu} = \rho h u^\mu u^\nu + p g^{\mu\nu} = (\rho + \rho\varepsilon + p)\, u^\mu u^\nu + p g^{\mu\nu}. \tag{2.23}$$

2.2.1.1 Low Mach number limit

The derivation of the low Mach number limit of this form of the equations is very similar to the derivation used for the Wilson formulation given in Appendix B. The equations obtained are as follows:

the continuity equation

$$\partial_0(\sqrt{\gamma}\rho W) + \partial_i\left(\sqrt{-g}\rho u^i\right) = 0; \tag{2.24}$$

the constrained momentum equation

$$\begin{aligned}\sqrt{\gamma}W\partial_0\left(h_0 u_j\right) = &-\sqrt{-g}u^i\partial_i\left(h_0 u_j\right) + \frac{1}{\rho}\sqrt{-g}\partial_j(p_0 + \pi) \\ &+ \sqrt{-g}h_0 u^\mu u^\nu \Gamma_{\nu j\mu} - \frac{1}{\rho}\sqrt{\gamma}W\zeta\partial_0\left(\frac{\pi}{\zeta}\right) - \frac{1}{\rho}\sqrt{-g}\zeta\partial_j\left(\frac{\pi}{\zeta}\right),\end{aligned} \tag{2.25}$$

where ζ is the relativistic counterpart of the β_0 variable found in the Newtonian low Mach number equations to ensure energy conservation, which we have relabelled in order to avoid confusion with the shift vector.

The energy equation is

$$\begin{aligned}\sqrt{\gamma}\rho W\partial_0(h_0 W) \approx &-\sqrt{-g}\rho u^i\partial_i(h_0 W) + \partial_0(\sqrt{\gamma}p_0) - \partial_i\left[\sqrt{-g}p_0\frac{\beta^i}{\alpha}\right] \\ &+ \sqrt{-g}\left(\rho h_0 u^\mu W + p_0\alpha g^{\mu 0}\right)\partial_\mu \ln\alpha - \sqrt{-g}\alpha\left(\rho h_0 u^\mu u^\nu + p_0 g^{\mu\nu}\right)\Gamma^0{}_{\nu\mu};\end{aligned} \tag{2.26}$$

and the velocity constraint

$$\partial_i\left(\sqrt{-g}\zeta_0 u^i\right) = \zeta_0\left(\tilde{S} - \frac{\sqrt{\gamma}W}{\rho p_\rho}\left[1 + \frac{p_T h_p}{h_T}\right]\partial_0 p_0\right). \tag{2.27}$$

These equations are significantly more complicated than those that shall be found in Section 2.2.3 using the Wilson formulation and are not directly comparable to the Newtonian low Mach number equations. Consequently, a lot more work would have to be done to implement these numerically and to test that they produce the correct results in the Newtonian limit.

2.2.2 Covariant 3+1 form

The Newtonian equations of Nonaka et al. (2010) are presented in a covariant form: by using the covariant (non-conservative) 3+1 form of the relativistic fluid equations, it is therefore hoped that this will produce equations which will have a sensible Newtonian limit that is directly comparable to the Newtonian equations. Alcubierre (2008) gives the equations of relativistic hydrodynamics in explicitly $3+1$ form as

$$\left(\partial_t - \beta^k \partial_k\right) D + {}^{(3)}\nabla_k \left(\alpha D v^k\right) = \alpha K D, \tag{2.28}$$

$$\left(\partial_t - \mathcal{L}_\beta\right) S_i + {}^{(3)}\nabla_k \left[\alpha \left(S_i v^k + \delta_i^k p\right)\right] = -\left(\mathcal{E} + D\right) {}^{(3)}\nabla_i \alpha + \alpha K S_i, \tag{2.29}$$

$$\left(\partial_t - \beta^k \partial_k\right) \mathcal{E} + {}^{(3)}\nabla_k \left[\alpha v^k \left(\mathcal{E} + p\right)\right] = \left(\mathcal{E} + p + D\right)\left(\alpha v^m v^n K_{mn} - v^m \partial_m \alpha\right) + \alpha K \left(\mathcal{E} + p\right). \tag{2.30}$$

Here α, β^i are the lapse and shift gauge variables, ${}^{(3)}\nabla_k$ is the covariant derivative associated with the spatial 3-metric γ_{ij}, $\mathcal{L}_\beta$ is the Lie derivative with respect to the shift vector, and K_{mn} is the extrinsic curvature with K its trace. The conserved variables are

$$\begin{aligned} D &= \rho_0 W, \\ S_i &= H W^2 v_i, \\ \mathcal{E} &= H W^2 - p - D, \end{aligned} \tag{2.31}$$

where $H = \varrho + p = \rho_0 \left(1 + \epsilon\right) + p$, p is the pressure, ρ_0 the specific rest mass density, ϵ the specific internal energy, ϱ the energy density, H the enthalpy and W is the Lorentz factor of the flow.

2.2.2.1 Low Mach number limit

To obtain the low Mach number limit of these equations, we make two fundamental assumptions. Firstly, that there is a reference state in which the metric and 'background' fluid quantities are time independent. We will use overbars, e.g. $\bar{p}$, to denote this. Secondly, we will assume that the full fluid quantities are perturbations from this reference state, so $p = \bar{p} + \pi$, and that the perturbation in the pressure is second order in the perturbation in the velocity: $\pi/\bar{p} = \mathcal{O}(V^2)$.

Details of the derivation can be found in Appendix A. We find that the momentum equation becomes

$$\left(\partial_t - \mathcal{L}_\beta\right) V_i + \alpha V^k {}^{(3)}\nabla_k V_i + \frac{{}^{(3)}\nabla_i \pi}{H} = -\frac{\varrho - \bar{\varrho}}{H} {}^{(3)}\nabla_i \alpha. \tag{2.32}$$

Taking the Newtonian limit sends the enthalpy H to the energy density ϱ, and the equation has the expected Low Mach form.

With the assumptions we have made, the energy equation reduces to

$$\left(\partial_t - \beta^k \partial_k\right)\left(HW^2\right) + {}^{(3)}\nabla_k \left[\alpha v^k H W^2\right] = -\beta^k \partial_k \bar{p} + H\left(-v^m \partial_m \alpha + \alpha K\right). \tag{2.33}$$

This equation is less straightforward to compare in the Newtonian limit, but after some manipulation we can obtain a suitable form for comparison which reduces to the correct Newtonian limit.

2.2.3 Wilson formulation

In this derivation, we start with the relativistic fluid equations in the form given in Section 1.4.3. The equations are then put into the Wilson formulation (Wilson, 1972; Font, 2008), where we define $D = \rho u^0$ and $U^\mu = u^\mu / u^0$ to get

$$\nabla_\mu \left(D U^\mu\right) = 0, \tag{2.34}$$

$$D h U^\nu \partial_\nu \left(u^0 U_\mu\right) + \partial_\mu p + U_\mu U^\nu \left(u^0\right)^2 \partial_\nu p = D h u^0 \Gamma_{\rho\nu\mu} U^\nu U^\rho, \tag{2.35}$$

$$\nabla_\mu \left(D h U^\mu\right) = u^0 U^\mu \nabla_\mu p. \tag{2.36}$$

The Wilson formulation does not typically perform as well as other formulations (e.g. the Valencia formulation) for flows with high Lorentz factor. However, we are investigating low Mach number systems, this shall not be an issue for us. The relativistic effects that we are interested here are those due to the strong graviational field rather than due to velocities approaching the speed of light.

The full derivation of the relativistic low Mach number equations using this formulation is given in detail in Appendix B; we shall just summarise them here. In our derivation, we have assumed that the spacetime is stationary, that the background state satisfies the TOV equations and have neglected terms that are quadratic in velocity (and are therefore $\mathcal{O}(\mathrm{Ma}^2)$). We obtain the following set of equations:

the continuity equation

$$\partial_t D = -\partial_i (D U^i) - D \Gamma^\mu{}_{\mu\nu} U^\nu; \tag{2.37}$$

the constrained momentum equation

$$\partial_t U_j = -U^i \partial_i U_j - U_j \frac{\mathrm{D} \ln u^0}{\mathrm{D}t} - \frac{1}{Dhu^0} \left(\partial_j p_0 + \zeta \partial_j \left[\frac{\pi}{\zeta} \right] \right) + \Gamma_{\rho\nu j} U^\nu U^\rho; \tag{2.38}$$

where we have defined the convective derivative $\frac{\mathrm{D}q}{\mathrm{D}t} = \partial_t q + U^i \partial_i q$.

The energy equation is

$$\partial_t (Dh) = -\partial_i (U^i Dh) + u^0 \frac{\mathrm{D}p_0}{\mathrm{D}t} - Dh\Gamma^\mu{}_{\mu\nu} U^\nu; \tag{2.39}$$

and the velocity constraint

$$\partial_i \left(\zeta U^i \right) = \zeta \left(S - \frac{1}{\bar{\Gamma}_1 p_0} \frac{\partial p_0}{\partial t} \right), \tag{2.40}$$

where

$$\zeta(t, r) = \zeta(0, t) \exp \left[\int_0^r \mathrm{d}r' \, \frac{1}{\bar{\Gamma}_1 p_0} \frac{\partial p_0}{\partial r'} \right], \tag{2.41}$$

and $S = -\Gamma^\mu{}_{\mu\nu} U^\nu$ are gravitational source terms.

Comparing these to the Newtonian equations (2.10) – (2.13), it can be seen that they are mostly very similar, but with the addition of the relativistic gravitational source terms. The $\zeta(t, r)$ term used here in the velocity constraint and momentum equation is the relativistic counterpart of the β_0 term in the Newtonian equations – we have renamed it in order to prevent potential confusion with the shift vector.

2.3 Numerical implementation of the relativistic low Mach number equations

In this section, we shall describe the numerical implementation of the relativistic low Mach number equations given above in Section 2.2.3, which were implemented by extending the Newtonian `pyro` code (Zingale, 2014). This code is a simplified, Python-based version of the full `MAESTRO` code (Nonaka et al., 2010); using this simplified code allowed us to focus on the implementation of the equations without worrying about the complexities of the numerical methods. In a weak planar gravitational field, the relativistic code produces qualitatively the same results as the original Newtonian code, indicating that the equations are behaving correctly in the Newtonian limit.

2.3.1 Numerical methodology

Rather than starting from scratch in our implementation of a solver for the relativistic low Mach number equations, the existing low Mach number `pyro` code was used as a base

so that we could take advantage of its multigrid solvers, high resolution shock capturing methods and overall framework. This was then built upon to implement a solver for our relativistic equations, using a numerical method adapted from that of the full `MAESTRO` code (Nonaka et al., 2010).

As in `MAESTRO`, a predictor corrector formalism is used to evolve the system. In the predictor step, an estimate of the expansion of the base state is calculated, followed by an estimate of the state at the new time. In the corrector step, this is used to compute a new base state update as well as the full state update. The reaction and gravitational source terms in the continuity and energy equations are included using Strang-splitting (Strang, 1968) in order to achieve second-order accuracy in time.

The main differences between our method and that of Nonaka et al. (2010) are:

- the use of gravitational source terms in addition to reactions when evolving D and Dh,
- the inclusion of gravitational source terms when evolving U^i,
- the addition of routines to deal with the metric and calculations of related quantities (e.g. Lorentz factor, Christoffel symbols).

Changes to specific steps of the numerical procedure are discussed in more depth in Appendix C.

2.3.1.1 Primitive variable recovery

In relativistic hydrodynamics, the transformation from the primitive variables (ρ, v^i, p) to the conserved variables is a simple algebraic relation with a closed-form solution. However, the inverse transformation requires solving a set of nonlinear equations with explicit dependence on the equation of state $p = p(\rho, \varepsilon)$ (Martí and Müller, 2015). In the Valencia formulation, the pressure of the physical state can be obtained by finding the root of the function

$$f(\bar{p}) = p(\rho_*(\bar{p}), \varepsilon_*(\bar{p})) - \bar{p}, \tag{2.42}$$

where we define

$$\rho_*(\bar{p}) = \frac{D}{W_*(\bar{p})}, \tag{2.43}$$

$$\varepsilon_*(\bar{p}) = \frac{\tau + D(1 - W_*(\bar{p})) + \bar{p}\left(1 - W_*(\bar{p})^2\right)}{DW_*(\bar{p})}, \tag{2.44}$$

$$W_*(\bar{p}) = \frac{1}{\sqrt{1 - v_*^i(\bar{p}) v_{*i}(\bar{p})}}, \tag{2.45}$$

$$v_*^i(\bar{p}) = \frac{S^i}{\tau + D + \bar{p}}, \tag{2.46}$$

and D, S^i and τ are the conserved variables, as given in Section 2.2.1. The root of (2.42) is found using a nonlinear root-finder, e.g. Newton-Raphson or Brent's method (Brent, 1973). The equation of state implemented in our code is

$$p = (\gamma - 1)\rho\varepsilon. \tag{2.47}$$

For the Wilson formulation (with fixed lapse function α and zero shift vector β^i) used in our implementation of the relativistic low Mach number equations, we can use an analagous technique to reclaim the primitive variables. From the momentum equation (2.38), we have $U^i = u^i/u^0$, which we can use to calculate the Lorentz factor W_* using (2.45). Using the fact that α is fixed and $W = \alpha u^0$, this will give us $u^0_* = W_*/\alpha$. This can then be used to find the density $\rho_* = D/u^0_*$. We can next find ε_* from

$$\varepsilon_*(\bar{p}) = \frac{Dh - \bar{p} - \rho_*(\bar{p})}{D}, \tag{2.48}$$

then use the equation of state to give us $p(\rho_*, \varepsilon_*)$. Finally, we can use this to find the root of (2.42).

2.3.2 Numerical results

We tested our equations by modelling several standard 2d systems and comparing the results to models produced using the general relativistic compressible solver `GRHydro` (Baiotti et al., 2005) from the Einstein Toolkit (Schnetter et al., 2004; Einstein Toolkit, 2010; Löffler et al., 2012). In order to test our methodology and compare it to existing work, we worked in the Newtonian regime. To our knowledge, there exist no previous derivations of the relativistic low Mach number equations, so there are no previous relativistic low Mach number results for us to compare against. By working in the weak gravity, Newtonian regime, we can compare against the results produced by `pyro` and `MAESTRO` to ensure that our implementation is consistent in the Newtonian limit.

To work in this limit, we used a $2 + 1$ metric that describes a Newtonian perturbation from Minkowski space (Reall, 2011),

$$ds^2 = -c^2\left(1 - \frac{2GM}{rc^2}\right)dt^2 + \left(1 + \frac{2GM}{rc^2}\right)\left(dx^2 + dr^2\right), \tag{2.49}$$

where G is the gravitational constant and M the mass of the star. Tests were performed with and without reactions.

The elliptic velocity constraint is solved in `pyro` and `MAESTRO` using multigrid methods. In `pyro`, the multigrid solver is cell-centred, whereas in `MAESTRO` the solver is nodal. Due to the stencil used by the cell-centred solver, at late times in the evolution some odd-even decoupling becomes noticeable in the solution. Nodal solvers do not suffer from such issues, however they are more difficult to implement numerically. We do not

believe that the odd-even decoupling is significant enough to affect our conclusions, so at this stage where we are testing the validity of our methodology, we believe the simpler cell-centred solver is sufficient.

The modelling in `GRHydro` was in 3+1 dimensions: we made the simulations effectively 2+1 by using very low resolution, uniform initial conditions and periodic boundary conditions in the redundant third spatial dimension such that they could be compared to the 2+1 `pyro` simulations.

2.3.2.1 Kelvin-Helmholtz instability

The Kelvin-Helmholtz instability (Helmholtz, 1868; Kelvin, 1910; Chandrasekhar, 1961a) is produced by the unstable growth of perturbations at the interface of two fluid flows and is driven by shearing motions. It was simulated by placing a layer of dense fluid on top of a less dense one, with a weak planar gravitational field acting in the vertical direction. The two layers of fluid are given a small initial horizontal velocity so that they move slowly in opposite directions, resulting in shearing at their interface. Typically, numerical tests are performed using a sharp transition at the shear interface, however in order for the solution to converge with increasing resolution, this interface should be smoothed so that spatial derivatives are finite (Robertson et al., 2010).

Using the initial conditions of McNally et al. (2012); Miczek et al. (2014), the initial horizontal velocity was given by

$$U^x = \begin{cases} U_1 - U_m \exp\left(\frac{y - y_{\text{center}}}{L_x}\right) & \text{if } y \leq y_{\text{center}} \\ U_2 + U_m \exp\left(\frac{-y + y_{\text{center}}}{L_x}\right) & \text{if } y > y_{\text{center}} \end{cases}, \tag{2.50}$$

where L_x is the horizontal extent of the numerical domain, U_1, U_2 are the velocities of the sections of fluid at infinity in the vertical direction and $U_m = \frac{1}{2}(U_2 - U_1)$. The initial conditions for the density are identical to those for the horizontal velocity, but with D exchanged for U.

Outflow boundary conditions were used for the vertical boundaries and periodic for the horizontal. The instability was seeded by making a small sinusoidal perturbation to the vertical velocity,

$$U^y = A \sin\left(\frac{4\pi x}{L_x}\right), \tag{2.51}$$

where A is some small constant ($A = 0.5U_1$ was used here).

The simulation after the Kelvin-Helmholtz instability has evolved 250 timesteps can be seen in Figure 2.2a for a run with resolution 512×512. The plot shows a passive advective scalar X that was initialised with the value $X = 1$ in the lower half of the domain and $X = 0$ in the upper half of the domain. Over the course of the simulation,

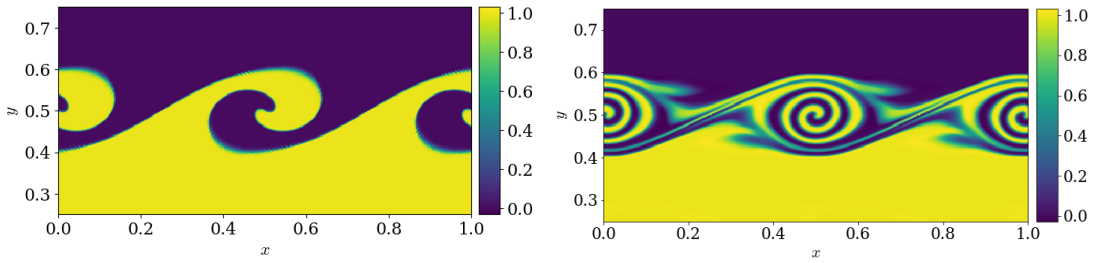

(A) 250 timesteps (B) 600 timesteps

FIGURE 2.2: Kelvin-Helmholtz instability simulation after 250 and 600 timesteps with resolution 512×512. The plot shows the passive advective scalar.

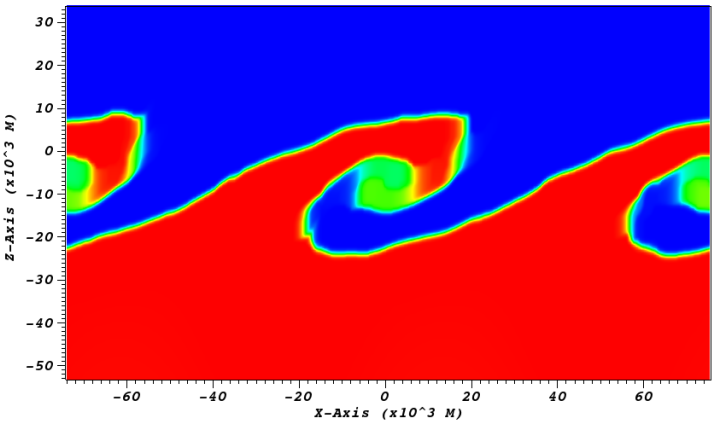

FIGURE 2.3: Kelvin-Helmholtz instability simulation produced using `GRHydro` after 90,000 timesteps. The plot shows the passive advective scalar.

shearing at the interface has produced the characteristic Kelvin-Helmholtz vortices. In Figure 2.2b, we show the same simulation evolved to 600 timesteps. The spirals are now more developed. It can be seen that even at this later stage the innermost spirals are very well resolved, however the boundary of the outermost spirals is a little smeared due to numerical dissipation.

The results produced by `pyro` can be compared to those produced by `GRHydro` in Figure 2.3. This figure is at a comparable stage in the evolution of the Kelvin-Helmholtz instability as Figure 2.2a, though as `GRHydro` is a fully compressible code it has taken many more timesteps to reach this point (90,000, so 360 times more than for the first `pyro` figure). As a result of the large number of timesteps required, the numerical errors at this point are much greater, leading to a larger amount of numerical dissipation. Consequently, the inner spiral has become significantly smeared. Clearly, this demonstrates how compressible codes such as `GRHydro` are ill-suited to modelling low Mach number flows.

2.3.2.2 Rayleigh-Taylor instability

Like the Kelvin-Helmholtz instability, the Rayleigh-Taylor instability (Rayleigh, 1900; Taylor, 1950; Chandrasekhar, 1961b) is produced by the unstable growth of perturbations at the interface of two fluids, however in this case the instability is produced by the acceleration of the fluids towards each other (e.g. due to a gravitational field). This was modelled by placing a layer of dense fluid on top of a less dense one, with a weak planar gravitational field acting in the vertical direction. Unlike in the Kelvin-Helmholtz example above, the fluid is initially stationary. The instability was then seeded by making a small sinusoidal perturbation of the vertical velocity:

$$U^y = A\cos\left(\frac{2\pi x}{L_x}\right)\exp\left(-\frac{(y-y_{\text{center}})^2}{\sigma^2}\right), \tag{2.52}$$

where L_x is the horizontal extent of the numerical domain, y_{center} is the vertical coordinate at the centre of the domain and A, σ are some small constants (0.1 was used here). The boundary between the fluids of different densities was smoothed to eradicate very high density gradients in the initial data, which can be difficult to deal with numerically, giving the initial density profile

$$D = D_1 + \frac{1}{2}(D_2 - D_1)\left(1 + \tanh\left[\frac{y - y_{\text{center}}}{0.5 y_{\text{smooth}}}\right]\right), \tag{2.53}$$

where $D_1 < D_2$. Outflow boundary conditions were used for the top and bottom boundaries and periodic for the horizontal boundaries.

The resulting instability after the simulation has evolved 1000 timesteps with resolution 512×512 can be seen in Figure 2.4. The plot of the passive advective scalar shows that the code has performed well in this case, producing the characteristic Rayleigh-Taylor spirals with very little smearing at the interface.

2.3.2.3 Gresho vortex

The Gresho vortex (Gresho and Chan, 1990; Liska and Wendroff, 2003) is a time independent solution of the hydrodynamics equations, with centrifugal forces balanced by pressure gradients. The initial angular velocity is given by

$$u_\phi(r) = \begin{cases} u_0 \frac{r}{R} & 0 \le r < R, \\ u_0\left(2 - \frac{r}{R}\right) & R \le r < 2R, \\ 0 & r \ge 2R, \end{cases} \tag{2.54}$$

where r is the radial distance from the centre of the vortex, and u_0 and R are parameters which determine that the maximum value of the angular velocity $u_{\phi,\max} = u_0$ occurs at

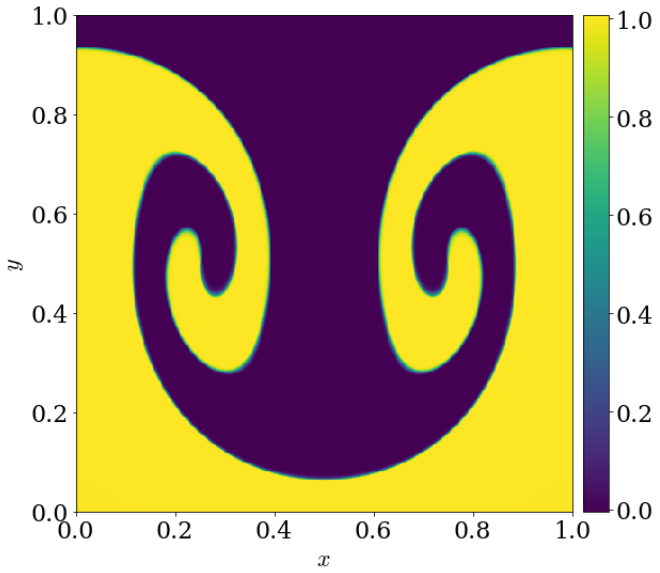

FIGURE 2.4: Rayleigh-Taylor simulation after 1000 timesteps with resolution 512×512. The plot shows the passive advective scalar.

$r = R$. The pressure gradient required to balance the vortex's centripetal force is

$$p(r) = \begin{cases} p_0 + \frac{u_0^2}{2}\left(\frac{r}{R}\right)^2 & 0 \leq r < R, \\ p_0 + u_0^2\left[\frac{1}{2}\left(\frac{r}{R}\right)^2 + 4\left(1 - \frac{r}{R} + \ln\frac{r}{R}\right)\right] & R \leq r < 2R, \\ p_0 + u_0^2\left(4\ln 2 - 2\right) & r \geq 2R, \end{cases} \tag{2.55}$$

where p_0 is a parameter that determines the maximum Mach number of the flow at $r = R$ (Miczek et al., 2014). For this test, we set $R = 0.2$, $u_0 = 0.1$ and $p_0 = 100$.

The Gresho vortex test is one that traditional compressible codes tend to perform poorly with compared to low Mach number codes. To evolve the simulation to the same point in time, compressible codes will have to carry out a much greater number of timesteps. With each timestep, a small amount of additional error will be produced. This error will build up, so that by the end of the simulation the larger number of timesteps will have resulted in a significantly less accurate solution.

The results for the low Mach number `pyro` simulation can be seen in Figure 2.5 for a test with a resolution of 512×512 after 660 timesteps. The initial data has been conserved fairly well, with the largest errors occurring where the speed in the vortex is greatest and at the edge of the vortex where it meets the static fluid. Elsewhere, the change in speed is negligible. In Figure 2.6, we show the relative change in the total kinetic energy over the course of the simulation. The total kinetic energy steadily decreases linearly with time for the majority of the simulation, the energy loss appearing to slightly slow down at the very end. At this point, it has only lost 0.25% of its initial energy.

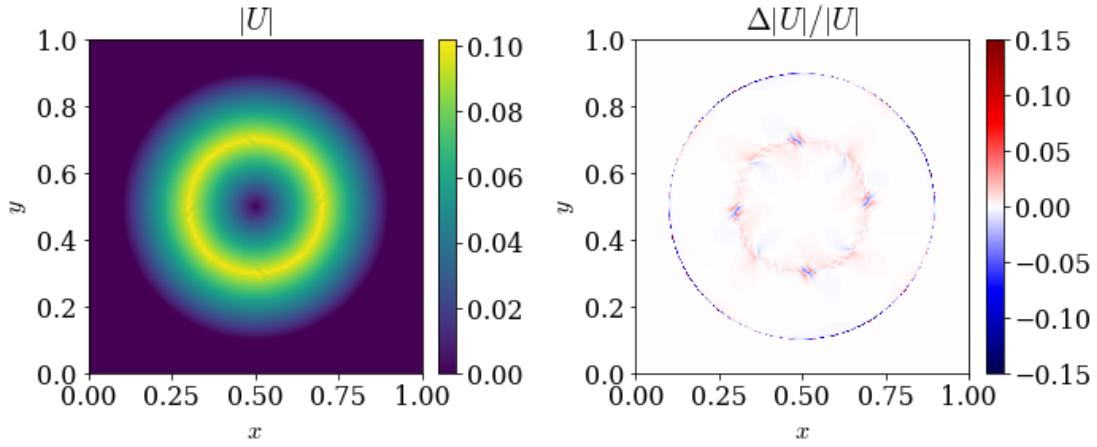

FIGURE 2.5: Gresho vortex after 660 timesteps. The plot shows the magnitude of the primitive velocity and the relative change in the velocity compared to the initial data. It can be seen that the initial data has been conserved to a high degree of accuracy, with the largest errors occurring where the speed is greatest and at the edge of the vortex.

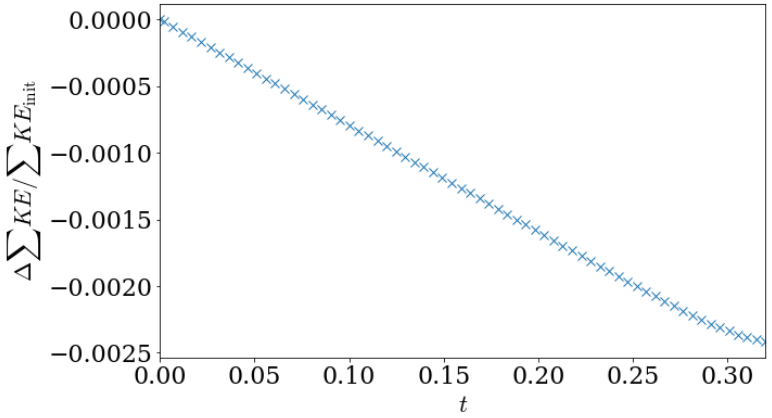

FIGURE 2.6: Change in the total kinetic energy for the Gresho vortex test over 660 timesteps. The kinetic energy decreases uniformly over the majority of the simulation. By the end, it has lost only 0.25% of its initial kinetic energy.

2.3.2.4 Reacting bubble rise

To test reactions in the code, we use the reacting bubble rise test found in the `MAESTRO` papers (Almgren et al., 2006b,a, 2008). In this test, a circular region of hot, underdense material was placed in a fluid which has been vertically stratified by a weak planar gravitational field (see Figure 2.7). A resolution of 512×512 was used. As can been seen in Figure 2.8, the bubble is convectively unstable, so rises. The bubble is deformed, its outer edges curling round underneath. In this example, as material in the bubble reaches the ignition temperature, it begins to burn, increasing the baryon fraction of burnt material. We use baryon fraction rather than mass fraction here as in nuclear fusion baryon number is conserved, whereas mass is not due to energy release (Townsley et al., 2016). In the plots, the top of the bubble appears to have some ridges; these

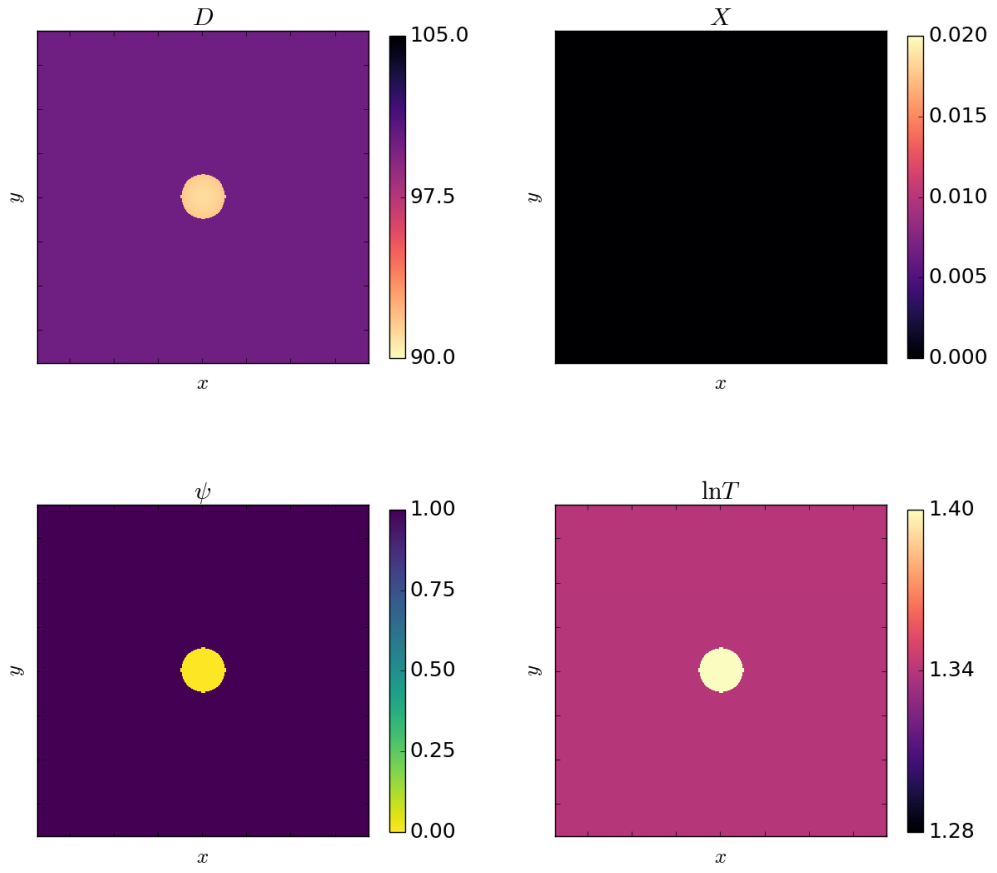

FIGURE 2.7: Reactive bubble simulation at initial time $t = 0$. The plots show the conserved density variable, $D = \rho u^0$, the baryon fraction of burnt material, X, a passive advective scalar, ψ, and the logarithm of the temperature, $\ln T$.

are believed to be numerical artefacts (e.g. because the initial data was discontinuous, the movement of the bubble produces Kelvin-Helmholtz instabilities at the interface (Robertson et al., 2010)), as they reduce in size as the resolution is increased. This was observed by repeating the test at a range of resolutions between 64×64 and 1024×1024. Some horizontal bands outside of the bubble can also be seen in the density and temperature plots. These are likewise thought to be purely numerical, resulting from certain horizontal averaging steps in the algorithm.

Comparing our results to the Newtonian simulations in the `MAESTRO` papers (Almgren et al., 2006b,a, 2008), we can see that they are qualitatively similar, with the same curling up of the bubble at the edges. However their simulations, which were performed at much higher resolution than ours, produced evolved bubbles with smooth tops, again suggesting that these are most likely numerical artefacts in our code.

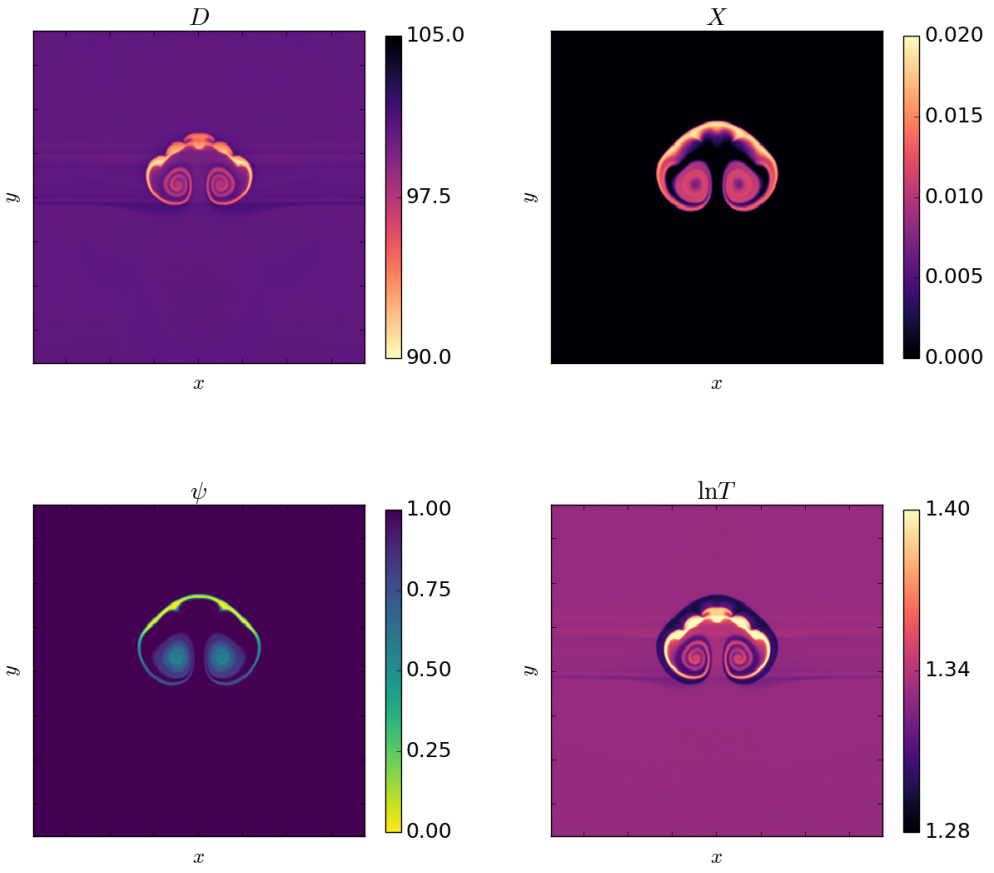

FIGURE 2.8: Reactive bubble simulation after it has been evolved through 500 timesteps. The plots show the conserved density variable, $D = \rho u^0$, the baryon fraction of burnt material, X, a passive advective scalar, ψ, and the logarithm of the temperature, $\ln T$.

2.4 Summary

In this chapter, we have described the development of a relativistic low Mach number formalism for modelling type I X-ray bursts and demonstrated its numerical implementation. We first discussed the advantages of sound-proof models when looking at systems where the characteristic velocity of the system is much smaller than the sound speed and various ways they can be implemented. We then summarised the derivation of the low Mach number equations in the Newtonian case, before presenting the derivation of the relativistic low Mach number equations. This was attempted using several different formulations of the general relativistic fluid equations; we concluded that the Wilson formulation was the most suitable as it produced a set of equations that most closely resembled the equations implemented in existing Newtonian codes such as `MAESTRO` (Nonaka et al., 2010) and `pyro` (Zingale, 2014). This therefore makes them easiest to compare and to implement numerically.

The rest of this chapter described how these equations were implemented numerically. Modelling of several test cases has shown that our code produces qualitatively the results that we would expect, indicating that the code and formalism are working correctly. As expected, our code has enabled us to gain a considerable speed up in simulating these systems when compared against conventional compressible solvers.

The numerical implementation presented here extended the Python-based `pyro` code. This code was chosen over the more accurate `MAESTRO` code for its relative simplicity, allowing us to focus on the implementation of the equations and additional requirements of a relativistic code. This choice was sufficient for the purposes of demonstrating the properties of the relativistic low Mach number equations, however this solver will not be suitable for conducting simulations of the scale we require to model X-ray bursts. In the future, we therefore intend to implement these equations in `MAESTRO` as well, enabling us to do much larger, more accurate simulations.

3 R3D2: the relativistic reactive Riemann problem

As described in Section 1.3, during an X-ray burst the burning front propagates as a deflagration through the ocean. On the scale of the entire star, the front can be treated as a discontinuity, with the burning reactions happening on very short timescales compared to the propagation of the front. In this case, we can approximate the system as a reactive Riemann problem. In this section, we shall describe work done to create an exact relativistic reactive Riemann solver (`R3D2`: Relativistic Reactive Riemann problem solver for Deflagrations and Detonations, DOI: 10.5281/zenodo.51096), building on work that has previously been done for the inert relativistic and the reactive Newtonian cases. We shall compare this to results gained from a compressible relativistic hydrodynamics code with a source term modelling the reactions, which itself was created by extending the fully compressible solver in `pyro` (Zingale, 2014). Some of this work is also described in Harpole and Hawke (2016).

3.1 The Riemann problem

As described in Section 1.5.1, the Riemann problem is defined as a system of hyperbolic conservation laws whose initial data $\boldsymbol{q}(0,x)$ consists of two states, $\boldsymbol{q}_L$, $\boldsymbol{q}_R$, separated by a single jump discontinuity (LeVeque, 2002),

$$\boldsymbol{q}(0,x) = \begin{cases} \boldsymbol{q}_L(0,x) & \text{if } x < 0, \\ \boldsymbol{q}_R(0,x) & \text{if } x > 0. \end{cases} \tag{3.1}$$

The solution of the Riemann problem involves computing the breakup of the initial discontinuity. This consists of a set of waves separating a set of constant states. These waves can be *simple waves*, which can take the form of nonlinear waves (such as shocks or rarefactions) and linear waves (such as contact discontinuities[1] and rotational linear waves in MHD). As shall be seen for the reactive Riemann problem, the solution may

[1] For fluids, across a *contact discontinuity* the density is discontinuous, but the pressure and normal velocity are continuous.

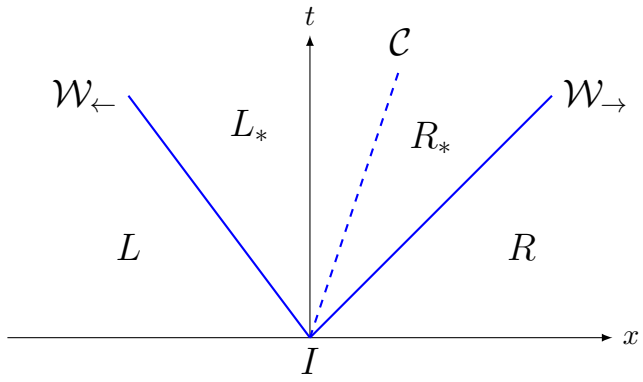

FIGURE 3.1: Wave pattern for the Riemann problem of a non-reactive fluid with initial left and right states L, R. The initial discontinuity decays into two nonlinear waves $\mathcal{W}_{\leftarrow}$, $\mathcal{W}_{\rightarrow}$, between which there are two new constant states L_*, R_* separated by a contact discontinuity $\mathcal{C}$.

also contain *compound waves* if the flux is non-convex[2]. These consist of multiple simple waves 'attached' to each other, moving as one. The exact nature of the wave pattern, including the number of waves in the final state, depends on the initial conditions of the problem. The solution is self-similar, such that it can be written as a function a single independent variable,

$$\xi = \frac{x}{t}. \tag{3.2}$$

Waves in the solution follow lines of constant ξ. In the case of a non-reactive fluid, the discontinuity will decay into two nonlinear waves moving in opposite directions. Between these two waves will be two new constant states $\boldsymbol{q}_L^*$, $\boldsymbol{q}_R^*$, separated by a contact discontinuity.

We can represent the time evolution of the Riemann problem in a non-reactive fluid with the initial state I as

$$I \rightarrow L\mathcal{W}_{\leftarrow}L_*\mathcal{C}R_*\mathcal{W}_{\rightarrow}R, \tag{3.3}$$

where $\mathcal{W}$ denotes a simple wave (a shock or a rarefaction) and $\mathcal{C}$ a contact discontinuity. The states are labelled L, L_*, R_* and R. The arrows represent the direction from which a fluid element enters the wave. This solution is illustrated in Figure 3.1.

The distinction between shock and rarefaction waves is defined as

$$\mathcal{W} = \begin{cases} \mathcal{R}, & p_b \leq p_a \\ \mathcal{S}, & p_b > p_a, \end{cases} \tag{3.4}$$

[2]The flux function $\boldsymbol{f}(\boldsymbol{q})$ is *convex* if $\boldsymbol{f}''(\boldsymbol{q})$ does not change sign in the domain over which $\boldsymbol{q}$ is defined. Consequently, the characteristic speed $\boldsymbol{f}'(\boldsymbol{q})$ varies monotonically with $\boldsymbol{q}$.

where the subscripts indicate quantities ahead of and behind the wave: $a \equiv L(R)$ and $b \equiv L_*(R_*)$ for $\mathcal{W}_\leftarrow(\mathcal{W}_\rightarrow)$. The initial discontinuity can decay in three distinct ways,

$$I \rightarrow \begin{cases} L\mathcal{S}_\leftarrow L_* \mathcal{C} R_* \mathcal{S}_\rightarrow R, & \text{if } p_L < p_* \text{ and } p_R < p_*, \\ L\mathcal{S}_\leftarrow L_* \mathcal{C} R_* \mathcal{R}_\rightarrow R, & \text{if } p_L < p_* \leq p_R, \\ L\mathcal{R}_\leftarrow L_* \mathcal{C} R_* \mathcal{R}_\rightarrow R, & \text{if } p_* \leq p_L \text{ and } p_* \leq p_R, \end{cases} \tag{3.5}$$

where $p_{L*} = p_{R*} = p_*$. Which of these three wave patterns the solution takes depends on the initial data.

3.1.0.1 Exact solution

For the Riemann problem of the inert non-relativistic Euler equations, it is possible to find an exact solution if the initial states $\boldsymbol{q}_L$ and $\boldsymbol{q}_R$ are known. As the pressure and normal velocity are constant across the contact wave, we can find the solution by solving an implicit equation for the pressure in the star states p_* (Toro, 1999),

$$f(p_*, \boldsymbol{q}_L, \boldsymbol{q}_R) \equiv f_L(p_*, \boldsymbol{q}_L) + f_R(p_*, \boldsymbol{q}_R) + \Delta v = 0, \tag{3.6}$$

where $\Delta v \equiv v_{x,R} - v_{x,L}$ is the difference across the wave of the normal velocity, v_x. This implicit equation is derived by considering the Rankine-Hugoniot conditions (previously described in 1.5.1.1). Denoting the left and right states L, R by S, the functions f_L, f_R are given by

$$f_S(p_*, \boldsymbol{q}_S) = \begin{cases} (p_* - p_S)\left(\frac{A_S}{p_* + B_S}\right)^{\frac{1}{2}} & \text{if } p_* > p_S \text{ (shock)}, \\ \frac{2c_{s,S}}{\gamma_S - 1}\left[\left(\frac{p_*}{p_S}\right)^{\frac{\gamma_S - 1}{2\gamma_S}} - 1\right] & \text{if } p_* \leq p_S \text{ (rarefaction)}, \end{cases} \tag{3.7}$$

where the constants A_S, B_S are given by

$$A_S = \frac{2}{(\gamma_S + 1)\rho_S}, \qquad B_S = \frac{\gamma_S - 1}{\gamma_S + 1} p_S. \tag{3.8}$$

The adiabatic index $\gamma_S = c_{p,S}/c_{V,S}$ is the ratio of specific heat capacities. It need not be the same in the left and right states – this solution holds even if the initial states have different equations of state. Once the pressure in the star states has been obtained, the other variables in the star states can then be calculated using the equation of state. The normal velocity in the star state $v_{x,*}$ is given by

$$v_{x,*} = \frac{1}{2}\left(v_{x,L} + v_{x,R}\right) + \frac{1}{2}\left[f_R\left(p_*\right) - f_L\left(p_*\right)\right]. \tag{3.9}$$

Finding this solution requires solving an implicit equation. This is often computationally expensive, so for many cases it is often preferable to use an approximate solver. Such solvers shall be described shortly in Section 3.2.

3.1.1 The reactive Riemann problem

In the reactive Riemann problem, we are interested in systems where the reactions happen 'instantly', i.e. where the species creation rate $\dot{\omega} \to \infty$. If we work in the frame of the flow, we see that the jump in the energy is given by

$$Q = \int_{X=0}^{X=1} \frac{\mathrm{d}X}{\rho} = q, \tag{3.10}$$

where X is the species mass fraction and q is the specific binding energy. Therefore, we can model the system by replacing the source terms that model the heat release by a change in the equation of state across the sharp, nonlinear wave that models the reaction. Specifically, the total internal energy e that specifies the equation of state will change to $e \to e - q$ when the reaction takes place. This is precisely the model considered by Zhang and Zheng (1989) when constructing their Newtonian reactive Riemann problem solution.

3.1.1.1 Detonations

A detonation is a discontinuous reactive wave across which the pressure increases. The equations are identical to those for shocks, but the interpretation changes. All 'known' (pre-shock) variables have the reactive equation of state. All 'unknown' (post-shock) variables use the inert equation of state – the reaction has taken place across the discontinuity.

It is possible that the resulting detonation wave is unstable. In general, detonations fall into two classes: unstable weak detonations and stable strong detonations. For the stable strong detonations, the characteristic waves impinge on the discontinuity from both sides (see Figure 3.2a). For the unstable weak detonations, the characteristics only enter the discontinuity on one side.

If the solution takes the form of an unstable weak detonation, the single detonation wave is replaced by a compound wave. The compound wave is composed of the fastest detonation wave that is stable, a Chapman-Jouget (CJ) detonation, and a rarefaction. This CJ detonation is where the characteristic waves are parallel to the discontinuity, propagating at the same speed. It can be thought of the 'weakest' possible strong detonation. As the post-detonation pressure will now no longer match the required post-wave pressure, the additional rarefaction wave is needed. This is illustrated in Figure 3.2b.

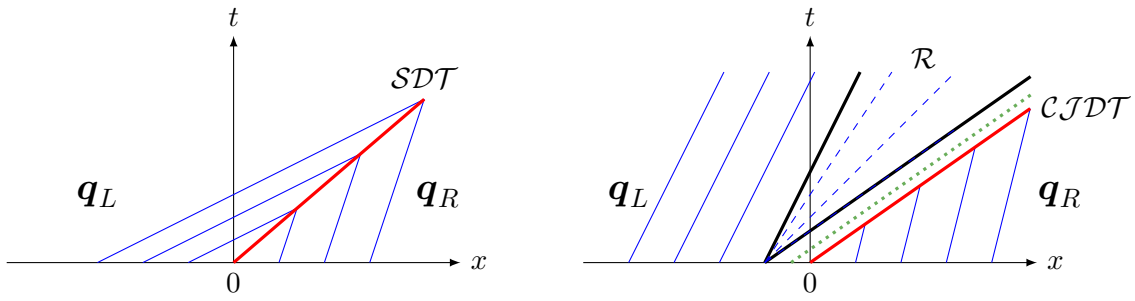

(A) Strong detonation (B) CJ detonation + rarefaction

FIGURE 3.2: Characteristics of detonations. For the strong detonation, characteristics (blue lines) from the initial left and right states meet, forming a strong detonation (thick red line).
For the unstable weak detonation, the characteristics impinge only on the right side of the discontinuity. It is replaced by a compound wave made up of a CJ detonation (red line) and a rarefaction (thick black lines and dashed blue lines). The dotted green line between them illustrates that the characteristics to the left of the CJ detonation are parallel to the discontinuity – in reality, the width of this region shrinks to zero and the detonation and rarefaction waves are attached together.

Using the calligraphic notation above used to represent the time evolution of the Riemann problem, we shall denote strong detonations by $\mathcal{SDT}$ and CJ detonations by $\mathcal{CJDT}$.

3.1.1.2 Deflagrations

A deflagration is a discontinuous reactive wave across which the pressure decreases. If the equation of state is convex and the reaction exothermic, then across a reactive discontinuous wave the pressure increases. Consequently, reactions cannot happen across a rarefaction wave and a discontinuity is required. However, this discontinuity will reduce the temperature along with the pressure. This means that, unless the material was already at the right temperature to react, any reaction across this wave would be unphysical. The solution for a deflagration therefore requires a compound wave.

The compound wave starts with an inert precursor shock which raises the temperature of the material to the ignition temperature. This follows the exact equations of the shock above (3.7). Next, there is a deflagration wave, across which the reaction takes place and the pressure drops. Again, this follows the shock equations for a rarefaction, but with the same interpretation as in the detonation case.

The deflagration wave need not be stable. As with detonations, deflagrations fall into two classes: stable weak deflagrations and unstable strong deflagrations. For the stable weak deflagrations, the characteristic waves from one side impinge on the discontinuity, but not the other (see Figure 3.3a). For the unstable strong deflagrations, neither set of characteristic waves impinge on the discontinuity.

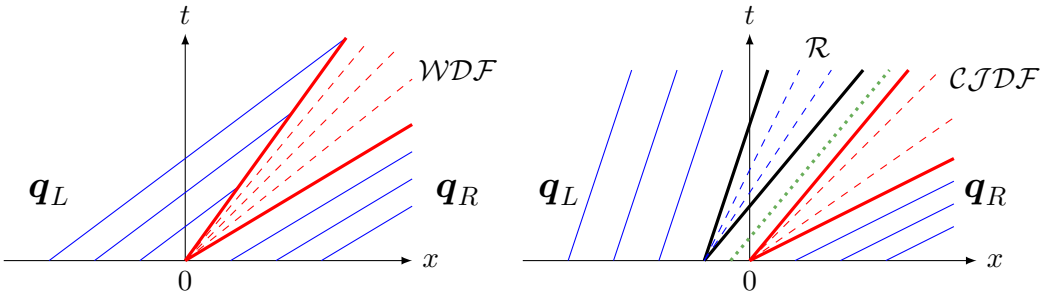

(A) Weak deflagration (B) CJ deflagration + rarefaction

FIGURE 3.3: Characteristics of deflagrations. For the weak deflagration, the characteristics (blue lines) impinge on the left side of the discontinuity only. To the right is a smooth transition zone (dashed red lines), on the left edge of which the reaction takes place.
In the case of an unstable strong deflagration, characteristics on either side of the discontinuity diverge. It is therefore replaced by a compound wave composed of a CJ deflagration (red lines) and a rarefaction (thick black lines and dashed blue lines). The dotted green line between them illustrates that the characteristics to the left of the CJ deflagration are parallel to the discontinuity – in reality, the width of this region shrinks to zero and the deflagration and rarefaction waves are attached together.

As for detonations, in the case of an unstable strong deflagration, the deflagration wave is replaced with a compound wave composed of a CJ deflagration (where the characteristics are parallel to the discontinuity) and a rarefaction wave (see Figure 3.3b).

In the calligraphic notation, we shall denote weak deflagrations by $\mathcal{WDF}$ and CJ deflagrations by $\mathcal{CJDF}$.

3.1.1.3 Pressure-volume plot

Another way to visualise the possible solutions for the reactive Riemann problem is to use a pressure-volume plot such as Figure 3.4 (Law, 2006). Here, we plot the dimensionless pressure $\hat{p} = p_*/p_S$ against the specific volume $\hat{v} = \frac{1}{\rho_*}/\frac{1}{\rho_S}$, such that the known state S corresponds to the point $(1, 1)$. From the Rankine-Hugoniot relations, which relate the known state to the unknown state, we obtain equations for the Rayleigh lines (solid red and blue)

$$\frac{p_* - p_S}{\frac{1}{\rho_*} - \frac{1}{\rho_S}} = -m^2, \tag{3.11}$$

where the mass flux $m = \rho_S v_S = \rho_* v_*$, and for the Hugoniot curve (solid black)

$$\frac{\gamma}{\gamma - 1}\left(\frac{p_*}{\rho_*} - \frac{p_S}{\rho_S}\right) - \frac{1}{2}\left(\frac{1}{\rho_*} + \frac{1}{\rho_S}\right)(p_* - p_S) = Q, \tag{3.12}$$

where Q is the reaction energy released per unit mass defined in (3.10).

The possible solutions are bounded by the intersections of the Rayleigh lines and the Hugoniot curve. The pressure jump across a detonation must therefore be somewhere

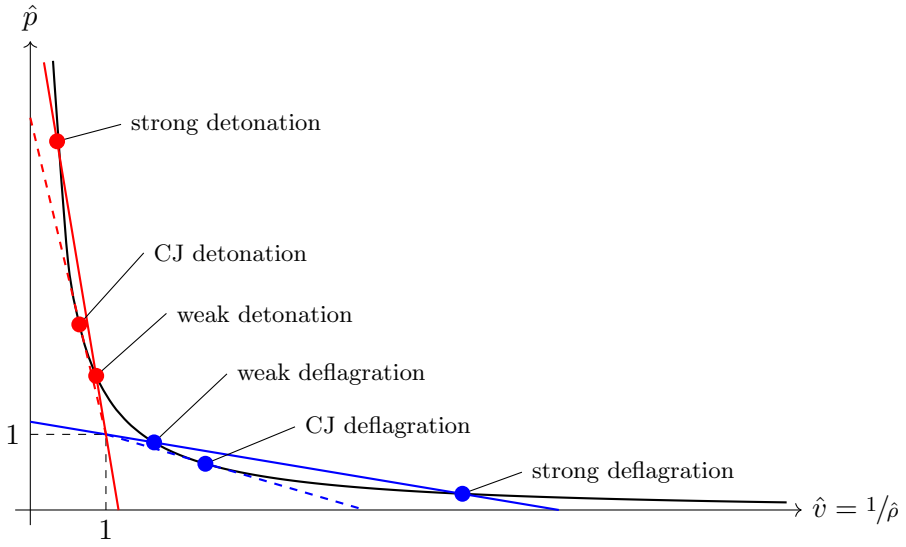

FIGURE 3.4: Solutions of the reactive Riemann problem, as bounded by the Rayleigh lines (solid red and blue lines) and the Hugoniot curve (solid black curve). The point $(1, 1)$, where $\hat{p} = \hat{v} = 1$, corresponds to the known state. The dashed lines indicate the tangents to the Hugoniot curve passing through the known state, which intersect with the curve at the CJ points.

between that of a weak detonation and a strong detonation (and similarly for deflagrations). The two tangent lines to the Hugoniot curve that pass through the known state intersect at the CJ points. These points separate strong detonations (deflagrations) from weak detonations (deflagrations), and occur where the flow in the unknown state is sonic, with velocity $v_* = c_{s,*}$.

3.1.2 The relativistic reactive Riemann problem

In Newtonian hydrodynamics, the solution to the Riemann problem depends only on the normal component of any vector quantities in the initial conditions. However, in relativistic systems, the Lorentz factor introduces a coupling between the normal and tangential components. As found by Pons et al. (2000); Rezzolla and Zanotti (2002) for the inert case, for high enough tangential velocities, the solution will smoothly transition from one wave pattern to another while maintaining the initial states otherwise unmodified. This is an entirely relativistic effect for which there is no Newtonian counterpart.

For the problem we are interested in (modelling neutron star oceans), we require general relativistic hydrodynamics. However, as described by Pons et al. (1998), we need only consider the simpler *special* relativistic Riemann problem. The Riemann problem for general relativistic fluids can be solved using special relativistic solvers following an appropriate coordinate transformation which transforms the coordinates at the interface to be locally Minkowskian. This follows from the equivalence principle, which states that

physical laws in a local inertial frame for an arbitrary spacetime have the same form as those of special relativity.

The introduction of this gauge transformation will change the component representation of the initial data and therefore change the detailed solution of the Riemann problem. However, the qualitative wave pattern is gauge invariant. Therefore transforming to the local inertial frame and using special relativistic solvers will allow us to determine the qualitative wave pattern with gravity.

In the following sections, we shall present the solutions for rarefactions and shocks, based on the equations given in Martí and Müller (1994); Pons et al. (2000); Martí and Müller (2003, 2015).

3.1.2.1 Rarefactions

Given the known state S ahead of the rarefaction wave, we wish to calculate the unknown state behind the wave. Across a rarefaction, the normal velocity v_x satisfies

$$\frac{\mathrm{d}v_x}{\mathrm{d}p} = \pm\frac{1}{\rho h W^2 c_s}\frac{1}{\sqrt{1+g\left(\xi_\pm, v_x, v_t\right)}}, \tag{3.13}$$

where v_t is the modulus of the tangential velocity, and we define the quantities

$$g(\xi_\pm, v_x, v_t) = \frac{v_t^2\left(\xi_\pm - 1\right)}{\left(1-\xi_\pm v_x\right)^2}, \tag{3.14}$$

and

$$\xi_\pm = \frac{v_x(1-c_s^2) \pm c_s\sqrt{(1-v^2)[1-v^2c_s^2-v_x^2(1-c_s)^2]}}{1-v^2c_s^2}. \tag{3.15}$$

The $+(-)$ sign corresponds to $S = R$ ($S = L$), and c_s is the local sound speed. The local speed $v = \sqrt{v_x^2+v_t^2}$ is used to calculate the Lorentz factor $W = (1-v^2)^{-1/2}$.

We also solve for the rest mass density and specific internal energy across the wave using

$$\frac{\mathrm{d}}{\mathrm{d}p}\begin{pmatrix}\rho\\ \epsilon\end{pmatrix} = \begin{pmatrix}\frac{1}{hc_s^2}\\ \frac{p}{\rho^2 h c_s^2}\end{pmatrix}. \tag{3.16}$$

To find the unknown state, we connect the known state ahead of the wave to the unknown state behind the wave by integrating (3.13). Using the fact that $hWv_t = \text{const}$ across a rarefaction wave, we can then calculate the tangential velocity as

$$v_t = h_S W_S v_{t,S}\left[\frac{1-v_x^2}{h^2+\left(h_S W_S v_{t,S}\right)^2}\right]^{\frac{1}{2}}, \tag{3.17}$$

where the S subscript denotes the value of the variable in state S. The relation for the tangential velocity holds across both continuous and discontinuous waves.

3.1.2.2 Shocks

Across a shock, the post-shock state can be found from the Taub adiabat (Thorne, 1973)

$$[\![h^2]\!] = \left(\frac{h_b}{\rho_b} + \frac{h_a}{\rho_a}\right)[\![p]\!], \tag{3.18}$$

where $[\![q]\!] = q_b - q_a$. In general, this is solved using the equation of state to obtain the post-shock enthalpy as a function of the post-shock pressure, $h = h(p)$.

From this we can compute the mass flux j across the shock

$$j(p) = \sqrt{\frac{p_S - p}{\frac{h_S^2 - h(p)^2}{p_S - p} - \frac{2h_S}{\rho_S}}}. \tag{3.19}$$

This gives the shock velocities

$$V_\pm(p) = \frac{\rho_S^2 W_S^2 v_{x,S} \pm j(p)^2\sqrt{1 + \rho_S^2 W_S^2\left(1 - v_{x,S}^2\right)/j(p)^2}}{\rho_S^2 W_S^2 + j(p)^2}. \tag{3.20}$$

Given the shock velocity we can compute the shock Lorentz factor $W_S = (1 - V_\pm^2)^{-1/2}$, from which the post-shock normal velocity is

$$v_x = \left(h_S W_S v_{x,S} \pm \frac{(p - p_S)}{j(p)\sqrt{1 - V_\pm(p)^2}}\right)\left[h_S W_S + (p - p_S)\left(\frac{1}{\rho_S W_S} \pm \frac{v_{x,S}}{j(p)\sqrt{1 - V_\pm(p)^2}}\right)\right]^{-1}. \tag{3.21}$$

The tangential velocity is the same as for the rarefaction wave in (3.17).

3.1.2.3 Pressure-volume plot

As we did in Section 3.1.1.3 for the Newtonian case, we can visualise the solutions of the relativistic reactive Riemann problem using a pressure-volume plot (Gao and Law, 2012). For relativistic flows, the Rayleigh and Hugoniot relations are given by

$$(\hat{p} - 1) = -u_S^2\left[\hat{v} - 1 + \frac{\gamma}{\gamma - 1}\left(\hat{p}\hat{v}^2 - 1\right) - \hat{q}\right], \tag{3.22}$$

$$\frac{\gamma + 1}{\gamma - 1}(\hat{p}\hat{v} - 1) - (\hat{p} - \hat{v}) - 2\hat{q} = \frac{\gamma}{\gamma - 1}\hat{p}(1 - \hat{v}^2) - \frac{\gamma}{(\gamma - 1)^2}(\hat{p}^2\hat{v}^2 - 1) + \left(\hat{p} + \frac{\gamma + 1}{\gamma - 1}\right)\hat{q} + \hat{q}^2, \tag{3.23}$$

where the normal 4-velocity for state S is $u_S = v_{x,S}W_S$, $\hat{p} = p_*/p_S$, $\hat{v} = \rho_S/\rho_*$ and $\hat{q} = \frac{\rho_S}{p_S}Q$. Comparing these to (3.11) and (3.12), the key difference is the appearance of the $\hat{q}$ term in the Rayleigh relation. This is because in the relativistic fluid equations, the reaction heat release term is present in the momentum conservation equation (through

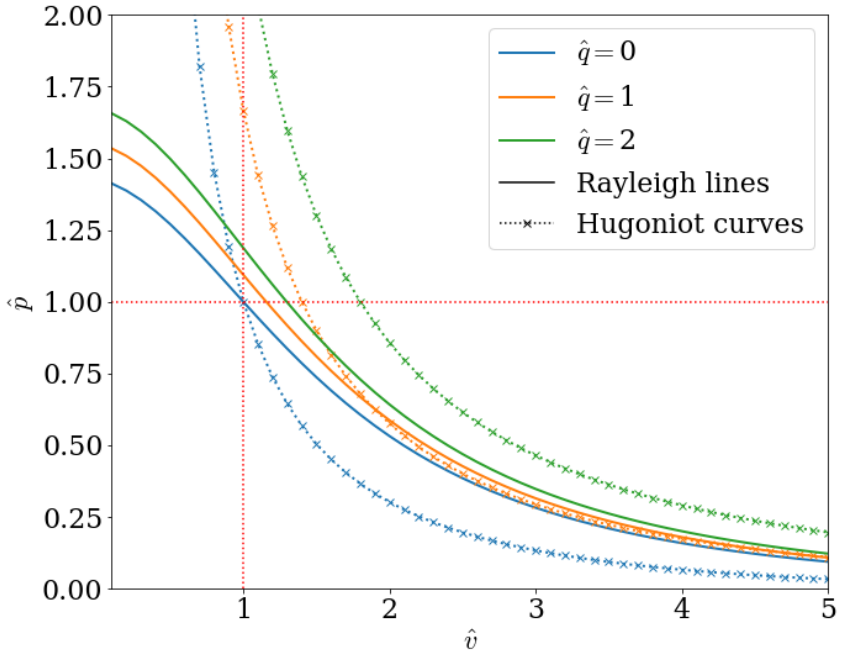

FIGURE 3.5: Pressure-volume plot for the relativistic reactive Riemann problem for normal 4-velocity $u_S = 0.35$ and various values of $\hat{q}$. Plotted are the Rayleigh lines and the Hugoniot curves. Unlike for the Newtonian case, both the Rayleigh lines and the Hugoniot curves are functions of functions of $\hat{q}$, with only the lines for $\hat{q}$ intersecting the (1,1) point.

the enthalpy). A consequence of this is that the Rayleigh lines are now a function of $\hat{q}$ and will no longer pass through the (1,1) point in the $\hat{p} - \hat{v}$ diagram if $\hat{q} \neq 0$.

This can be seen in Figure 3.5, where the Rayleigh lines and Hugoniot curves are plotted for various values of $\hat{q}$. As $\hat{q}$ increases, both the Rayleigh lines and Hugoniot curves move rightwards, away from the (1,1) point. By varying $\hat{q}$, the possible solutions of the problem changes. For $\hat{q} = 0$, there is a single intersection of the curves at (1,1): only weak deflagrations are a valid solution. For $\hat{q} = 1$, there are two intersections, indicating that both weak and strong deflagrations are valid for the range of $\hat{p}$ and $\hat{v}$ where the Rayleigh line is above the Hugoniot curve. For $\hat{q} = 2$, there are no intersections. There are therefore no valid deflagration solutions for this initial data.

In the Newtonian Rayleigh relation (3.11), the normal velocity appears in the mass flux term. In the corresponding relativistic relation (3.22), the tangential velocity v_t appears due to the introduction of the Lorentz factor in the normal 4-velocity. This means that the Rayleigh curve is now a function of v_t, and its intersection with the Hugoniot curve will change as v_t. This is illustrated in Figure 3.6, where the Rayleigh lines are plotted for a range of tangential velocities. For $v_t \lesssim 0.65$, there is only a single intersection of the Rayleigh line and the Hugoniot curve – only weak deflagrations are possible for this system. For $0.65 \lesssim v_t \lesssim 0.9$, the lines intersect twice, such that both weak and strong deflagrations are valid solutions for the range of $\hat{p}$ and $\hat{v}$ where the Rayleigh line

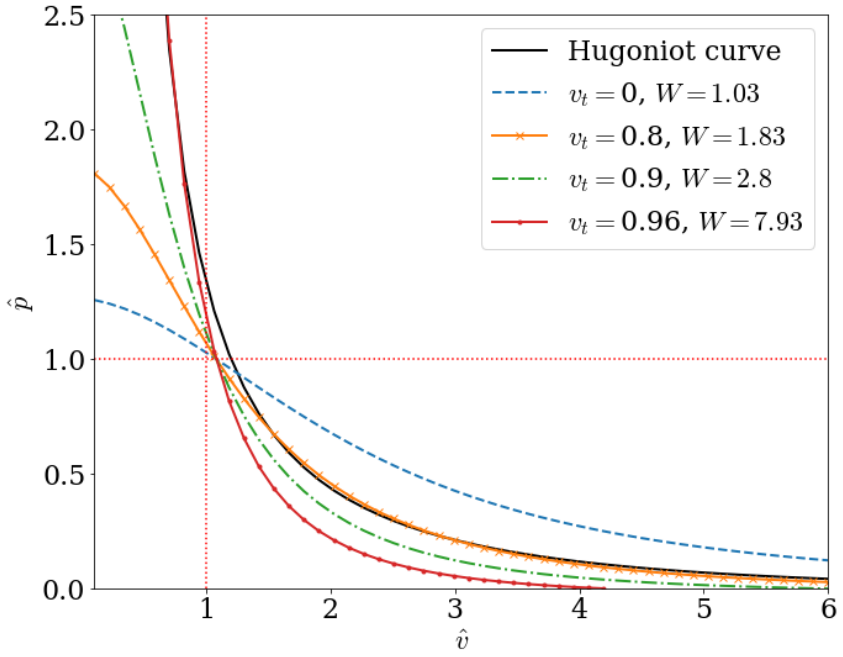

FIGURE 3.6: Pressure-volume plot for the relativistic reactive Riemann problem for normal 3-velocity $v_{x,S} = 0.25$ and $\hat{q} = 0.5$. Plotted are the Hugoniot curve and Rayleigh lines for various values of tangential velocity v_t. Unlike for the Newtonian case, the coupling of the tangential velocity via the Lorentz factor means that the range of possible solutions changes with the tangential velocity.

is above the Hugoniot curve. For $v_t \gtrsim 0.9$, there are no intersections of the curves in the bottom right deflagration quadrant, and so no deflagration solutions exist for the problem. However, it can be seen for the curve with $v_t = 0.96$ that if the tangential velocity is increased still further, the lines will now intersect in the detonation quadrant of the diagram. Weak detonations are therefore now valid solutions of the problem. As the Lorentz factor $W \to \infty$, the Rayleigh and Hugoniot curves still only intersect once. There are therefore no possible strong detonation solutions for this system for any value of the tangential velocity.

Another feature of the relativistic Rayleigh curves that can be seen in this plot is that they become less straight as W increases. In the Newtonian problem, the Rayleigh curves are straight lines. In the relativistic problem, the curvature of the curves increases with the Lorentz factor. It is this that changes their intersection with the Hugoniot curve and results in the changes in the range of possible solutions.

3.2 Approximate solvers

For fully compressible codes, we can solve the evolution equations using Godunov's method (Godunov, 1959), where intercell fluxes are treated as the solutions to local

Riemann problems (Toro et al., 1994; Toro, 1999). We then evolve the system by solving these Riemann problems.

Godunov's method can be summarised as follows for the hyperbolic system of PDEs

$$\partial_t \boldsymbol{q} + \partial_x \left(\boldsymbol{f}(\boldsymbol{q}) \right) = 0, \tag{3.24}$$

where $\boldsymbol{f}(\boldsymbol{q})$ is the flux vector. The intercell flux $\boldsymbol{f}_{i+1/2}^{\,n}$ at time $t = t^n$ is calculated by solving the Riemann problem to find the state at the cell boundary, $\boldsymbol{q}_*(\bar{\boldsymbol{q}}_i^{\,n}, \bar{\boldsymbol{q}}_{i+1}^{\,n})$, using the cell-averaged variables, $\bar{\boldsymbol{q}}$, on either side of the boundary as initial data. The intercell flux is then

$$\bar{\boldsymbol{f}}_{i+1/2}^{\,n} = \boldsymbol{f}\left(\boldsymbol{q}_*(\bar{\boldsymbol{q}}_i^{\,n}, \bar{\boldsymbol{q}}_{i+1}^{\,n}) \right). \tag{3.25}$$

Using this, the solution at time $t = t^{n+1}$ is given by

$$\boldsymbol{q}_i^{n+1} = \boldsymbol{q}_i^n - \frac{\Delta t}{\Delta x} \left(\boldsymbol{f}_{i+1/2}^{\,n} - \boldsymbol{f}_{i-1/2}^{\,n} \right), \tag{3.26}$$

where Δt is the timestep and Δx is the cell spacing. In the vast majority of cases, it is too computationally expensive and unnecessary to calculate the exact solution of the Riemann problem each time: we therefore use approximate solvers. This method is conservative and first order; there also exist higher order methods based on this approach known as *high-resolution shock-capturing* (HRSC) methods, such as the MUSCL-type solvers described in Section 1.5.2.

We implemented two approximate solvers, the RHLLE and RHLLC methods, in a relativistic reactive compressible solver that was built by extending the existing Newtonian compressible solver in `pyro` (Zingale, 2014). This compressible solver performs reasonably well for systems with moderate Mach number, however breaks down for high Mach number Ma $\gtrsim 10$. This could be improved by implementing alternative methods to deal with the stiff burning source term, for example the modified fractional step method of Helzel et al. (2000). Such methods resolve the sub-grid structure of the burning region, rather than smearing out the reaction across the entire cell.

3.2.1 RHLLE

An approximate solution for the relativistic Riemann problem can be found using the relativistic HLLE method, as described in Martí and Müller (2003, 2015) and derived by Schneider et al. (1993) based on the method of Harten, Lax and van Leer (Harten et al., 1983). Like this original Newtonian HLL method, it is positively conservative in the sense that it preserves initially positive densities, energies and pressures. It is based

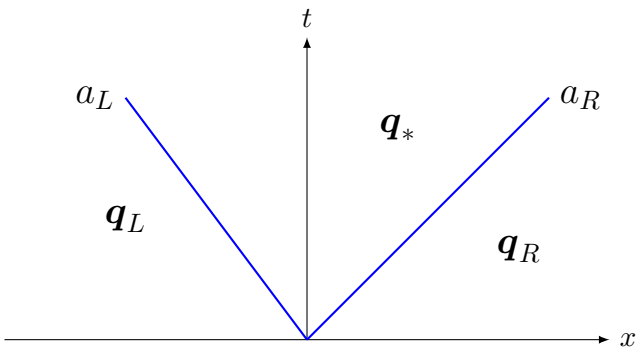

FIGURE 3.7: Wave pattern for the Riemann problem used in the HLLE method.

on an approximate solution with a single intermediate state:

$$\boldsymbol{q}^{\text{HLLE}}\left(\frac{x}{t};\boldsymbol{q}_L,\boldsymbol{u}_R\right)=\begin{cases}\boldsymbol{q}_L & \text{for } x<a_Lt,\\ \boldsymbol{q}_* & \text{for } a_Lt\le x\le a_Rt,\\ \boldsymbol{q}_R & \text{for } x>a_Rt,\end{cases} \tag{3.27}$$

where a_L, a_R are lower and upper bounds for the smallest and largest signal velocities and $\boldsymbol{q}=(D,S^i,\tau)$ is the state vector of conserved variables. The intermediate state is the average of the Riemann solution between the slowest and fastest signals and is given by

$$\boldsymbol{q}_*=\frac{a_R\boldsymbol{q}_R-a_L\boldsymbol{q}_L-\boldsymbol{f}(\boldsymbol{q}_R)+\boldsymbol{f}(\boldsymbol{q}_L)}{a_R-a_L}, \tag{3.28}$$

and the numerical flux vector by

$$\hat{\boldsymbol{f}}^{\text{HLLE}}=\frac{a_R^+\boldsymbol{f}(\boldsymbol{q}_L)-a_L^-\boldsymbol{f}(\boldsymbol{q}_R)+a_R^+a_L^-(\boldsymbol{q}_R-\boldsymbol{q}_L)}{a_R^+-a_L^-}, \tag{3.29}$$

where

$$a_L^-=\min\{0,a_L\},\qquad a_R^+=\max\{0,a_R\}. \tag{3.30}$$

The limiting characteristic wavespeeds are given by

$$a_L=\frac{\bar{v}-\bar{c}_s}{1-\bar{v}\bar{c}_s},\qquad a_R=\frac{\bar{v}+\bar{c}_s}{1+\bar{v}\bar{c}_s}, \tag{3.31}$$

where $\bar{v}$, $\bar{c}_s$ are the averages (arithmetic or Roe-like) between the left and right states of the fluid velocity and sound speed.

3.2.2 RHLLC

The HLLC method was originally developed by Toro et al. (1994), and builds on the HLLE method by including the contact surface in the wave pattern. This means it can better resolve contact discontinuities.

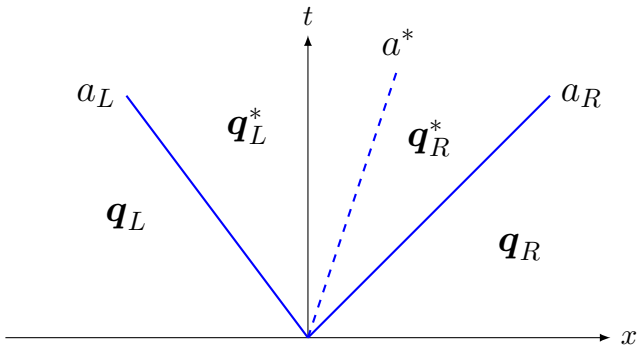

FIGURE 3.8: Wave pattern for the Riemann problem used in the HLLC method.

Mignone and Bodo (2005) extended the HLLC method to relativistic fluids. The full scheme shall be presented here as the notation and state vector used in the paper is different from that used here and there is a sign error in Equation (17). The solution is given by

$$\boldsymbol{q}^{\mathrm{HLLC}}\left(\frac{x}{t};\boldsymbol{q}_L,\boldsymbol{q}_R\right)=\begin{cases}\boldsymbol{q}_L & \text{for } x<a_Lt,\\ \boldsymbol{q}_L^* & \text{for } a_Lt\le x\le a^*t,\\ \boldsymbol{q}_R^* & \text{for } a^*t\le x\le a_Rt,\\ \boldsymbol{q}_R & \text{for } x>a_Rt,\end{cases} \tag{3.32}$$

where a^* is the speed of the contact discontinuity separating the left and right intermediate states L^*, R^*. The corresponding intercell numerical flux is given by

$$\hat{\boldsymbol{f}}^{\mathrm{HLLC}}=\begin{cases}\boldsymbol{f}_L & \text{for } x<a_Lt,\\ \boldsymbol{f}_L^* & \text{for } a_Lt\le x\le a^*t,\\ \boldsymbol{f}_R^* & \text{for } a^*t\le x\le a_Rt,\\ \boldsymbol{f}_R & \text{for } x>a_Rt,\end{cases}, \tag{3.33}$$

where the intermediate state fluxes $\boldsymbol{f}_L^*$ and $\boldsymbol{f}_R^*$ can be expressed using the Rankine-Hugoniot jump conditions as

$$a(\boldsymbol{q}^*-\boldsymbol{q})=\boldsymbol{f}^*-\boldsymbol{f}, \tag{3.34}$$

where all quantities are on the same side (the L or R subscript shall henceforth be dropped for clarity).

If a_L and a_R are given, then (3.34) represents a system of $2n$ equations (where n is the number of components of $\boldsymbol{q}$) for the $4n+1$ unknowns $\boldsymbol{q}_L^*, \boldsymbol{q}_R^*, \boldsymbol{f}_L^*, \boldsymbol{f}_R^*$ and a^*. We can obtain three additional constraints from requirements that the pressure and normal velocity are continuous across the contact wave: $v_{x,R}^*=v_{x,L}^*$ and $p_R^*=p_L^*$, and that $a^*=v_{x,R}^*=v_{x,L}^*$. This brings the total of equations to $2n+3$, which is still not enough to solve the system. Further assumptions must therefore be made on the intermediate

state fluxes $\boldsymbol{f}^*$. We assume that they can be written as

$$\boldsymbol{f}^* = \begin{pmatrix} D^* v_x^* \\ S_x^* v_x^* + p^* \\ S_y^* v_x^* \\ S_x^* - D^* v_x^* \end{pmatrix}. \tag{3.35}$$

Consequently, both $\boldsymbol{q}^*$ and $\boldsymbol{f}^*$ are expressed in terms of the five unknowns $D^*, v_x^*, S_y^*, \tau^*$ and p^*. The normal components of the momentum in the star region, $S_{x,L}^*$ and $S_{x,R}^*$, are not independent, as for consistency we require that $S_x^* = (\tau^* + p^* + D^*)v_x^*$. This gives 11 equations with 11 unknowns, so the system can now be solved.

Writing out (3.34) explicitly gives the following set of equations

$$\begin{aligned} D^*(a - a^*) &= D(a - v_x), \\ S_x^*(a - a^*) &= S_x(a - v_x) + p^* - p, \\ S_y^*(a - a^*) &= S_y(a - v_x), \\ \tau^*(a - a^*) &= \tau(a - v_x) + p^* a^* - p v_x. \end{aligned} \tag{3.36}$$

Substituting the second and last of these into the consistency condition for S_x^*, we get the following expression for a^* in terms of p^*

$$(A + a p^*)a^* = B + p^*, \tag{3.37}$$

where $A = (\tau + D)a - S_x$ and $B = S_x(a - v_x) - p$.

If we impose the condition $p_R^* = p_L^*$ across the contact discontinuity, a quadratic equation can be found for a^*,

$$f_E^{\mathrm{HLLE}}(a^*)^2 - (E^{\mathrm{HLLE}} + f_{S_x}^{\mathrm{HLLE}})a^* + S_x^{\mathrm{HLLE}} = 0, \tag{3.38}$$

where f_E^{HLLE} and $f_{S_x}^{\mathrm{HLLE}}$ are the energy and momentum components of the HLLE flux (note the HLLE flux used here is slightly different from the one above and uses a_L and a_R rather than a_L^- and a_R^+). E^{HLLE} and S_x^{HLLE} are the energy and momentum components of the HLLE state vector, where the energy $E = \tau + D$. Only the root with the minus sign is physically acceptable, since a^* must lie in the range $(-c, c)$.

To solve the system, a^* is found using (3.38), then this is used to calculate p^* from (3.37). These can then be used to calculate the components of $\boldsymbol{q}^*$ and therefore $\boldsymbol{f}^*$ using (3.36).

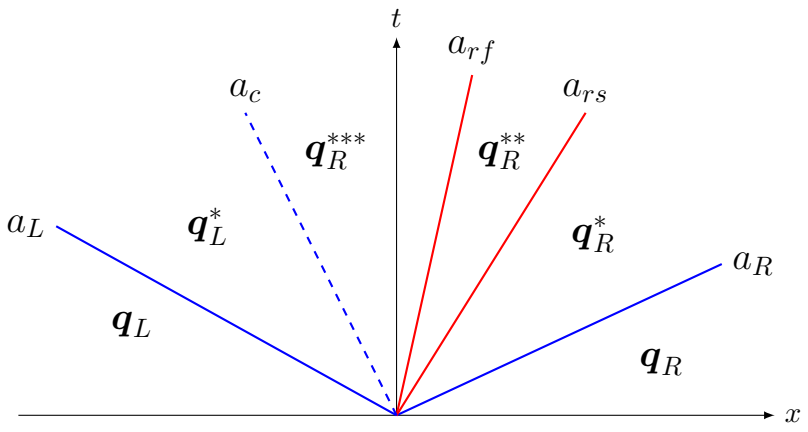

FIGURE 3.9: Wave pattern for the reactive Riemann problem, adapted from (Beccantini and Studer, 2010). a_c is the contact wave speed, a_{rf} the reaction rarefaction wave speed and a_{rs} the reaction shock wave speed.

3.2.3 Reactions

If reactions are included, we obtain a source term such that the system of equations to be solved is now given by

$$\frac{\partial \boldsymbol{q}}{\partial t} + \frac{\partial \boldsymbol{f}^{(i)}(\boldsymbol{q})}{\partial x^i} = \boldsymbol{S}. \tag{3.39}$$

The state vector of conservative variables is extended to account for the evolution of the baryon fraction X,

$$\boldsymbol{q} = \begin{pmatrix} D \\ S^j \\ \tau \\ DX \end{pmatrix}, \tag{3.40}$$

and the fluxes are extended to produce

$$\boldsymbol{f}^{(i)} = \begin{pmatrix} Dv^i \\ S^j v^i + p\delta^{ij} \\ \tau v^i + pv^i \\ DXv^i \end{pmatrix}. \tag{3.41}$$

The source term is given by

$$\boldsymbol{S} = \begin{pmatrix} 0 \\ 0 \\ Dq\dot{\omega} \\ D\dot{\omega} \end{pmatrix}, \tag{3.42}$$

where q is the specific energy release of burning.

In the HLLC method, we obtain the additional jump condition

$$DX^*(a - a^*) = DX(a - v_x^*). \tag{3.43}$$

Despite the fact that introducing burning adds an additional burning wave to the wave pattern (see Figure 3.9), with high enough resolution this shall be resolved without changing the numerical method to include this wave (just as HLLE will resolve contact discontinuities properly at high enough resolution).

3.3 Numerical results

The following section contains results of simulations performed for a range of 1d systems using the exact and compressible solvers `R3D2` and `pyro`, and for a 2d system using the compressible solver.

3.3.1 1d simulations using the exact solver

`R3D2` solves the reactive relativistic Riemann problem by solving the equations given in Section 3.1.1 and Section 3.1.2 to find the characteristic waves and the intermediate states. To do this, it makes use of root finders and ODE solvers from the `numpy` and `scipy` Python libraries. Given the solution's wave pattern, wavespeeds and states, the full solution can be plotted.

3.3.1.1 Inert relativistic Riemann problem

The effect of tangential velocity in the inert relativistic Riemann problem can be seen in Figure 3.10. In this example, the final state is made up of a left-going rarefaction wave, a contact wave and a right-going shock wave: $\mathcal{R}_{\leftarrow}\,\mathcal{C}\,\mathcal{S}_{\rightarrow}$. We can see that the introduction of a tangential velocity changes the intermediate wave states. In the case with a non-zero right state tangential velocity $v_t(L, R) = (0, 0.9)$, the density and pressure in the intermediate states is increased, whereas in the case with non-zero tangential velocities for both states $v_t(L, R) = (0.9, 0.9)$, they are decreased. For both cases, the normal velocity in the intermediate states is reduced.

3.3.1.2 Reactive relativistic Riemann problem

The wave pattern of the solution to the reactive Riemann problem can be very sensitive to initial conditions. In Figure 3.11, the effect on a detonation of changing the energy of reaction q can be seen. For $q = 0.1$, the solution is a single left-going strong detonation wave, $\mathcal{SDT}_{\leftarrow}$. Decreasing to $q < 0.1$ causes a contact wave and a right-going shock wave to be produced, $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}\,\mathcal{S}_{\rightarrow}$, as the decrease in reaction energy increases the pressure and density in the intermediate star state. Increasing to $q > 0.1$ produces a contact

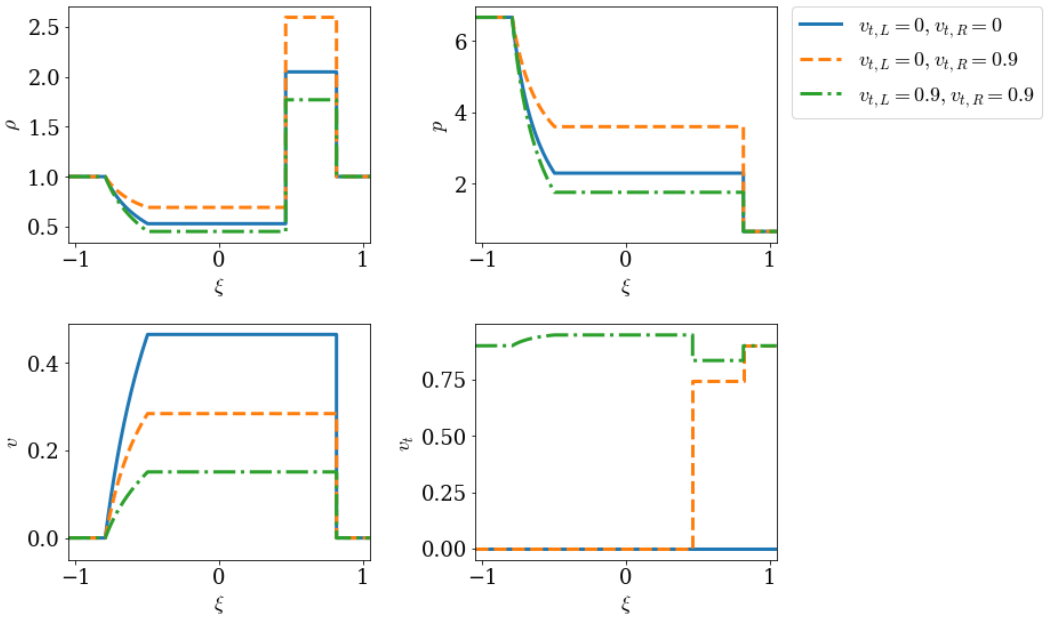

FIGURE 3.10: Exact solution of the inert relativistic Riemann problem for different initial tangential velocities, calculated using `R3D2`. All cases share the initial conditions $(\rho, v_x, \varepsilon)_L = (1, 0, 10)$, $(\rho, v_x, \varepsilon)_R = (1, 0, 1)$. For the blue curves, $v_t(L, R) = (0, 0)$, for the orange dashed curves, $v_t(L, R) = (0, 0.9)$, and for the green dot-dashed curves, $v_t(L, R) = (0.9, 0.9)$. The final state consists of a left-going rarefaction wave, a contact wave and a right-going shock wave: $\mathcal{R}_{\leftarrow}\,\mathcal{C}\,\mathcal{S}_{\rightarrow}$. Increasing the tangential velocity alters the intermediate state, reducing the normal velocity.

wave and a right-going rarefaction wave, $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}\,\mathcal{R}_{\rightarrow}$, as the increase in reaction energy decreases the pressure, density and velocity in the intermediate star state.

Similarly, varying the reaction energy q also changes the wave pattern in the case of a deflagration, as seen in Figure 3.12. For $q = 0.1$, the solution is a single left-going deflagration wave, $\mathcal{CJDF}_{\leftarrow}\,\mathcal{R}_{\leftarrow}$. For $q < 0.1$, a right-going shock wave is produced, $\mathcal{CJDF}_{\leftarrow}\,\mathcal{R}_{\leftarrow}\,\mathcal{C}\,\mathcal{S}_{\rightarrow}$, and for $q > 0.1$, a right-going rarefaction wave is produced, $\mathcal{CJDF}_{\leftarrow}\,\mathcal{R}_{\leftarrow}\,\mathcal{C}\,\mathcal{R}_{\rightarrow}$.

As can be seen in Figure 3.13, varying the reaction energy can also cause the reactive wave to transition. In this example, increasing the reaction energy causes the left-going reactive wave in the system to transition from a CJ deflagration to a weak deflagration, then on to a CJ detonation.

3.3.1.3 Tangential velocity

As described above, Pons et al. (2000); Rezzolla and Zanotti (2002) found that sufficiently high tangential velocities can change the wave pattern in the inert relativistic Riemann problem. Here we investigate this for the reactive relativistic system to see whether it could be relevant for X-ray bursts on rapidly rotating neutron stars: given

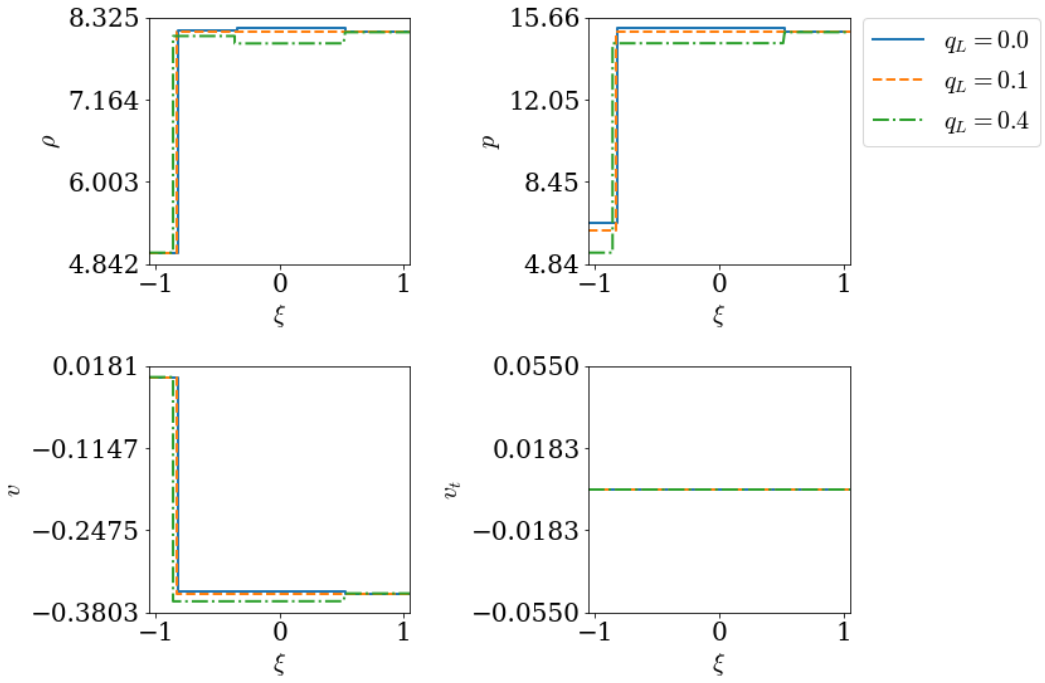

FIGURE 3.11: Exact solution of a strong detonation wave for different values of $q = (0.0, 0.1, 0.4)$, corresponding to the solid blue, dashed orange and dash-dotted green curves respectively. The initial conditions are $(\rho, v_x, v_t, \varepsilon)_L = (5, 0, 0, 2)$, $(\rho, v_x, v_t, \varepsilon)_R = (8.114, -0.349, 0, 2.773)$. For $q = 0.1$, the solution is $\mathcal{SDT}_{\leftarrow}$. Varying q introduces a contact wave and a right-going wave. Decreasing to $q < 0.1$ instead produces $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}\,\mathcal{S}_{\rightarrow}$, and increasing to $q > 0.1$ produces $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}\,\mathcal{R}_{\rightarrow}$.

the conditions in the ocean, would be possible for a fast enough tangential velocity to develop that could induce a deflagration to detonation transition.

As seen in Figure 3.14, it was found that varying the tangential velocity did indeed change the wave pattern. For $v_{t,L} = 0$, the solution consists of a left-going strong detonation, a contact wave and a right-going rarefaction, $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}\,\mathcal{R}_{\rightarrow}$. As the tangential velocity increases, the pressure and density in the intermediate star state increases until at $v_{t,L} \simeq 0.70541$ the rarefaction disappears and there is just a strong detonation and contact wave, $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}$. For higher tangential velocities, there is then a right-going shock wave, $\mathcal{SDT}_{\leftarrow}\,\mathcal{C}\,\mathcal{S}_{\rightarrow}$.

In the Figure 3.14, a transition is induced in the non-reactive wave by introducing a non-zero tangential velocity. As shown in Figure 3.15, we can also induce a transition in the reactive wave by introducing a non-zero tangential velocity. Increasing the tangential velocity in the right state from zero, we can observe the reactive wave to transition from a CJ deflagration to a weak deflagration, to a CJ detonation then finally on to a strong detonation.

We can define a critical tangential velocity at which the wave pattern transitions between these different states. This critical tangential velocity is plotted against the density ρ_L in Figure 3.16. The solver that was used here to find the critical velocity breaks down

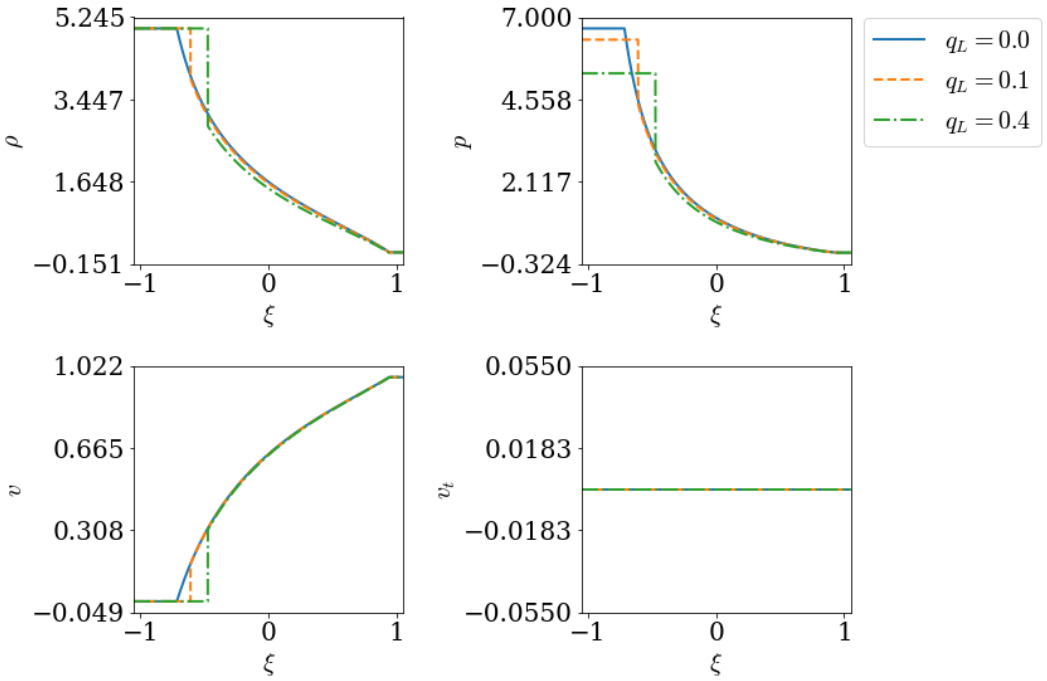

FIGURE 3.12: Exact solution of a deflagration wave for different values of $q = (0.0, 0.1, 0.4)$, corresponding to the solid blue, dashed orange and dash-dotted green curves respectively. The initial conditions are $(\rho, v_x, v_t, \varepsilon)_L = (5, 0, 0, 2.0)$, $(\rho, v_x, v_t, \varepsilon)_R = (0.1, 0.973, 0, 0.149)$. For $q = 0.1$, the solution is $\mathcal{CJDF}_{\leftarrow}\mathcal{R}_{\leftarrow}$. As for the detonation, varying q introduces a contact wave and a right-going wave. For $q < 0.1$, the solution becomes $\mathcal{CJDF}_{\leftarrow}\mathcal{R}_{\leftarrow}\mathcal{CS}_{\rightarrow}$, and for $q > 0.1$ it is $\mathcal{CJDF}_{\leftarrow}\mathcal{R}_{\leftarrow}\mathcal{CR}_{\rightarrow}$.

for $v_{t,crit} \lesssim 0.06$, but it can be seen that it tends towards $v_{t,crit} = 0$ at $\rho \sim 7$. At this point, the pressure in the star state $p_* = p_R$, such that the right-going wave transitions from a rarefaction wave to a shock wave.

From Figure 3.16, we can see that unless the system with zero tangential velocity is already very near the deflagration to detonation transition point, the tangential velocity required to induce a transition is close to the speed of light. In neutron star oceans, the flame is very unlikely to come close to the deflagration to detonation transition point and the star does not rotate rapidly enough for such high tangential velocities to be obtained. Consequently, this effect is not relevant for these systems.

3.3.1.4 Pressure-velocity plots

Another way to to look at the Riemann problem is by plotting the solution in pressure-velocity, p-v, phase space. The velocity plotted here is the normal velocity v_x. Pressure-velocity plots can be thought of a graphical way of solving the Riemann problem to find the wave pattern and intermediate state. The Hugoniot curves of the left and right states are plotted. Their point of intersection determines the intermediate state. By considering the change in pressure and velocity along the lines joining the left, intermediate and right states, the nature of the waves can be determined. Moving from the

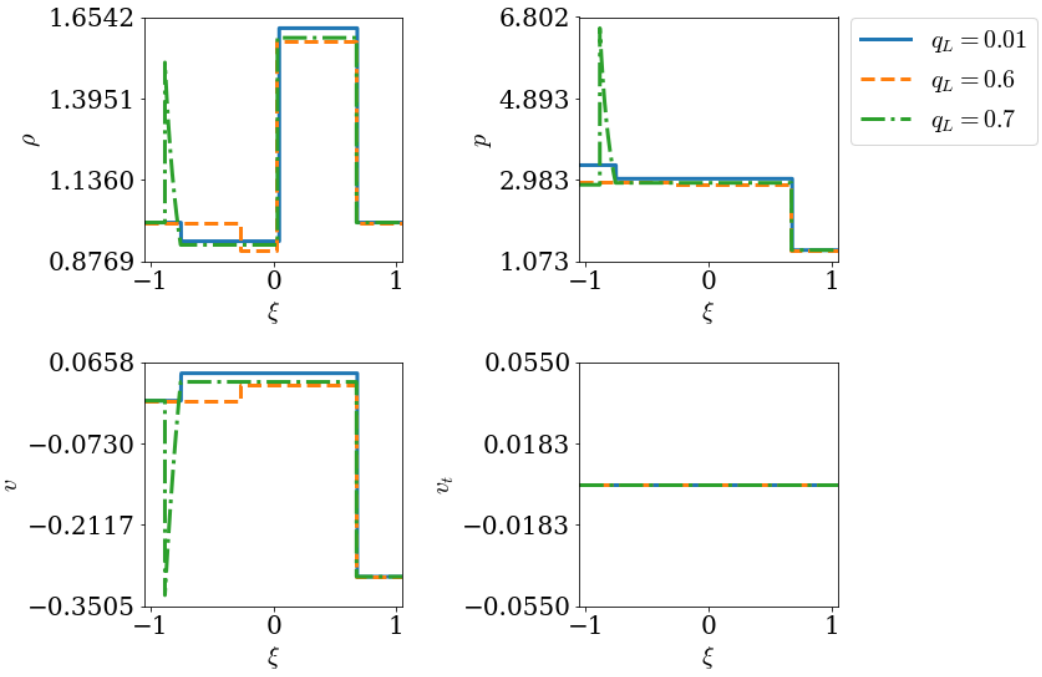

FIGURE 3.13: Deflagration to detonation transition for different values of $q = (0.01, 0.6, 0.7)$, corresponding to the solid blue, dashed orange and dotted green curves respectively. The initial conditions are $(\rho, v_x, v_t, \varepsilon)_L = (1, 0, 0, 5.0)$, $(\rho, v_x, v_t, \varepsilon)_R = (1, -0.3, 0, 2)$. For $q = 0.01$, the solution is a CJ deflagration, $(\mathcal{CJDF}_{\leftarrow}\mathcal{R}_{\leftarrow})\mathcal{CS}_{\rightarrow}$. As q is increased to $q = 0.6$, the CJ deflagration becomes a weak deflagration, $\mathcal{WDF}_{\leftarrow}\mathcal{CS}_{\rightarrow}$, and at $q = 0.7$ this has transitioned to a CJ detonation, $(\mathcal{CJDT}_{\leftarrow}\mathcal{R}_{\leftarrow})\mathcal{CS}_{\rightarrow}$.

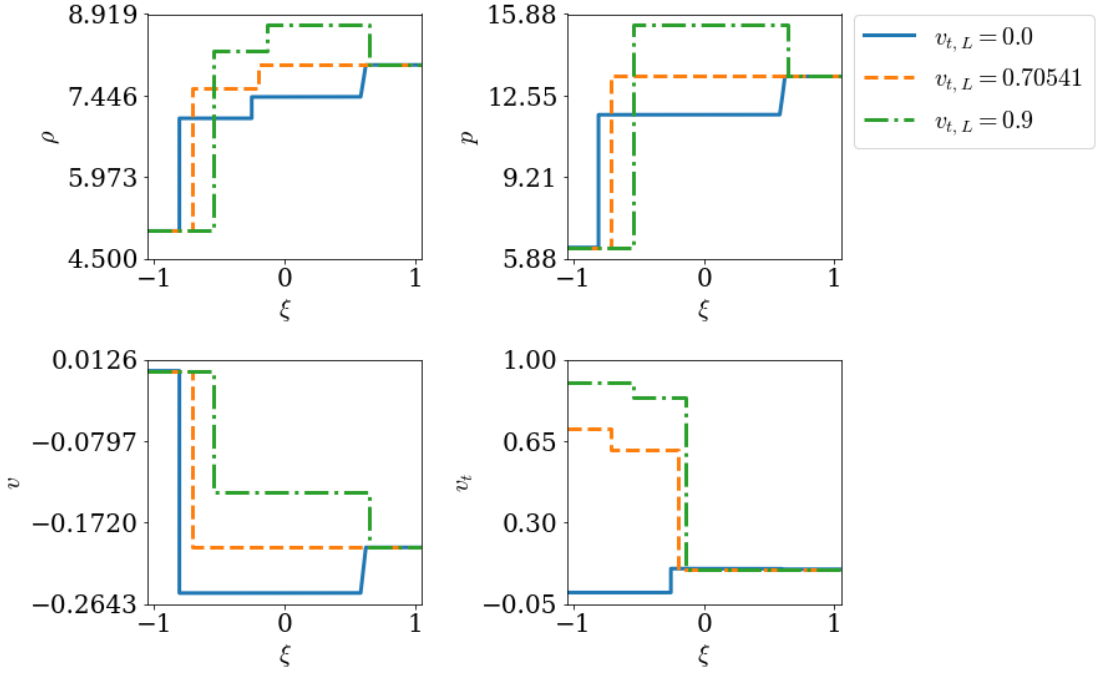

FIGURE 3.14: Exact solution of a strong detonation wave for different values of $v_{t,L} = (0.0, 0.70541, 0.9)$, corresponding to the solid blue, dashed orange and dash-dotted green lines respectively. The initial conditions are $(\rho, v_x, \varepsilon)_L = (5, 0, 2.0)$, $(\rho, v_x, v_t, \varepsilon)_R = (8, -0.2, 0.0, 2.5)$. For $v_{t,L} = 0$, the solution is $\mathcal{SDT}_{\leftarrow}\mathcal{CR}_{\rightarrow}$. At $v_{t,L} \simeq 0.70541$, the rarefaction disappears to produce $\mathcal{SDT}_{\leftarrow}\mathcal{C}$. At even higher tangential velocities, the solution becomes $\mathcal{SDT}_{\leftarrow}\mathcal{CS}_{\rightarrow}$

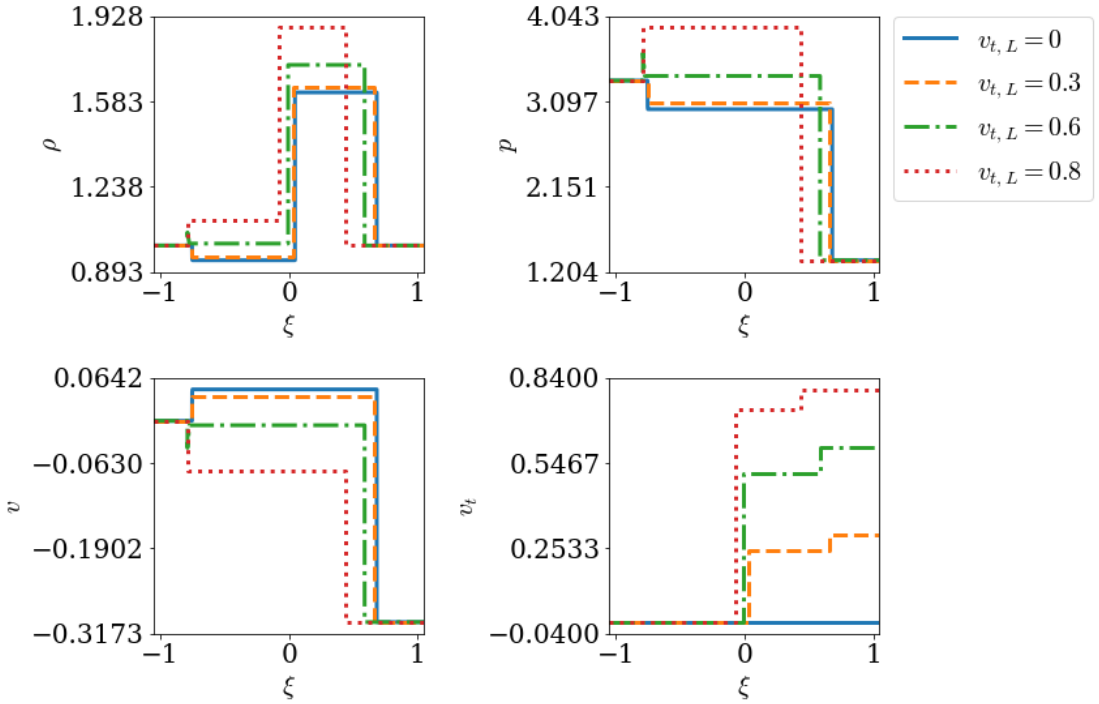

FIGURE 3.15: Deflagration to detonation transition for different values of $v_{t,L} = (0.0, 0.3, 0.6, 0.8)$, corresponding to the solid blue, dashed orange, dash-dotted green and dotted red lines respectively. The initial conditions are $(\rho, v_x, v_t, \varepsilon)_L = (1, 0, 0, 5.0)$, $(\rho, v_x, v_t, \varepsilon)_R = (1, -0.3, v_t, 2)$. For $v_{t,L} = 0$, the solution is $(\mathcal{CJDF}_{\leftarrow}\mathcal{R}_{\leftarrow})\mathcal{C}\mathcal{S}_{\rightarrow}$. At $v_{t,L} = 0.3$, the CJ deflagration has transitioned to a weak deflagration, producing the solution $\mathcal{WDF}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow}$. At $v_t = 0.6$, the deflagration has transitioned to a CJ detonation, $(\mathcal{CJDT}_{\leftarrow}\mathcal{R}_{\leftarrow})\mathcal{C}\mathcal{S}_{\rightarrow}$, then at $v_t = 0.8$, this has transitioned to a strong detonation, $\mathcal{SDT}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow}$.

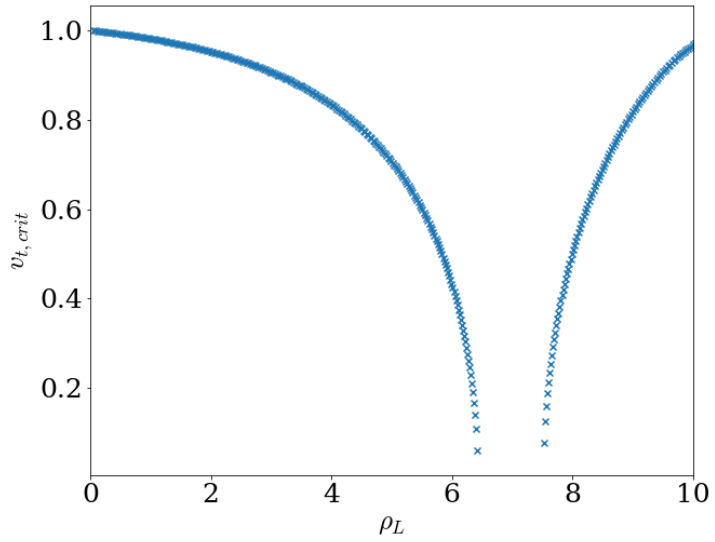

FIGURE 3.16: Plot of the critical tangential velocity for different values of ρ_L. The initial conditions are $(v_x, \varepsilon)_L = (0, 2.0)$, $(\rho, v_x, v_t, \varepsilon)_R = (8, -0.2, 0.0, 2.5)$. The method used to find the critical tangential velocity has failed to converge for $v_{t,crit} \lesssim 0.06$, hence the lack of points in this region.

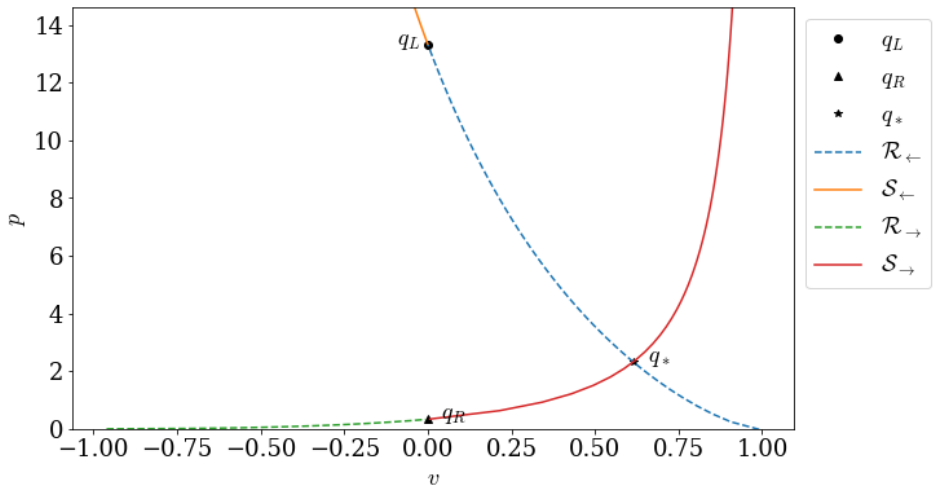

FIGURE 3.17: Pressure-velocity plot for the inert relativistic Riemann problem with initial conditions $(\rho, v_x, v_t, \varepsilon)_L = (10, 0, 0, 2)$, $(\rho, v_x, v_t, \varepsilon)_R = (2, 0, 0, 0.5)$ and solution $\mathcal{R}_\leftarrow \mathcal{C} \mathcal{S}_\rightarrow$.

left/right state to the intermediate state, a positive change in pressure along the line connecting the states corresponds to a shock (and a negative change to a rarefaction).

Inert relativistic Riemann problem

Figure 3.17 shows the pressure-velocity plot for an inert problem with the solution $\mathcal{R}_\leftarrow \mathcal{C} \mathcal{S}_\rightarrow$. The plot shows that the left and right states $\boldsymbol{q}_L$, $\boldsymbol{q}_R$ are connected by a left-going rarefaction (dashed blue line) and a right-going shock (solid red line). The two curves intersect at the intermediate star state $\boldsymbol{q}_*$.

Reactive relativistic Riemann problem

Figure 3.18 shows the plot for a reactive problem whose solution consists of a compound CJ deflagration-rarefaction wave preceded by a shock: $\mathcal{R}_\rightarrow \mathcal{C}\mathcal{J}\mathcal{D}\mathcal{F}_\rightarrow \mathcal{S}_\rightarrow$. Beginning with the unburnt fluid in the right state $\boldsymbol{q}_R$, the precursor shock wave (solid orange line) raises the temperature of the fluid until it reaches the lower ignition point, marked i, and the gas ignites. There is then a discontinuity in the phase space as the burning reaction increases the pressure and velocity of the gas to that of the upper ignition point, which intersects the curve with the equation of state of the burnt gas (solid green line). This point is intersected by the Chapman-Jouget deflagration wave, which then joins to the rarefaction wave (dashed blue line) at the Chapman-Jouget point CJ.

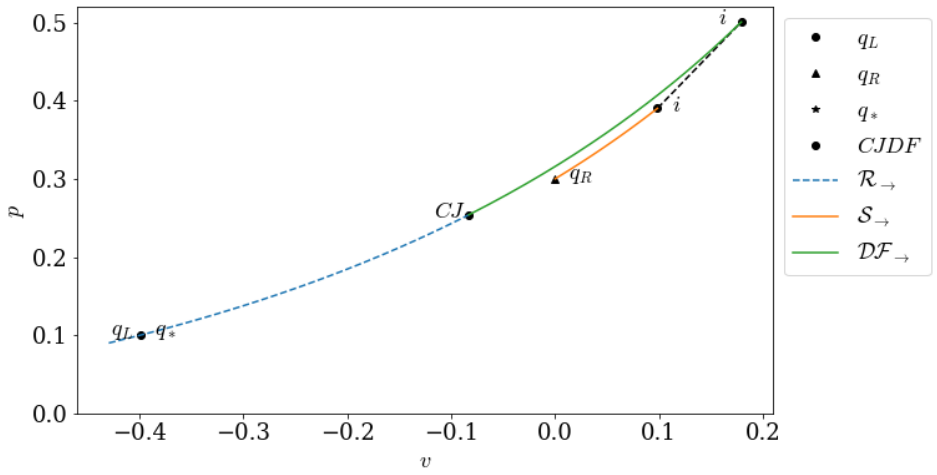

FIGURE 3.18: Pressure-velocity plot for the reactive relativistic Riemann problem with initial conditions $(\rho, v_x, v_t, \varepsilon)_L = (0.2432, -0.3992, 0, 0.6168)$, $(\rho, v_x, v_t, \varepsilon)_R = (0.5, 0, 0, 1)$ and solution $\mathcal{R}_\rightarrow \mathcal{CJDF}_\rightarrow \mathcal{S}_\rightarrow$. In the initial data, the burnt fluid is on the left and the unburnt fluid on the right.

3.3.2 1d simulations using the compressible solver

In the previous section, we demonstrated the performance of the exact solver `R3D2`. In this section, we shall use this as a benchmark to test the reactive relativistic compressible solver implemented in `pyro` and investigate how accurate it is.

Inert relativistic Riemann problem

Figure 3.19 shows the solutions for the inert problem with solution $\mathcal{R}_\leftarrow \mathcal{C} \mathcal{S}_\rightarrow$. It can be seen that the compressible solver performs well, with a slight smoothing of the pressure and normal velocity at the tail of the rarefaction. This is to be expected: as described in Section 1.5.2.5, the use of slope limiters in finite volume methods leads to clipping at extrema.

Reactive relativistic Riemann problem

Figure 3.20 shows the solutions for the reactive problem, where the initial conditions have produced a left-going deflagration wave and a right-going shock wave: $\mathcal{CJDF}_\leftarrow \mathcal{R}_\leftarrow \mathcal{C} \mathcal{S}_\rightarrow$. The solution found by the compressible solver follows the exact solution fairly closely away from the waves, but there is a lot of smoothing and some overshooting at the discontinuities. The compressible solver has also failed to capture the CJ deflagration on the left hand side, with the solution instead taking the form of an unstable strong deflagration. Given that the solver performed much better for the previous inert problem, this indicates that the significant reduction in accuracy is largely due to the introduction of the reactive source terms. Stiff source terms such as these are often difficult to model

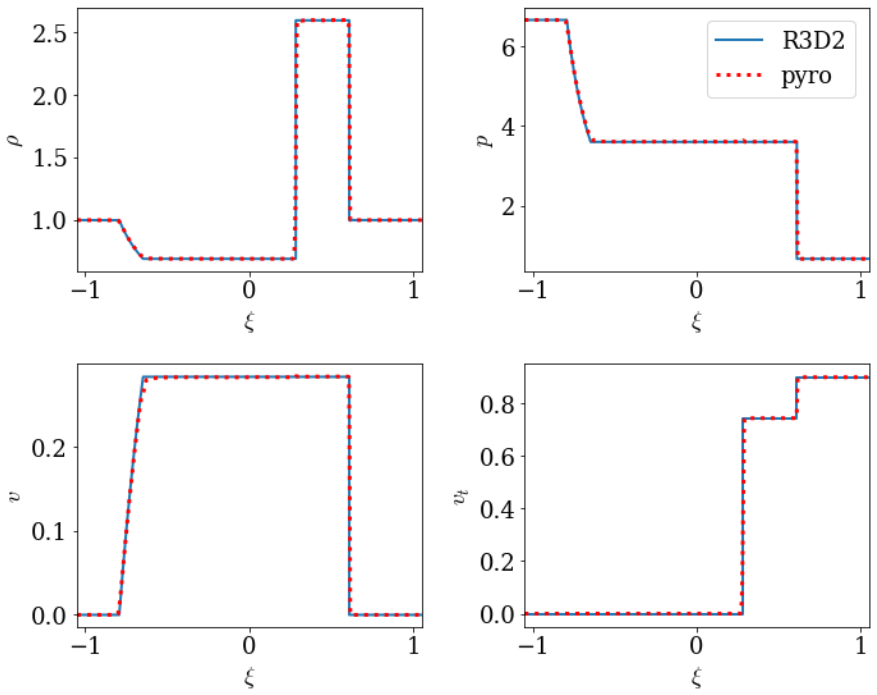

FIGURE 3.19: Solution of the inert relativistic Riemann problem with the initial conditions $(\rho, v_x, v_t, \varepsilon)_L = (1, 0, 0, 10)$, $(\rho, v_x, v_t, \varepsilon)_R = (1, 0, 0.9, 1)$ and solution $\mathcal{R}_{\leftarrow}\mathcal{C}\mathcal{S}_{\rightarrow}$. The solid blue curve shows the solution found using the exact solver, the dotted red curve using the compressible solver with a resolution of 256 in the x-direction.

using standard schemes, especially if the reaction zone is underresolved and the region where the source terms become non-negligible is smeared out across a larger region of the grid.

RHLLE vs RHLLC

The effect of using the RHLLC solver rather than the RHLLE solver can be seen in Figure 3.21 for a test with resolution of 256 in the x-direction. It can be seen that the RHLLE solver produces some smearing of the data across the contact wave. This smearing decreases as the resolution is increased. As expected, the RHLLE solver performs about as well as the RHLLC solver at resolving the shock and rarefaction waves.

3.3.3 2d simulations using the compressible solver

3.3.3.1 Bubble in shock tube

The compressible solver can also be used to model 2d systems. Here we use it to investigate the case of a bubble in shock tube. In this system, a hot bubble of unburnt fluid was placed in a tube of cooler, burnt fluid. A resolution of 512×512 was used. A shock wave travelling at relativistic speed was then sent along the tube in order to

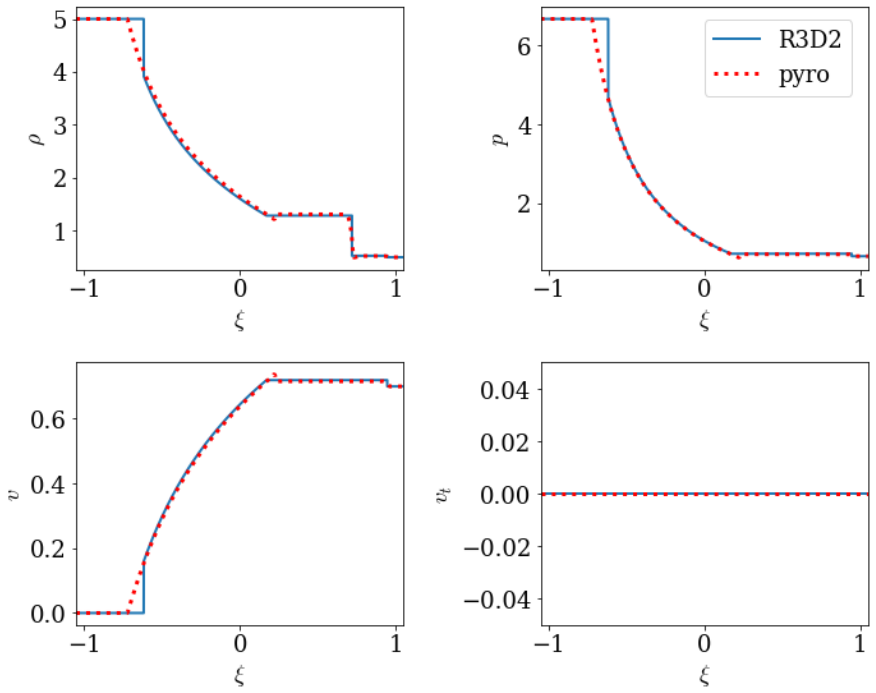

FIGURE 3.20: Solution of the reactive relativistic Riemann problem with the initial conditions $(\rho, v_x, v_t, \varepsilon)_L = (5, 0, 0, 2.1)$, $(\rho, v_x, v_t, \varepsilon)_R = (0.5, 0.7, 0, 2)$. In the initial data, the unburnt fluid is on the left and the burnt fluid on the right. The solid blue curve shows the solution found using the exact solver, the dotted red curve using the compressible solver with a resolution of 256 in the x-direction.

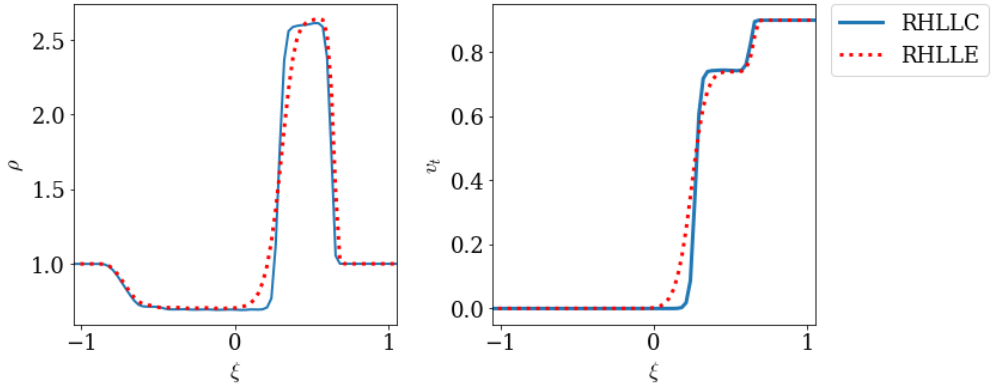

FIGURE 3.21: Solution of the inert relativistic Riemann problem using the RHLLE and RHLLC solvers in `pyro` with a resolution of 256 in the x-direction. The initial conditions were $(\rho, v_x, v_t, \varepsilon)_L = (1, 0, 0, 10)$, $(\rho, v_x, v_t, \varepsilon)_R = (1, 0, 0.9, 1)$. It can be seen that there is some smearing at the contact discontinuity by the RHLLE solver.

investigate the behaviour of the bubble under the impact of the wave. The test was carried out for both the inert and reactive cases so as to determine which effects were due to the shock and which due to the burning reactions. The code was able to cope with systems containing shocks with Mach number Ma $\lesssim 5$, however at Ma $= 10$ the code was unstable at the shock front for the resolutions tested.

Among the quantities plotted for these systems is the Schlieren number. Plots of this quantity help to clearly illustrate density discontinuities, so are particularly useful for systems with shock waves. The Schlieren number $\mathcal{S}$ is calculated as

$$\mathcal{S} = \left|\nabla^2 \rho\right| . \tag{3.44}$$

Figure 3.22 shows the system just before the shock wave with Ma $= 5$ makes contact with the bubble. In Figure 3.23, we plot the bubble after the shock wave has passed through, with the inert case shown in the top half of each plot and the reactive case in the lower half. Comparing the plots before and after the shock wave has passed through, it can be seen that the shock wave deforms the bubble, increasing its density and changing its shape to that of a disc whose edges have curled inwards. Comparing the inert and reactive bubbles, the density and Schlieren plots are qualitatively very similar. However, looking at the species creation rate and the baryon fraction, it is clear that the burning has changed the composition of the bubble. Comparing the density and species baryon fraction plots in the reactive case, it can be seen that ahead of the remaining unburnt fluid is a disc of burnt fluid with lower density – this is the fluid that was originally in the bubble that has now burnt due to the impact of the shock wave. Behind this disc, the unburnt fluid is still burning – the species creation rate peaks in this region with the lowest baryon fraction of burnt material.

The Schlieren plot illustrates the different waves produced by the wave interacting with the bubble. The left-most wavefront was produced when the wave first hit the left side of the bubble, the next wavefront when the wave exited the right side of the bubble.

3.3.3.2 Spinning bubble

In order to investigate the performance of the code for a more complex system, the previous system was modified to give the bubble an initial angular momentum of $\omega = 0.5$. The bubble is rotating clockwise in the 2d plane of the simulation. The results of this test can be seen in Figure 3.24. It was found that in this case the bubble no longer ignited; this could be due to a number of reasons. As can be seen, the structure of the bubble post-shock is significantly distorted compared to the previous non-spinning case. The interaction of the spinning bubble with the shock front has caused many smaller vortices to be created, which in turn have broken up the structure of the bubble. Focusing of reflected shocks is reduced, lowering the peak temperature such that the

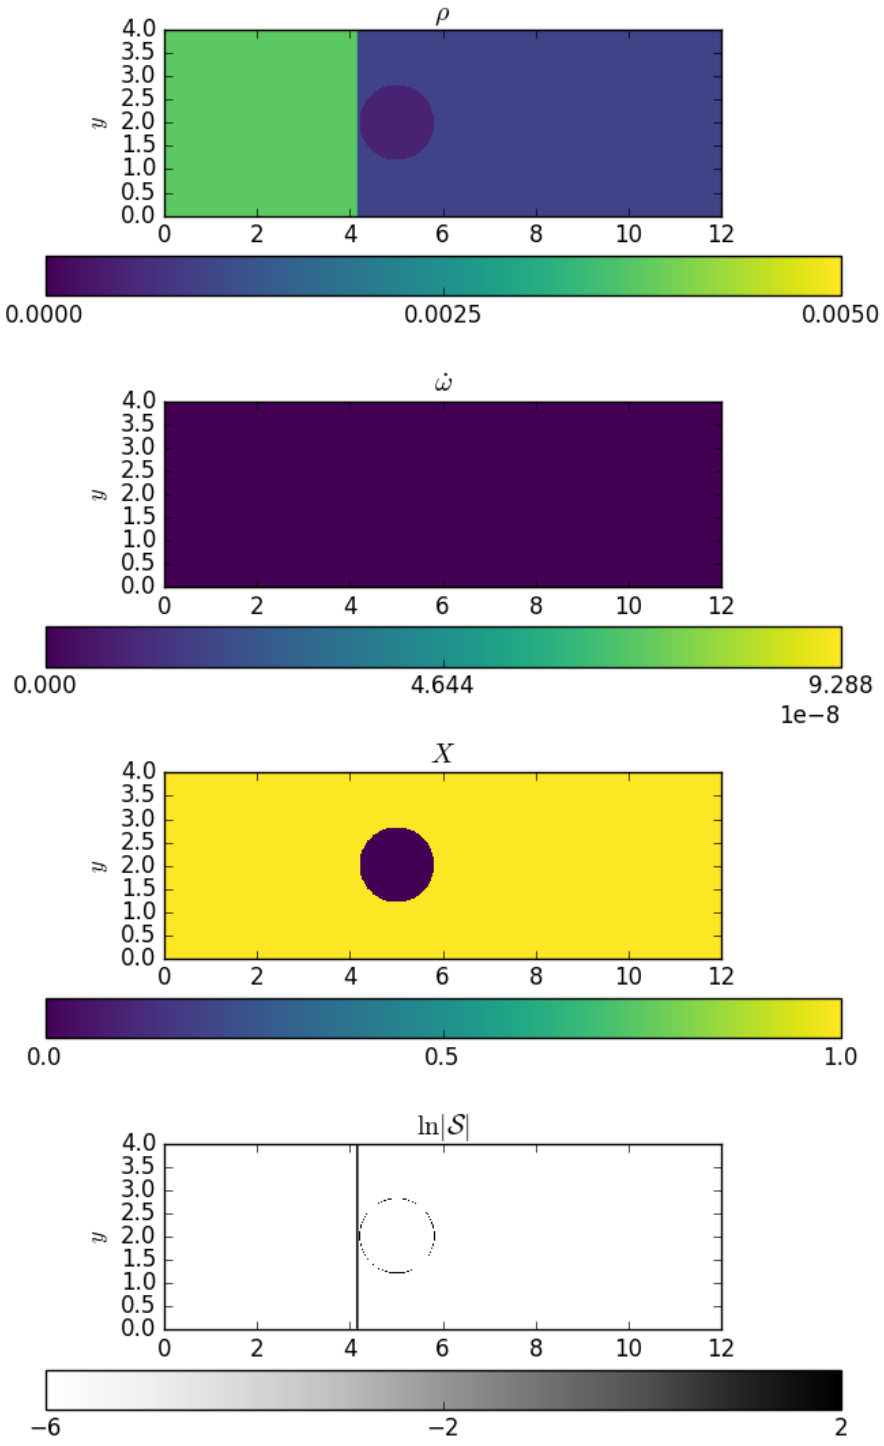

FIGURE 3.22: Special relativistic shock with Ma = 5 before the shock wave passes through the bubble. Plots from top to bottom are of density, species creation rate, baryon fraction of burnt material and the log Schlieren number. Note that the scales in all plots are arbitrary (with the exception of the baryon fraction which is a dimensionless quantity).

ignition temperature is not reached. The initial rotation of the bubble in the plane of the simulation accounts for the asymmetry of the shocked bubble.

Although this is an ideal fluid simulation, it is possible that numerical viscosity here is acting to mimic physical viscosity and producing turbulence. If this is so, then this simulation demonstrates the effects described in Section 1.4.2, where turbulent eddies can seriously disrupt burning regimes. Here, the formation of turbulent eddies has prevented ignition from being reached entirely – the peak species creation rate has been reduced by three orders of magnitude from ~ 13.4 to ~ 0.01. This test was repeated for a range of resolutions between 64×64 and 1024×1024 in order to determine whether this effect was due to insufficient resolution. As the resolution was increased, the focusing and peak temperature did increase, however not at a rate that indicates the peak temperature would ever approach the peak temperature seen in the previous case of the static bubble, even if a much greater resolution that presented here was used. We therefore believe that the disruption of the burning regime is a genuine effect, and not a consequence of insufficient resolution.

3.4 Summary

In this chapter, we have looked at the relativistic reactive Riemann problem. We began by looking at the theory of the Riemann problem, and how an exact solution can be obtained for the inert, reactive and relativistic reactive cases. This theory was then used to construct an exact solver of the relativistic reactive Riemann problem, `R3D2`. The performance of this solver was compared against a fully compressible relativistic solver implemented in `pyro`, where burning was accounted for by including reactive source terms. Several tests of the two codes were performed, showing that their results were generally consistent with each other and produced qualitatively the features expected.

In neutron star oceans, it is conceivable that a burning front moving latitudinally may encounter a region with a fast tangential velocity produced by the star's rotation. In Section 3.3.1.3, we investigated whether a fast tangential velocity may be able to cause a deflagration wave to transition to a detonation. Unlike for the Newtonian Riemann problem where only the velocity perpendicular to the interface is relevant, in the relativistic case the tangential velocity can become significant through the Lorentz factor. It was found that such a transition is possible, but only for tangential velocities that are a significant fraction of the speed of light or for systems already on the verge of transitioning. Consequently, it is unlikely that such a transition would occur for a burning front in a neutron star ocean.

Relativistic deflagrations also occur in other astrophysical systems such as supernovae and gamma ray bursts. However, given the combination of extreme conditions required for the transition to take place, it seems unlikely that this transition would occur in

these systems either. As we've seen, not only does there need to be some process which generates an extremely high, relativistic tangential velocity (preferably over an extended period of time in order to maximise the chances of the reaction occurring), but the fluid must also be very hot and on the verge of reacting (without already having done so). The fluid must also be dense enough that burning in the fluid with zero tangential velocity propagates as a deflagration rather than a detonation.

Core-collapse supernovae involve explosive burning of fast moving material. The progenitor stars can have very high rotation rates, so may even be possible for a fast tangential velocity to develop. However, even taking an optimistic estimate of the surface rotational velocity of the progenitor of $\sim 300\,\mathrm{km\,s^{-1}}$, this is still only a fraction of a percent of the speed of light. The burning is also believed to propagate as a detonation from the start, so there would be no opportunity for it to transition.

Black hole accretion disks are another system where hot, burning material moving at high speeds can be found. Instabilities in the inner accretion disk are believed to be a source of gamma ray bursts (Perna et al., 2006). If the accretion disk was rotating fast enough, and unreacted material from the black hole jet were to fall back onto the disk, the necessary conditions could potentially be reached for the transition to occur. However, this is again very unlikely given how extreme the conditions need to be.

In Section 3.3.2, we saw that the accuracy of the compressible solver was significantly reduced by the introduction of the stiff reactive source terms. This is likely to be a result of the way we have implemented source terms in the code. We used the first order operator splitting method for its simplicity and stability, however for stiff source terms this leads to significant errors, especially if the reaction zone is underresolved such that it is smeared across a larger region than is physical. For improved accuracy for reactive problems, we could use a higher order method for the source terms such as Strang splitting (LeVeque, 2002) or the modified fractional step method of Helzel et al. (2000).

In Section 3.3.3.2, the simulation of a special relativistic shock passing through an initially spinning bubble found that the burning was disrupted by the formation of turbulent eddies, which prevented ignition from being reached. This is consistent with the standard theory of turbulent combustion where turbulent eddies broaden flames, lowering the temperature in the reaction zone such that the reaction rate decreases (Klein, 1998; Peters, 2000). Tests were repeated for a range of resolutions to investigate whether this result was due to insufficient resolution: these suggested that the disruption of the burning observed is likely to be a genuine effect, and even a test with much greater resolution would be unable to reach the same peak temperature as seen in the case of the static bubble. It would be of interest to investigate this system further using a more accurate model to investigate further whether this observed turbulence is a real physical effect or a numerical artefact.

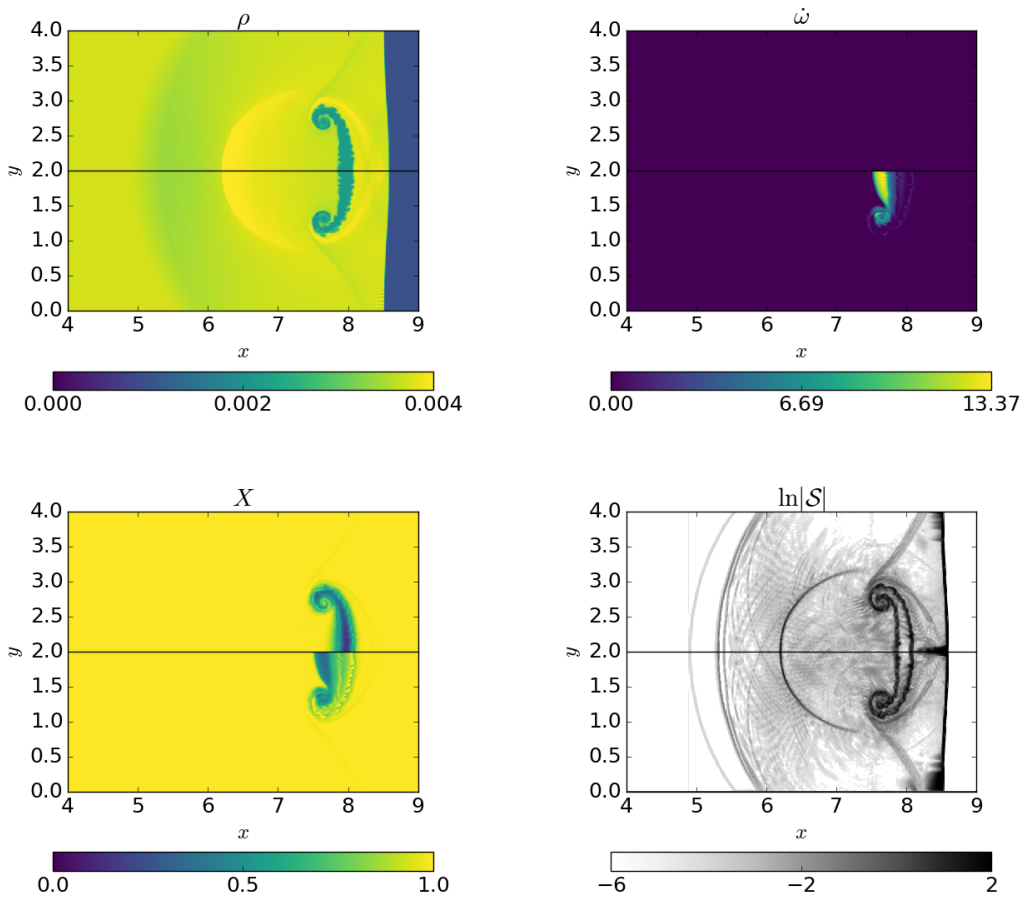

FIGURE 3.23: Special relativistic shock with Ma = 5 after the shock wave has passed through the bubble. Note that we have reduced the limits of the x-axis so as to focus on the region around the bubble, rather than plotting the full domain as in Figure 3.22. The top half of each plot shows the inert case, the bottom half the reactive case.

It can be seen that the density ρ and log Schlieren number $\ln|\mathcal{S}|$ plots appear to be unchanged in the reactive case. On the other hand, the species reaction rate $\dot{\omega}$ shows that the material to the left of the bubble is burning. This corresponds to the region of low species fraction X. In the inert case, the material in the bubble remains unburnt, however in the reactive case much of this has reacted, and the unburnt material has been pushed to the rear side of the bubble.

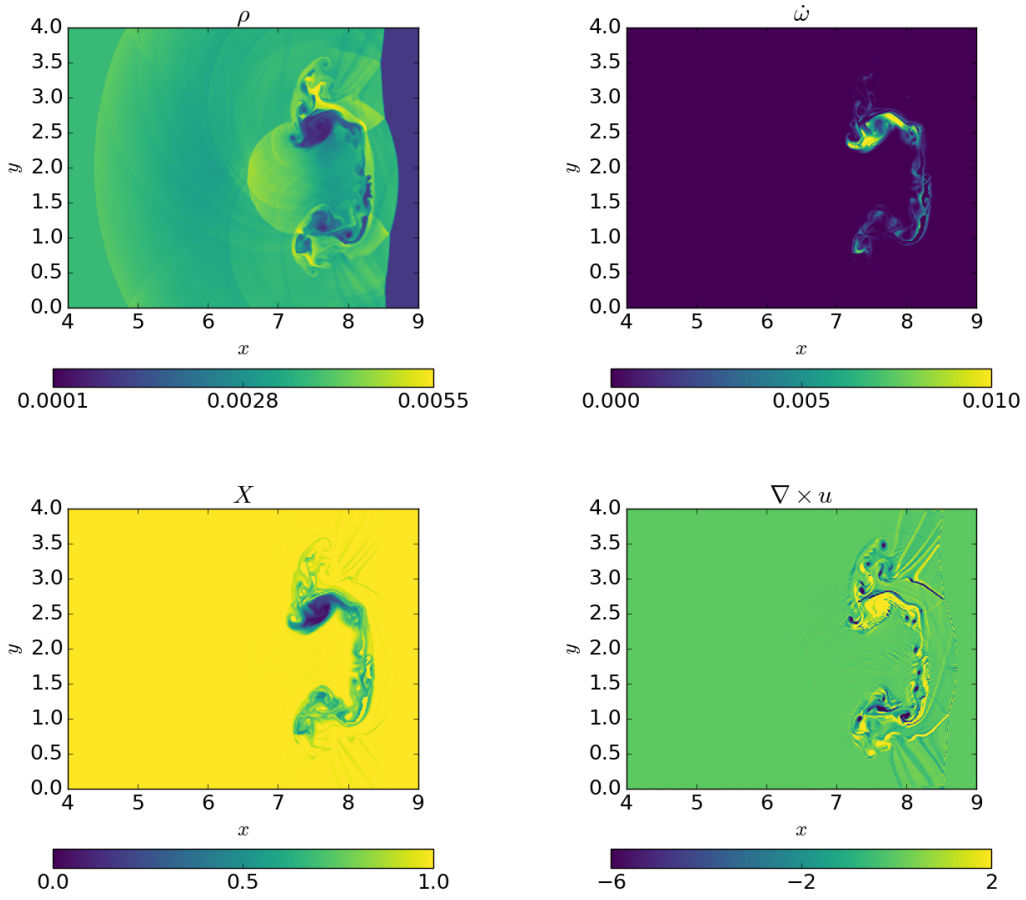

FIGURE 3.24: Special relativistic reactive shock with Ma = 5 after it has passed through an initially spinning bubble. The bottom right plot shows the vorticity.

Comparing to the non-spinning bubble in Figure 3.23, we can see that the bubble has become distorted and broken up by the vortices. The peak reaction rate $\dot{\omega}$ is much lower: the vortices have inhibited burning. This can also be seen in the plot of the species fraction X: in the non-spinning case, the minimum mass fraction was higher as material in the bubble had burnt, and the unburnt material had been pushed to the rear side of the bubble. For the spinning bubble, there is now a region centred at $\sim (7.5, 2.5)$ of completely unburnt material corresponding to a region of low density. The unburnt material largely corresponds to the regions of lower density which indicate material initially inside the bubble. Given the low reaction rate, it is possible that a large part of the increase in species fraction in these regions may have come from the turbulent mixing with the burnt material outside the bubble rather than burning taking place.

4 Shallow water equations

The propagation of the burning front is a multiscale problem: the behaviour on large scales is dominated by the Coriolis force (Spitkovsky et al., 2002), whereas on small scales turbulent burning reactions dominate (Malone et al., 2011). Another large scale effect that needs to be considered in general relativity is that rotating systems experience frame dragging effects and will deform at the equator (Paschalidis et al., 2017). There have not yet been any studies of how these may affect the propagation of fronts in the ocean, but it is possible that they may provide mechanisms that explain observed features such as multiple peaks and burst oscillations. We therefore believe it will be important to include these large scale effects in our model. In order to accurately model the whole system, it will be necessary to use a scheme that captures the physics at both large and small scales. In the previous sections, we have described work to better understand the small scale physics; in this section, we shall instead concentrate on the large scale physics.

On the scale of the entire neutron star, the ocean itself is relatively shallow: its depth is only $\sim 10\,\mathrm{m}$, whereas the neutron star's radius is $\sim 10\,\mathrm{km}$. In meteorology, the ocean and atmosphere are often modelled using the shallow water equations, an approximation derived from integrating the (spatially) 3d fluid equations in the vertical direction to get a 2d set of equations. This reduced dimensionality reduces the computational cost of evolving the system considerably. It therefore seems reasonable that such an approach could be useful to model the large scale physics of our problem. This approach has been used previously in the context of neutron star atmospheres by Spitkovsky et al. (2002) in their paper modelling the effects of the Coriolis force on the propagation of the burning front and by Inogamov and Sunyaev (2010) to model the spread of matter over the star's surface during accretion. Both groups used the Newtonian shallow water equations.

The reduced dimensionality of the shallow water model may also help avoid vertical boundary issues. As described in Section 1.5.2.8, for the low Mach and compressible models, care must be taken with the vertical boundary conditions. In the horizontal directions, it is possible to simply use periodic or outflow boundary conditions. However, this is not appropriate in the vertical direction, where it is necessary to ensure that hydrostatic equilibrium is properly enforced.

Due to the highly stratified nature of the neutron star ocean (see Figure 1.2) and the extreme temperature sensitivity of the nuclear reactions, the nuclear burning processes in type I X-ray bursts are confined to a thin shell (Schwarzschild and Härm, 1965; Yoon et al., 2004; Galloway and Keek, 2017). This confinement produces the thin-shell instability, leading to thermonuclear runaway. This motivates a multilayer shallow water model, such that the burning layer can be accurately resolved.

In this chapter, we shall describe our derivation of the relativistic shallow water equations and how we include multiple layers and reaction terms. We shall then present a numerical implementation of the single layer Newtonian and relativistic shallow water equations.

4.1 Newtonian shallow water equations

The shallow water equations are an approximation for modelling the flow of a fluid with a free surface under gravity. They are derived by assuming that the depth of the fluid is small with respect to the wave length or curvature of the top free surface, or equivalently that the pressure distribution is given by hydrostatic equilibrium and is a result of assuming that the vertical acceleration of the fluid has a negligible effect on the pressure (Toro, 2001).

To derive the shallow water equations, we consider fluid in a channel of unit width, assuming that the vertical velocity of the fluid is negligible and that the horizontal fluid velocity $u(t,x)$ is roughly constant across the cross section of the channel (LeVeque, 2002). The gravitational field is taken to be constant and act in the vertical direction, with its acceleration given by g. Assuming that the fluid is incompressible with constant density $\bar{\rho}$ and depth $h(t,x)$, the total mass in the section of the channel between $[x_1, x_2]$ at time t is

$$\int_{x_1}^{x_2} \mathrm{d}x\, \bar{\rho} h(t,x). \tag{4.1}$$

The momentum density is given by $\bar{\rho}u(t,x)$; integrating this vertically gives the mass flux $\bar{\rho}u(t,x)h(t,x)$. Substituting this into the continuity equation, we find the evolution equation for the fluid depth

$$\partial_t h + \partial_x(hu) = 0. \tag{4.2}$$

Looking at the momentum equation and substituting for the mass flux, we obtain

$$\partial_t(\bar{\rho}hu) + \partial_x(\bar{\rho}hu^2 + p) = 0. \tag{4.3}$$

Assuming that the fluid is in hydrostatic equilibrium, with the pressure at distance $h-y$ below the surface given by $\bar{\rho}g(h-y)$, the total pressure is found by integrating vertically to get

$$p = \frac{1}{2}\bar{\rho}gh^2. \tag{4.4}$$

Substituting into the momentum equation, we find

$$\partial_t(hu) + \partial_x\left(hu^2 + \frac{1}{2}gh^2\right) = 0. \tag{4.5}$$

This can be trivially extended to a 3d domain to give

$$\partial_t h + \partial_i\left(hu_i\right) = 0, \tag{4.6}$$

$$\partial_t\left(hu_i\right) + \partial_j\left(hu_iu_j + \frac{1}{2}gh^2\delta_{ij}\right) = 0. \tag{4.7}$$

With the view to deriving a general relativistic form of these equations, it will be useful to give these equations in a form which does not assume a constant gravitational acceleration. Instead of using the fluid depth h as a primitive variable, we can instead use the geopotential Φ (St-cyr et al., 2008). For a flat spacetime with constant gravitational field, $\Phi = gh$. This gives us the following alternative form of the Newtonian shallow water equations

$$\partial_t\Phi + \partial_i\left(\Phi u_i\right) = 0, \tag{4.8}$$

$$\partial_t\left(\Phi u_i\right) + \partial_j\left(\Phi u_iu_j + \frac{1}{2}\Phi^2\delta_{ij}\right) = 0. \tag{4.9}$$

4.2 General relativistic shallow water equations

When deriving the Newtonian shallow water equations, we integrated the 2d fluid equations in the vertical direction to obtain their integral average. For the relativistic case, this is not so straightforward: for a non-static spacetime, what exactly we define to be the vertical direction becomes less obvious (i.e. rather than just radial lines from the centre of the star, it might now make more sense to follow the geodesics defined by the timelike Killing vectors).

We shall begin with the fully compressible GR hydrodynamics equations:

$$\nabla_\mu(\rho u^\mu) = 0, \tag{4.10}$$

$$u^\nu\nabla_\nu u_\mu + \frac{1}{\rho h}h_\mu{}^\nu\nabla_\nu p = 0, \tag{4.11}$$

where $h_\mu{}^\nu = u_\mu u^\nu + \delta_\mu{}^\nu$ is the projection operator, and $h = 1 + \varepsilon + \frac{p}{\rho}$ now represents the fluid enthalpy. Note that we only need the continuity and momentum equations – by assuming (the relativistic equivalent of) hydrostatic equilibrium, we no longer need the energy equation.

4.2.1 Spherical symmetry

For this first derivation, we shall assume that the background spacetime is spherically symmetric. For a static, spherically symmetric spacetime, Birkhoff's theorem tells us that the metric can be written as

$$ds^2 = -e^{-2\Phi(r)}dt^2 + e^{2\Psi(r)}dr^2 + r^2 d\Omega^2, \tag{4.12}$$

where $\Phi(r)$ and $\Psi(r)$ are functions of the radius only. The Tolman-Oppenheimer-Volkoff (TOV) equation for the gradient of the pressure is then given by

$$\nabla_\mu p - \rho h \nabla_\mu \Phi = 0. \tag{4.13}$$

Assuming a gamma-law equation of state, we can write ρh in terms of the vertically averaged density $\bar{\rho}$ and the pressure p as

$$\rho h = \bar{\rho} + \frac{\gamma}{\gamma - 1} p. \tag{4.14}$$

We then find that

$$\frac{\gamma - 1}{\gamma} \nabla_\mu \left(\ln \left[\bar{\rho} + \frac{\gamma}{\gamma - 1} p \right] \right) - \nabla_\mu \Phi = 0. \tag{4.15}$$

Substituting in the vertically averaged density, the continuity equation can be written as

$$\nabla_\mu (\bar{\rho} u^\mu \Phi) = 0 \qquad \Rightarrow \qquad \nabla_\mu (\Phi u^\mu) = 0. \tag{4.16}$$

Using the TOV equation (4.13), we can substitute for the second term in the Euler equation to get

$$u^\nu \nabla_\nu u_\mu + h_\mu{}^\nu \nabla_\nu \Phi = 0. \tag{4.17}$$

Multiplying by Φ, we can find this in conservative form

$$\nabla_\nu \left(\Phi u^\nu u_\mu \right) + \frac{1}{2} h_\mu{}^\nu \nabla_\nu \Phi^2 = 0. \tag{4.18}$$

We can now write the equations in the hyperbolic form used before, with the primitive variables, conserved variables and fluxes given by

$$\begin{aligned} \boldsymbol{w} &= (\Phi, v_j), \\ \boldsymbol{q}(\boldsymbol{w}) &= \left(\Phi W, \Phi W^2 v_j \right) \\ &= (D, S_j), \\ \boldsymbol{f}^{(i)}(\boldsymbol{w}) &= \left(\Phi W v^i, \Phi W^2 v_j v^i + \frac{1}{2} \Phi^2 \delta^i{}_j \right) \\ &= \left(D v^i, S_j v^i + \frac{D^2}{2W^2} \delta^i{}_j \right). \end{aligned} \tag{4.19}$$

As we have assumed the background spacetime to be spherically symmetric, the components of the shift vector $\beta^i = 0$. The non-trivial components of the stress energy tensor used in the source terms are given by

$$\tilde{T}^{tt} = \Phi g^{tt}, \qquad \tilde{T}^{ij} = \frac{1}{2}\Phi^2 \delta^{ij}. \tag{4.20}$$

Integrating the TOV equation (4.13), we find that the pressure p is

$$p(r) = \frac{\gamma - 1}{\gamma}\left(Ae^{\frac{\gamma}{\gamma-1}\Phi} - \bar{\rho}\right), \tag{4.21}$$

where A is some constant that can be defined by setting e.g. the surface pressure $p(h) = p_h$.

We also find that ρh is

$$\rho h = \bar{\rho} A e^{\frac{\gamma}{\gamma-1}\Phi}. \tag{4.22}$$

By fixing e.g. the surface pressure, we can obtain a value for the constant A, and this therefore allows us to convert between the shallow water equations' conserved variables and those of the compressible equations.

4.2.2 Axisymmetry

In the previous section, an expression was derived for the pressure gradient assuming the background metric to be spherically symmetric. Such a metric describes the spacetime that would be produced by a spherical, non-rotating neutron star. However, we are interested in the effects of the Coriolis force on the propagation of the burning front, which requires that the background metric is at the very least axisymmetric to account for the rotation of the neutron star. In this section we shall therefore find an expression for the pressure gradient assuming axisymmetry, using the integral of motion approach given in Section 3.4.3 of Gourgoulhon (2010).

For a spherical star rotating at angular velocity Ω, the fluid 4-velocity is given by

$$\boldsymbol{u} = u^t \boldsymbol{k}, \tag{4.23}$$

where

$$\boldsymbol{k} = \boldsymbol{\xi} + \Omega \boldsymbol{\chi} \tag{4.24}$$

is a Killing vector if Ω is constant, and where $\boldsymbol{\xi} = \boldsymbol{\partial}_t$ and $\boldsymbol{\chi} = \boldsymbol{\partial}_\phi$ are Killing vectors generating the timelike and axisymmetric symmetries of the system. Exploiting the symmetries of the system, the fluid equation of motion can be written as

$$\boldsymbol{\nabla} \ln\left(-h\boldsymbol{u}\cdot\boldsymbol{k}\right) - \frac{\boldsymbol{u}\cdot\boldsymbol{\chi}}{\boldsymbol{u}\cdot\boldsymbol{k}}\boldsymbol{\nabla}\Omega = 0, \tag{4.25}$$

where h is the fluid enthalpy. If we assume that Ω is constant, this gives us the first integral of motion

$$\ln\left(-h\boldsymbol{u}\cdot\boldsymbol{k}\right) = \text{const.} \tag{4.26}$$

Given that $W = \alpha u^t$ and $u_t = \boldsymbol{u}\cdot\boldsymbol{\xi}$, we find that

$$\boldsymbol{u}\cdot\boldsymbol{k} = u_t = g_{tt}u^t = -\alpha^2\frac{W}{\alpha} = -\frac{\alpha}{W}. \tag{4.27}$$

The first integral therefore becomes

$$\ln\left(\frac{h}{\sqrt{{}^{(3)}\gamma}}\right) = \text{const.} \qquad \Rightarrow \qquad \frac{h}{\sqrt{{}^{(3)}\gamma}} = A, \tag{4.28}$$

where $\sqrt{{}^{(3)}\gamma}$ is the square root of the determinant of the spatial 3-metric γ_{ij} and A is some constant.

We can use this to find an expression for the pressure using the definition of the enthalpy from the ideal gas equation of state,

$$h = 1 + \frac{\gamma p}{(\gamma-1)\rho}, \tag{4.29}$$

where γ is the adiabatic index.

Substituting this into (4.28), we find

$$p = \frac{(\gamma-1)\rho}{\gamma}\left(A\sqrt{{}^{(3)}\gamma} - 1\right). \tag{4.30}$$

For the shallow water equations, we assume that the density ρ is constant. Taking the gradient, we then find an expression for the gradient of the pressure which is proportional to the gradient of the spatial metric determinant

$$\boldsymbol{\nabla}p = \frac{(\gamma-1)\bar{\rho}}{\gamma}A\boldsymbol{\nabla}\sqrt{{}^{(3)}\gamma}. \tag{4.31}$$

4.2.3 Layer interaction

If there are multiple layers of fluid, extra terms must be included to account for interaction between the layers. For a simple system with no heat or mass transfer, buoyancy terms must be included so that the motion of the layers is coupled together. If there is some heat or mass transfer, friction and mass transfer terms must also be included.

In the multilayer shallow water model of Spitkovsky et al. (2002); Bouchot et al. (2010), the buoyancy terms are included in the momentum equation for layer j as

$$\partial_t(h_j u_j) + \partial_x(h_j u_j^2 + g h_j^2/2) + g h_j \partial_x \left(\sum_{k>j} h_k + \sum_{k<j} \frac{\rho_k}{\rho_j} h_k \right) = 0. \tag{4.32}$$

These terms are simply the gradient of the pressure, $\partial_x p = \partial_x(\rho g h) = \rho g \partial_x h$, where $\rho = \bar{\rho} h$ and so appears as just h_j in the above equation. Given that for the Newtonian system $\Phi = gh$, in the relativistic system this term will become $\Phi_j \nabla \Phi_k$.

The multilayer relativistic shallow water equations then become

$$\nabla_\mu \left(\Phi^L u^{L\mu} \right) = 0, \tag{4.33}$$

$$\nabla_\mu \left(\Phi^L u^L_\nu u^{L\mu} \right) + \frac{1}{2} \delta^i{}_\nu \nabla_i \left(\Phi^L \right)^2 + \Phi^L \delta^i{}_\nu \nabla_i \left(\sum_{M>L} \Phi^M + \sum_{M<L} \frac{\rho^M}{\rho^L} \Phi^M \right) = 0, \tag{4.34}$$

where the $L, M...$ superscript indicates the layer.

If we instead include heating and friction terms, the Newtonian equations become

$$\partial_t(g h_j) + \boldsymbol{\nabla} \cdot (g h_j \boldsymbol{u}_j) + g \left((Q_{j-1} - Q_j) + \frac{\rho_{j+1}}{\rho_j} (Q_{j+1} - Q_j) \right) = 0 \tag{4.35}$$

$$\partial_t(g h_j u_j) + g \partial_x(h_j u_j^2 + g h_j^2/2) + g(Q_j + \mu_f) \left((u_j - u_{j-1}) + \frac{\rho_{j+1}}{\rho_j} (u_j - u_{j+1}) \right) = 0, \tag{4.36}$$

where Q_j is the mass flux into the layer and μ_f is the friction. Note that if $j - 1 < 0$ or $j + 1 > n_{\text{layers}}$, then the pair of terms involving these quantities are ignored.

The presence of the g coefficient of the heating and friction terms complicates matters when it comes to adding this to the relativistic equations, not least because it remains a challenge to concretely define heat transport in relativistic fluids from first principles (Andersson and Lopez-Monsalvo, 2011). We can mitigate this by using a qualitative model and redefining Q and μ: instead of being mass flux (or flux of h) and the friction coefficient of h, we redefine them as the flux of $\Phi = gh$ and the friction coefficient of Φ, such that in the Newtonian limit

$$\tilde{Q} = gQ, \quad \tilde{\mu}_f = g\mu_f. \tag{4.37}$$

Including the heating and friction terms as well, the multilayer relativistic shallow water equations then become

$$\nabla_\mu \left(\Phi^L u^{L\mu}\right) + \left(\tilde{Q}^{L-1} - \tilde{Q}^L\right) + \frac{\rho^{L+1}}{\rho^L}\left(\tilde{Q}^{L+1} - \tilde{Q}^L\right) = 0, \tag{4.38}$$

$$\nabla_\mu \left(\Phi^L u^L_\nu u^{L\mu}\right) + \frac{1}{2}\delta^i{}_\nu \nabla_i \left(\Phi^L\right)^2 + \Phi^L \delta^i{}_\nu \nabla_i \left(\sum_{M>L} \Phi^M + \sum_{M<L} \frac{\rho^M}{\rho^L}\Phi^M\right) + \left(\tilde{Q}^L + \tilde{\mu}_f\right)\left(\left[u^L_\nu - u^{L-1}_\nu\right] + \frac{\rho^{L+1}}{\rho^L}\left[u^L_\nu - u^{L+1}_\nu\right]\right) = 0. \tag{4.39}$$

4.2.4 Burning

In order to include reactions in the relativistic shallow water equations, we shall first add reaction terms to the relativistic compressible equations. To model burning, we shall use the simple 2-species model of Cumming and Bildsten (2000); Spitkovsky et al. (2002) which simulates the triple alpha reaction. Let us assume that the helium abundance Y is given by

$$\frac{\mathrm{d}Y}{\mathrm{d}t} = -\frac{\epsilon_{3\alpha}}{E_{\mathrm{He}}}, \tag{4.40}$$

where $\epsilon_{3\alpha}$ is the heating rate per unit mass due to the triple alpha reaction, given by Cumming and Bildsten (2000) as

$$\epsilon_{3\alpha} = 5.3 \times 10^{21}\ \mathrm{erg\,g^{-1}\,s^{-1}} \frac{\rho_5^2 Y^3}{T_8^3} \exp\left(\frac{-44}{T_8}\right), \tag{4.41}$$

where ρ_5 is the density in units of $10^5\,\mathrm{g\,cm^{-3}}$ and T_8 is the temperature in units of $10^8\,\mathrm{K}$. E_{He} is the energy release per unit mass of helium burning during the triple alpha process. As shown in Ooyama (1969); Spitkovsky et al. (2002), the height of the top fluid layer represents the temperature. An increase in the temperature of the fluid due to heating corresponds to addition of fluid to a shallow water layer and therefore an increase in its height. An adiabatic temperature rise corresponds to a decrease in the pressure, which in the shallow water equations also corresponds to an increase in height.

Defining the species creation rate as $\frac{\mathrm{d}Y}{\mathrm{d}t} = \dot{Y}$, in the relativistic case the continuity equation gives

$$\nabla_\mu \left(\rho u^\mu Y\right) = \rho \dot{Y}, \tag{4.42}$$

such that we obtain the new conserved variable and corresponding flux

$$\zeta = \rho W Y, \qquad f_\zeta = \rho u^i Y = \zeta\left(v^i - \frac{\beta^i}{\alpha}\right). \tag{4.43}$$

The energy equation must also be amended to account for the heating produced during burning. Starting from the compressible energy equation

$$\nabla_\mu(\rho h u^\mu) - u^\mu \nabla_\mu p = \rho H, \tag{4.44}$$

where H is the heating rate per unit mass (which for our case will be equal to $\epsilon_{3\alpha}$). Writing this in terms of the conserved variables and performing the 3+1 decomposition, we get

$$\partial_t(\sqrt{\gamma}\tau) + \partial_i\left(\sqrt{-g}\left[\tau\left(v^i - \frac{\beta^i}{\alpha}\right) + pv^i\right]\right) = \sqrt{\gamma}DH. \tag{4.45}$$

In order to obtain the reactive relativistic shallow water equations, we can introduce the mass fraction X to the primitive state vector and $\Phi XW = DX$ to the conserved state vector (as we did for the first of the reactive continuity equations), such that

$$\nabla_\mu\left(\Phi X u^\mu\right) = \Phi\dot{X}. \tag{4.46}$$

As the shallow water equations do not evolve any energy or temperature-like quantities, no heating term is required.

The previous interlayer mass transfer approach makes sense if we assume that the ocean is highly stratified, with heavier burnt material sinking below the product layer on a timescale much less than that of the burning (such that they can be treated as instantaneous) and that each layer has its own distinct composition with no (or at least negligible) mixing. The layers then will represent the physical stratified layers of the ocean. However, in reality it is likely that these layers will not be quite so distinct, and so using instead a species mass fraction to represent the ocean composition is more appropriate and makes it a lot easier to link the shallow water and compressible equations. The mass transfer term is essentially just redefining the layer boundaries, shifting the geopotential up/down but not necessarily representing real physics.

Adding the evolution of the species mass fraction to the multilayer shallow system, we can obtain the reactive relativistic multilayer shallow water equations. This extends (4.38) with the addition of

$$\nabla_\mu\left(\Phi^L X^L u^{L\mu}\right) = \Phi^L \dot{X}^L. \tag{4.47}$$

4.2.5 Characteristics of the relativistic equations

In the numerical implementation below, we shall be using the inert relativistic shallow water equations with a constant gravitational field in the vertical direction. Before we carry out any tests, we shall investigate the characteristics of these equations in order to

determine their physical regime of validity. This is necessary as the equations themselves do not enforce that the speeds remain less than $c = 1$. As described in Section 1.5.1, by looking at the eigenvalues of the Jacobian, we can determine the wavespeeds of the characteristics of the Riemann problem for the system. This will tell us if we need to place limits on other quantities such as the geopotential in order for the speeds in the system to remain subluminal and therefore physical.

For the relativistic system with a constant gravitational field in the vertical direction, we can set $\Phi = gh$ in the general relativistic equations (4.19), giving the conserved variables and fluxes

$$\begin{aligned} \boldsymbol{q} &= \left(hW, hW^2 v_j\right) \\ &= \left(D, S_j\right), \\ \boldsymbol{f}^{(i)} &= \left(hWv_i, hW^2 v_j v_i + \frac{1}{2}gh^2\delta_{ij}\right) \\ &= \left(Dv_i, S_j v_i + \frac{gD^2}{2W^2}\delta_{ij}\right). \end{aligned} \tag{4.48}$$

In the solution of the Riemann problem, the wavespeeds can be calculated from the eigenvalues of the Jacobian $J(\boldsymbol{q}) = \partial \boldsymbol{f}(\boldsymbol{q})/\partial \boldsymbol{q}$. Looking at the 1d case where $v = v_i$, we obtain the eigenvalues

$$\lambda_\pm = \frac{1}{2}v\left(\Phi v^2 - \Phi + 2\right) \pm \frac{1}{2}\sqrt{\Phi}(1 - v^2)\sqrt{\Phi v^2 + 4}. \tag{4.49}$$

These are clearly significantly more complex than the eigenvalues of the Newtonian shallow water equations, $\lambda_\pm = v \pm \sqrt{\Phi}$. In order to gain some physical insight, we can consider various limits of the eigenvalues. Taking the Newtonian limit (small Φ and v), we find

$$\lim_{\substack{\Phi \to 0 \\ v \to 0}} \lambda_\pm = v \pm \sqrt{\Phi}, \tag{4.50}$$

which is the same as the advective wavespeed in the Newtonian case. Taking the limit as $v \to 0$, we find the acoustic speed to be

$$\lim_{v \to 0} \lambda_\pm = \pm\sqrt{\Phi}, \tag{4.51}$$

which again is the same as for the Newtonian system. This limit has the important consequence for our numerical tests that in order for the wavespeeds to be physical (i.e. less than the speed of light $c = 1$), we will need to make sure that $\Phi < 1$.

To investigate the behaviour of the system as the speed becomes relativistic, we can take the limit as the Lorentz factor $W \to \infty$, or equivalently as $W^{-2} \to 0$, which is

$$\lim_{W^{-2} \to 0} \lambda_\pm = 1. \tag{4.52}$$

The wave speed therefore tends to $c = 1$ as the fluid speed becomes highly relativistic. This demonstrates that the relativistic shallow water equations given above are consistent with special relativity, as a Riemann problem with physical, subluminal initial data shall produce a solution with physical, subluminal waves.

4.3 Numerical implementation

We implemented a solver for the Newtonian and relativistic shallow water equations by extending the `CASTRO` code (Almgren et al., 2010), a compressible solver for astrophysical flows that is built on top of the `AMReX` adaptive mesh refinement framework (Zhang et al., 2016). In the following sections, we demonstrate the results of 1d and 2d dam break tests performed to validate this implementation.

The Newtonian solver solves the standard shallow water equations in their conservative form $\partial_t \boldsymbol{q} + \partial_i \boldsymbol{f}^{(i)} = 0$, where

$$\boldsymbol{q} = \begin{pmatrix} h \\ hv_x \\ hv_y \end{pmatrix} \qquad \text{and} \qquad \boldsymbol{f}^{(i)} = \begin{pmatrix} hv_i \\ hv_x v_i + \frac{1}{2} g h^2 \delta_{ix} \\ hv_y v_i + \frac{1}{2} g h^2 \delta_{iy} \end{pmatrix}. \tag{4.53}$$

The relativistic solver solves the equations given in (4.48). Note that because we no longer evolve any thermodynamic variables, reclaiming the primitive variables for the conserved variables used here is much more straightforward than for the relativistic Euler equations. The Lorentz factor can be found from

$$W = \sqrt{1 + \frac{S_i S^i}{D^2}}, \tag{4.54}$$

allowing the fluid height and 3-velocity to be calculated from D and S_j.

The spatial integration is carried out using the finite volume method with the MC flux limiter (as described in Section 1.5.2) and HLLE solver (see Section 3.2.1). The time integration is carried out using the method of lines with RK3. The shallow water solvers were implemented by modifying the existing hydrodynamics routines so that they used the above fluxes, adjusting the conservative/primitive variables conversion routines and adjusting the HLL solver to use the correct signal speed. For the relativistic solver, care had to be taken to ensure that the primitive variables were properly calculated (in a few parts of the original code it was found that the trivial variable recovery possible for Newtonian systems was assumed – this had to be changed so that it called the conserved to primitive routine instead).

4.3.1 Dam break problem

The dam break problem is a special case of the Riemann problem where $v_L = v_R = 0$, so called because it models what would happen if a dam separating two bodies of water at different heights were to suddenly burst (LeVeque, 2002). It is the shallow water equivalent of the shock tube problem, with initial data for the height $h = h(t, x)$ and fluid velocity $v = v(t, x)$ given by

$$h(0, x) = \begin{cases} h_l & \text{if } x < 0, \\ h_r & \text{if } x > 0, \end{cases} \qquad v(0, x) = 0. \tag{4.55}$$

In the sections below, a 1d test shall be performed with an initial discontinuity at $x = 0$.

A 2d problem can be created by producing a circular dam. There exists no analytic solution for the circular dam, however the problem can be reduced to a 1d problem in the radial direction with geometric source terms (Alcrudo and Garcia Navarro, 1993). The 2d solution can therefore be compared against a high resolution 1d solution.

4.3.1.1 Newtonian dam break

The analytic solution for the dam break problem with non-zero initial water height on both sides was first derived by Stoker (1957). Given the scarcity of this solution in the literature, its derivation is outlined in Appendix D.

The results of the 1d dam breaking test can be seen in Figure 4.1. The initial conditions used were $h_l = 11$, $h_r = 1$ and $g = 1$. The simulation was run at two resolutions to demonstrate the convergence of the solver. The results match up with the exact solution very well, particularly for the high resolution run. In the low resolution run, the locations of the head and tail of the rarefaction fan are slightly over-/underestimated (as can be seen in the inset plot) and there is a small amount of smearing at the shock front. This will largely be a result of the slope limiters used in the finite volume solver, as discussed in Section 1.5.2.5.

In the high resolution run, there are some small oscillations in the velocity in the region behind the shock front. In the low resolution case, these have been smeared out, however they have been able to grow in the higher resolution run. They are likely to have occured due to the large height difference in the initial data, which have produced a very strong shock. This is supported by the fact that when the height of the left state was reduced from $h_l = 11$ to $h_l = 2$, these oscillations were no longer present.

Figure 4.2 shows the results of the 2d radial dam breaking test. The initial conditions of LeVeque (2002) were used: a dam of initial radius $r = 0.5$ and fluid height $h = 2$ inside

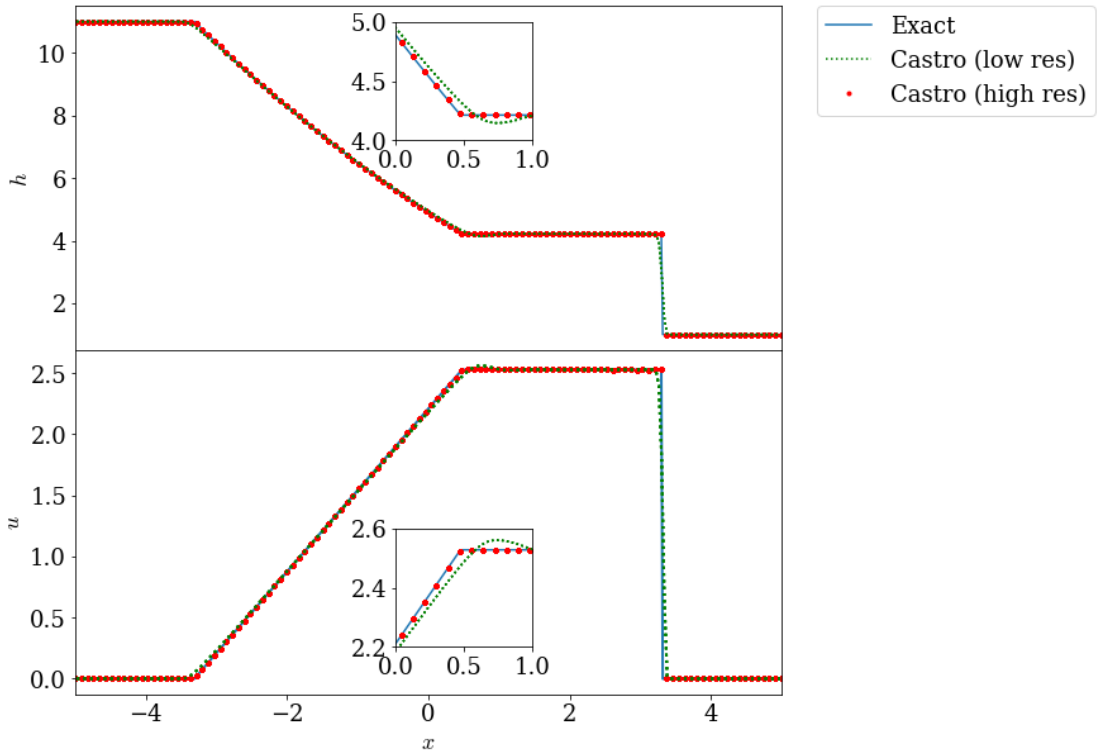

FIGURE 4.1: Results of the Newtonian dam breaking test with $h_l = 11$ and $h_r = 1$ at time $t = 1$. The simulation was carried out on a the domain $x = [-5, 5]$, $y = [0, 0.15625]$, with 256 grid points in the x-direction and 4 in the y-direction for the low resolution run and 4096 in the x-direction and 64 in the y-direction for the high resolution run, making it effectively 1d. The plot shows the fluid height and velocity. The inset plots demonstrate the over/undershooting in the low resolution run at the tail of the rarefaction.

the dam, and $h = 1$ outside. As mentioned above, there is no exact solution for this system that we can compare the results against: instead, the results shall be compared against a high resolution simulation of the reduced 1d system

$$\partial_t \boldsymbol{q} + \partial_r \boldsymbol{f} = \boldsymbol{S}, \tag{4.56}$$

where

$$\boldsymbol{q} = \begin{pmatrix} h \\ wh \end{pmatrix}, \qquad \boldsymbol{f} = \begin{pmatrix} wh \\ w^2 h + \frac{1}{2} g h^2 \end{pmatrix}, \qquad \boldsymbol{S} = \begin{pmatrix} -wh/r \\ -w^2 h/r \end{pmatrix}. \tag{4.57}$$

$w = \sqrt{u^2 + v^2}$ is the radial velocity and r the radial distance from the centre of the domain.

Comparing the 1d and 2d solutions, it can be seen that they match up very well for the high resolution run. In particular, the high resolution 2d simulation has been able to capture the shock fronts well, with no noticeable smearing or overshooting. The smearing of the shock fronts is clearly visible in the low resolution run, however as expected this is much reduced as the resolution is increased, demonstrating the convergence of the solver.

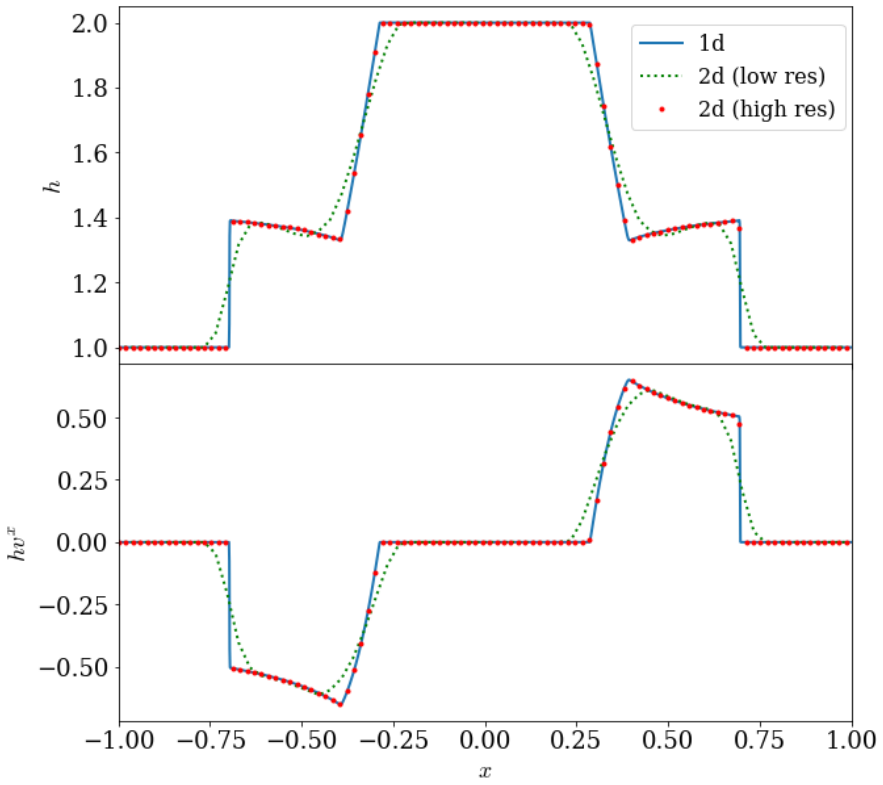

FIGURE 4.2: Results of the Newtonian radial dam breaking test with $h_{\text{in}} = 2$ and $h_{\text{out}} = 1$ and the initial dam at $r = 0.5$ at time $t = 0.15$. The 2d simulation was carried out on a the domain $x = [-1, 1]$, $y = [-1, 1]$, with 64×64 grid points for the low resolution run and 1028×1028 for the high resolution run. The plot shows the fluid height and momentum in the x-direction at $y = 0$.

4.3.1.2 Relativistic dam break

The results of the 1d relativistic dam breaking test with $h_l = 0.9$ and $h_r = 0.05$ at time $t = 0.4$ can be seen in Figure 4.3. Unfortunately, for this case we do not have an exact solution to compare against, however we can compare the low resolution run against the much sharper high resolution run. The low resolution run is more diffusive, with smearing at the shock front and at the head and tail of the rarefaction. The shock front has propagated at a significant fraction of the speed of light, with the shocked fluid in the intermediate region having a high, relativistic speed of 0.635.

The radial dam breaking test results with initial data $h_{\text{in}} = 0.9$ and $h_{\text{out}} = 0.2$ at time $t = 0.5$ can be seen in Figures 4.4 and 4.5. As for the Newtonian case, we have no exact solution against which to compare the results, but can instead compare against a high resolution 1d solution with radial source terms. The state vector, flux and source terms for the 1d simulation are given by

$$\boldsymbol{q} = \begin{pmatrix} \Phi W \\ \Phi W^2 w \end{pmatrix}, \qquad \boldsymbol{f} = \begin{pmatrix} \Phi W w \\ \Phi W^2 w^2 + \frac{1}{2}\Phi^2 \end{pmatrix}, \qquad \boldsymbol{S} = \begin{pmatrix} -\Phi W / r \\ -\Phi W^2 w / r \end{pmatrix}, \tag{4.58}$$

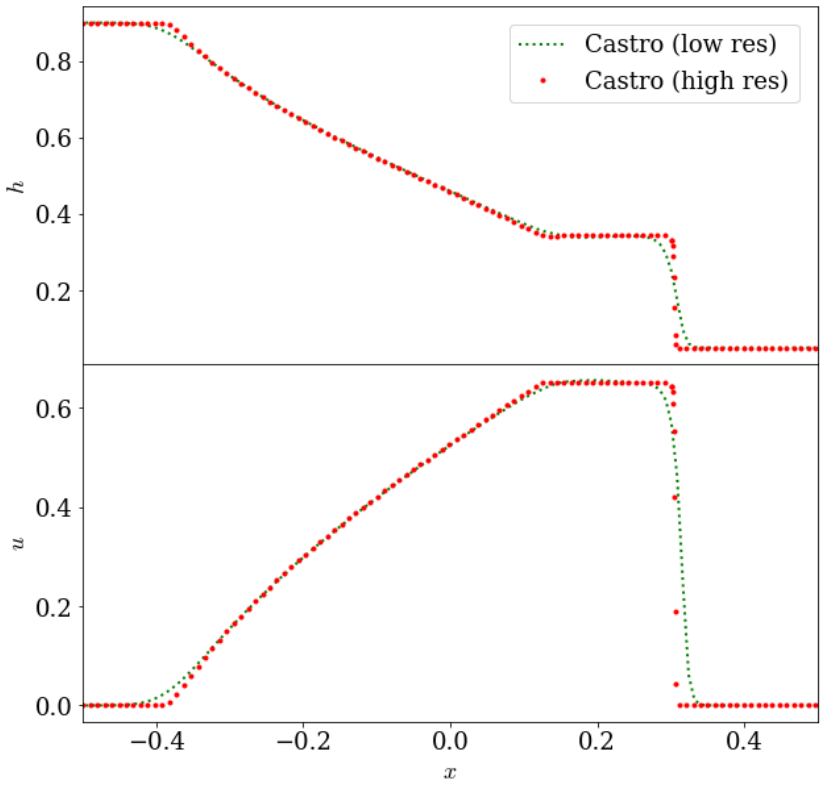

FIGURE 4.3: Results of the relativistic dam breaking test with $h_l = 0.9$ and $h_r = 0.05$ at time $t = 0.4$. The simulation was carried out on a the domain $x = [-0.5, 0.5]$, $y = [0, 0.03125]$, with 128 grid points in the x-direction and 4 in the y-direction for the low resolution run and 1024 in the x-direction and 32 in the y-direction for the high resolution run, making it effectively 1d. The plot shows the fluid height and primitive 3-velocity in the x-direction.

where $w = \sqrt{v_x^2 + v_y^2}$ is the primitive radial velocity and r the radial distance from the centre of the domain.

It can be seen that the high resolution 2d simulation matches up well with the 1d simulation, however there is some smearing at the discontinuities as a result of clipping by the slope limiters. As can be seen, this smearing is significantly reduced in the high resolution run, indicating that the solution is converging towards the correct solution.

Although we have used different initial data for this relativistic test than for the previous Newtonian test in Figure 4.2 in order to ensure the shock speeds remain within the physical regime, we can compare the tests in the two regimes qualitatively. The profiles of the conserved variables are mostly similar, however there are a couple of noticeable differences. The curves on either side of the central peak in the conserved height variable $D = hW$ are noticeably curved in the relativistic case. For the Newtonian system, these were straight, linear functions of distance. This is due to the nonlinear coupling of the wave speed and the Lorentz factor. Another consequence of this nonlinear coupling is that the region around the maximum/minimum conserved speed variable $S_x = hW^2 v_x$ is

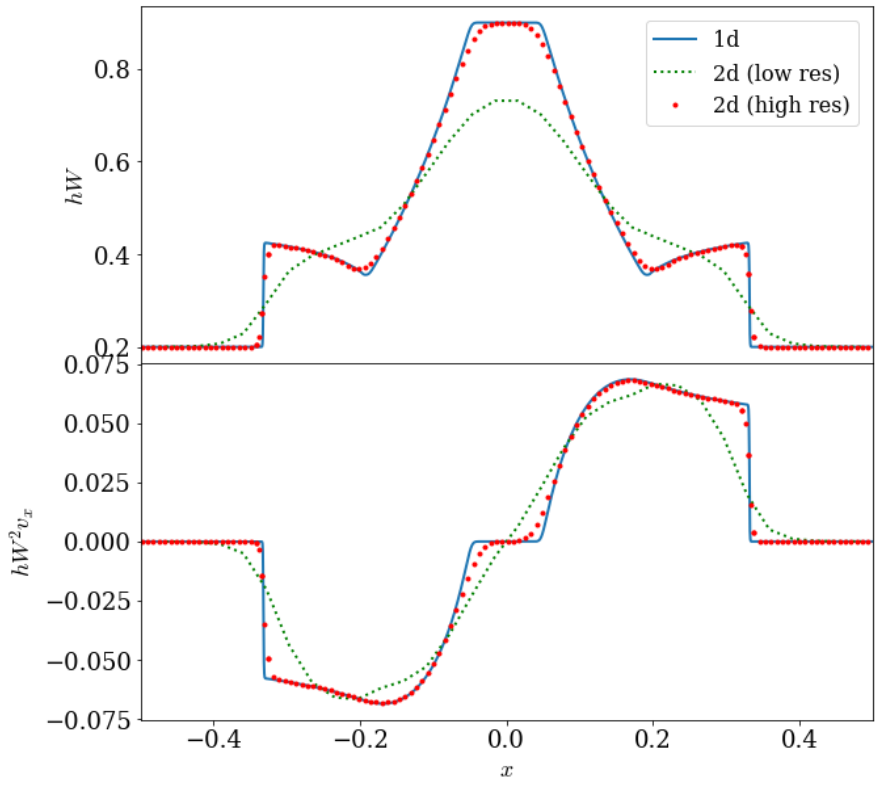

FIGURE 4.4: Results of the relativistic radial dam breaking test with $h_{\text{in}} = 0.9$ and $h_{\text{out}} = 0.2$ and initial dam radius of $r = 0.2$ at time $t = 0.5$ compared against a high resolution (2048 grid points) 1d simulation with geometric source terms. The simulation was carried out on a the domain $x = [-0.5, 0.5]$, $y = [-0.5, 0.5]$, with 32×32 grid points for the low resolution run and 256×256 for the high resolution run. The plot shows the conserved variables hW and $hW^2 v_x$.

significantly rounded for the relativistic system. This was a sharp peak for the Newtonian case.

Tangential velocity

In Section 3.3.1.3, we found that a fast tangential velocity could change the wave pattern in the relativistic Riemann problem. This was investigated by repeating the 1d dam breaking test with an additional tangential velocity of $v_t = 0.5$, $v_t = 0.7$ and $v_t = 0.9$. The results of this can be seen in Figure 4.6. With the addition of a non-zero tangential velocity, in order for the wavespeeds to remain subluminal the initial height of the fluid had to be reduced to ensure that $D = hW < 1$ for even the most relativistic case, where the Lorentz factor $W \sim 2.3$. The initial heights used were $h_l = 0.41$, $h_r = 0.01$.

As the primitive normal velocity in the intermediate star state was already reasonably high for the case with zero tangential velocity ($v_x = 0.561$), the effects of introducing a tangential velocity to this system are quite significant. As the tangential velocity is increased, the shock speed is decreased, such that the shock front has propagated less

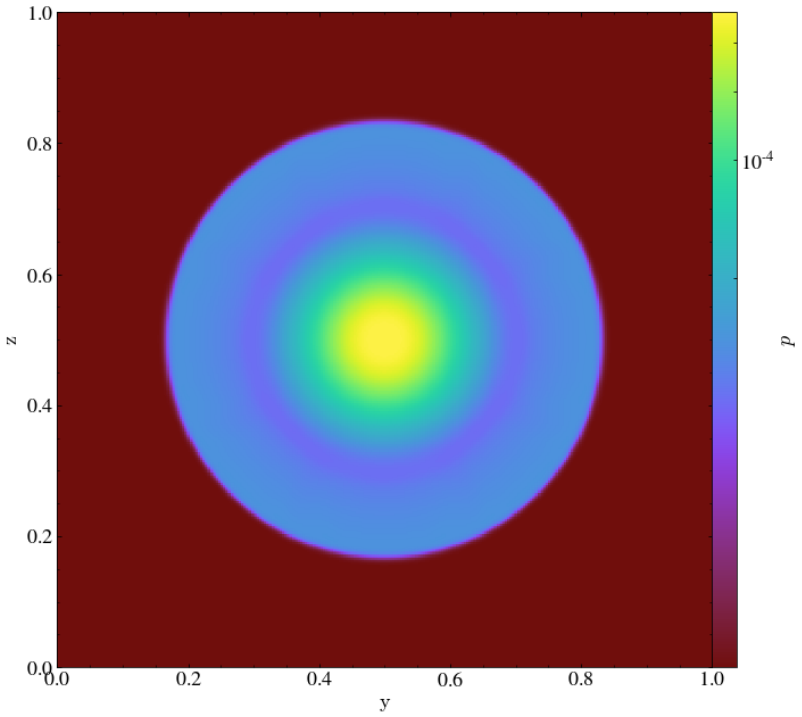

FIGURE 4.5: Results of the relativistic radial dam breaking test with $h_{\text{in}} = 0.9$ and $h_{\text{out}} = 0.2$ and initial dam radius of $r = 0.2$ at time $t = 0.5$. The simulation was carried out on a the domain $x = [-0.5, 0.5]$, $y = [-0.5, 0.5]$, with 256×256 grid points. The plot shows the fluid pressure.

far in the final state. We can interpret this physically by thinking of the inertia of the system. If the fluid is already moving at high speed in the tangential direction, then it will take a greater force in order to get it to change direction. Consequently, as the tangential velocity is increased, the shock speed in the normal direction decreases.

The conserved variable hW increases monotonically with increased tangential velocity. As v_t is increased, this increases the energy of the intermediate state, increasing hW in the region. Given the nonlinear coupling between v_t and the Lorentz factor, this change is nonlinear, with hW in the $v_t = 0.9$ case significantly greater than it is for the runs with smaller tangential velocities.

As the inset plot shows, increasing v_t causes the height in the intermediate state to increase, however this is a very small change, even for $v_t = 0.9$. The normal velocity in the intermediate state is reduced. This is in accordance with the findings of Section 3.3.1.3, where the magnitude of the normal velocity in the intermediate state was reduced with increased tangential velocity, and is what would be expected for a relativistic system where the overall speed of the flow cannot exceed the speed of light. For the case with the highest tangential velocity, $v_t = 0.9$, the inset plot shows that in the intermediate state there are some small oscillations present in the height. This is likely due to the

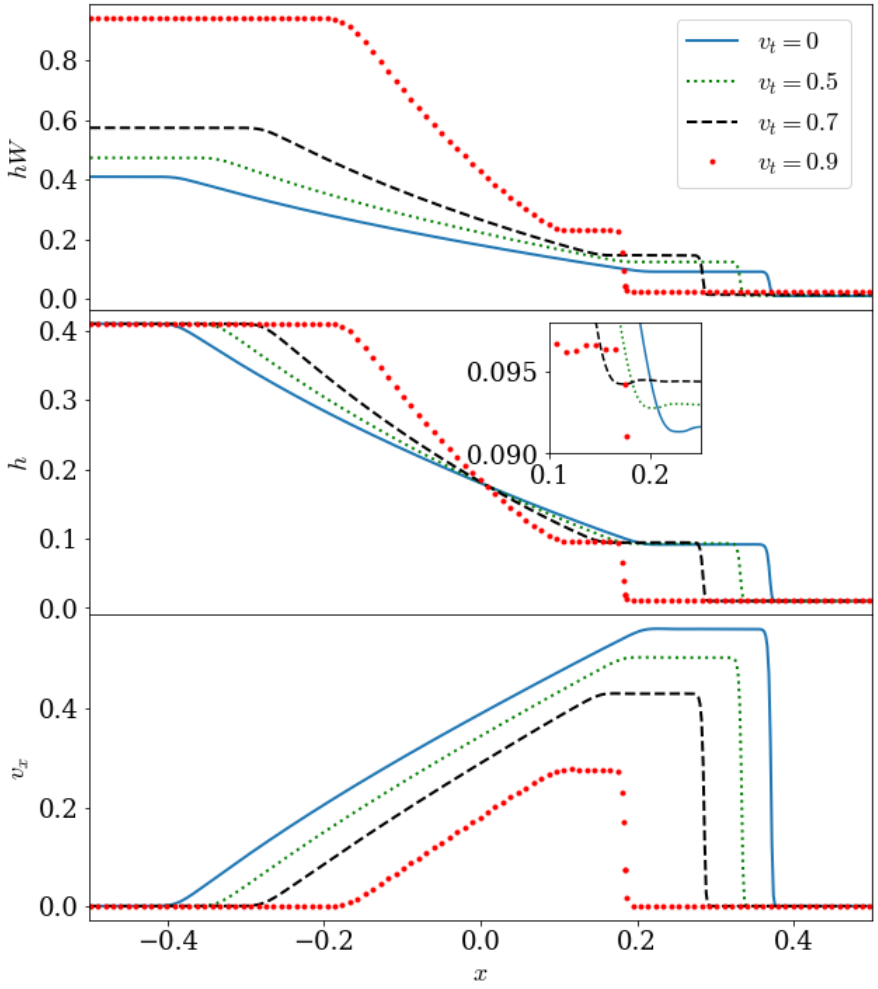

FIGURE 4.6: The effects of introducing a non-zero tangential velocity in the relativistic dam breaking test. The previous 1d dam breaking test was repeated with the addition of a non-zero tangential velocity v_t. A resolution of 512 was used in the x-direction. The plots show the conserved variable $D = hW$, the height and the normal velocity at time $t = 0.6$. Increasing v_t increases hW, decreases the shock speed and changes the height and normal velocity in the intermediate state. The inset plot shows a close up of the height in the intermediate state.

high relativistic velocities involved in this test. The inset plot also shows that there is a small undershoot at the tail of the rarefaction for all values of v_t.

4.4 Summary

In this chapter, we looked at the shallow water equations, considering how they may be extended to model relativistic fluids and include multiple layers and reactions. We began by looking at the motivation behind the shallow water equations and how they can be derived from the compressible fluid equations. This derivation was then repeated for the general relativistic fluid equations to obtain a set of general relativistic shallow

water equations. In order to compare these with the Newtonian equations, a non-traditional form of the Newtonian equations was used which evolved the geopotential of the fluid layer rather than the fluid height. This was done in order for it to be possible to model a varying gravitational field in the general relativistic case (the traditional form of the equations assumes constant gravitational acceleration g). The tests of the numerical implementation presented here all use a constant graviational field, however this is something that we hope to implement in the future. The shallow water equations were then extended to allow for multiple layers of fluid, adding buoyancy terms, and adding an equation for the evolution of the species mass fraction(s) so that reactions could be modelled.

We then implemented the relativistic form of the single layer equations in the `CASTRO` hydrodynamics code. Results of dam breaking tests in 1d and 2d were presented, demonstrating the performance of the code. For the Newtonian 1d dam breaking test, the code was shown to reproduce analytic results. The 2d circular dam breaking tests were shown to match up well with high resolution 1d dam breaking simulations with geometric source terms. The effects of a non-zero tangential velocity were investigated for the relativistic 1d dam breaking problem. As found in the previous chapter, it was shown that varying the tangential velocity changes the shock speed and intermediate state.

The numerical implementation of the relativistic shallow water equations presented here can be considered special relativistic in the horizontal plane of the domain, where we have used the Minkowski metric. However, a key component of the shallow water equations is the inclusion of a gravitational field in the vertical direction. This gravitational field means that overall the metric for our system is non-Minkowski, and we can consider the implementation to be general relativistic for a system with a constant gravitational field. In the future, we intend to extend this to the full general relativistic shallow water equations, allowing for a gravitational field that may vary in time and/or across the physical domain to be studied. This will allow us to use a Kerr- or Schwarzschild-like metric to investigate the effects of frame dragging and the oblateness of the neutron star on X-ray burst propagation.

The numerical implementation we have demonstrated includes only the single layer, inert shallow water equations. In the future, we hope to extend this to multiple layers and reactions using the formulation presented in this chapter.

5 Multiscale methods

In previous chapters, we have looked at modelling the small scale features of the burning front using the low Mach number approximation, and at modelling the large scale, Coriolis-force dominated front propagation using the shallow water equations. In this chapter, we shall combine these different models in order to produce a multiscale model which incorporates the physics from the smallest centimetre scales up to the scale of the entire neutron star.

In standard mesh refinement techniques (Berger and Oliger, 1984), the same physical model is used at all levels of the mesh refinement hierarchy. However, by using different physical models at different resolutions, we will be able to better capture the physics at different length scales: a model that works well at one length scale is unlikely to be the best model at length scales 6 orders of magnitude smaller.

Adaptive mesh and algorithm refinement (AMAR) (Weinan, 2011) builds on the standard adaptive mesh refinement (AMR) algorithm, adding the possibility for different physical models or numerical solvers to be used at different levels of the grid hierarchy. This can offer significant computational speed up, with cheaper models used on coarser grids where accuracy is less important. It can also help model systems where there exists a separation of scales, with some processes operating at larger scales only, and some at smaller scales only. Allowing for different physical models at these different scales can help to better capture the physics of these processes.

In our system, we wish to model scales from $\sim 10\,\mathrm{km}$ down to $\sim 1\,\mathrm{cm}$. This will allow us to capture both the large scale Coriolis forces and the small scale nuclear reactions (Malone et al., 2011). Assuming the coarsest grids to have a resolution of $\sim 10\,\mathrm{m}$, we estimate that around 10 levels of mesh refinement will be required, giving an increase in resolution of a factor of $2^{10} \sim 10^3$ at the finest level. The coarsest grid shall cover the whole domain and evolve the system using the single layer shallow water equations. This shall then be refined to grids which will evolve the multilayer shallow water equations. The next levels of refinement, around the 1 m scale, will evolve the system using the fully compressible fluid equations. The finest grids shall cover only the part of the domain at the flame front and its immediate vicinity, and will evolve this using the low Mach number equations.

To couple these different grids together, it will be necessary to communicate the data between them where they overlap and intersect. As the grids have different resolutions, mesh refinement methods must be used to increase or decrease the resolution. On different grids we will be using different physical models and so will use different sets of conserved variables. Consequently, we will need to develop techniques to consistently convert variables between the models.

At this point, it may appear that the intermediate compressible model is unnecessary: why not use a multiscale model comprised of the shallow water and low Mach number models alone? Whilst the low Mach number model will require much fewer timesteps than a compressible simulation of the same spatial resolution, it is a lot more expensive to evolve a single timestep of the low Mach number model since it requires computationally time-consuming projection methods and elliptic solves. This trade off means that on the coarser grids with fewer grid points and a less restrictive timestep condition, it is more efficient to use the faster compressible model. As suggested by Heyl (2004), there could be large scale modes in the ocean that interact with the small scale features in the flame front as they pass through. In order to get proper coupling between these small and large scale features, it is necessary to have sufficient compressible levels, as described by Motheau et al. (2018).

In this chapter, we shall begin by looking at adaptive mesh refinement and adaptive mesh and algorithm refinement. We shall then describe the additional numerical methods required to implement the multiscale model, before showing the results of tests performed on the numerical implementation which demonstrate its performance and limitations. It is the overall plan of this work to include the low Mach number, compressible and shallow water models. However, in the numerical implementation of the multiscale scheme presented in this work we shall focus on a large scale model, coupling the compressible and shallow water models only.

5.1 Adaptive mesh refinement

Adaptive mesh refinement is a numerical technique used to efficiently model large physical domains whilst retaining the high resolution required to resolve features such as boundary layers or regions with steep gradients, shocks or discontinuities. It was first developed by Berger and Oliger (1984) for the solution of hyperbolic partial differential equations. This was later developed by Berger and Colella (1989) to produce a conservative scheme for modelling gas dynamics, which was then extended to the 3d case by Bell et al. (1994). There are now a number of high level AMR frameworks, including `AMReX` (Zhang et al., 2016), `Cactus` (Goodale et al., 2003), `Chombo` (Colella et al., 2013), `Enzo` (O'Shea et al., 2004; Bryan et al., 2014), `FLASH` (Fryxell et al., 2000) and `Paramesh` (MacNeice et al., 2000) (see Dubey et al. (2014) for a review).

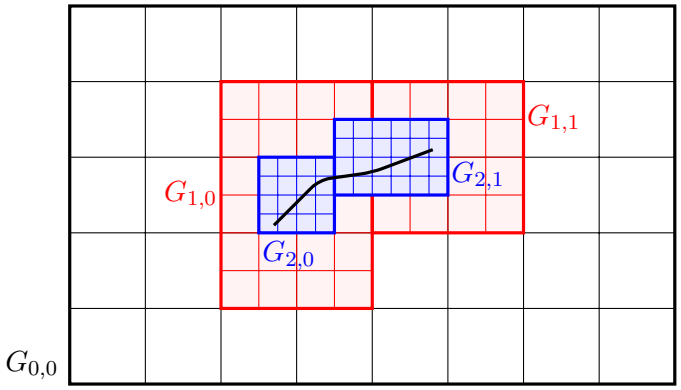

FIGURE 5.1: Example of a 2d AMR grid hierarchy with 3 levels of mesh refinement and refinement ratio $r_\ell = 2$. The entire domain is covered by the coarsest level 0 black grid, $\ell_0 = G_{0,0}$. The black line indicates a feature of interest that requires higher resolution. This feature is enclosed within the red level 1 grids $\ell_1 = \bigcup_i G_{1,i}$ and the blue level 2 grids $\ell_2 = \bigcup_i G_{2,i}$.

Several different approaches to AMR exist, including unstructured grids, mesh distortion and point-wise structured refinement; we shall use the original block-structured approach of Berger and Oliger (1984). The basic idea of block-structured AMR is to cover the entire domain with a coarse grid, then adaptively place finer grids on regions where higher resolution is needed. Finer grids still can be added until the required resolution is reached. If the problem being solved is time dependent, the locations of the regions requiring higher resolution are likely to move with time. Consequently, the grids must be updated and adapted throughout the course of the simulation.

In the sections below, we shall outline the main principles of block-structured AMR. Further details can be found in Bell and Day (2011).

5.1.1 Grid hierarchy

The set of grids of increasing resolution is known as the *grid hierarchy*. An example of a grid hierarchy for a 2d problem can be seen in Figure 5.1. The entire domain is covered by the coarsest grid, $G_{0,0}$. In this example, there is a feature of interest (indicated by the thick black line), around which we desire higher resolution. We therefore place higher resolution grids in the vicinity of this feature so that it is completely enclosed within grids of the finest resolution. The ratio of the resolutions between each successive level of grids is known as the *refinement ratio*, $r_\ell = \Delta x_\ell / \Delta x_{\ell+1}$ – this is typically 2. Grids must satisfy 'proper nesting requirements': grids at finer levels must start and end at the corners of cells in the next coarsest level, and there must be at least $n_{\text{buf}} > 0$ level ℓ_j cells between the edge of grids on level ℓ_{j+1} and any uncovered ℓ_{j-1} cells, unless the level ℓ_{j+1} grid is touching the boundary of the physical domain. An important consequence of this is that interfaces between adjacent levels coincide with cell boundaries on both

levels. As can be seen in the figure, a fine grid can cross a coarser grid boundary and still be properly nested.

Note that it is not necessary for all the grids in a particular level to be connected. If there were another feature of interest elsewhere in the domain, it would be possible for finer grids to be created to cover that feature, with none of these grids intersecting with or connected to the existing finer grids.

As the solution evolves with time, the feature of interest may move outside of the region covered by the finer grid, or change in a way that it is no longer necessary to model the feature at such high a resolution. At discrete intervals between timesteps, the simulation will therefore assess the current state of the grids and *regrid* if necessary. Regridding involves investigating the current cells to see if they satisfy some tagging criteria set by the user. Tagged cells are then grouped into rectangular boxes, and new grids generated or old grids destroyed where required.

5.1.2 Synchronisation

When evolving the system, finer grids are typically evolved using a shorter timestep, $\Delta t_\ell = \Delta t_0 / r^\ell$. For every timestep of the coarse grid, grids on level ℓ are evolved by r^ℓ timesteps. This is known as *subcycling*. For each subcycled timestep, ghost data will be required at the interfaces of the fine grids with their adjacent coarser grids. It will therefore be necessary to map data between the coarse and fine grids. Data will also need to be mapped from coarser grids to finer grids when new grids are created during regridding. Data on the fine grid is mapped to the coarse grid using a restriction routine, and mapped from the coarse grid to the fine grid using an interpolation routine. The restriction and interpolation operations can take the form of simple averaging and linear interpolation, or involve more elaborate higher order reconstruction routines depending on the behaviour of the system and the accuracy required.

The second step of synchronisation between grids is *refluxing*. This corrects for the difference in coarse and fine grid fluxes at the fine grid boundaries to ensure conservation. The correction to the coarse cell on level ℓ at the fine grid boundary is made by taking the fluxes on level $\ell + 1$ averaged over the level ℓ timestep and spatially averaged over the level ℓ cell face.

5.1.3 Error estimation and regridding

Error estimation routines indicate areas where a higher solution accuracy is required. Such areas are then tagged for further refinement. The method used to estimate the error can be very simple, e.g. searching for cells where the velocity exceeds some value, or can be more sophisticated, e.g. calculating the error on the coarse grid using Richardson

extrapolation. In Richardson extrapolation, the data on the grid is coarsened then integrated for a timestep, then compared to the result of integrating first then coarsening. The method used for error estimation is ultimately determined by the properties of the system to be modelled, the computational constraints of the problem and the level of accuracy desired.

Once the cells requiring refinement have been tagged, these cells are clustered together to form rectangular boxes and refined to generate the new grids. The size and number of these new finer grids is determined by the number of buffer cells chosen to surround each tagged cell and the grid efficiency $\eta =$ (number of cells tagged)/(number of cells in new grids). Regridding must be done simultaneously at all finer levels to ensure that the new grid hierarchy does not violate proper nesting requirements. It is often not necessary to regrid every timestep: provided a sufficient number of buffer cells have been chosen, regions requiring higher resolution should not move outside of the finer grids over the course of a single timestep. As regridding is quite a costly operation, it is therefore a more optimal solution to instead regrid every integer $n > 1$ timesteps.

5.1.4 Reconstruction

Prolongation and restriction methods communicate the data between the different grid resolutions. We shall outline these methods below for a 2d grid with a refinement ratio $r = 2$; they can be trivially extended to higher dimensions.

Prolongation

Prolongation reconstructs variables on the fine grid from those on the coarse grid. This is done by first calculating the (limited) slopes in the cells in the x and y directions, S_x and S_y, then interpolating. Given the coarse state at the centre of cell $(2i, 2j)$ in the fine grid coordinate system $\boldsymbol{q}_{\mathrm{c},2i,2j}$, we then calculate the fine grid states within this cell $\boldsymbol{q}_{\mathrm{f},2i,2j}$ (as shown by Figure 5.2) as:

$$\begin{aligned}
\boldsymbol{q}_{\mathrm{f},2i,2j} &= \boldsymbol{q}_{\mathrm{c},2i,2j} - \frac{1}{4}(S_x\Delta x + S_y\Delta y),\\
\boldsymbol{q}_{\mathrm{f},2i+1,2j} &= \boldsymbol{q}_{\mathrm{c},2i,2j} + \frac{1}{4}(S_x\Delta x - S_y\Delta y),\\
\boldsymbol{q}_{\mathrm{f},2i,2j+1} &= \boldsymbol{q}_{\mathrm{c},2i,2j} + \frac{1}{4}(-S_x\Delta x + S_y\Delta y),\\
\boldsymbol{q}_{\mathrm{f},2i+1,2j+1} &= \boldsymbol{q}_{\mathrm{c},2i,2j} + \frac{1}{4}(S_x\Delta x + S_y\Delta y).
\end{aligned} \tag{5.1}$$

Δx and Δy are the grid spacings of the coarse grid.

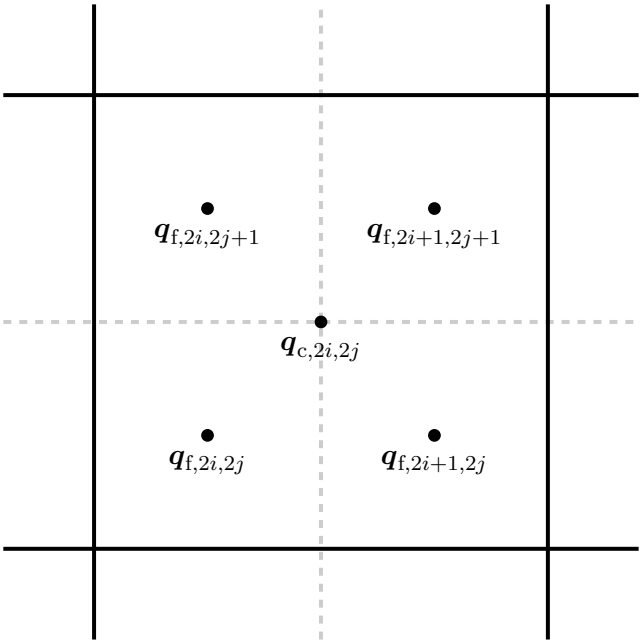

FIGURE 5.2: Coarse grid (solid black) with respect to fine grid (dashed grey). Note that the cells are labelled with respect to the position of the bottom left hand corner, but the cell averaged states are assigned to live at the centres of the cells.

Restriction

Restriction interpolates the variables on the fine grid to the coarse grid. To obtain the variables on the coarse grid, we interpolate, taking the average of its four neighbouring fine grid cells:

$$\boldsymbol{q}_{\mathrm{c},2i,2j} = \frac{1}{4}\left(\boldsymbol{q}_{\mathrm{f},2i,2j} + \boldsymbol{q}_{\mathrm{f},2i+1,2j} + \boldsymbol{q}_{\mathrm{f},2i,2j+1} + \boldsymbol{q}_{\mathrm{f},2i+1,2j+1}\right). \tag{5.2}$$

5.1.5 AMR algorithm

An outline of the block-structured AMR algorithm for a system with 2 levels and the refinement ratio r is as follows:

While time $t_c < t_{\mathrm{end}}$:

- Regrid (if it is time)
- Advance coarse grids in time $t_c \to t_c + \Delta t_c$
- Advance fine grids in time r times
- Synchronise coarse and fine data

This can then be extended recursively to an arbitrary number of levels of refinement.

5.2 Adaptive mesh and algorithm refinement

As described above, in standard AMR, the same physical model is evolved on all grids. All grids will therefore store and evolve the same sets of physical variables. In adaptive mesh and algorithm refinement, this is no longer the case. In order to better capture the physics at the smaller scales, a different physical model may be evolved or a more accurate numerical method may be used on the finer grids. AMAR is a relatively new technique that is starting to be used more widely for systems involving physics across multiple scales. It has been used to model a number of different physical systems. Below are listed a few recent examples.

- Garcia et al. (1999) used AMAR to embed a particle method within a continuum model, using a strongly coupled scheme where both methods simultaneously evaluate different regions of flow and continuously exchange information across the interface (Weinan, 2011).
- In the context of modelling CO_2 injected into aquifers, Møyner and Nilsen (2017); Becker et al. (2017) coupled multilayer vertical equilibrium models (which are similar to the shallow water equations and found by integrating all quantities vertically) with the compressible equations to produce a hybrid model. At large scales away from the gas inlet, the coarser vertical equilibrium models are used; at finer scales near the inlet, the compressible equations are used. At the interface, details in the vertical direction are reconstructed from the imposed equilibrium pressure distribution.
- A hybrid low Mach number/compressible scheme has recently been developed by Motheau et al. (2018) and implemented in their `LAMBDA` code. Built using routines from `CASTRO` and `MAESTRO` with `AMReX` for the AMR, this model uses a compressible solver for the coarse grids and a low Mach number solver for the fine grids. This allows them to capture small amplitude, large wavelength acoustic waves which would be filtered out by a low Mach-only model.
- Xu et al. (2016) used a multiscale model to investigate the propagation of defects in solids, using a nonlocal peridynamics model near the discontinuity, and a local PDE model away from the discontinuity. In the immediate vicinity of the defect, discontinuous Galerkin methods (which are better suited to modelling systems with discontinuities) were used. Away from the defect, cheaper continuous Galerkin methods were used instead.

The introduction of additional physical models and numerical solvers adds some additional complexity to the AMAR algorithm. The various physical models may evolve different sets of physical variables. The finer grids may be defined using a different

coordinate system, such that coarse grid boundaries no longer coincide with fine grid boundaries. Solvers may involve non-local calculations requiring data defined on multiple grids and levels. As a result, synchronisation of data between the different grids becomes much more complex (especially if the code is run in parallel). Standard restriction and interpolation routines may need to be modified in order to convert between the different sets of physical variables in a consistent way and synchronise data across grids with different coordinate systems.

It is likely that the fine grid model will have different timestep restrictions to the coarse grid model, such that $\Delta t_c \neq r\Delta t_f$. It may even be the case that the system is evolved with 'inverse' subcycling, with coarse grids evolved more frequently than finer grids. For example, for a model comprised of compressible hydrodynamics on the coarse grid and the low Mach number model on the finer grids (such as that of Motheau et al. (2018)), the CFL condition for the compressible grids will be much more restrictive than that for the low Mach number grids: $\Delta t_{\mathrm{LM}}/\Delta t_{\mathrm{comp}} \sim 1/\mathrm{Ma}$. The timesteps for the compressible grids will therefore be much shorter than for the low Mach number grids, and for each advancement of the coarsest low Mach number grid, the finest compressible grid will be advanced $\mathcal{O}(1/\mathrm{Ma})$ times.

These additional complexities are some of the main reasons why AMAR has not been used more widely in the past, as they often require significant modifications to existing AMR frameworks. Of the examples listed above, our implementation is clearly most similar to the work of Motheau et al. (2018): our work can be thought of as a complementary approach to modelling similar astrophysical systems. However, rather than coupling the compressible fluid equations to the low Mach number equations to better model the small scale physics, we have chosen to focus on the opposite end of the scale, coupling the compressible fluid equations to the shallow water equations so we can better model the large scale physics. In the future, it would be interesting to join these two approaches together, producing a single model that can investigate the physics across the entire range of scales relevant for modelling X-ray bursts.

5.3 Analytic multiscale methods

In problems such as those we wish to investigate, it is often the case that the state variables can be split into two: a set of slow variables describing slower processes that happen on longer timescales, and a set of fast variables describing faster processes that occur on shorter timescales. The processes that we are more interested in are described by the slow variables, so it would be useful to extract these from the full equations to obtain evolution equations for the slow variables alone. An analytic technique used in multiscale modelling to obtain a model for these slow variables is to introduce a small parameter ε to explicitly separate them from the fast variables, and use a multiscale

expansion to match terms of the same timescale (Weinan, 2011). This small parameter essentially rescales the slow variables to make them $\mathcal{O}(1)$, while reducing the fast variables to be $\mathcal{O}(\varepsilon)$. As we shall demonstrate in the next couple of sections, this technique provides an alternative way to derive the low Mach number and shallow water equations which is consistent with the traditional approaches used in previous chapters.

5.3.1 Low Mach number equations

To illustrate this method, we shall first use it to derive the Newtonian low Mach number equations, given above in (2.10)–(2.12). In this system, the small parameter will be the Mach number. To begin with, we shall first consider the momentum equation

$$D_t v_i + \frac{1}{\rho}\partial_i p = -\partial_i \phi, \tag{5.3}$$

where $D_t = \partial_t + v^i\partial_i$ is the material derivative. If we non-dimensionalise this, setting all variables to be the product of a dimensionless variable $\hat{q}$ and a reference parameter q_r, such that $q = \hat{q}q_r$, then we obtain the dimensionless momentum equation

$$D_t v_i + \frac{1}{\mathrm{Ma}^2}\frac{1}{\rho}\partial_i p = -\frac{1}{\mathrm{Fr}^2}\partial_i \phi, \tag{5.4}$$

where $\mathrm{Fr} = v_r/\sqrt{g_r \ell_r}$ is the (reference) Froude number. We now introduce the small parameter $\varepsilon \sim \mathrm{Ma} \sim \mathrm{Fr}$, such that the dimensionless momentum equation becomes

$$D_t v_i + \varepsilon^{-2}\frac{1}{\rho}\partial_i p = -\varepsilon^{-2}\partial_i \phi. \tag{5.5}$$

If we assume that all variables can be decomposed into a slow variable q_0 and a fast variable q_1 like $q = q_0 + \varepsilon q_1$, then the different scales will separate at each order of ε. This also works for the time variable, which we can decompose into a slow timescale τ and a fast timescale $\tilde{t}$ such that $\tilde{t} = \varepsilon\tau$ and $D_t \to D_\tau + \varepsilon D_{\tilde{t}}$. Looking at the order-by-order expansion, the $\mathcal{O}(\varepsilon^{-2})$ terms give us an expression for the divergence of the slow timescale pressure,

$$\frac{1}{\rho_0}\partial_i p_0 = -\partial_i \phi_0, \tag{5.6}$$

the $\mathcal{O}(\varepsilon^{-1})$ terms give

$$\frac{1}{\rho_0}\partial_i p_1 - \frac{1}{\rho_0^2}\rho_1 \partial_i p_0 = -\partial_i \phi_1, \tag{5.7}$$

and the $\mathcal{O}(\varepsilon^0)$ terms give an evolution equation for the slow timescale velocity,

$$D_\tau (v_0)_i + \frac{1}{\rho_0}\partial_i p_2 - \frac{1}{\rho_1}\partial_i p_1 = 0. \tag{5.8}$$

Repeating this analysis for the continuity equation and taking the leading order part, we get an evolution equation for the slow timescale density,

$$D_\tau \rho_0 + \rho_0 \partial_i (v_0)^i = 0. \tag{5.9}$$

Similarly, taking the leading order part of the energy equation (and assuming an ideal gas equation of state) gives us an evolution equation for the slow timescale pressure,

$$D_\tau p_0 + \gamma p_0 \partial_i (v_0)^i = 0. \tag{5.10}$$

Taking this last equation gives us a divergence constraint for the slow timescale velocity,

$$\partial_i (v_0)^i = -\frac{1}{\gamma p_0} D_\tau p_0. \tag{5.11}$$

Just as we did in our original derivation, we have obtained this divergence constraint from the energy equation and the equation of state. If we impose that the slow timescale pressure p_0 is constant in time such that $D_\tau p_0 = 0$, then we regain the divergence constraint of incompressible hydrodynamics for the slow timescale velocity.

Given that in our problem the gravitational field will be constant in time so $\phi = \phi_0$, and taking (5.6), we find that the slow timescale pressure must be in hydrostatic equilibrium,

$$\partial_i p_0 = -\rho_0 \partial_i \phi. \tag{5.12}$$

As $\phi_1 = 0$, (5.7) now tells us that p_1 is a linear function of p_0 (i.e. it is not a linearly independent variable),

$$\partial_i p_1 = \frac{\rho_1}{\rho_0} \partial_i p_0 = -\rho_1 \partial_i \phi. \tag{5.13}$$

This is consistent with the 'traditional' derivation of the low Mach number equations, where we find that lowest order pressure perturbations of relevance to the low Mach number approximation are at second order in Mach number.

Finally, from the momentum equation (5.8), we find the familiar low Mach number expression containing the derivative of the second order pressure perturbation,

$$D_\tau (v_0)_i + \frac{1}{\rho_0} \partial_i p_2 = -\partial_i \phi. \tag{5.14}$$

Comparing this to the Newtonian low Mach number momentum equation (2.11), we see that we can identify $p_2 = \pi$. The β_0 compressibility factor of (2.11) is absent in this multiscale derivation – as in the original derivation (see Section B.4), it is still necessary to add this term in 'by hand' by carefully considering the energy conservation along particle paths.

5.3.2 Shallow water equations

Another system of equations we can derive using the multiscale approach are the Newtonian shallow water equations. In this case, instead of looking at slow and fast variables, we can instead think of small scale and large scale variables describing processes that occur at shorter and longer length scales. We shall use the small parameter to rescale the vertical direction under the assumption that this is much smaller in extent than the horizontal directions. To do this, we define the rescaled vertical coordinate $Z = z/\varepsilon$ and let I, J denote the horizontal coordinates, such that we can now write the continuity and momentum equations as

$$\partial_t \rho + \partial_I \left(\rho v^I\right) + \varepsilon^{-1} \partial_Z \left(\rho v^Z\right) = 0, \tag{5.15}$$

$$\partial_t (\rho v_J) + \partial_I \left(\rho v_J v^I + \delta^I_J p\right) + \varepsilon^{-1} \partial_Z \left(\rho v_J v^Z\right) = -\rho \partial_J \phi, \tag{5.16}$$

$$\partial_t (\rho v_Z) + \partial_I \left(\rho v_Z v^I\right) + \varepsilon^{-1} \partial_Z \left(\rho v_Z v^Z + p\right) = -\rho \partial_Z \phi. \tag{5.17}$$

We next decompose the variables into a large scale part q_0 and a small scale part q_1, $q = q_0(t, x^I) + \varepsilon q_1(t, x^I, Z)$. Substituting this in and keeping terms up to $\mathcal{O}(\varepsilon^0)$, we get

$$\partial_t \rho_0 + \partial_I \left(\rho_0 v_0^I\right) + \varepsilon^{-1} \partial_Z \left(\varepsilon \left[\rho_1 v_0^Z + \rho_0 v_1^Z\right]\right) = 0, \tag{5.18}$$

$$\begin{aligned} &\partial_t \left(\rho_0 v_{0J}\right) + \partial_I \left(\rho_0 v_{0J} v_0^I + \delta^I_J p_0\right) + \\ &\qquad \varepsilon^{-1} \partial_Z \left(\varepsilon \left[\rho_1 v_{0J} v_0^Z + \rho_0 v_{1J} v_0^Z + \rho_0 v_{0J} v_1^Z\right]\right) = -\rho_0 \partial_J \phi_0, \end{aligned} \tag{5.19}$$

$$\begin{aligned} &\partial_t \left(\rho_0 v_{0Z}\right) + \partial_I \left(\rho_0 v_{0Z} v_0^I\right) + \\ &\qquad \varepsilon^{-1} \partial_Z \left(\varepsilon \left[\rho_1 v_{0Z} v_0^Z + \rho_0 v_{1Z} v_0^Z + \rho_0 v_{0Z} v_1^Z + p_1\right]\right) = 0, \end{aligned} \tag{5.20}$$

where we have enforced the fact that the small scale variables are constant in the vertical direction, $\partial_Z q_0 = 0$. We now impose a fixed boundary at the bottom, such that $v_Z = 0$ at this boundary. As v_{0Z} is a function of time and the horizontal directions only, $v_{0Z} = v_{0Z}(t, x^I)$, this means that $v_{0Z} = 0$ everywhere. This reduces our equations to

$$\partial_t \rho_0 + \partial_I \left(\rho_0 v_0^I\right) + \partial_Z \left(\rho_0 v_1^Z\right) = 0, \tag{5.21}$$

$$\partial_t \left(\rho_0 v_{0J}\right) + \partial_I \left(\rho_0 v_{0J} v_0^I + \delta^I_J p_0\right) + \partial_Z \left(\rho_0 v_{0J} v_1^Z\right) = -\rho_0 \partial_J \phi_0. \tag{5.22}$$

Finally, we can integrate the equations in the vertical direction, denoting vertically integrated quantities by angle brackets to obtain

$$\partial_t \langle \rho_0 \rangle + \partial_I \left\langle \rho_0 v_0^I \right\rangle = 0, \tag{5.23}$$

$$\partial_t \left\langle \rho_0 v_{0J} \right\rangle + \partial_I \left\langle \rho_0 v_{0J} v_0^I + \delta^I_J p_0 \right\rangle = - \left\langle \rho_0 \partial_J \phi_0 \right\rangle . \tag{5.24}$$

If we define the gravitational potential to be $\partial_j \phi = g \delta^z_j$, the $\langle \rho_0 \partial_J \phi_0 \rangle$ term then vanishes. By definition, ρ_0 is constant in Z, so using the equation for hydrostatic equilibrium

$$p_0(z) = g \rho_0 z, \tag{5.25}$$

we get

$$\langle p_0 \rangle = \frac{1}{2} g \rho_0 H^2, \tag{5.26}$$

where $H = H(t, x^I)$ is the fluid height. Again, using the fact that ρ_0 is constant in Z, we find that its vertical integral must be $\langle \rho_0 \rangle = \rho_0 H$. Substituting this in and letting $u^I = \langle v_0^I \rangle / H$, we obtain the shallow water equations in their standard form

$$\partial_t H + \partial_I \left(H u_I \right) = 0, \tag{5.27}$$

$$\partial_t \left(H u_J \right) + \partial_I \left(H u_J u^I + \frac{1}{2} g H^2 \delta_J^I \right) = 0. \tag{5.28}$$

5.3.3 Limitations of the multiscale approach

From the above multiscale derivations of the low Mach number and shallow water equations, we can gain insight into a number of limitations of the multiscale approach. With the introduction of the small parameter ε, we have explicitly separated out the slow and fast timescales for the low Mach number equations, and the vertical and horizontal spatial scales in the shallow water equations. By separating the scales, we have assumed that there is no interaction of features at the different scales. For the problems we are interested in, this is likely to be a valid assumption for most of the behaviour we wish to model, however it is important to keep this in mind when e.g. designing the tagging criteria for the mesh refinement so that features are not lost when passed between the two scales in the numerical implementation.

However, this may be more problematic when modelling other systems. The lack of interaction between scales is particularly limiting for systems where there are resonances between features at different scales. If this is the case, the modelled solution will pass through these resonances, and will be unable to be 'locked in' to them, as may be the case in the true physical solution. As shall be described further in Section 6, an example of an astrophysical system with resonances between features at different scales which the multiscale approach would have difficultly capturing is the *p*-*g* coupling instability.

In splitting the scales, we have also assumed that the higher order terms grow sublinearly as a function of the fast timescale, i.e. $q_1(\tau, \tilde{t})/\tau \to 0$ as $\tau \to \infty$. This is a necessary condition to ensure that the leading order terms remain dominant over the time period of interest. The multiscale approach will fail to accurately model systems where this is not the case.

5.4 Numerical methods

In the following sections we shall describe the numerical methods used to implement our multiscale model. For simplicity, we shall concentrate only on a system with two

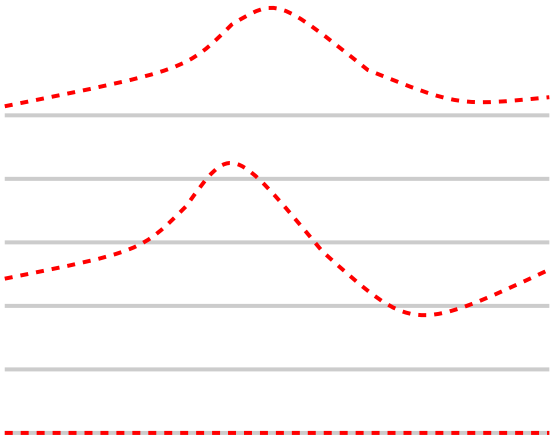

FIGURE 5.3: Equally spaced layers of the compressible grid (solid grey) embedded in the shallow water grid with variably spaced layers (dashed red). Note that the bottom layers of the shallow water and compressible grids coincide and that the top layer of the compressible grid must be beneath the fluid surface.

levels of refinement which evolves the shallow water equations on the coarse grid and the compressible equations on the fine grid(s).

In going between the shallow water and compressible models, the problem is different in the vertical and horizontal directions. In the horizontal directions, the compressible grid cells match up in a uniform, consistent way with the shallow water ones: the square shallow water cells are split into four smaller square cells (see Figure 5.2). However, the problem is more difficult in the vertical direction. The compressible grid layers are uniformly spaced in height, but for the multilayer shallow water system, the layers are not uniformly spaced and each one can intersect several compressible layers (see Figure 5.3). Note that the surface shallow water layer must be above the top compressible layer – the compressible cells must be contained within the fluid (i.e. they should not extend into the atmosphere above the fluid surface). In the single layer shallow water model, we only model the top surface layer.

In the multiscale model, at the level where we convert from one model to another we have two grids of the same resolution but with the two different sets of conserved variables of the models to be converted between. The conversion between the models is therefore always done between grids of equal resolution. As described above, for the shallow water to compressible case, this still requires some interpolation and/or refinement in the vertical direction, which shall be outlined below.

5.4.1 Shallow water to compressible

Single layer shallow water

In order to simplify the code, the program treats all models as if they operate over the same number of dimensions. In the single layer case, this means that we store the 1d (2d) shallow water variables on a 2d (3d) grid. All values in the extra dimension should be the same, however over the course of a timestep small variations may introduced. When converting the variables back to the compressible grids, we must therefore average the shallow water quantities in the vertical direction to smooth out these small variations. As shall be shown in the examples below, care must be taken when performing this operation. In `AMReX`, each grid may be divided up into multiple patches which may then be assigned to different parallel processes. Averaging in the vertical direction may therefore require averaging variables across multiple parallel processes.

Multilayer shallow water

In the multilayer case, we will need to interpolate the data in the vertical direction in order to represent it on the compressible grid in addition to converting it from the shallow water to the compressible variables. Each of the shallow water layers are defined as surfaces of equal pressure, with $p = 0$ at the surface. Each layer is at constant density, which is defined in the initial data of the problem.

The first step of the shallow water to compressible conversion is to obtain the height of each layer and reclaim the 3-velocities from the conserved variables. The height is calculated from the geopotential as

$$h(\Phi) = \frac{2}{1 - e^{-2\Phi}} - R, \tag{5.29}$$

where R is the radius of the star at its crust at the base of the ocean. The next step is to interpolate ρ, p and v_i in the vertical direction to obtain those quantities at each of the layers of the compressible grid. The height of the compressible layer on which the state is to be determined, h_{C}, is calculated. For each (x, y) grid point, the heights of the shallow water layers at that point are calculated, starting from the top layer, until one is reached that is below h_{C}. Calling this layer ℓ, we have now found the two shallow water layers above and below the compressible layer at this coarse grid point: ℓ and $\ell - 1$. The state at h_{C} will then be given by

$$\boldsymbol{q}_{\mathrm{C}} = (1 - \alpha)\boldsymbol{q}_{\mathrm{C}}(\boldsymbol{q}_{\mathrm{SW},\ell-1}) + \alpha\boldsymbol{q}_{\mathrm{C}}(\boldsymbol{q}_{\mathrm{SW},\ell}), \tag{5.30}$$

where α is the 'layer fraction', representing where the compressible layer is with respect to its neighbouring shallow water layers and is given by

$$\alpha = \frac{h_{\mathrm{C}} - h_{\ell-1}}{h_{\ell} - h_{\ell-1}}. \tag{5.31}$$

Once we have interpolated the primitive variables to the compressible grid, the conserved compressible variables are calculated. At the vertical boundaries, it is necessary to enforce hydrostatic equilibrium at this point to ensure that there is a pressure gradient; as discussed in Section 1.5.2.8, if the standard outflow conditions are used, this can lead to unphysical mass transfer and instabilities.

Note that we have chosen to interpolate the primitive variables rather than the conserved variables as it was found to be more stable. It was found that the conserved variables can vary with height in a significantly nonlinear way, such that their linear interpolation can lead to large inaccuracies.

5.4.2 Compressible to shallow water

Single layer shallow water

When calculating the shallow water variables in the single layer case, we use the compressible variables at the base of the ocean. As in the shallow water to compressible case, this is complicated by the fact that there may be multiple patches in the vertical direction. As (for the sake of simplicity) we treat the shallow water model as a fully 2d (3d) model (rather than the 1d (2d) model it represents physically), grid cells may have to look at the cells in a patch 'below' in order to find the state at the base of the ocean (rather than at the base of that patch). To do this, we create a new grid containing the values of the compressible grid at the base of the ocean and synchronise across patches on different processes before beginning the compressible to shallow water variable conversion.

Multilayer shallow water

In the multilayer case, the first step of the compressible to shallow water conversion is to perform the conserved to primitive calculation for the compressible variables in order to obtain p, ρ and v_i. The next step is to use the pressure to calculate the heights of the shallow water layers at each (x, y) point and find where they intersect the compressible grid. This can be done using simple linear interpolation or by using the equation for hydrostatic equilibrium to more accurately model the pressure profile between layers. Once the heights have been obtained, the horizontal 3-velocities v_i are calculated on the shallow water grid using interpolation in the vertical direction.

Finally, the shallow water variables are calculated on the shallow water grid. The geopotential is calculated from the layer heights by inverting (5.29); from this and the 3-velocities the conserved shallow water variables are obtained.

5.5 Numerical implementation

The multiscale model was implemented by extending `CASTRO` and `AMReX`. The numerical implementation of the shallow water equations was demonstrated above in Section 4.3. Here, we shall first demonstrate the compressible solver for Newtonian and relativistic systems, then present tests demonstrating the performance and current limitations of the multiscale solver.

5.5.1 Compressible solver

The original `CASTRO` code was designed for studying Newtonian compressible hydrodynamics. We have extended it to special relativistic hydrodynamics. Its performance shall be demonstrated in the following tests.

As for the shallow water solver, the spatial integration is carried out using the finite volume method with the MC flux slope limiter and HLLE solver, and the time integration is carried out using the method of lines with RK3. The compressible solver solves the standard special relativistic hydrodynamics equations in conservative form with gravitational source terms. The conserved state vector, flux and sources are given by

$$\boldsymbol{q} = \begin{pmatrix} \rho W \\ \rho h W^2 v_j \\ \rho h W^2 - p - \rho W \end{pmatrix}, \qquad \boldsymbol{f}^{(i)} = \begin{pmatrix} \rho W v_i \\ \rho h W^2 v_j v_i + \delta^i_j p \\ (\rho h W^2 - \rho W) v_i \end{pmatrix}, \qquad \boldsymbol{s} = \begin{pmatrix} -\rho W g \\ 0 \\ -\rho W v_z g \end{pmatrix}, \tag{5.32}$$

where v_j are the components of the primitive velocity and v_z is the component in the positive vertical direction.

5.5.1.1 Sod shock tube

The Sod shock tube (Sod, 1978) is a standard test used to demonstrate hydrodynamics codes' performance when resolving shock fronts. It is also one of the few problems in special relativistic hydrodynamics for which an exact solution is known. It takes the form of a Riemann problem: the initial data consists of a left and right state, separated by a discontinuity. We shall use the initial data given in Table 5.1 in both our Newtonian and special relativistic tests.

	ρ	v^x	v^y	p
left	1	0	0	1
right	0.125	0	0	0.1

TABLE 5.1: Initial conditions for the Sod shock tube (Sod, 1978). The discontinuity is located at $x = 0$ and the adiabatic index is $\gamma = 1.4$.

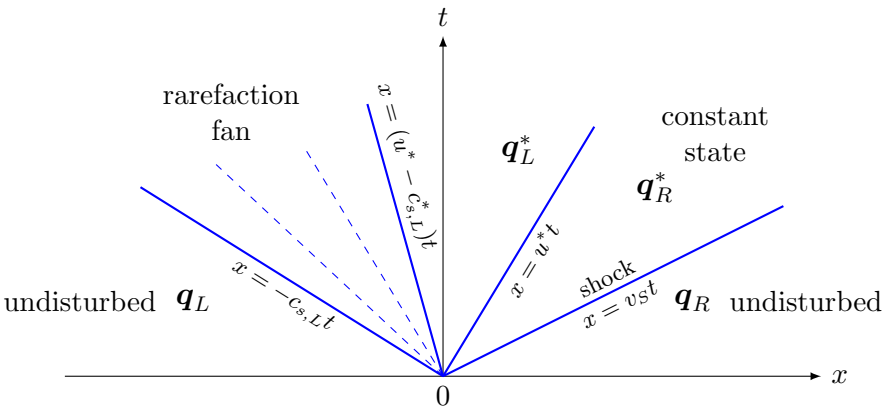

FIGURE 5.4: Wave structure of the shock tube problem with initial states $\boldsymbol{q}_L$ and $\boldsymbol{q}_R$, separated by a discontinuity at $x = 0$. At time $t > 0$, the solution consists of a left-going rarefaction, two intermediate 'star' states of constant pressure and velocity separated by a contact discontinuity, and a right-going shock wave with speed v_S.

Newtonian shock tube

The exact solution of the shock tube is given by Sod (1978); Toro (1999). The structure of the solution can be seen in Figure 5.4: it consists of a left-going rarefaction, two intermediate star states of constant pressure and velocity separated by a contact discontinuity, and a right-going shock.

The results of the test, performed using the original Newtonian `CASTRO` code, can be seen in Figure 5.5. The exact solution was calculated using the method given by Toro (1999). The numerical solution matches the exact solution very well for the high resolution run, with very little smearing at the discontinuities. The low resolution run is clearly more diffusive at the discontinuities, as would be expected for the finite volume methods used. There is also a small overshoot at the tail of the rarefaction. Comparing the low and high resolution runs to the exact solution, we can see that the results are converging to the correct solution.

Relativistic shock tube

The shock tube test was repeated for the relativistic compressible equations. The exact solution was calculated using the `R3D2` code. The results can be seen in Figure 5.6. It can be seen that the discontinuities are resolved fairly well for the high resolution run,

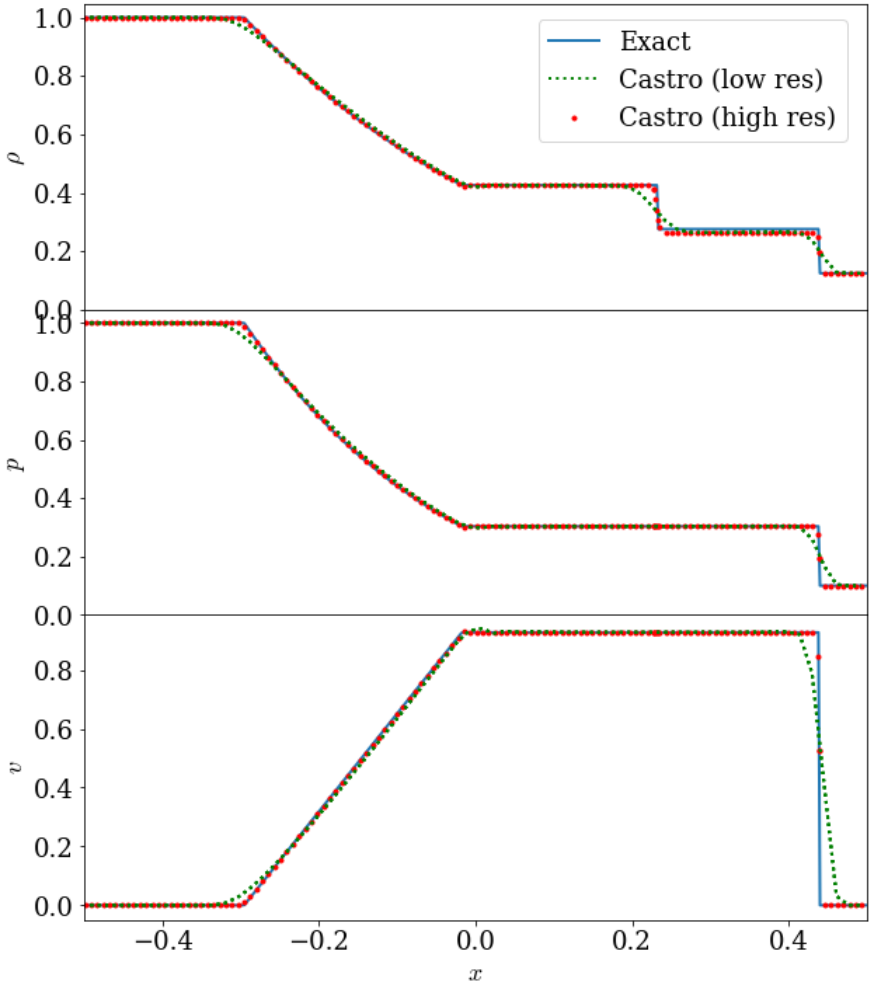

FIGURE 5.5: Results of the Newtonian shock tube test using the initial conditions in Table 5.1 at time $t = 0.25$. The simulation was carried out on the domain $x = [-0.5, 0.5]$, $y = [0, 0.03125]$, with 64 grid points in the x-direction and 4 in the y-direction for the low resolution run, and 512×16 for the high resolution run, making it effectively 1d. Plotted are the density ρ, the pressure p and the velocity in the x-direction v.

however there is a small amount of smearing, especially at the contact discontinuity. This contrasts with the Newtonian simulation in Figure 5.5, which was able to capture the contact discontinuity more clearly, despite being lower resolution than the high resolution relativistic run. Otherwise, the solution matches up very well with the exact solution. Comparing the low resolution and high resolution runs, we can see that the results are converging to the exact solution.

The additional error in the relativistic simulation is likely to be a result of the signal speeds used in the HLL solver. For the Newtonian case, we estimate these from the speed of sound in the fluid. In the relativistic case, we instead set these to the speed of light. This is overly conservative, and as this test shows has made our model particularly diffusive at the contact discontinuity. This could be improved by using a better estimate

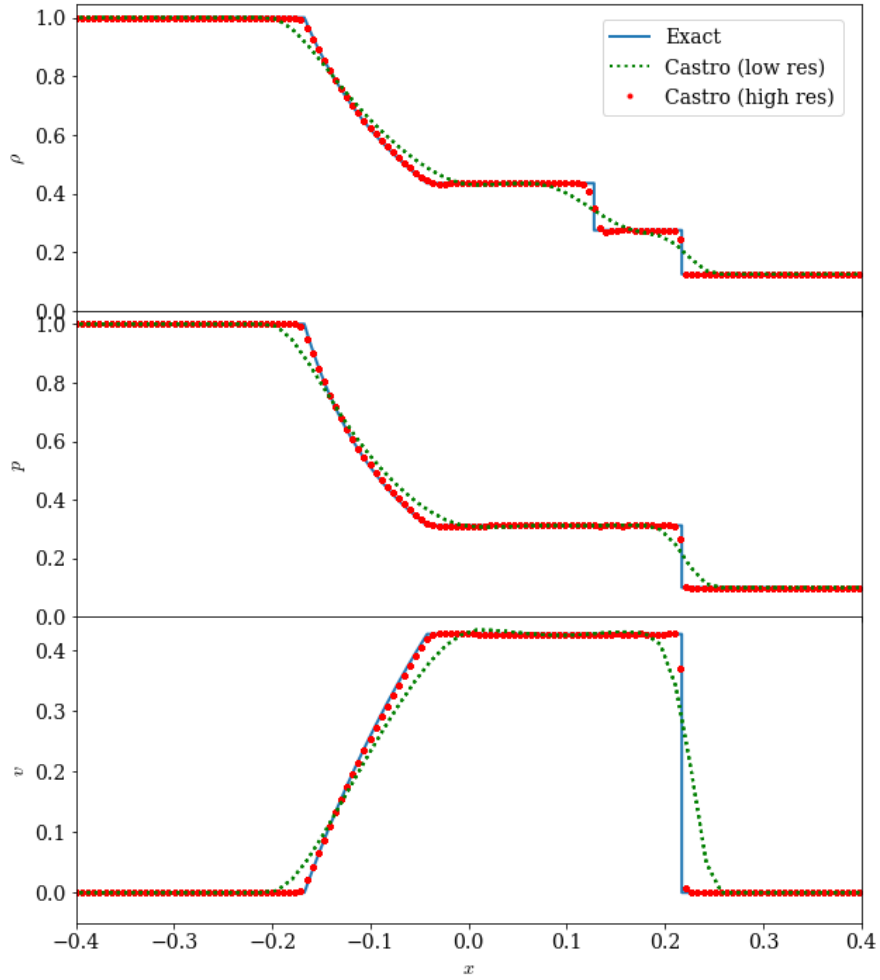

FIGURE 5.6: Results of the relativistic shock tube test using the initial conditions in Table 5.1 at time $t = 0.3$. The simulation was carried out on the domain $x = [-0.5, 0.5]$, $y = [0, 0.03125]$, with 64 grid points in the x-direction and 4 in the y-direction for the low resolution run and 512×16 for the high resolution run, making it effectively 1d. Plotted are the primitive density ρ, the pressure p and the 3-velocity in the x-direction v. The exact solution was calculated using the `R3D2` code.

of the signal speeds, or by implementing the HLLC solver described in Section 3.2.2, which would allow the contact discontinuities to be better resolved.

5.5.1.2 Blast wave

The blast wave is a Riemann problem with a large initial pressure jump. Like the Sod shock tube, exact solutions exist for both the Newtonian and special relativistic cases (Centrella and Wilson, 1984).

We shall use the initial data of Winkler and Norman (1986); Martí and Müller (2003) given in Table 5.2.

	ρ	v^x	v^y	v^z	p
left	1	0	0	0	1000
right	1	0	0	0	0.01

TABLE 5.2: Initial conditions for the blast wave (Winkler and Norman, 1986; Martí and Müller, 2003). The discontinuity is located at $x = 0$ and the adiabatic index is $\gamma = 5/3$.

Newtonian blast wave

The results of the Newtonian blast wave simulation can be seen in Figure 5.7. The approximate solution matches the exact solution very well, with the discontinuities very sharply resolved with little smearing for the high resolution run. There is a very small overshoot in the velocity at the tail of the rarefaction and in the density on the left side of the blast zone. The low resolution run is a lot more diffusive, with some oscillations in the velocity behind the tail of the rarefaction. The density in the blast zone is also underestimated, with the profile of the blast zone significantly rounded – the lower resolution used in this run is not sufficient to resolve this feature. Comparing the low and high resolution runs, it is clear that the solution is converging to the correct solution as the resolution increases.

Relativistic blast wave

The results from the relativistic blast wave test can be seen in Figure 5.8. In order to demonstrate the improvements gained from increased resolution, the test was run without AMR on a coarse grid with 128 points in the x direction, and run with AMR and 4 levels of mesh refinement. As can be seen, in the lower resolution run without AMR, the blast wave has been significantly smeared due to clipping from the slope limiters. There is also some smearing at the head and tail of the rarefaction. As would be expected, this clipping has been significantly reduced in the higher resolution run with AMR.

There are significant oscillations in the variables either side of the waves in the central region. These are clearly unphysical, as the waves produced by the initial discontinuity should not yet have reached these regions and so they should have remained in their initial state. The oscillations appear to be greatest to the right of the shock wave, however they also appear to the left of the rarefaction. The inset plot in the density panel shows that the magnitude of these oscillations is actually greater in the higher resolution run. In the lower resolution case, these oscillations have been somewhat smeared out, slowing their growth. In this test, a CFL number of 0.5 was used. Reducing this to 0.2 and increasing the number of points in the y-direction so that they were the same as in the x-direction did not appear to have any affect on the appearance of these oscillations.

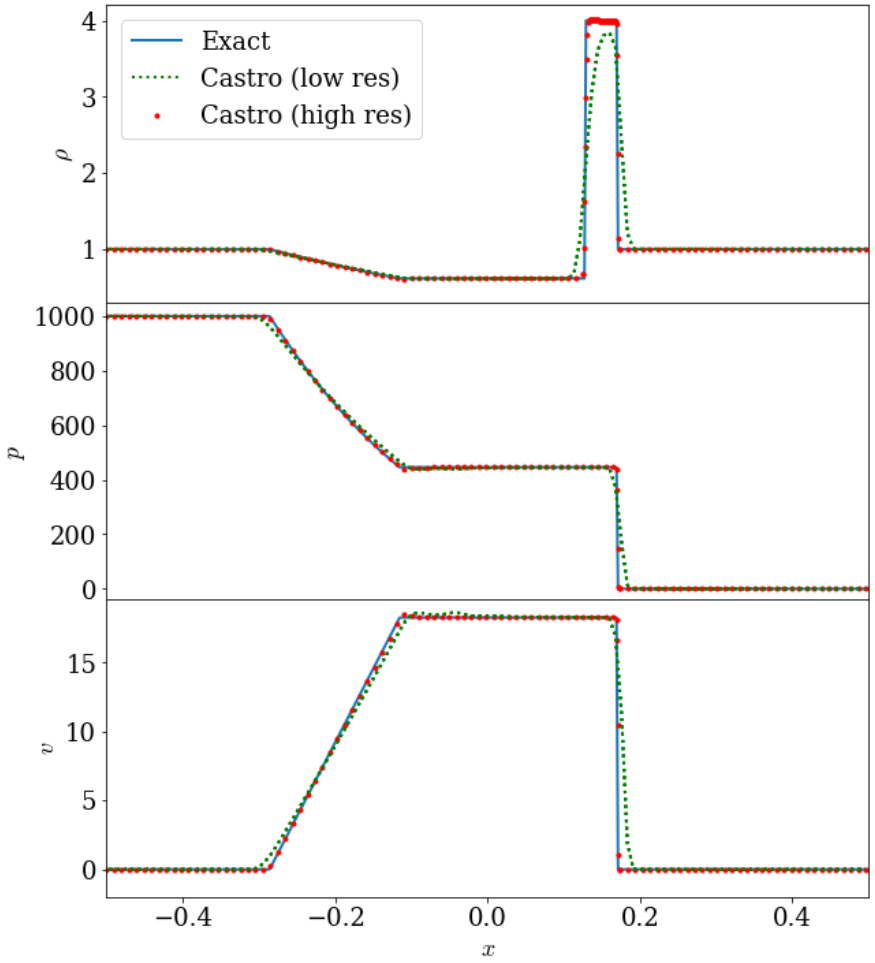

FIGURE 5.7: Results of the Newtonian blast wave test using the initial conditions in Table 5.2 at time $t = 0.007$. The simulation was carried out on the domain $x = [-0.5, 0.5]$, $y = [0, 0.03125]$, with 128 grid points in the x-direction and 4 in the y-direction for the low resolution test and 1024×32 for the high resolution test, making it effectively 1d. Plotted are the density ρ, the pressure p and the velocity in the x-direction v.

Looking at the plot, it appears that the density in the blast wave has been significantly underestimated. This is especially true for the lower resolution run. In the exact solution, the density in the blast wave is 10.4, whereas in the simulation the highest value of the density in the blast zone is 3.1. However, on closer inspection, this may not be as bad as it seems. As can be seen in the exact solution, the blast zone is extremely narrow. The two points in the higher resolution run nearest the blast zone lie either side of the peak. Looking at the peak density in consecutive timesteps, it can be seen to increase and decrease as these points move closer and further away from the location of the peak. The significant underestimation of the density may therefore be largely a result of lack of resolution of the blast zone, rather than being entirely due to clipping from the slope limiters.

The rarefaction wave is captured reasonably accurately, however the blast wave has

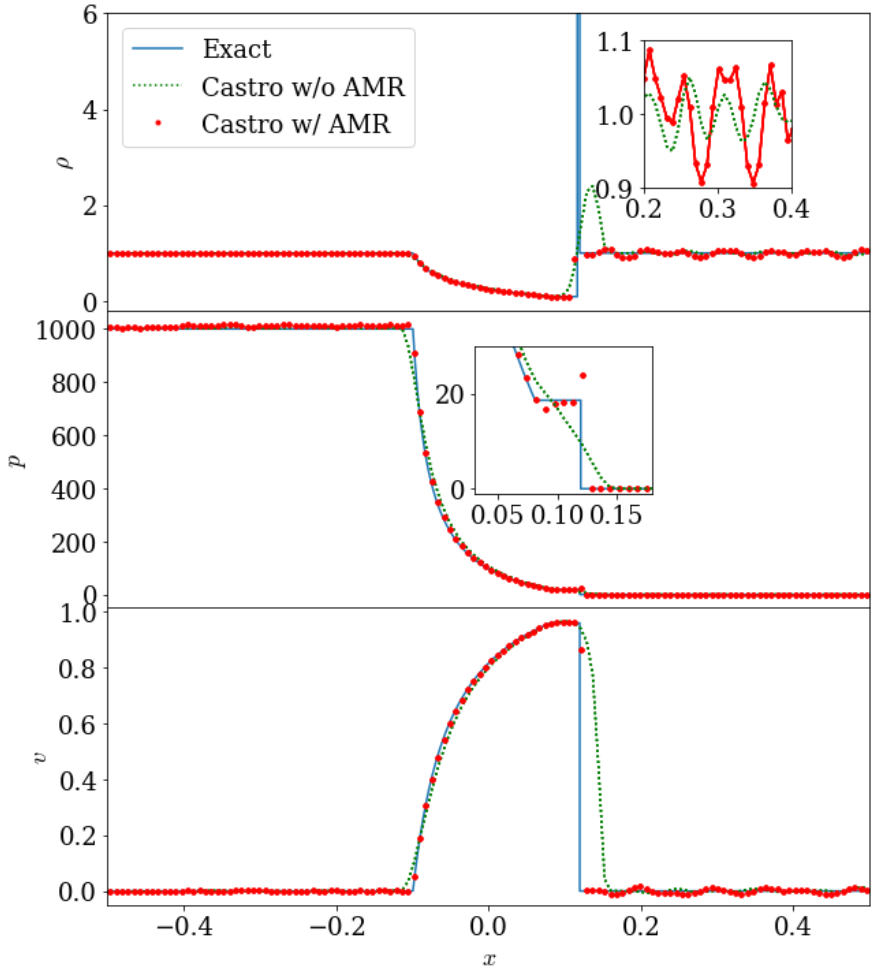

FIGURE 5.8: Results of the relativistic blast wave test using the initial conditions in Table 5.2 at time $t = 0.12$. The simulation was carried out on the domain $x = [-0.5, 0.5]$, $y = [0, 0.03125]$, with 128 grid points in the x-direction and 4 in the y-direction on the coarsest grid, making it effectively 1d. The results of two runs are shown: the dotted green line shows the results from a run without AMR, the red dots the results from a run with 4 levels of mesh refinement. Plotted are the primitive density ρ, the pressure p and the 3-velocity in the x-direction v. The exact solution was calculated using the `R3D2` code. The inset plots show close ups of the oscillations in the density and of the pressure discontinuity around the shock front.

become smeared in all variables. The lack of accuracy in this example is perhaps unsurprising given that the blast wave is highly relativistic, with its normal 3-velocity approaching the speed of light. This test therefore demonstrates the limitations of the code and shows that results are unlikely to be reliable for highly relativistic systems.

5.5.2 Multiscale solver

In our implementation of the multiscale solver coupling the compressible and shallow water models, we begin with the simplest possible hydrodynamics model (Newtonian

incompressible hydrodynamics) and gradually build up the complexity to a 3d relativistic compressible system.

When implementing the multiscale system, the x-axis was defined to be the vertical direction in order to make it easier to write code that could run in both 2d and 3d. This is reflected in the plots below, where the vertical axis is labelled as the x-axis. In all of the multiscale tests below, the shallow water model was used for the coarsest level, and the compressible model for all finer levels (unless otherwise specified).

5.5.2.1 Incompressible + shallow water

The simplest possible version of the multiscale solver is that of 2d incompressible hydrodynamics coupled to the 1d shallow water equations at large scales. In this system, the shallow water equations track the surface of the fluid. When reconstructing the 2d system, the pressure is calculated using the equation for hydrostatic equilibrium. If the height of the fluid at the surface is h, the pressure at a height of x above the bottom boundary is

$$p(x) = \frac{g}{2}(h - x)^2 . \tag{5.33}$$

Given the pressure gradient in the vertical direction, the horizontal motion of the fluid in the 2d incompressible system varies with depth, with fluid at the bottom moving faster than that at the top. The multiscale approximation assumes that any variation in the vertical direction (other than pressure) is negligible, so in order for the solution to be accurate, any 2d behaviour such as this must be covered by the incompressible grids. This can be achieved by suitable tagging criteria which check for a non-zero velocity or gradient of a non-pressure variable in the vertical direction.

Newtonian 1d

Figure 5.9 shows the fluid pressure at time $t = 0.1$ for the Newtonian 1d dam breaking simulation with three levels of mesh refinement. In the shock front near the top of the plot there is a small kink – this is a numerical error resulting from the AMR patches. When the data is transferred between the incompressible and shallow water models, vertical averaging and flooring operations are carried out to translate the data from the 2d (3d) incompressible model to the 1d (2d) shallow water one. However, for this first incompressible multiscale implementation, this operation was only carried out locally for each patch. This has led to the error in the shock front location, which grows as the simulation progresses, and highlights the importance of doing this calculating globally, across all patches in the vertical direction. Unfortunately, such an operation is much more computationally expensive as it requires communication between multiple processes and cannot be parallelised. The improved calculation of the vertical average

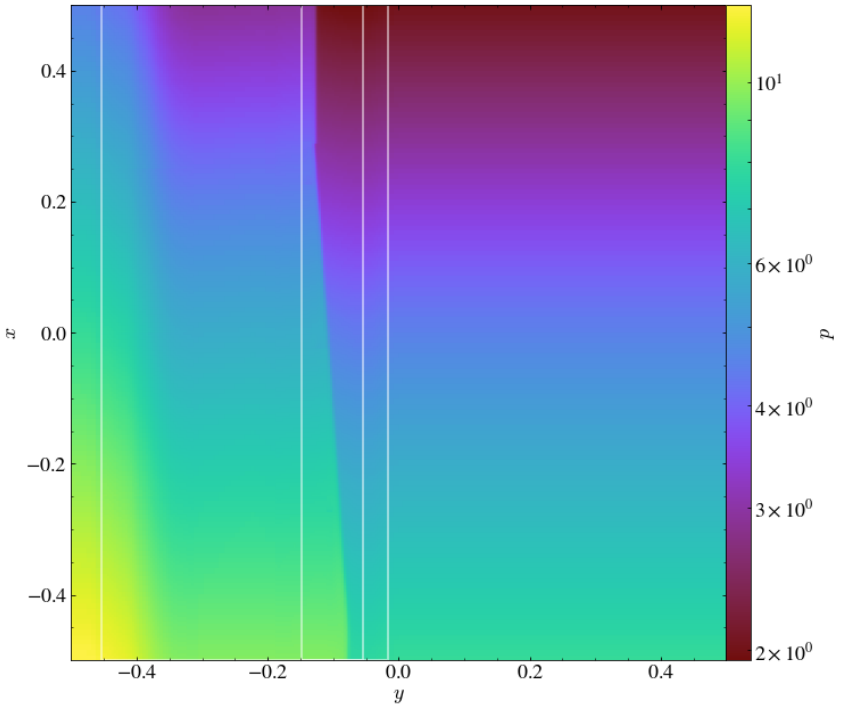

FIGURE 5.9: Newtonian multiscale 1d dam breaking at time $t = 0.1$ with $h_l = 5$ and $h_r = 4$. The plot shows the fluid pressure, with the outlines of the finer grids overlaid. The shock front has remained sharp, however there is a small kink near the top of the plot due to a numerical error resulting from the AMR patches.

across multiple patches was implemented in the next multiscale model presented in the following section. Looking ahead to the results of the special relativistic test in Figure 5.12, these numerical artefacts from the patches have disappeared – the shock front is now smooth.

Newtonian 2d

In Figure 5.10 we plot the results of the 2d radial dam breaking test. This test was performed using the Newtonian incompressible multiscale code with three levels of refinement (one shallow water grid and two incompressible grids). In the plot, this is compared against the same test run using the incompressible code only and the 1d shallow water code with geometric source terms. As expected, the shallow water model produces the sharpest profile – the profile of the central peak is noticeably more rounded for the incompressible and multiscale models. There is some mismatch in the location of the shock front and height of the shocked fluid between the shallow water model and the other two, with the shock front in the shallow water model having a slightly lower shock speed. This indicates that there is some small error in our model e.g. when converting between heights and pressures.

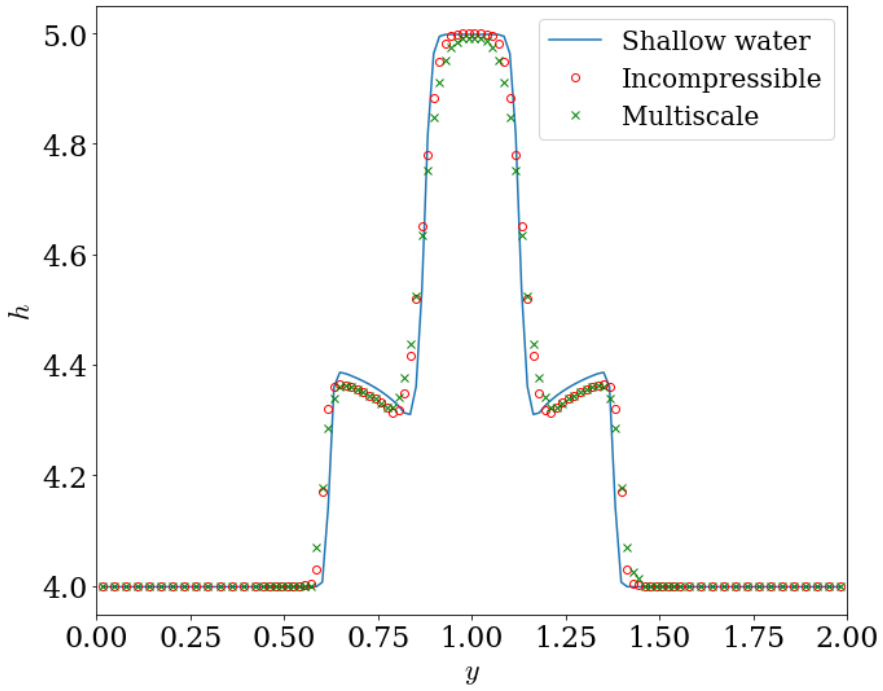

FIGURE 5.10: Newtonian multiscale 2d dam breaking at time $t = 0.07$ with $h_{\text{in}} = 5$ and $h_{\text{out}} = 4$. The shallow water and incompressible grids had resolution 256×256. The multiscale model had 3 levels of mesh refinement, with the finest grid at the same resolution as the shallow water and incompressible grids.

Relativistic 1d

The results of the relativistic incompressible multiscale 1d dam breaking test can be seen in Figure 5.11. The shock has been sharply resolved with little smearing. There is some smearing at the head and tail of the rarefaction, but the profile is otherwise in line with what we expect. Figure 5.12 shows the pressure profile for the entire domain used in this test. Comparing to Figure 5.9, we can see that the shock front is much straighter, thanks to averaging the state across all patches in the vertical direction rather than across individual patches.

Relativistic 2d

Figure 5.13 shows the results of the relativistic incompressible radial dam breaking test; plotted is the fluid height at time $t = 0.6$, with the outlines of the finer grids overlaid. The test was run with three levels of mesh refinement consisting of a single coarse shallow water grid and two finer incompressible grids. The grids remained fixed over the course of the simulation.

In Figure 5.14, this is plotted against the results produced using the incompressible model only and the shallow water model only. As to be expected, the multiscale and

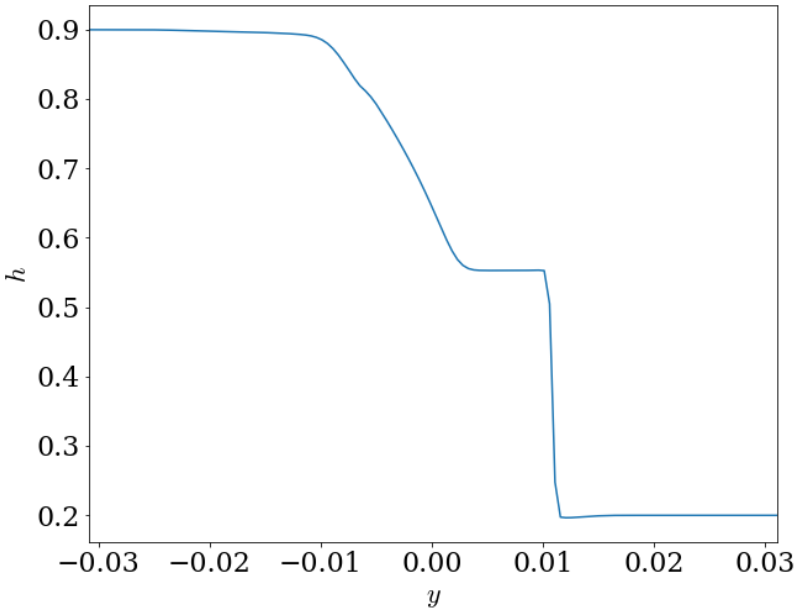

FIGURE 5.11: Results of the relativistic incompressible multiscale 1d dam breaking test with 3 levels of mesh refinement at time $t = 0.048$ with $h_l = 0.9$ and $h_r = 0.2$. The coarse grid resolution was 128×128. The plot shows the fluid height.

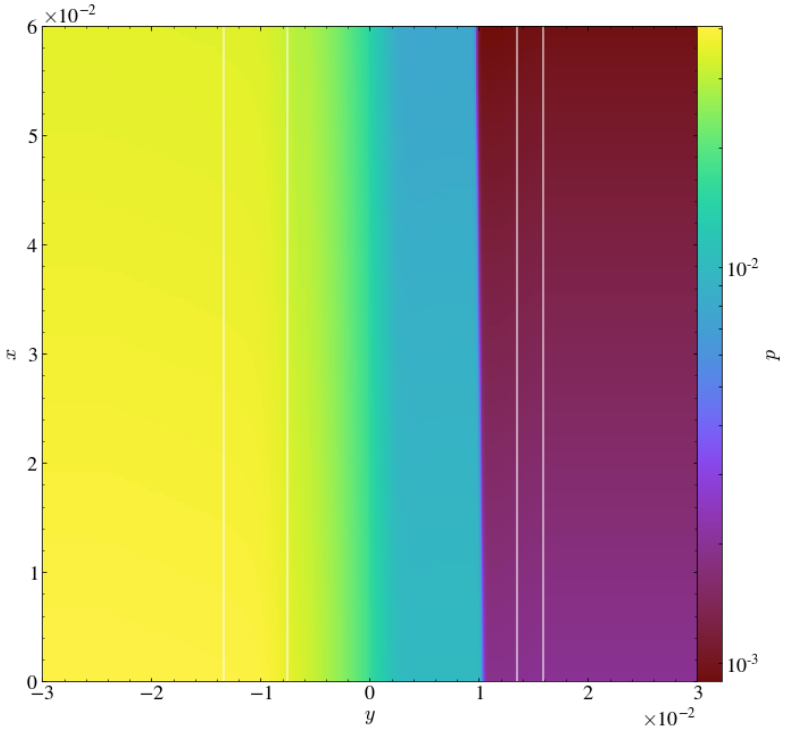

FIGURE 5.12: Results of the relativistic incompressible multiscale 1d dam breaking test with 3 levels of mesh refinement at time $t = 0.048$ with $h_l = 0.9$ and $h_r = 0.2$. The coarse grid resolution was 128×128. The plot shows the fluid pressure, with the outlines of the finer grids overlaid.

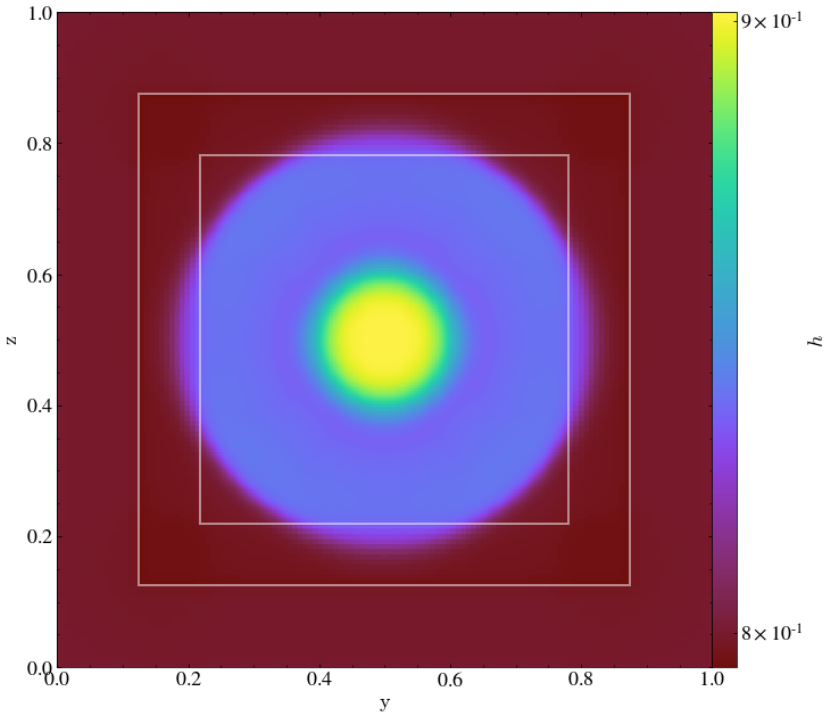

FIGURE 5.13: Results of the relativistic incompressible multiscale radial dam breaking test with 3 levels of mesh refinement after 205 (coarse grid) timesteps at time $t = 0.6$ with an initial dam radius of $r = 0.2$, $h_{\text{in}} = 0.9$, $h_{\text{out}} = 0.8$ and $g = 0.1$. The coarse grid resolution was $8 \times 64 \times 64$. The plot shows the fluid height, with the outlines of the finer incompressible levels overlaid.

incompressible models agree very well, especially in the central region where the multiscale model is using the incompressible equations and has the same resolution on its finest grid as the incompressible model. The outer shock front is more smeared for the multiscale model than it is for the incompressible run: this reflects the fact that the multiscale model is using a coarser grid here.

There is a small error in the height at the interface of the shallow water and incompressible models. This error formed at the start of the simulation, however did not significantly grow in magnitude. It suggests that there is some mismatch in the conversion of data between the two models.

The shallow water model is considerably sharper than the other two models. The locations of the shocks calculated using the shallow water solver do not quite agree with the other two. This is to be expected: physically they are different models, so should predict different wavespeeds. All that we require for the purposes of our multiscale model are that the models agree 'well enough' in the matching region where we convert between the two that we can capture the features we are interested in to some desired level of accuracy.

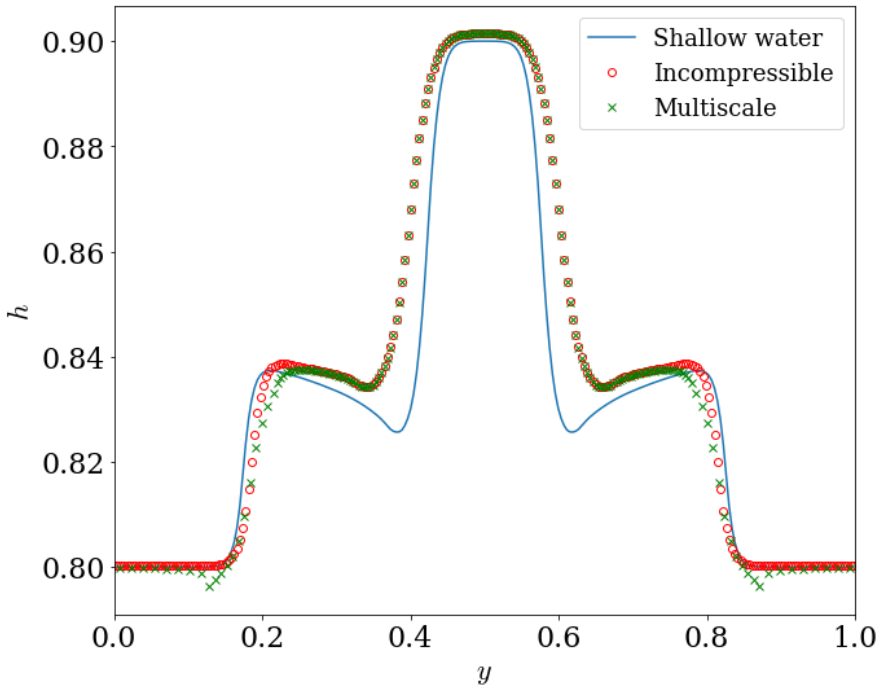

FIGURE 5.14: Comparison of the relativistic radial dam breaking test for the multiscale, incompressible and shallow water models. The finest level of the multiscale run had the same resolution as the incompressible and shallow water runs. The plot shows the fluid height.

5.5.2.2 Compressible + shallow water

Adding compressibility to the system makes it significantly more challenging as it is much harder to reconstruct the pressure and density profile of the fluid from the shallow water variables. However, we can simplify this task greatly by assuming the fluid to be adiabatic at the shallow water/compressible interface. This means that the thermodynamic variables can be expressed as functions of the height only. An adiabatic fluid is defined by

$$p^{1-\gamma}T^{\gamma} = \text{constant}. \tag{5.34}$$

Using the ideal gas equation of state, this means we can write the density as a function of the pressure only

$$\rho(p) = Kp^{\gamma}, \tag{5.35}$$

where K is some constant that must be given. Using the equation for hydrostatic equilibrium

$$\frac{\mathrm{d}p}{\mathrm{d}x} = -g\rho(x), \tag{5.36}$$

we can now integrate this to get the pressure as a function of depth. Assuming that the pressure vanishes at the surface, so $p(x = h) = 0$, we then get

$$p(x;h) = \left(\frac{\gamma-1}{\gamma}gK(h-x)\right)^{\frac{\gamma}{\gamma-1}}. \tag{5.37}$$

Inverting this, we obtain the expression for the fluid height as a function of pressure at some depth x

$$h(p;x) = x + \frac{1}{gK}\frac{\gamma}{\gamma-1}p(x)^{\frac{\gamma-1}{\gamma}}. \tag{5.38}$$

Similar functions can be obtained for the density as a function of depth and height as a function of density.

Note that the assumption of adiabaticity is only required at the shallow water/compressible interface – it is entirely possible for the fluid to be non-adiabatic in the region(s) covered by the compressible patches away from this interface. As shall be seen in the tests in this section (and shall be discussed further at the end of the chapter), the assumption of adiabaticity is clearly insufficiently accurate for all but the smoothest of features, leading to the development of instabilities at the model interface. It has been used here due to its simplicity, however in the future we hope to use a more accurate model or develop better ways of matching the data at the model interface.

When converting from the compressible to the shallow water levels, we are taking 2d (3d) data and converting it to 1d (2d) data. In order to do this, it is necessary to take some kind of average of the data in the vertical direction. As found in the previous section, this is complicated by the fact that there may be multiple patches in the vertical direction, so in order to take the average across all data in the vertical direction, there needs to be some communication of data between patches. As this will require syncing of the parallel threads and/or processes, this will slow down the performance of the code.

In the following sections, we shall present tests performed on the multiscale model made up of the relativistic compressible and relativistic shallow water models.

Dam breaking

The dam breaking tests of Section 4.3.1 were repeated with the multiscale solver. The tagging criteria were the magnitude of the pressure gradient and the primitive speed, with cells exceeding given critical values tagged for refinement. Figure 5.15 shows the primitive y-velocity for the 1d dam breaking test for a run with 3 levels of mesh refinement at time $t = 0.05$. It can be seen that the shock front has been properly captured, remaining relatively sharp as the simulation has evolved.

The state appears to transition smoothly across the level interfaces, with no visible artefacts in any of the variables. Regridding has ensured that the region of the simulation where the state evolves remains in the section covered by the finest grid, well away from the level boundaries. This test therefore demonstrates that the multiscale model is able to maintain the static state of the unperturbed fluid without the development of errors at the level interfaces.

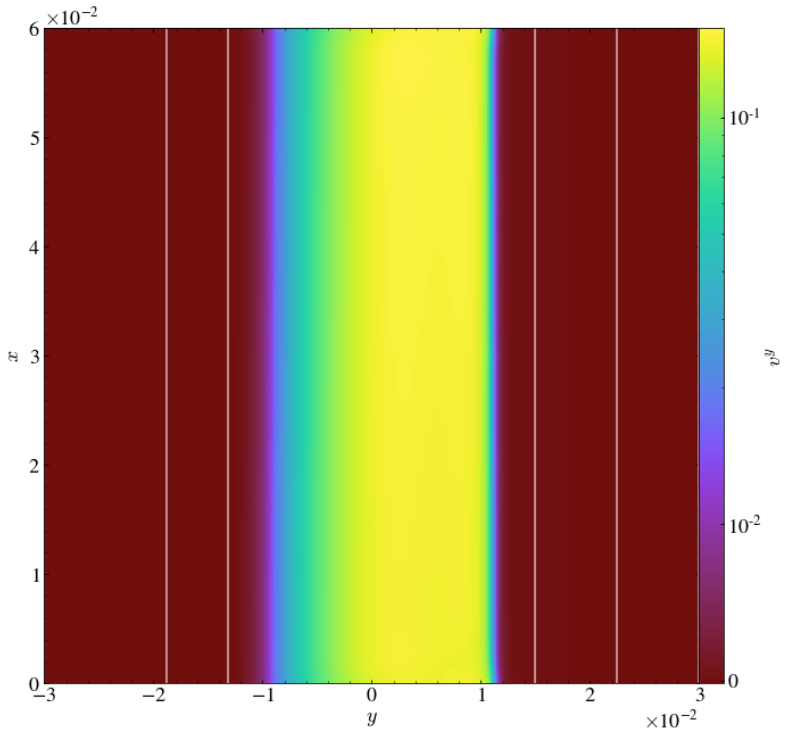

FIGURE 5.15: Results of the compressible multiscale 1d dam breaking test with 3 levels of mesh refinement at time $t = 0.05$. The coarse grid resolution was 64×64 and the initial data was $h_l = 0.9$, $h_r = 0.4$. The plot shows the primitive y-velocity, with the outlines of the finer levels overlaid. It can be seen that the shock front has remained sharp.

Figure 5.16 shows the height for the circular dam breaking test with three levels of mesh refinement at time $t = 0.5$. The coarse grid resolution was $8 \times 64 \times 64$, and the grids remained fixed throughout the simulation.

The multiscale circular dam break solution was compared against similar compressible and shallow water simulations (see Figure 5.17). These simulations had the same resolution as the finest multiscale grid. It can be seen that the multiscale simulation lines up very well with the compressible simulation in the central part of the domain where it is using the compressible model. As was found in the incompressible test, there is a small error in the height at the shallow water/compressible model interface. This error first developed early on in the simulation, however did not significantly grow in magnitude. It most likely corresponds to a mismatch in the conversion of variables between the two models, suggesting that e.g. our assumption of adiabaticity is not valid for this system. This error could be reduced by employing some kind of numerical smoothing techniques at the interface or by improving the scheme converting variables between the compressible and shallow water models.

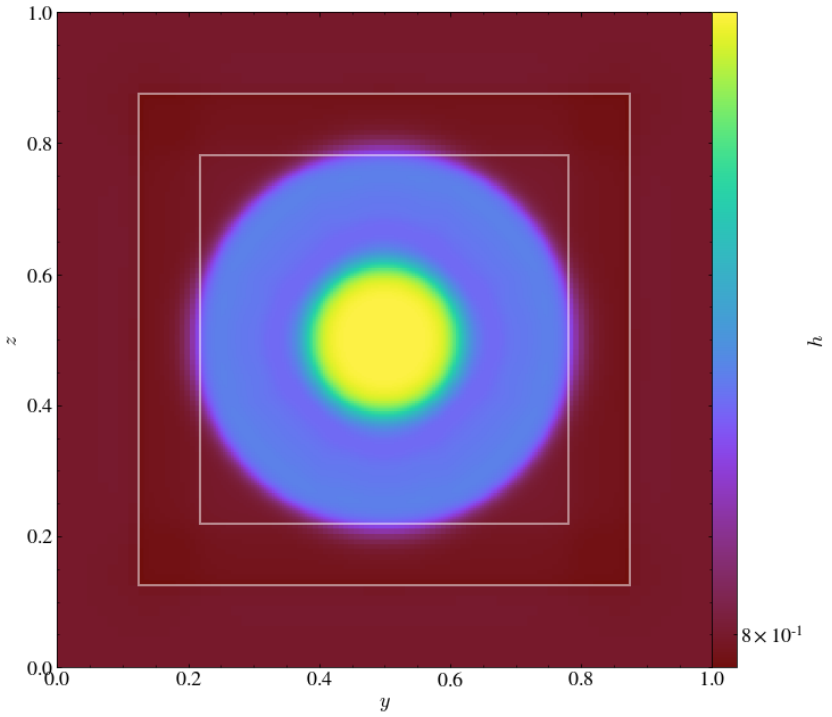

FIGURE 5.16: Results of the relativistic multiscale circular dam breaking test at time $t = 0.5$ with initial dam radius of $r = 0.2$, $h_{\text{in}} = 0.9$ and $h_{\text{out}} = 0.8$. 3 levels of mesh refinement were used, with the locations of the grids fixed throughout the simulation. The coarsest grid had resolution $8 \times 64 \times 64$. The plot shows the fluid height, with the outlines of the finer levels overlaid.

Fixed mesh refinement

In the dam breaking tests, we showed that the multiscale system produced results largely consistent with simulations produced using a single model. However, in these tests large features were intentionally kept away from the model interfaces. In this section, we shall investigate how well the multiscale model copes when such features pass through these interfaces.

Figure 5.18 shows the results of the previous 1d dam breaking test with fixed mesh refinement. This was performed using the compressible model only and using the multiscale model. Both tests had three levels of mesh refinement, with the coarsest grid having a resolution of 64×64. The finest grids were centred around the initial discontinuity at $x = 0$. The grid hierarchy was calculated at the start of the simulation, and remained unaltered after that.

It can be seen that for the compressible-only model, the shock fronts were able to pass through the level interfaces without producing any noticeable errors. However, this is clearly not the case for the multiscale model, where large errors have developed at the boundaries of both the shallow water/compressible interface and the finer compressible/-compressible interface. These errors grew over time, first beginning to develop when the

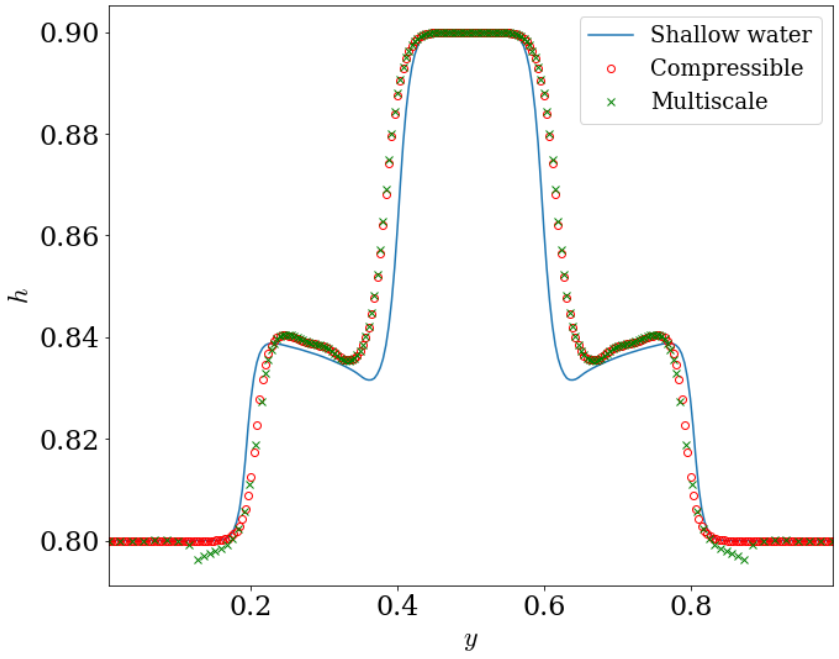

FIGURE 5.17: Height profile of the multiscale relativistic circular dam breaking simulation (green crosses) with corresponding compressible (red circles) and shallow water (blue line) simulations. The shallow water simulation produced the sharpest result, preserving the profile of the shock fronts with minimal smearing. The compressible case was more diffusive, with some smearing at the discontinuities. The multiscale results line up well with the compressible simulation in the central part of the regime where it is using the compressible model, however there are small errors at the shallow water/compressible interfaces.

shock fronts first crossed the shallow water/compressible interface. It is likely that the errors occurred as a result of reflections from the model interface. As the shock front passed through the interface, small errors in the conversion of quantities between the two models produced reflected waves. These waves interacted with the incoming waves, causing the errors to grow greater in magnitude. This is further demonstrated by the high resolution multiscale test, where again three levels of mesh refinement were used, but the coarsest grid had a resolution of 256×256. The instabilities in the central region are worse at this higher resolution: this is consistent with them being produced by the interaction of waves produced by numerical errors at the level interfaces.

Despite the development of significant instabilities in the central part of the domain, the shock fronts have been able to propagate through the interfaces between the two models, and their locations are not too dissimilar from the locations produced by the compressible-only model. The shock fronts themselves are fairly well resolved (especially for the higher resolution run), and there is no overshooting in the region preceding the shocks.

This test demonstrates that this implementation of the multiscale model is unable to

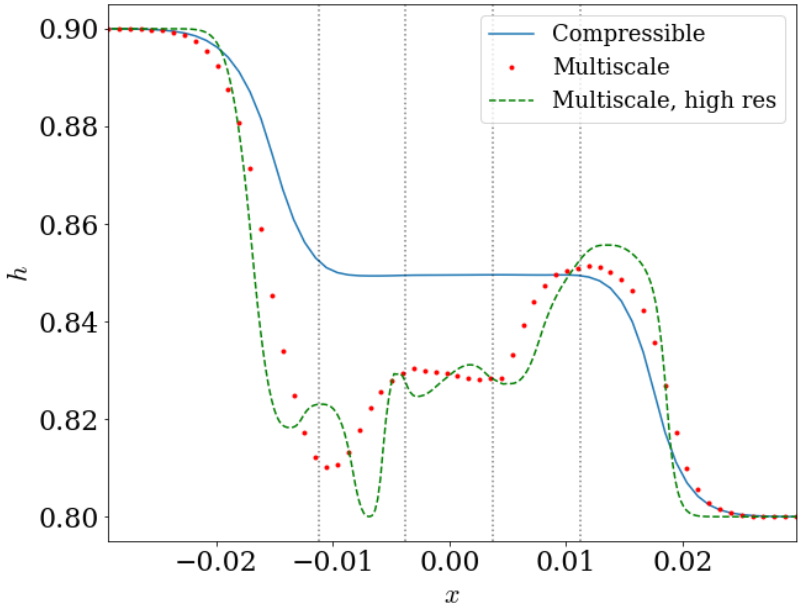

FIGURE 5.18: Results of the 1d dam breaking test with fixed mesh refinement performed for compressible only and multiscale with coarse grid resolution 64×64. Also shown is the result for a higher resolution multiscale run with coarse grid resolution 256×256. The plot shows the height at time $t = 0.09$, with the outlines of the finer levels overlaid. For the multiscale runs, the shock fronts have passed through the shallow water/compressible interfaces, and their locations line up with the location of the shock fronts in the compressible case. However, significant instabilities have developed in the shallow water region behind the shocks.

cope with features with large gradients (such as a shock front) passing through the interfaces between models without the development of instabilities in the finer grids. However, the shock fronts themselves were able to propagate through the interfaces largely unchanged, suggesting that the results of this test could be much improved if the development of instabilities at the model interfaces could be suppressed. This could be done by introducing some artificial smoothing in this region (e.g. similar to the slope limiting techniques used in MUSCL-type schemes to prevent the development of spurious oscillations described in Section 1.5.2.2), or by improving the physical model used to convert quantities between the compressible and shallow water models.

Figure 5.19 shows the results of a second test with fixed mesh refinement, however the initial data used here was much smoother – the height was initialised with a Gaussian profile. This smoother feature is much more like the sort of features we would expect our model to be able to represent when modelling neutron star oceans. We would wish sharper features to be resolved in finer detail, so would ensure that these are always covered by finer grids and more accurate physical models. The features that would pass through model boundaries would be larger scale, smoother features of the sort we look at in this test.

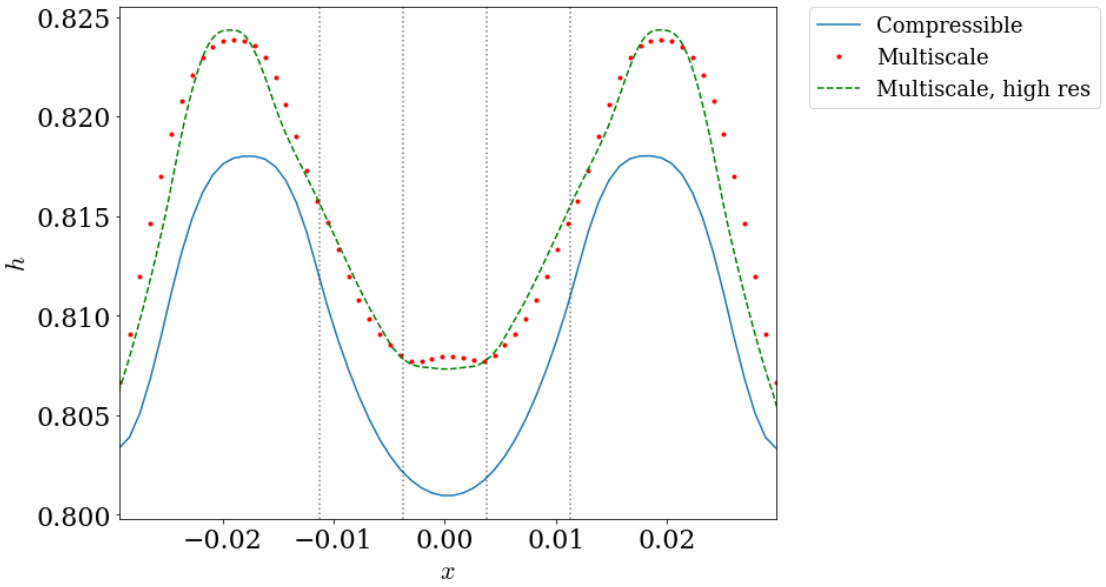

FIGURE 5.19: Fixed mesh refinement with an initially Gaussian height profile at time $t = 0.1$. The coarse grid resolution was 64×64 for the compressible and low resolution multiscale runs, and 256×256 for the high resolution multiscale run. Unlike for the dam breaking test, the smooth feature is able to pass through the level boundaries without the development of significant errors that destabilise the code. The finer level at the centre of the domain is outlined.

The height was initialised with a Gaussian profile at the centre of the domain with $\sigma = 0.1$ and peak height 10% greater than that of the undisturbed fluid. Again, three levels of mesh refinement were used, with the finer levels centred around the initial peak at the centre of the domain. The test was again performed for a compressible-only run and a multiscale run, both of which had a coarse grid resolution of 64×64. It was also repeated for a higher resolution multiscale run, with a coarse grid resolution of 256×256.

Unlike in the previous test, it can be seen that the results of this test were stable, qualitatively correct and converging. The fluid was able to propagate through the level interfaces without the development of instabilities in the central region. This is likely to be because the gradient in the state as it crossed the boundaries was much smaller in this test, so as the feature passed through, any reflected waves produced due to errors at the boundary would have been much smaller in magnitude. Consequently, these will have been suppressed before they could interact with incoming waves and grow to become instabilities. This is supported by the good agreement of the low and high resolution multiscale runs: even at the higher resolution, instabilities of the form seen in the previous test have not developed.

As we saw with the shock fronts in the dam breaking test, the locations of the fluid peaks produced by the multiscale model largely agree with those in the compressible model. However, there is a significant difference in the height of these peaks.

This test demonstrates that qualitatively the multiscale model can cope reasonably well with smooth features passing through model interfaces, however significant quantitative errors are produced in the state of the fluid in these features. More work is therefore required to minimise these errors in order for the multiscale scheme to be consistent with single model schemes.

Gresho vortex

Here we perform a variation of the Gresho vortex test described in Section 2.3.2.3. In the following tests, we set $R = 0.2$, $u_0 = 0.1$ and $p_0 = 0.01$.

The Gresho vortex test is one where traditional compressible codes tend to perform poorly compared to low Mach number codes. This test therefore demonstrates one of the advantages of including the low Mach number model in our multiscale scheme to model the physics at the smallest scales.

Unlike in the previous low Mach number tests, for these tests the solution will not be time independent. The fluid is initially assumed to be adiabatic and in hydrostatic equilibrium. Any pressure differences are therefore the result of differences in the height of the fluid – regions with a higher pressure at the base of the ocean will have a greater fluid height in the multiscale data. This height gradient means that over time, we expect the fluid in the vortex to flow outwards under the effects of gravity. However, the overall vortex structure should still be preserved, as should the total energy (the sum of the gravitational potential energy and the kinetic energy).

The height profile of the vortex at time $t = 1$ for a multiscale test with 2 levels of mesh refinement can be seen in Figure 5.20. As expected, the vortex structure has spread out so that the height minimum has moved from a radius of $r = 0.2$ from the centre of the vortex to $r \approx 0.3$. The symmetry of the vortex has been preserved reasonably well, though the square nature of the finer grid has distorted the profile on the outside of the vortex so that it is no longer circular. As seen in the circular dam breaking tests, there is a small error in the height at the model interface, however this has not grown significantly large in magnitude.

Figure 5.21 compares the results of the multiscale test to similar runs for shallow water-only and compressible-only models. All three tests used the same resolutions and grid hierarchies. Looking at the multiscale test, the errors at the model interface can be seen more clearly in this plot – as it appeared from the previous plot, these are relatively small in magnitude.

As seen in the Gaussian test, the locations of the extrema in the multiscale run agree well with those in the compressible-only run, however the magnitude of these extrema are significantly different. This can be somewhat explained by this plot by looking at

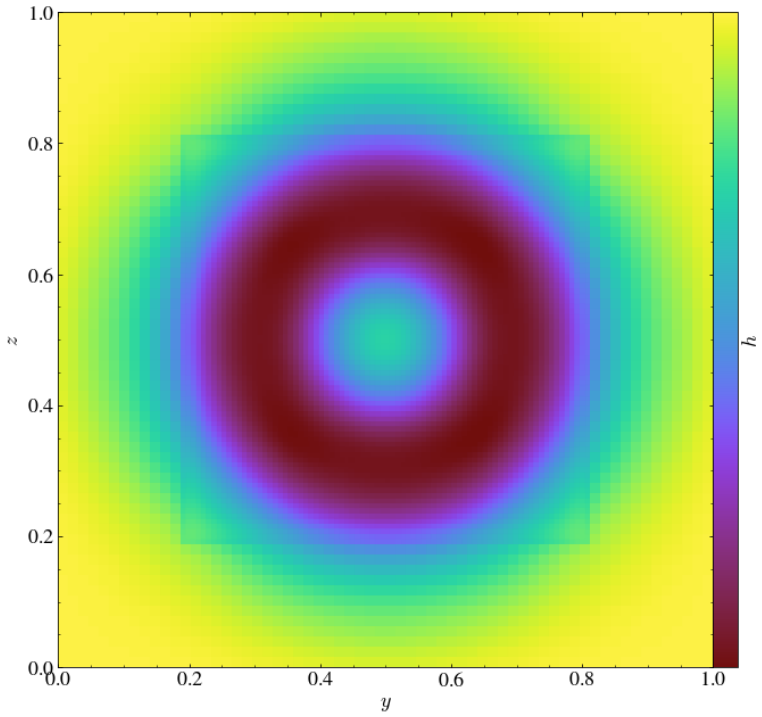

FIGURE 5.20: Results of the Gresho vortex test at $t = 1$ using two levels of mesh refinement. The coarsest grid has a resolution of $8 \times 64 \times 64$. The plot shows the fluid height on a slice perpendicular to the vertical axis. The central region covered by the finer compressible grid can be identified by the small height error at the boundary. In the initial data, the height was at a minimum $r = 0.2$ from the centre of the vortex – the vortex has spread out so that the height minimum is now at $r \approx 0.3$.

the outer parts of the domain. Here, the height profile produced by the multiscale model closely follows that produced by the shallow water model. This is to be expected, as the multiscale model is using the shallow water model in this part of the domain. This reduces the height of the fluid in this region, leading to a reduction in the height in the compressible region as well.

In Section 2.3.2.3, we found that the low Mach number solver was able to conserve the total kinetic energy in the Gresho vortex test very well. In Figure 5.22, we look at how well energy is conserved by the three different models in this test. As this version of the Gresho vortex test is not time independent, we would expect for the total energy (the sum of the gravitational potential energy and the kinetic energy) to be conserved, rather than just the kinetic energy. In order to calculate the total gravitational potential energy across the domain, for each cell in the horizontal plane we integrated the gravitational potential energy over the fluid height, then summed this over the entire grid:

$$U_{\text{total}} = \sum_{i,j} U_{i,j} = \sum_{i,j} \Delta y_{i,j} \Delta z_{i,j} \int_0^{h_{i,j}} \mathrm{d}x \, g\rho(x; h_{i,j}) x, \tag{5.39}$$

where $\Delta y_{i,j}$ and $\Delta z_{i,j}$ are the dimensions of the cell at (i, j) and $\rho(x; h_{i,j})$ is the fluid

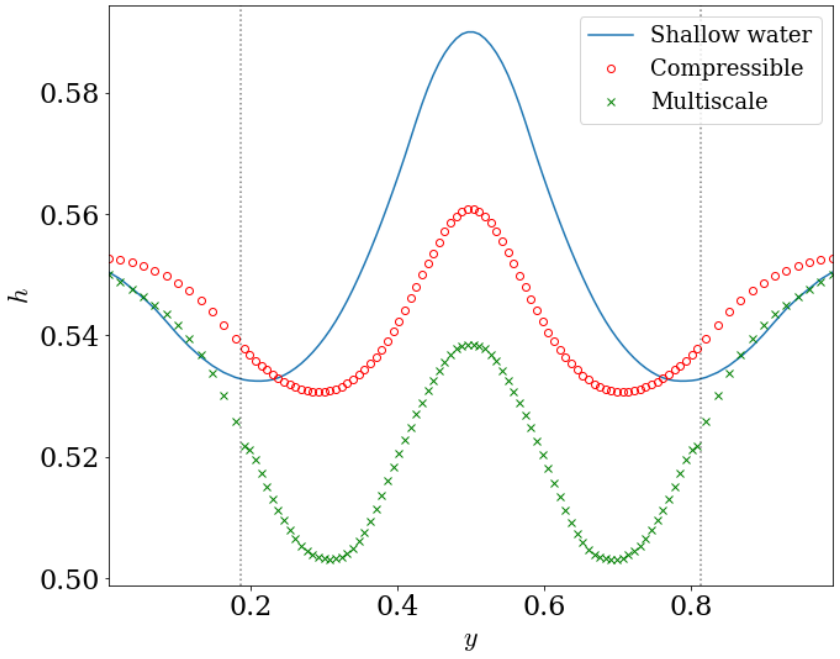

FIGURE 5.21: Results of the Gresho vortex test at $t = 1$ for the shallow water, compressible and multiscale models with 2 levels of mesh refinement. The plot shows the fluid height along the y-axis through the centre of the simulation's x-z domain. The coarsest grid has a resolution of $8 \times 64 \times 64$. The vertical dashed lines indicate the boundaries of the central fine grid. The locations of the minima in the multiscale model line up with those in the compressible model, however their magnitude is significantly reduced. This is likely to be due to a mismatch in the heights of the shallow water and compressible models at their interface.

density at height x from the ocean floor, calculated using (5.37) from the assumption of adiabaticity. The kinetic energy was calculated similarly, integrating the density in the vertical direction up to the height at (i, j) to obtain the total mass of the fluid in the column of fluid, multiplying this by $\frac{1}{2}v_{i,j}^2$, then again summing over the entire grid. Both vertical integrals were carried out numerically.

The plot shows evolution of the change in energy up to the time $t = 1$. By this point, we can observe that the changes in energy for the 3 different models are clearly distinct. In order to judge how concerned we should be about the overall conservation of energy, it is important to compare this time to the dynamical timescale for the system: a decrease of 10% in total energy over several hundred dynamical timescales would indicate that the energy is lost at a much slower rate than for the same decrease over a single dynamical timescale. For this system, we can estimate the dynamical timescale from the maximum speed of the vortex, u_0, and its diameter, $4R$: $t_{\text{dyn}} \sim \frac{4R}{u_0} \sim 8$. This plot therefore shows the change in energy over $\sim \frac{1}{8}t_{\text{dyn}}$, so given that the energy in the multiscale model has already decreased by $\sim 6\%$, this demonstrates very poor energy conservation.

From the plot, it can be seen that the shallow water-only model (solid lines) conserves

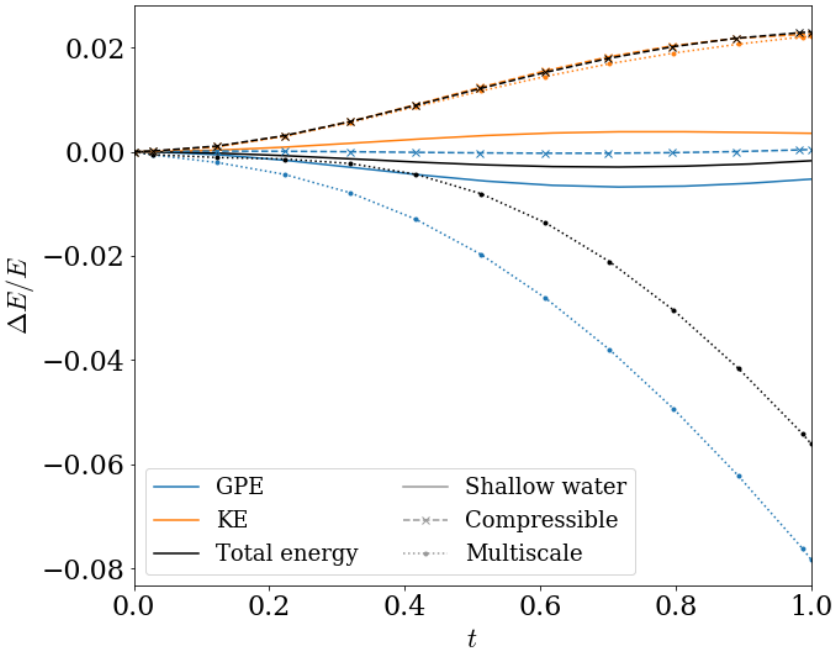

FIGURE 5.22: Energy conservation of Gresho vortex test. Plotted are the relative changes in the total energy (gravitational potential energy + kinetic energy), the gravitational potential energy (GPE) and the kinetic energy (KE). The line markers indicate the results for the 3 different tests. It can be seen that the shallow water model has the best energy conservation, the multiscale model the poorest.

the total energy very well, with the decrease in gravitational potential energy balancing the increase in kinetic energy. The compressible-only model (dashed lines with crosses) performs less well, with the gravitational potential energy of the system not changing significantly over the course of the simulation and therefore not balancing off against the increase in kinetic energy. Looking back at Figure 5.21, this is not surprising – the height profile of the fluid in the compressible test is considerably flatter than that of the other two tests. Consequently, the change in total energy is completely dominated by the increase in kinetic energy. The multiscale test (dotted lines) has the poorest energy conservation. The change in kinetic energy is comparable to that of the compressible-only model, however this is not enough to balance off against the much greater decrease in gravitational potential energy. This large decrease in gravitational potential energy is most likely a result of the underestimate of the height.

This test is an example of a system where the shallow water and compressible models we have implemented do not agree particularly well. Despite this, our multiscale model has performed reasonably well, with the resulting height profile largely consistent with the models it is using in each of the domains. The errors are largely a result of the disagreements between the two models we have tried to couple.

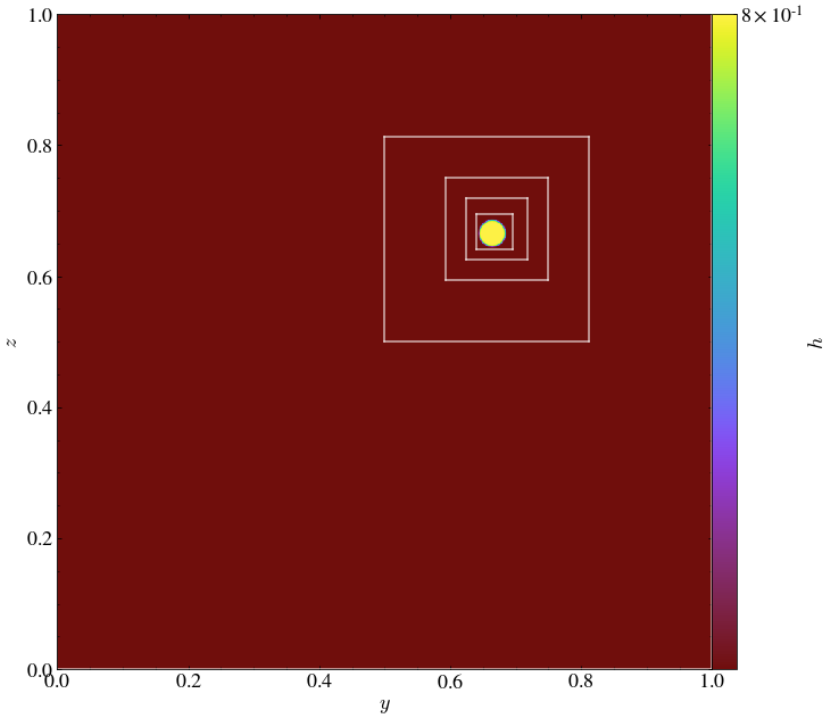

FIGURE 5.23: Initial data used for the scaling test. A small cylinder of fluid was initialised with a height of $h = 0.8$, with the surrounding fluid height $h = 0.78$. Overlaid are the outlines of the finer grids for a run with 5 levels of mesh refinement.

Scaling

The aim of the multiscale model is to be able to capture the physics across a wide range of physical scales in a computationally efficient way. It is therefore the aim for it to be faster to run a simulation using the multiscale model than using a compressible-only model. In this section, we show the results of a test devised to investigate this. The initial data can be seen in Figure 5.23. A small cylinder of fluid with a height of $h = 0.8$ was placed in the top right of the domain, with the surrounding fluid having a height of $h = 0.78$. The cylinder's radius was approximately the width of two coarse grid cells. This data was evolved to the time $t = 0.0002$, approximately one coarse grid timestep. This test was repeated with increasing numbers of grids for the shallow water-only model, compressible-only model and the multiscale model. For the multiscale model, all grids evolved the shallow water model with the exception of the finest, which used the compressible model.

This test was devised to try and represent the sort of grid hierarchy we might use when modelling X-ray bursts. In our models, the total domain would be $\sim 10\,\text{km}$ across, with the flame front only spanning a few metres. The compressible grids would only cover the front and the surrounding area, $\lesssim 1\%$ of the total domain.

N levels	1	2	3	4	5
SWE	0.15	0.22	0.59	3.56	35.72
Compressible	1.21	1.68	4.17	26.61	306.78
Multiscale	0.15	1.34	2.91	23.94	275.69

TABLE 5.3: Time taken for the scaling test with coarse grid resolution $8 \times 64 \times 64$ to reach $t = 0.0002$ for various models and maximum number of AMR levels. The time is quoted in seconds without the initialisation time for simulations run using 6 MPI processes. All grids in the multiscale tests were shallow water grids, except the finest which used the compressible model.

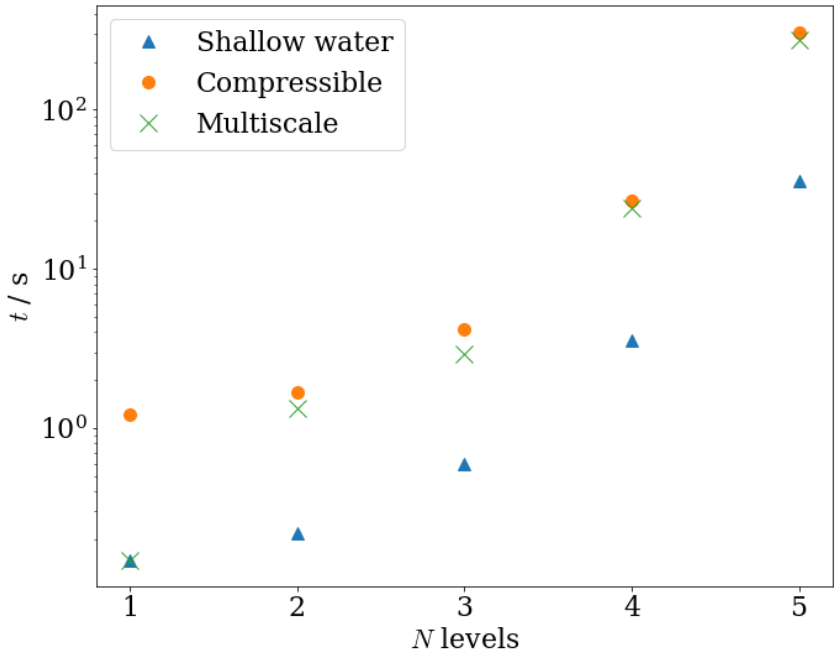

FIGURE 5.24: Plot showing the total CPU time taken for various runs of the 1d dam breaking test against the maximum refinement level used in the simulation. The plot shows runs using the multiscale model, the compressible model only and the shallow water model only.

The total CPU time taken to run each test is listed in Table 5.3 and plotted in Figure 5.24. As would be expected, the shallow water-only model is by far the fastest, consistently more than 8 times faster than the compressible-only model. The multiscale model is faster than the compressible-only model, though not by as much as might be hoped. Despite the fact that the finest compressible grid only covers a very small part of the domain, it appears that the multiscale calculations are still completely dominated by the evolution of this grid, and suggests that there are significant overheads incurred in converting between the two models. The performance of the multiscale model could be improved by optimising this conversion such that these overheads are reduced – little work has been made in optimising the code at this point, so there is almost certainly significant room for improvement.

Matching data at the model interface

As demonstrated in the tests in this section, the multiscale model presented here is unable to accurately model the propagation of features across the model interface unless the features are very smooth (as in the Gaussian test of Section 5.5.2.2), with significant instabilities developing at the interface when the features first begin to cross it. As has been suggested, these errors are likely arise from the failure to match the fluid properties produced by the two models at the interface to sufficient accuracy.

One way that the errors could be reduced is by dropping our assumption of adiabaticity at the interface. This was used in order to simplify the conversion of variables between the two models, however it may not be sufficiently accurate to use in our multiscale scheme for modelling non-smooth features.

Another way that errors could be reduced is by introducing numerical techniques that minimise the non-physical waves produced by the Riemann solvers across the cell boundaries at the interface of the models when the models do not match up sufficiently well. One such technique is the Ghost Fluid Method, first developed by Fedkiw et al. (1999) and later extended to relativistic fluids by Millmore (2010); Muddle (2015). This method was developed in order to eliminate nonphysical oscillations produced at material interfaces when solving the Euler equations. The model interface in our multiscale model is equivalent to the material interface considered in the original Ghost Fluid Method: although the real, physical system on either side of the interface is the same, the different models used to describe the fluid on each side of the interface means that they in effect act as if they are in fact different materials with different physical properties. It should be noted that the original Ghost Fluid Method shows unphysical solutions in certain cases, but modified Ghost Fluid Methods such as those used by Muddle (2015) can alleviate these problems.

5.6 Summary

In this chapter, we looked at combining the physical models described in previous chapters to produce a multiscale model that captures the dominant physics across a wide range of scales. In the first section, we presented an overview of standard adaptive mesh refinement techniques and described how they may be extended to include different physical models at different scales using adaptive mesh and algorithm refinement. In the following section, we looked at the derivation of the low Mach number and shallow water equations using the analytic multiscale approach of decomposing the equations into variables associated with the problems' fast and slow timescales. Next, we described the numerical methods used implement the multiscale model, in particular how information was converted between the two models. We then presented tests of a multiscale

solver which couples the shallow water and compressible models for both Newtonian and relativistic systems.

The next step in developing our multiscale model is to include the low Mach number equations, applying them to the finest grids. There has already been some work on a hybrid low Mach number/compressible model by Motheau et al. (2018), who have implemented a Newtonian model built on the `CASTRO` and `MAESTRO` codes to create the new hybrid `LAMBDA` code. It is expected that from a modelling perspective, coupling the low Mach number and compressible models should be relatively simple compared to the shallow water/compressible coupling, as both models evolve largely the same sets of variables. However, for this system it will be important to ensure that care is taken at the interface to filter out any sound waves and ensure that any features entering grids covered by the low Mach number model are travelling at suitably subsonic speeds. The main difficulty here which has prevented this from being included in this work is that `CASTRO` uses the C++ interface of `AMReX`, however `MAESTRO` still uses the older Fortran interface (a C++ version is currently being developed). This makes the implementation of the coupling of the two models much less straightforward. In the future, we intend to incorporate some of the routines from the `LAMBDA` code in the code in this work to add the (relativistic) low Mach number equations to our model.

In the multiscale model implemented here, we have used a metric with a constant gravitational acceleration, such that the system can be considered special relativistic in the horizontal plane of the domain. We intend to extend this in the future to the full general relativistic shallow water and compressible equations, allowing us to model the physical system we are most interested in: neutron star oceans. To model X-ray bursts, we shall also evolve the species fraction(s) and include heating terms in order to model burning.

Another improvement that could be made to our model would be to use higher order numerical schemes in order to obtain better accuracy. In our current model, we have used low order schemes as they are simpler to implement and are more stable. However, as seen in the tests above, these methods are highly diffusive and high resolution is necessary in order to obtain high accuracy. This is not optimal, as higher resolution simulations are more computationally expensive to run. A better way to obtain higher accuracy is to use higher order numerical methods in the first place. For example, we could use a higher order WENO scheme for the reconstruction, rather than the HRSC methods used here which reduce to first order at nonlinear features (as described in Section 1.5.2.5). Likewise, we could use a higher order Runge-Kutta method for the time integration instead of the third order method we have used.

The fixed mesh refinement tests in Section 5.5.2.2 demonstrated that our current model is vulnerable to the development of errors at level interfaces as features pass through them. If the features include large pressure gradients, the errors can grow so large that they destabilise the code. As described in Section 1.5.2.7, one reason why these errors

occur because our scheme is not well-balanced – the discrete equations we evolve do not satisfy the discrete form of the equations of hydrostatic equilibrium. In our simulations of X-ray bursts, there should not be any features with large gradients passing through level interfaces. As the flame propagates as a deflagration rather than a detonation, we do not expect there to be discontinuities present in the solution. In addition to this, regridding should ensure that any features with large gradients remain well away from the level boundaries. However, smoother features should be able to pass between the levels. The errors from the scheme not being well-balanced grow over time, so if we are to model any significant time period, it is likely that we will need to take steps to make sure that our scheme is well-balanced.

When translating between the shallow water and compressible models, we assumed that the fluid was adiabatic and in hydrostatic equilibrium. As shown in the radial dam breaking and Gresho vortex tests, this led to errors at the model interfaces. This suggests that this assumption is not sufficiently valid for this system. The accuracy could therefore be improved by finding a better way of 'translating' the data between the two models. As described in Section 5.5.2.2, one way to do this would be to use the Ghost Fluid Method of Fedkiw et al. (1999); Millmore (2010); Muddle (2015). Care must be taken here – as seen in the scaling tests in Section 5.5.2.2, there is a numerical cost involved in these data translation routines. It is necessary to ensure that the additional costs of translating the data between state vectors, grids etc. does not end up outweighing the gains of the AMAR approach. Additional complications will arise here when we extend our model to general relativity and will need to use the general relativistic equivalent of the hydrostatic equilibrium equations – the TOV equations – which are likely to be significantly more complex for all but the most simple spacetimes.

As described in Section 5.2, in previous work AMAR has been used to produce schemes with multiple numerical models as well as multiple physical models. This is an approach that we could use in our model, using higher order numerical solvers on the finest grids and cheaper, lower order methods on the coarser grids. This would give us improved accuracy on the finest scales, potentially meaning that fewer levels of refinement are required in order to obtain the desired accuracy and therefore leading to improved computational efficiency overall. At the same time, using lower order methods on the coarsest grids would enable us to keep calculations on these grids as cheap as possible, so that the expensive higher order calculations are only performed where needed.

6 Discussion & conclusions

In this thesis, we have looked at the physics of type I X-ray bursts in neutron star oceans and have developed new techniques to model these systems numerically. In particular, we have extended the existing low Mach number and shallow water approximations to general relativistic fluids, combining these models to produce a multiscale model which is computationally efficient whilst capturing the relevant physics across a wide range of scales.

In Chapter 1, a literature review of X-ray bursts was presented. It was found that X-ray bursts have a very rich phenomenology, with an array of different observed features. The physics in neutron star oceans is very extreme compared to that on Earth, and as such it is difficult to ascertain the dominant processes that drive the propagation of burning fronts during bursts. There have been many models proposed to try and explain the observed features, but currently there does not exist a universal model that can explain all observations. A few studies have included general relativistic effects, however no simulations of the propagation of the burning front have been carried out using the fully general relativistic fluid equations. In the rest of this thesis, we therefore developed tools that will allow general relativistic effects to be included in simulations of bursts.

Chapter 2 described the derivation of the general relativistic low Mach number equations, extending the existing Newtonian low Mach number equations to the general relativistic case. After considering a range of formulations of the general relativistic equations, the Wilson formulation was selected in order that the resulting relativistic equations would resemble the Newtonian equations, and therefore making it easier to compare them and implement numerically. This numerical implementation was then described, and was carried out by adapting the existing Newtonian `pyro` code. Some standard tests were performed on this code, with the metric set to model a weak, planar gravitational field. The code produced qualitatively the results that would be expected for this regime.

Chapter 3 described work carried out to model relativistic detonations and deflagrations. An exact Riemann solver was developed to investigate relativistic detonations and deflagrations, and in particular to see if the effects of tangential velocities could be important for deflagrations in neutron star oceans. The code for this solver, `R3D2`, has been published in Harpole and Hawke (2016). It was found that a fast tangential

velocity can cause a deflagration to transition into a detonation. However, the velocity required is so great that unless the state is already on the verge of transitioning, this process is unlikely to occur in any real, physical system, and certainly would not occur in neutron star oceans. A compressible solver, built by extending the compressible solver in `pyro` to reactive relativistic systems, was used to investigate 2d systems.

In Chapter 4, we derived the general relativistic shallow water equations. This model was developed in order to capture the large scale effects of the Coriolis force on the propagation of the burning fronts. The equations were extended to a multi-layer model by including layer interaction terms, and burning was also included by evolving the species fraction(s). A numerical implementation of the relativistic shallow water equations was demonstrated.

Chapter 5 described how the models of the previous sections may be combined to produce a multiscale model, with the small scale physics captured by the low Mach number equations and the large scale physics by the shallow water equations. We first looked at standard adaptive mesh refinement, and described how this can be extended to include algorithm refinement. Analytic multiscale methods were investigated and shown to be an alternative approach to deriving the low Mach number equations. The numerical implementation of a relativistic multiscale model coupling the shallow water equations and compressible hydrodynamics was presented.

In this thesis, we developed several models to capture the physics across a range of scales and implemented these in a number of codes. We extended the Newtonian `pyro` code to model relativistic low Mach number systems. While the numerical implementation was more complex than for conventional compressible codes, the low Mach number approximation allowed us the gain considerable computational speed up when modelling flows in this regime. We developed an exact solver for the reactive relativistic Riemann problem, `R3D2`, allowing us to investigate the behaviour of detonations and deflagrations in relativistic fluids. We implemented relativistic shallow water equations in `Castro`. This was shown to be far cheaper to evolve than the full relativistic compressible system, enabling us to model the large scale flow of (shallow) relativistic fluids. Finally, we coupled this relativistic shallow water model with the relativistic compressible equations to produce a hybrid multiscale model. This was challenging to implement, given that we had to adapt an existing AMR code that assumed the same physical model was to be used on all levels of refinement. In the tests presented in this thesis, our current implementation of this multiscale model did not demonstrate the considerable speed up we might have hoped for. However, given how much faster the shallow water model is to evaluate than the compressible model, we believe that with some optimisation this could potentially offer a huge computational speed up when applied to modelling flames in the neutron star ocean, allowing us to model a much larger fraction of the neutron star surface than is currently possible using existing approaches.

The overall plan for this project is to improve and extend the current computational model such that simulations can be run on larger physical domains and include more of the relevant physics. As mentioned above, this would include adding the low Mach number equations to the multiscale model to capture the physics at the smallest scales. Eventually, will be able to create a model on the scale of the whole neutron star ocean that can produce phenomena consistent with X-ray burst observations, and which therefore can help shed some light on the underlying physics.

A key part of the physics that we would like to investigate are the effects of rotation. In general relativity, rotating systems experience frame dragging effects and will deform at the equator (Paschalidis et al., 2017). There have not yet been any studies of how these may affect the propagation of fronts in the ocean, but it is possible that they may provide mechanisms that explain multiple peaks and burst oscillations. In order to study rotation, we will need to use a stationary reference state (rather than the current static one); this will considerably complicate our equations and their numerical implementation.

To our knowledge, this is the first time such a model coupling the shallow water and compressible equations has been used. We believe that this model and the general multiscale approach has many other potential applications and could be of great use in astrophysical modelling, which often involves systems with physics across a wide range of scales.

For example, in core collapse supernovae there are many processes operating across a large range of length scales. The dominant physics acting at the smallest and largest scales is very different, and so is best captured using different physical models. At the smallest scales, the physics is dominated by turbulent nuclear burning in the detonation front as it propagates outwards. However, the entire detonation front extends over a region several orders of magnitudes greater in size, and in order to accurately model its propagation, a sizeable fraction of the domain over which it spreads must be modelled so that instabilities that manifest on larger scales such as the standing accretion shock instability (SASI) (Blondin et al., 2003) can be captured. A multiscale approach would couple the small scale and large scale models, allowing for the physics to be described using the most physically appropriate and most computationally efficient model at each scale.

Another potential application of multiscale methods would be for binary neutron star mergers. A current challenge when modelling these mergers is to calculate the saturation point of the magneto-rotational instability (MRI) (Pessah and Goodman, 2009; Müller, 2016; Sawai et al., 2013; Sawai and Yamada, 2016) and dynamo effects (Mösta et al., 2015). Existing simulations have had difficulties with the convergence of these quantities, with the saturation point continuing to grow in magnitude as the resolution is increased.

The multiscale method, with its ability to capture different physics at different scales, could help in solving this problem.

Convection in stellar interiors is another area where a multiscale approach could be useful. In massive stars, convection occurs in the core. It is a large scale process which mixes material in the core, taking the products of burning out from the centre and replacing them with fresh fuel. Convection within stars is also highly turbulent, however given the large length scales over which it operates, current computing resources are nowhere near powerful enough to fully resolve this turbulence down to the smallest microscopic scales. Current computational modelling therefore typically involves the use of significantly simplified models, such as mixing length theory (MLT) and Reynolds-averaged Navier-Stokes (RANS) (Arnett et al., 2015). The convective flow is very subsonic, so is best modelled using solvers suited to low Mach number flows (Kupka and Muthsam, 2017). The boundary of this convective core with the radiative region is particularly difficult to model, however is of great importance for the evolution of the star (Meakin and Arnett, 2007). It is very thin compared to the extent of the core, and is believed to be dominated by turbulent processes. The physics of this zone determines how efficiently material can move in and out of the core, and therefore controls how much fuel the core has access to before it runs out. It also determines how much of the burnt material can move out to the outer layers of the star where it can then affect quantities we are able to directly measure (e.g. metallic absorption lines). A multiscale model would allow us to capture both the large scale, low Mach number convective motion in the entire core and the small scale turbulent processes necessary to effectively model the physics in the boundary zone.

Beyond these examples, there are many other astrophysical systems featuring physical processes which operate over a large range of scales where multiscale methods could be useful. However, as discussed in Section 5.3.3, there are some systems where multiscale methods of the form presented here would not be appropriate as they would be unable to capture physics key to the system. An example of a process that would be poorly described by a multiscale model is stellar mode coupling. In stars, it is possible for oscillatory modes to become coupled and subsequently grow in amplitude unstably. One such example is the *p*-*g* coupling instability where high frequency acoustic waves and low frequency internal gravity waves can become coupled and grow to large energies in very short timescales. This instability could be important in e.g. the evolution of neutron star binaries (Weinberg et al., 2013), where it could significantly influence the gravitational wave signal and electromagnetic emission as they merge. It affects the tidal deformability of the star and therefore will be important for the calculation of the Love number. The ratio of the frequencies of these two modes is typically $\omega_p/\omega_g \gg 1$. As multiscale methods rely on the assumption that features on different scales do not interact, they would be unsuitable for modelling a system such as this without some modifications.

Despite the limitations of multiscale methods when modelling systems where resonant behaviour is important, as has been shown in this thesis, multiscale methods offer many significant advantages over traditional modelling techniques. There is currently significant work being done in this area in the field of engineering, including the development of new analytic and numerical tools. There is a lot of potential for the application of multiscale methods to astrophysical systems, which will make the large scale, high resolution modelling of a wide range of systems computationally feasible.

A Details of the derivation of the relativistic low Mach number equations in the 3+1 formulation

In this appendix, we shall outline the derivation of the relativistic low Mach number equations in the 3+1 formulation. We begin with the covariant 3+1 fluid equations given by (2.28) – (2.30).

Combining the energy (2.30) and continuity (2.28) equations and expanding, we have

$$\left(\partial_t - \beta^k\partial_k\right)\left(HW^2\right) + {}^{(3)}\nabla_k\left[\alpha v^k HW^2\right] = \left(\partial_t - \beta^k\partial_k\right)p + HW^2\left(\alpha v^m v^n K_{mn} - v^m\partial_m\alpha + \alpha K\right). \quad \text{(A.1)}$$

We can also expand the momentum equation (2.29) as

$$HW^2\left[\left(\partial_t - \beta^k\partial_k\right)v_i + \alpha v^{k\,(3)}\nabla_k v_i\right] + {}^{(3)}\nabla_i p + v_i\left[\left(\partial_t - \beta^k\partial_k\right)\left(HW^2\right) + {}^{(3)}\nabla_k\left[\alpha v^k HW^2\right]\right] = \left(HW^2 + p\right){}^{(3)}\nabla_i\alpha + \alpha K S_i, \quad \text{(A.2)}$$

which after further manipulations gives

$$\left(\partial_t - \mathcal{L}_\beta\right)v_i + \alpha v^{k\,(3)}\nabla_k v_i + \frac{{}^{(3)}\nabla_i p}{HW^2} = -\frac{v_i}{HW^2}\left[\left(\partial_t - \beta^k\partial_k\right)p\right] - v_i\left(\alpha v^m v^n K_{mn} - v^m\partial_m\alpha\right) - \frac{\varrho + pW^2v^2}{HW^2}{}^{(3)}\nabla_i\alpha. \quad \text{(A.3)}$$

To obtain the low Mach number limit of these equations, we make two fundamental assumptions. Firstly, we assume that there is a reference state in which the metric and 'background' fluid quantities are time independent. We will use overbars, e.g. $\bar{p}$, to denote this. Secondly, we will assume that the full fluid quantities are perturbations from this reference state, so $p = \bar{p} + \pi$, and that the perturbation in the pressure is second order in the perturbation in the velocity: $\pi/\bar{p} = \mathcal{O}(V^2)$.

We find that the reference state obeys the momentum equation,

$$-\mathcal{L}_\beta \bar{v}_i + \alpha \bar{v}^{k\,(3)}\nabla_k \bar{v}_i + \frac{{}^{(3)}\nabla_i \bar{p}}{\bar{H}\bar{W}^2} = -\frac{\bar{v}_i}{\bar{H}\bar{W}^2}\left[-\beta^k \partial_k \bar{p}\right] - \bar{v}_i\left(\alpha \bar{v}^m \bar{v}^n K_{mn} - \bar{v}^m \partial_m \alpha\right) - \frac{\bar{\varrho} + \bar{p}\bar{W}^2\bar{v}^2}{\bar{H}\bar{W}^2}{}^{(3)}\nabla_i \alpha. \quad \text{(A.4)}$$

If we replace the gradient of the reference state pressure in the full momentum equation with this result, we get the expression

$$\begin{aligned}\left(\partial_t - \mathcal{L}_\beta\right) v_i + \frac{\bar{H}\bar{W}^2}{HW^2}\mathcal{L}_\beta \bar{v}_i + \alpha v^{k\,(3)}\nabla_k v_i - \frac{\bar{H}\bar{W}^2}{HW^2}\alpha \bar{v}^{k\,(3)}\nabla_k \bar{v}_i + \frac{{}^{(3)}\nabla_i \pi}{HW^2} = \\ \frac{(v_i - \bar{v}_i)}{HW^2}\left[\beta^k \partial_k \bar{p}\right] - \frac{v_i}{HW^2}\left[\left(\partial_t - \beta^k \partial_k\right)\pi\right] - \\ v_i\left(\alpha v^m v^n K_{mn} - v^m \partial_m \alpha\right) + \frac{\bar{H}\bar{W}^2}{HW^2}\bar{v}_i\left(\alpha \bar{v}^m \bar{v}^n K_{mn} - \bar{v}^m \partial_m \alpha\right) \\ - \frac{\varrho - \bar{\varrho} + pW^2 v^2 - \bar{p}\bar{W}^2\bar{v}^2}{HW^2}{}^{(3)}\nabla_i \alpha. \quad \text{(A.5)}\end{aligned}$$

Let us write $v_i = \bar{v}_i + V_i$ and assume that V_i is small, so that terms quadratic in V_i can be dropped. We can also write $H = \bar{H} + \Theta$ to get

$$\mathcal{H} = 1 - \frac{\bar{H}\bar{W}^2}{HW^2} = \frac{\Theta}{H}\left(1 - 2\bar{W}^2 \bar{v}^i V_i\right) + 2\bar{W}^2 \bar{v}^i V_i + \mathcal{O}(V^2). \quad \text{(A.6)}$$

Using the fact that the background quantities are time independent, $\partial_t \bar{v}_i = 0$, after some manipulation the full momentum equation becomes

$$\begin{aligned}\left(\partial_t - \mathcal{L}_\beta\right) V_i - \mathcal{H}\mathcal{L}_\beta \bar{v}_i + \alpha\left(V^{k\,(3)}\nabla_k V_i + \bar{v}^{k\,(3)}\nabla_k V_i + V^{k\,(3)}\nabla_k \bar{v}_i + \mathcal{H}\bar{v}^{k\,(3)}\nabla_k \bar{v}_i\right) + \frac{{}^{(3)}\nabla_i \pi}{HW^2} = \\ \frac{V_i}{HW^2}\beta^k \partial_k \bar{p} - \frac{\bar{v}_i + V_i}{HW^2}\left[\left(\partial_t - \beta^k \partial_k\right)\pi\right] - \\ \left(\mathcal{H}\bar{v}_i + V_i\right)\left(\alpha \bar{v}^m \bar{v}^n K_{mn} - \bar{v}^m \partial_m \alpha\right) - \bar{v}_i\left(2\alpha V^m \bar{v}^n K_{mn} - V^m \partial_m \alpha\right) - \\ \frac{(\varrho - \bar{\varrho}) + \left(\pi \bar{W}^2 \bar{v}^2 + 2\left(\bar{p} + \pi\right)\bar{W}^2 \bar{v}^k V_k\right)}{HW^2}{}^{(3)}\nabla_i \alpha. \quad \text{(A.7)}\end{aligned}$$

If we restrict this to the case where the reference state is in hydrostatic equilibrium (so $\bar{v}_i \to 0$ and $\bar{W} \to 1$), we find the gradient of the reference state pressure to be

$$\frac{{}^{(3)}\nabla_i \bar{p}}{\bar{H}} = -\frac{\bar{\varrho}}{\bar{H}}{}^{(3)}\nabla_i \alpha. \quad \text{(A.8)}$$

As we assume the metric, and hence the lapse, to be fixed and given by the initial data for a particular problem, this allows us to replace derivatives of the background pressure with known terms.

If we also assume that the perturbation in the velocity is small then the full state Lorentz factor $W \simeq 1$, as the Lorentz factor is quadratic in the velocity, and the full momentum

equation reduces to

$$\left(\partial_t - \mathcal{L}_\beta\right) V_i + \alpha V^k {}^{(3)}\nabla_k V_i + \frac{{}^{(3)}\nabla_i \pi}{H} = -\frac{V_i}{H}\left[\left(\partial_t - \beta^k \partial_k\right) p\right] - \frac{\left(\varrho - \bar{\varrho}\right)}{H} {}^{(3)}\nabla_i \alpha, \quad \text{(A.9)}$$

where we have dropped all terms that are quadratic in the velocity. Finally, we argue that the pressure derivative terms on the right hand side are, thanks to the assumptions of hydrostatic equilibrium for the reference state, at most first order in the velocity. This gives the final form

$$\left(\partial_t - \mathcal{L}_\beta\right) V_i + \alpha V^k {}^{(3)}\nabla_k V_i + \frac{{}^{(3)}\nabla_i \pi}{H} = -\frac{\varrho - \bar{\varrho}}{H} {}^{(3)}\nabla_i \alpha. \quad \text{(A.10)}$$

Taking the Newtonian limit of this sends the enthalpy H to the energy density ϱ, and the equation has the expected low Mach form.

With the assumptions we have made, the energy equation reduces to

$$\left(\partial_t - \beta^k \partial_k\right)\left(HW^2\right) + {}^{(3)}\nabla_k\left[\alpha v^k H W^2\right] = -\beta^k \partial_k \bar{p} + H\left(-v^m \partial_m \alpha + \alpha K\right). \quad \text{(A.11)}$$

We need to keep the Lorentz factor terms inside the derivatives on the left hand side, as they may be first order in the velocity. This equation is less straightforward to compare in the Newtonian limit. The trace of the extrinsic curvature (the final term) would vanish in any standard Newtonian limit. The derivative of the lapse would, however, not vanish. We instead replace it with the pressure derivative in (A.8), and obtain the correct Newtonian limit as desired.

B Details of the derivation of the relativistic low Mach number equations in the Wilson formulation

In this appendix, we shall give a detailed derivation of the low Mach number equations derived from the Wilson formulation of the relativistic hydrodynamics equations, (2.34) – (2.36).

B.1 Continuity equation

Taking (1.22) and substituting in D and U^μ, we find

$$\nabla_\mu(DU^\mu) = 0. \tag{B.1}$$

Expanding out the covariant derivative and performing the 3+1 split into the time and spatial parts, we get

$$\partial_t D + \partial_i(DU^i) = -D\Gamma^\mu{}_{\mu\nu}U^\nu. \tag{B.2}$$

When implementing the equations numerically, it will be necessary to have expressions for the evolution of the base state conserved density $\bar{D}$ and the perturbed state conserved density, $D' = D - \bar{D}$. The base state conserved density evolution equation can be defined by laterally averaging (averaging over the plane parallel to the stellar surface, perpendicular to the radial direction) over the full continuity equation, giving

$$\partial_t \bar{D} + \partial_r\left(\bar{D}\bar{U}\right) = -\overline{D\Gamma^\mu{}_{\mu\nu}U^\nu}, \tag{B.3}$$

where we have defined $\bar{U} = \bar{U}(t, r)$ as the base state velocity such that the full velocity field is given by $U^i = \bar{U}(e_r)^i + U'^i$. This can be subtracted from the full density evolution equation to obtain the evolution of the perturbational conserved density:

$$\partial_t D' + \partial_i\left(D'U^i + \bar{D}U'^i\right) = -\left(D\Gamma^\mu{}_{\mu\nu}U^\nu\right)'. \tag{B.4}$$

We capture the changes that can occur due to significant convective overturning by imposing the constraint that the lateral average of the perturbational conserved density $\overline{D'} = 0$ for all time, which gives

$$\partial_t \overline{D'} + \partial_i \eta_D = 0, \tag{B.5}$$

where we have defined $\eta_D = \overline{D' U^r}$.

B.2 Momentum equation

Starting with (1.24) and substituting in D and U^μ,

$$DhU^\nu \partial_\nu (u^0 U_\mu) + \partial_\mu p + U_\mu U^\nu (u^0)^2 \partial_\nu p = Dhu^0 \Gamma_{\rho\nu\mu} U^\nu U^\rho. \tag{B.6}$$

Splitting the time and spatial parts of the first term and taking the j-th spatial component,

$$\partial_t U_j + U^i \partial_i U_j + \frac{1}{u^0} U_j \frac{\mathrm{D}u^0}{\mathrm{D}t} + \frac{1}{Dhu^0} \partial_j p + \frac{u^0}{Dh} U_j U^\nu \partial_\nu p = \Gamma_{\rho\nu j} U^\nu U^\rho. \tag{B.7}$$

As described in Vasil et al. (2013), the Newtonian equivalent of this equation does not conserve energy. It is therefore necessary to incorporate a constraint to ensure that energy is conserved. The form of the constrained equation can be found using Lagrangian analysis.

The Lagrangian density for relativistic fluids is given by Rezzolla and Zanotti (2013) as

$$\begin{aligned} \mathcal{L}_{\text{fluid}} &= \sqrt{-g} \left(R - 2\kappa\rho \left[1 + f + \frac{1}{2} \lambda g_{\mu\nu} u^\mu u^\nu \right] \right) \\ &= \sqrt{-g} \left(R - 2\frac{\kappa D}{u^0} \left[h - sT - \frac{p u^0}{D} + \frac{1}{2} \lambda g_{\mu\nu} U^\mu U^\nu (u^0)^2 \right] \right), \end{aligned} \tag{B.8}$$

where R is the Ricci scalar, $\kappa = 8\pi G$ is the standard gravitational coupling constant, λ is a Lagrange multiplier enforcing the constraint on the 4-velocity of the fluid $u^\mu u_\mu = -1$ and $f = \epsilon - sT = h - 1 - sT - \frac{pu^0}{D}$ is the specific Helmholtz free energy. We define h to be the specific enthalpy, ϵ the specific internal energy, s the specific entropy and T the temperature. The action is given by

$$S = \int \mathrm{d}^4 x \, \mathcal{L}_{\text{fluid}}(g, U^\mu, D, s, \lambda). \tag{B.9}$$

From Hamilton's principle and assuming that perturbations to the metric $\delta g^{\mu\nu}$ vanish (as we are considering fixed spacetimes),

$$\delta S = \int \mathrm{d}^4x \left(\frac{\partial \mathcal{L}}{\partial U^\mu}\delta U^\mu + \frac{\partial \mathcal{L}}{\partial D}\delta D + \frac{\partial \mathcal{L}}{\partial s}\delta s + \frac{\partial \mathcal{L}}{\partial \lambda}\delta\lambda \right). \tag{B.10}$$

Conservation of mass and entropy are given by

$$\nabla_\mu(DU^\mu) = 0, \qquad U^\mu \nabla_\mu s = 0. \tag{B.11}$$

Varying these equations, letting $D \to D + \delta D$, $s \to s + \delta s$ etc., where $\delta D/D$, $\delta s/s$ etc. $\ll 1$, and only keeping terms up to first order in perturbations, we find

$$\nabla_\mu(D\delta U^\mu + \delta D U^\mu) = 0, \tag{B.12}$$

$$\delta U^\mu \nabla_\mu s + U^\mu \nabla_\mu \delta s = 0. \tag{B.13}$$

Varying the constraint $\lambda(U^\mu U_\mu - 1) = 0$, it can be shown that consistency requires $\delta\lambda = 0$.

These expressions are very similar to those found in the Newtonian case, so we can expect they can be parametrised by ξ^μ, an infinitesimal displacement vector, in a similar way,

$$\begin{aligned} \delta U^\nu &= U^\mu \nabla_\mu \xi^\nu - \xi^\mu \nabla_\mu U^\nu = -\mathcal{L}_\xi U^\mu, \\ \delta D &= -\nabla_\mu(D\xi^\mu) = -\mathcal{L}_\xi D - D\nabla_\mu \xi^\mu, \\ \delta s &= -\xi^\mu \nabla_\mu s = -\mathcal{L}_\xi s. \end{aligned} \tag{B.14}$$

Substituting these into equations (B.12) and (B.13), we find that they are consistent and therefore that the parametrisation is valid. We can also see why $\delta g = 0$: we know that $\nabla_\mu g = 0$ so, contracting with u^μ,

$$(U^\mu + \delta U^\mu)\nabla_\mu(g + \delta g) = 0 \qquad \Rightarrow \qquad \delta U^\mu \nabla_\mu g + U^\mu \nabla_\mu \delta g = 0. \tag{B.15}$$

It can then be parametrised similarly to δs:

$$\delta g = -\mathcal{L}_\xi g = -\xi^\mu \nabla_\mu g = 0. \tag{B.16}$$

We can now substitute the parametrised perturbations from (B.14) into the perturbation of the action (B.10) and integrate by parts to obtain

$$\begin{aligned}\delta S &= \int \mathrm{d}^4x \left(\frac{\partial \mathcal{L}}{\partial U^\nu}[U^\mu \nabla_\mu \xi^\nu - \xi^\mu \nabla_\mu U^\nu] + \frac{\partial \mathcal{L}}{\partial D}[-\nabla_\mu (D\xi^\mu)] + \frac{\partial \mathcal{L}}{\partial s}[-\xi^\mu \nabla_\mu s] \right) \\ &= \left[\frac{\partial \mathcal{L}}{\partial U^\mu} U^\mu \xi^\nu \right]_{\partial \mathcal{M}} - \int \mathrm{d}^4x\, \xi^\nu \nabla_\mu \left(\frac{\partial \mathcal{L}}{\partial U^\nu} U^\mu \right) \\ &\quad - \int \mathrm{d}^4x\, \xi^\nu (\nabla_\nu U^\mu) \frac{\partial \mathcal{L}}{\partial U^\mu} - \left[\frac{\partial \mathcal{L}}{\partial D} D\xi^\nu \right]_{\partial \mathcal{M}} \\ &\quad + \int \mathrm{d}^4x\, \xi^\mu D \nabla_\mu \left(\frac{\partial \mathcal{L}}{\partial D} \right) - \int \mathrm{d}^4x\, \xi^\mu (\nabla_\mu s) \frac{\partial \mathcal{L}}{\partial s}. \end{aligned} \tag{B.17}$$

Assuming that the boundary terms vanish, then as δS must vanish for any infinitesimal ξ^μ, we find

$$\nabla_\mu \left(\frac{\partial \mathcal{L}}{\partial U^\nu} U^\mu \right) + \frac{\partial \mathcal{L}}{\partial U^\mu} \nabla_\nu U^\mu - D \nabla_\nu \left(\frac{\partial \mathcal{L}}{\partial D} \right) + \frac{\partial \mathcal{L}}{\partial s} \nabla_\nu s = 0, \tag{B.18}$$

which is very similar to the Newtonian equation. We wish to impose the constraint that the pressure perturbations remain small, so

$$\mathcal{C} = 2\sqrt{-g}\kappa [p(D, s) - p_0(x^\mu)] = 0. \tag{B.19}$$

For the system under this constraint, the Lagrangian density becomes

$$\mathcal{L} \to \mathcal{L} - \mu \mathcal{C}, \tag{B.20}$$

where μ is a Lagrange multiplier. Substituting into (B.18), we get

$$- D \nabla_\nu \left(\mu \frac{\partial \mathcal{C}}{\partial D} \right) + \mu \frac{\partial \mathcal{C}}{\partial s} \nabla_\nu s = 0. \tag{B.21}$$

Using

$$\frac{\partial \mathcal{C}}{\partial D} = 2\sqrt{-g}\kappa \left. \frac{\partial p}{\partial D} \right|_s, \qquad \frac{\partial \mathcal{C}}{\partial s} = 2\sqrt{-g}\kappa \left. \frac{\partial p}{\partial s} \right|_D, \tag{B.22}$$

we get

$$\begin{aligned} 0 &= -D \nabla_\nu \left(2\mu\sqrt{-g}\kappa \left. \frac{\partial p}{\partial D} \right|_s \right) + 2\mu\sqrt{-g}\kappa \left. \frac{\partial p}{\partial s} \right|_D \nabla_\nu s \\ &= -\nabla_\nu \left(2\mu\sqrt{-g}\kappa p \Gamma_1 \right) + 2\mu\sqrt{-g}\kappa \nabla_\nu p, \end{aligned} \tag{B.23}$$

where

$$\Gamma_1 = \left. \frac{\partial \ln p}{\partial \ln \rho} \right|_s = \left. \frac{\partial \ln p}{\partial \ln D} \right|_s \tag{B.24}$$

is the first adiabatic exponent. Letting $p' = 2\mu\sqrt{-g}\kappa p\Gamma_1$ and replacing $p \to p_0$ from the constraint equation (B.19),

$$-\nabla_\nu p' + \frac{p'}{\Gamma_1 p_0}\nabla_\nu p_0 = 0. \tag{B.25}$$

We can define the generalised pressure as the Legendre transform of $\mathcal{L}$ with respect to D,

$$\mathcal{P} = \mathcal{L} - D\frac{\partial \mathcal{L}}{\partial D}. \tag{B.26}$$

Calculating this for the constrained Lagrangian density (with the unconstrained Lagrangian density taking the form of (B.8)), we obtain

$$\mathcal{P} = \sqrt{-g}\left(R + 2\kappa p + 2\kappa\lambda\Gamma_1 p\right). \tag{B.27}$$

Interpreting the curvature term to act as a sort of background pressure, we find

$$\mathcal{P} = \sqrt{-g}R + 2\sqrt{-g}\kappa\left(p + \lambda\Gamma_1 p\right) = p_0 + p', \tag{B.28}$$

so $p' = \pi$, the perturbation to the background pressure.

We can now create a constrained version of the momentum equation by adding the constraint (B.25) (divided by Dhu^0 to match the existing momentum terms) and replacing $p \to p_0$ using (B.19):

$$\begin{aligned}\partial_t U_j + U^i\partial_i U_j + U_j\frac{\mathrm{D}\ln u^0}{\mathrm{D}t} + \frac{1}{Dhu^0}\partial_j p_0 + \frac{u^0}{Dh}U_j U^\nu\partial_\nu p_0 \\ = \Gamma_{\rho\nu j}U^\nu U^\rho - \frac{1}{Dhu^0}\partial_j p' + \frac{p'}{\Gamma_1 p_0 Dhu^0}\partial_j p_0.\end{aligned} \tag{B.29}$$

We can check the consistency of this approach by using the generalised pressure to calculate the energy-momentum tensor. From (Andersson and Comer, 2007), this is given by

$$T^{\mu\nu} = g^{\lambda\nu}\left(\mathcal{P}\delta^\mu_{\ \lambda} + n^\mu\mu_\lambda\right), \tag{B.30}$$

where $n^\mu = nu^\mu = nU^\mu u^0$ is the particle 4-current, $\mu_\nu = \mu u_\nu = \mu U_\nu u^0$ is the momentum covector and $n\mu = \rho h = \frac{Dh}{u^0}$. We can therefore rewrite this in the Wilson formulation as

$$T^{\mu\nu} = \mathcal{P}g^{\mu\nu} + DhU^\mu U^\nu u^0. \tag{B.31}$$

Substituting in the expression found for the generalised pressure, we indeed regain the form of the energy-momentum for a relativistic fluid in thermodynamic equilibrium,

$$T^{\mu\nu} = (p_0 + \pi)g^{\mu\nu} + DhU^\mu U^\nu u^0 = pg^{\mu\nu} + DhU^\mu U^\nu u^0. \tag{B.32}$$

B.2.1 Comparison to Newtonian case

In the Newtonian case (Vasil et al., 2013), the velocity evolution equation is modified by the constraint as

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\frac{1}{\rho}\boldsymbol{\nabla} p' + \boldsymbol{a} \qquad \rightarrow \qquad \frac{\partial \boldsymbol{u}}{\partial t} = -\frac{\beta_0}{\rho}\boldsymbol{\nabla}\left(\frac{p'}{\beta_0}\right) + \boldsymbol{a}, \tag{B.33}$$

where the acceleration

$$\boldsymbol{a} = -\boldsymbol{u}\cdot\boldsymbol{\nabla}\boldsymbol{u} + \frac{\rho-\rho_0}{\rho}\boldsymbol{g}. \tag{B.34}$$

It would therefore be reasonable to assume that the relativistic version can be written with a similar integrating factor, such that the constraint terms in (B.25) can be written as

$$-\xi\nabla_\mu\left(\frac{\pi}{\xi}\right). \tag{B.35}$$

Expanding this out and using the definition of ξ, we find

$$\begin{aligned} &\nabla_\mu\pi - \frac{\pi}{\xi}\nabla_\mu\xi = \nabla_\mu\pi - \frac{\pi}{2\sqrt{-g}\kappa\Gamma_1 p_0}\nabla_\mu p_0 \\ \Rightarrow \qquad &\frac{1}{\xi}\nabla_\mu\xi = \frac{1}{2\sqrt{-g}\kappa\Gamma_1 p_0}\nabla_\mu p_0. \end{aligned} \tag{B.36}$$

Integrating, this becomes

$$\xi = A\exp\left(\frac{\ln p_0}{2\sqrt{-g}\kappa\Gamma_1}\right), \tag{B.37}$$

for some constant A which we can set to 1 without loss of generality. The constrained momentum equation then becomes

$$\begin{aligned} \partial_t U_j + U^i\partial_i U_j + U_j\frac{\mathrm{D}\ln u^0}{\mathrm{D}t} + \frac{1}{Dhu^0}\partial_j p_0 + \frac{u^0}{Dh}U_j U^\nu\partial_\nu p_0 \\ = \Gamma_{\rho\nu j}U^\nu U^\rho - \frac{\xi}{Dhu^0}\partial_j\left(\frac{\pi}{\xi}\right). \end{aligned} \tag{B.38}$$

We can also repeat this for the other pressure derivative term to give

$$\begin{aligned} \partial_t U_j = &- U^i\partial_i U_j - U_j\frac{\mathrm{D}\ln u^0}{\mathrm{D}t} + \Gamma_{\rho\nu j}U^\nu U^\rho \\ &- \frac{1}{Dhu^0}\left(g^\nu{}_j + (u^0)^2 U_j U^\nu\right)\left(\partial_\nu p_0 + \xi\partial_\nu\left[\frac{\pi}{\xi}\right]\right). \end{aligned} \tag{B.39}$$

Terms that are $\sim U_j U^i$ are $\mathcal{O}(\mathrm{Ma}^2)$ and so negligible compared to the other terms. We can therefore discard these, leaving us with

$$\partial_t U_j = -U^i\partial_i U_j - U_j\frac{\mathrm{D}\ln u^0}{\mathrm{D}t} + \Gamma_{\rho\nu j}U^\nu U^\rho - \frac{1}{Dhu^0}\left(\partial_j p_0 + \xi\partial_j\left[\frac{\pi}{\xi}\right]\right). \tag{B.40}$$

B.3 Energy equation

Taking the energy equation (1.23), we find

$$\nabla_\mu \left(u^\mu \rho h\right) = u^\mu \nabla_\mu p. \tag{B.41}$$

Substituting in D and U^μ and expanding out the time and spatial derivatives, this becomes

$$\partial_t(Dh) + \partial_i(U^i Dh) = u^0 \frac{\mathrm{D}p}{\mathrm{D}t} - Dh\Gamma^\mu{}_{\mu\nu}U^\nu, \tag{B.42}$$

where we have defined the convective derivative $\frac{\mathrm{D}p}{\mathrm{D}t} = \partial_t p + U^i \partial_i p$. In the low Mach number approximation, we assume that the total pressure is close to the base pressure, so we can replace the pressure in this term with the base pressure to give

$$\partial_t(Dh) + \partial_i(U^i Dh) = u^0 \frac{\mathrm{D}\bar{p}}{\mathrm{D}t} - Dh\Gamma^\mu{}_{\mu\nu}U^\nu. \tag{B.43}$$

The equation for the base state can be found by averaging laterally to get

$$\partial_t(\overline{Dh}) + \partial_r\left(\bar{U}\overline{Dh}\right) = u^0 \overline{\frac{\mathrm{D}\bar{p}}{\mathrm{D}t}} - \overline{Dh\Gamma^\mu{}_{\mu\nu}U^\nu}, \tag{B.44}$$

where if

$$\begin{aligned} \frac{\mathrm{D}\bar{p}}{\mathrm{D}t} &= \partial_t \bar{p} + \bar{U}\partial_r \bar{p} + U'^r \partial_r \bar{p} \\ &= \psi + U'^r \partial_r \bar{p}, \end{aligned} \tag{B.45}$$

then the lateral average is $\overline{\frac{\mathrm{D}\bar{p}}{\mathrm{D}t}} = \psi$. We therefore find the evolution of the base conserved enthalpy Dh to be

$$\partial_t(\overline{Dh}) + \partial_r\left(\bar{U}\overline{Dh}\right) = u^0\psi - \overline{Dh\Gamma^\mu{}_{\mu\nu}U^\nu}. \tag{B.46}$$

Subtracting from the equation for the full state, we find the evolution of the perturbed state to be

$$\partial_t(Dh)' + \partial_i\left(U^i(Dh)' + U'^i\overline{Dh}\right) = u^0 U'^r \partial_r \bar{p} - \left(Dh\Gamma^\mu{}_{\mu\nu}U^\nu\right)'. \tag{B.47}$$

B.4 Velocity constraint

The velocity constraint casts the equation of state $p_0 = p_0(D, T)$ into an elliptic constraint on the velocity field.

Differentiating the equation of state along particle paths,

$$\frac{\mathrm{D}p}{\mathrm{D}t} = p_D \frac{\mathrm{D}D}{\mathrm{D}t} + p_T \frac{\mathrm{D}T}{\mathrm{D}t}, \tag{B.48}$$

where we have defined $p_D = \partial p/\partial D|_T$ and $p_T = \partial p/\partial T|_D$. Rearranging,

$$\frac{\mathrm{D}D}{\mathrm{D}t} = \frac{1}{p_D}\left(\frac{\mathrm{D}p}{\mathrm{D}t} - p_T \frac{\mathrm{D}T}{\mathrm{D}t}\right). \tag{B.49}$$

Writing the energy equation, (B.43), in the form of convective derivatives, we find

$$\partial_i U^i + \frac{1}{D}\frac{\mathrm{D}D}{\mathrm{D}t} + \frac{1}{h}\frac{\mathrm{D}h}{\mathrm{D}t} = \frac{u^0}{Dh}\frac{\mathrm{D}p}{\mathrm{D}t} - \Gamma^{\mu}{}_{\mu\nu}U^{\nu}. \tag{B.50}$$

Similarly, the continuity equation can be expressed as

$$\frac{\mathrm{D}D}{\mathrm{D}t} + D\partial_i U^i + D\Gamma^{\mu}{}_{\mu\nu}U^{\nu} = 0. \tag{B.51}$$

Expanding out the convective derivative of the enthalpy h and rearranging,

$$\frac{\mathrm{D}T}{\mathrm{D}t} = \frac{h}{h_T}\left(\frac{u^0}{Dh}\frac{\mathrm{D}p}{\mathrm{D}t} - \Gamma^{\mu}{}_{\mu\nu}U^{\nu} - \partial_i U^i - \frac{1}{D}\frac{\mathrm{D}D}{\mathrm{D}t} - \frac{h_p}{h}\frac{\mathrm{D}p}{\mathrm{D}t}\right). \tag{B.52}$$

Substituting this into the expression above for the convective derivative of the density and grouping terms:

$$\frac{\mathrm{D}D}{\mathrm{D}t}\left(1 - \frac{p_T h}{p_D h_T D}\right) = \frac{1}{p_D}\left(\frac{\mathrm{D}p}{\mathrm{D}t}\left[1 - \frac{p_T u^0}{h_T D} + \frac{p_T h_p}{h_T}\right] + \frac{p_T h}{h_T}\Gamma^{\mu}{}_{\mu\nu}U^{\nu} + \frac{p_T h}{h_T}\partial_i U^i\right). \tag{B.53}$$

Given that the enthalpy is given by

$$h = 1 + \varepsilon + p/\rho = 1 + \varepsilon + \frac{pu^0}{D}, \tag{B.54}$$

and that by definition $h_p = u^0/D$, the coefficient of $\mathrm{D}p/\mathrm{D}t$ just becomes $1/p_D$. Substituting this back into the continuity equation and cancelling terms, we obtain

$$\partial_i U^i + \frac{1}{p_D D}\frac{\mathrm{D}p}{\mathrm{D}t} = -\Gamma^{\mu}{}_{\mu\nu}U^{\nu}. \tag{B.55}$$

We can define the first adiabatic exponent as

$$\Gamma_1 = \left.\frac{\partial \ln p}{\partial \ln \rho}\right|_s, \tag{B.56}$$

such that

$$\frac{1}{p_D D} = \frac{1}{\Gamma_1 p}. \tag{B.57}$$

The velocity constraint then becomes

$$\partial_i U^i + \frac{1}{\Gamma_1 p}\frac{\mathrm{D}p}{\mathrm{D}t} = S, \tag{B.58}$$

where we have written the curvature term on the right hand side as the source term S. In order to express this as an elliptic equation, we now seek some function $\zeta(t, r)$ such that

$$\frac{1}{\zeta}\partial_i\left(\zeta U^i\right) = \partial_i U^i + \frac{1}{\Gamma_1 \bar{p}} U^i \partial_i \bar{p}. \tag{B.59}$$

Expanding the left hand side and cancelling repeated terms, we get

$$\frac{1}{\zeta}\partial_i \zeta = \frac{1}{\Gamma_1 \bar{p}}\partial_i \bar{p}. \tag{B.60}$$

Laterally averaging, such that $\partial_i \to \partial_r$ and $\Gamma_1 \to \bar{\Gamma}_1$, and indicating differentiation with respect to r by $'$, this reduces to

$$\frac{\zeta'}{\zeta} = \frac{1}{\bar{\Gamma}_1}\frac{\bar{p}'}{\bar{p}}. \tag{B.61}$$

Integrating this from 0 to r, we get

$$\zeta(t, r) = \zeta(0, t)\exp\left[\int_0^r \mathrm{d}r' \, \frac{1}{\bar{\Gamma}_1 \bar{p}}\frac{\partial \bar{p}}{\partial r}\right]. \tag{B.62}$$

Substituting this back in, we can now write the velocity constraint as

$$\partial_i\left(\zeta U^i\right) = \zeta\left(S - \frac{1}{\bar{\Gamma}_1 \bar{p}}\frac{\partial \bar{p}}{\partial t}\right). \tag{B.63}$$

Laterally averaging this, we obtain the velocity constraint for the base state velocity $\bar{u}$,

$$\partial_i\left(\zeta \bar{U}^i\right) = \zeta\left(\overline{S} - \frac{1}{\bar{\Gamma}_1 \bar{p}}\frac{\partial \bar{p}}{\partial t}\right), \tag{B.64}$$

then subtracting this from the constraint for the full state, we get the constraint on the velocity perturbation,

$$\partial_i\left(\zeta U''^i\right) = \zeta\left(S - \overline{S}\right). \tag{B.65}$$

In the Newtonian low Mach number codes `MAESTRO` (Nonaka et al., 2014) and `pyro` (Zingale, 2014), the velocity constraint is enforced by solving an elliptic equation. We can derive this for our equations by first decomposing the velocity field (which at this stage does not obey the constraint) into a part that does obey the constraint, $U^{i,d}$, and another part proportional to the gradient of some scalar, ϕ:

$$U^{i,*} = U^{i,d} + \frac{\zeta}{Dhu^0}\partial^i \phi, \tag{B.66}$$

where the coefficient of the gradient is suggested from the form of the momentum equation, and

$$\partial_i \left(\zeta U^{i,d}\right) = \zeta \left(S - \frac{1}{\bar{\Gamma}_1 \bar{p}} \frac{\partial \bar{p}}{\partial t}\right). \tag{B.67}$$

Multiplying the full velocity field by ζ and taking the divergence, we find

$$\partial_i \left(\zeta U^{i,*}\right) = \zeta \left(S - \frac{1}{\bar{\Gamma}_1 \bar{p}} \frac{\partial \bar{p}}{\partial t}\right) + \partial_i \left(\frac{\zeta^2}{Dhu^0} \partial^i \phi\right). \tag{B.68}$$

B.5 Hydrostatic equilibrium

In order to derive suitable initial conditions for the base state pressure, we need to find the relativistic form of the equation for hydrostatic equilibrium. If we assume the initial metric and base state variables to be functions of r only, we can find this by deriving the TOV equation of the system (Thorne, 1977).

Taking the r-component of the equation for the conservation of the energy-momentum tensor, $\nabla_\mu T^{\mu\nu} = 0$, we find

$$0 = \partial_r T^{rr} + T^{rr} \Gamma^\nu{}_{r\nu} + T^{tt} \Gamma^r{}_{tt} + T^{rr} \Gamma^r{}_{rr} + T^{\theta\theta} \Gamma^r{}_{\theta\theta} + T^{\phi\phi} \Gamma^r{}_{\phi\phi}. \tag{B.69}$$

Let the metric take the form of a weak planar gravitational field, $g_{\mu\nu}(t,x,y,r) = \mathrm{diag}\left(-c^2\left(1 - \frac{2GM}{rc^2}\right), \left(1 + \frac{2GM}{rc^2}\right)(1,1,1)\right)$, where G is the gravitational constant and M the mass of the star. Letting $r = R + r'$, where R is the radius of the neutron star, r' the vertical distance from the surface and $r' \ll R$, expanding the g_{tt} component, we get

$$1 - \frac{2GM}{rc^2} = 1 - \frac{2GM}{(R+r')\,c^2} = 1 - \frac{2GM}{Rc^2\left(1 + \frac{r'}{R}\right)} \approx 1 - \left(\frac{2GM}{Rc^2}\left[1 - \frac{r'}{R}\right]\right). \tag{B.70}$$

Defining $g \equiv GM/R$, the corresponding non-zero Christoffel symbols are

$$\begin{aligned} \Gamma^t{}_{tr} = \Gamma^t{}_{rt} = \frac{g}{R\alpha^2}, \qquad \Gamma^r{}_{tt} = \frac{g\alpha^2}{c^2 R}, \qquad \Gamma^r{}_{xx} = \Gamma^r{}_{yy} = \frac{g}{R\alpha^2} \\ \Gamma^r{}_{rr} = \Gamma^x{}_{xr} = \Gamma^x{}_{rx} = \Gamma^y{}_{yr} = \Gamma^y{}_{ry} = -\frac{g}{R\alpha^2}. \end{aligned} \tag{B.71}$$

The non-zero components of the energy-momentum tensor for a perfect fluid and this metric are $T^{tt} = -\rho h c^2 g^{tt}$ and $T^{rr} = p g^{rr} = p\alpha^2/c^2$; substituting this into the r-th component of the equation for the conservation of energy-momentum, we obtain

$$0 = \partial_\nu T^{\nu r} + T^{\mu r} \Gamma^\nu{}_{\mu\nu} + T^{\nu\mu} \Gamma^r{}_{\mu\nu} = \frac{1}{c^2} \partial_r \left(p\alpha^2\right) + \frac{g}{Rc^2}\left(\rho h c^2 - 3p\right). \tag{B.72}$$

B.6 Newtonian limit

To check that the low Mach number equations derived above are of the same form as those in Almgren et al. (2006a) in the Newtonian limit, we assume the weak planar metric of the previous section and let $D \to \rho$, $U^i \to u^i$, $u_0 \to 1$ and $\alpha \to 1$.

Doing this for the continuity equation, we find

$$\partial_t \rho = -\partial_i \left(\rho u^i\right), \tag{B.73}$$

which is the Newtonian continuity equation, as desired. Looking next at the energy equation and using

$$U^\nu \Gamma^\mu_{\mu\nu} = \Gamma^t{}_{tr} U^r = \frac{g}{Rc^2\alpha^2} U^r \approx \frac{g}{Rc^2} u^r, \tag{B.74}$$

we get

$$\partial_t(\rho h) + \partial_i \left(u^i \rho h\right) = \frac{\mathrm{D}p}{\mathrm{D}t} - \frac{\rho h g}{Rc^2} u^r, \tag{B.75}$$

which again coincides with the Newtonian case. Looking at the momentum equation,

$$\partial_t U_j = -U^i \partial_i U_j - U_j \frac{\mathrm{D} \ln u^0}{\mathrm{D}t} + \Gamma_{\rho\nu j} U^\nu U^\rho - \frac{1}{Dhu^0} \left(\partial_j p_0 + \xi \partial_j \left[\frac{\pi}{\xi}\right] \right). \tag{B.76}$$

The last term will disappear, and the Christoffel term will reduce to $\Gamma^t{}_{tr} U_t U^t = -\frac{g}{Rc^2}$, such that we get

$$\partial_t u_j = -u^i \partial_i u_j - \frac{g}{Rc^2} - \frac{1}{\rho h} \left(\partial_j p_0 + \xi \partial_j \left[\frac{\pi}{\xi}\right] \right). \tag{B.77}$$

Using the equation for hydrostatic equilibrium, we can write $\partial_j p_0 = -\rho \partial_j \phi$, where ϕ is the scalar gravitational potential. In the low Mach number limit, $h \to 1$, so we again reclaim the Newtonian equation,

$$\partial_t u_j = -u^i \partial_i u_j - \frac{\xi}{\rho} \partial_j \left(\frac{\pi}{\xi}\right) - \frac{g}{Rc^2} + \partial_j \phi. \tag{B.78}$$

C Changes to the MAESTRO numerical method

In order to implement the relativistic equations, some changes had to be made to the numerical method used in MAESTRO. Here we shall describe in more detail the main departures from the time advancement algorithm described in Appendix A of Nonaka et al. (2010), and shall refer back to specific steps of their algorithm.

React State

react_state$[D^{\rm in}, (Dh)^{\rm in}, X^{\rm in}, T^{\rm in}, (\Gamma^{\mu}_{\mu\nu}U^{\nu}Dh)^{\rm in}, (\Gamma^{\mu}_{\mu\nu}U^{\nu}D)^{\rm in}, p_0^{\rm in}] \rightarrow [D^{\rm out}, (Dh)^{\rm out}, X^{\rm out}, T^{\rm out}, (D\dot{\omega})^{\rm out}]$ evolves the density and enthalpy through the half timestep $\Delta t/2$. We first evolve the species fraction and temperature according to

$$\frac{\mathrm{d}X}{\mathrm{d}t} = \dot{\omega}, \qquad \frac{\mathrm{d}T}{\mathrm{d}t} = \frac{1}{h_T}\left(-h_X\dot{\omega} + H\right). \tag{C.1}$$

We then define

$$(D\dot{\omega})^{\rm out} = \frac{D^{\rm out}\left(X^{\rm out} - X^{\rm in}\right)}{\Delta t/2}, \tag{C.2}$$

$$(Dh)^{\rm out} = (Dh)^{\rm in} + \frac{\Delta t}{2}\left(-\Gamma^{\mu}_{\mu\nu}U^{\nu}Dh\right)^{\rm in} + \frac{\Delta t}{2}\left(u^0\frac{\mathrm{D}p_0}{\mathrm{D}t}\right)^{\rm in} + \frac{\Delta t}{2}\left(DH_{\rm nuc}\right)^{\rm out}, \tag{C.3}$$

where we have treated the curvature and pressure terms like external heat sources. In our metric, if we assume that $\partial_t p_0$ is small, then $\frac{\mathrm{D}p_0}{\mathrm{D}t}$ is given by

$$\frac{\mathrm{D}p_0}{\mathrm{D}t} = \partial_t p_0 + U^r\partial_r p_0 \approx U^r\partial_r p_0. \tag{C.4}$$

The density also needs to be updated as, unlike in the MAESTRO case, our continuity equation has gravitational source terms:

$$D^{\rm out} = D^{\rm in} + \frac{\Delta t}{2}\left(-\Gamma^{\mu}_{\mu\nu}U^{\nu}D\right)^{\rm in}. \tag{C.5}$$

We update the temperature using the equation of state, $T^{\rm out} = T(D^{\rm out}, (Dh)^{\rm out}, X^{\rm out})$.

Enforce TOV

enforce_tov$[p_0^{\rm in}, (Dh)_0^{\rm in}] \rightarrow [p_0^{\rm out}]$ updates the base state pressure through Δt in

time. This replaces the **Enforce HSE** function in MAESTRO: instead we enforce that that base state pressure obeys the relativistic TOV equations by updating the base pressure as

$$p_{0,j+1}^{\text{out}} = p_{0,j}^{\text{in}} + \frac{\Delta r}{2}\left(\partial_r p_{0,j+1} + \partial_r p_{0,j}\right), \tag{C.6}$$

where $\partial_r p_0$ is found from the TOV equations.

Calculate u^0

calc_u0$[U^{i,\text{in}}] \rightarrow [u^{0,\text{out}}]$ calculates u^0 by first calculating the Lorentz factor from the spatial velocity:

$$W = \frac{1}{\sqrt{1 - g_{ij}U^iU^j}}, \qquad u^0 = \frac{W}{\alpha}, \tag{C.7}$$

where we have used the fact that in our metric the shift factor $\beta^i = 0$ so the metric component $g_{0i} = 0$.

Predict interface states

D_states$[D_{ij}^{\text{in}}, u_{ij}, v_{ij}] \rightarrow [D_{i+1/2,j}^{\text{out}}, D_{i,j+1/2}^{\text{out}}]$ predicts the cell-centred density to the edges. The equation for the density (without source terms) takes the conservative form

$$\partial_t D = -\partial_i(\rho U^i). \tag{C.8}$$

A piecewise linear prediction to the interface is given by

$$\begin{aligned} D_{i+1/2,j}^{n+1/2} &= D_{i,j}^n + \frac{\Delta x}{2}\left.\frac{\partial D}{\partial x}\right|_{i,j} + \frac{\Delta t}{2}\left.\frac{\partial D}{\partial t}\right|_{i,j} \\ &= D_{i,j}^n + \frac{\Delta x}{2}\left.\frac{\partial D}{\partial x}\right|_{i,j} + \frac{\Delta t}{2}\left(-\frac{\partial (DU^x)}{\partial x} - \frac{\partial (DU^y)}{\partial y}\right) \\ &= D_{i,j}^n + \frac{\Delta x}{2}\left(1 - \frac{\Delta t}{\Delta x}U^x\right)\frac{\partial D}{\partial x} - \underbrace{\frac{\Delta t}{2}D\frac{\partial U^x}{\partial x}}_{\text{non-advective}} - \underbrace{\frac{\Delta t}{2}\frac{\partial (DU^y)}{\partial y}}_{\text{transverse}}. \end{aligned}$$

psi_states$[\psi_{ij}^{\text{in}}, u_{ij}, v_{ij}] \rightarrow [\psi_{i+1/2,j}^{\text{out}}, \psi_{i,j+1/2}^{\text{out}}]$ predicts the cell-centred passive advective scalar to the edges. The equation for the scalar (without source terms) takes the advective form

$$\partial_t \psi = -U^i \partial_i \psi. \tag{C.9}$$

A piecewise linear prediction to the interface is given by

$$\begin{aligned}\psi^{n+1/2}_{i+1/2,j} &= \psi^n_{i,j} + \frac{\Delta x}{2}\left.\frac{\partial\psi}{\partial x}\right|_{i,j} + \frac{\Delta t}{2}\left.\frac{\partial\psi}{\partial t}\right|_{i,j} \\ &= \psi^n_{i,j} + \frac{\Delta x}{2}\left.\frac{\partial\psi}{\partial x}\right|_{i,j} + \frac{\Delta t}{2}\left(-U^x\frac{\partial\psi}{\partial x} - U^y\frac{\partial\psi}{\partial y}\right) \\ &= \psi^n_{i,j} + \frac{\Delta x}{2}\left(1 - \frac{\Delta t}{\Delta x}U^x\right)\frac{\partial\psi}{\partial x} - \underbrace{\frac{\Delta t}{2}U^y\frac{\partial\psi}{\partial y}}_{\text{transverse}}.\end{aligned}$$

Steps 2C & 6C

Compute provisional base state velocity forcing by discretising the base velocity component of the momentum equation,

$$\partial_t U_0 = -U_0\partial_r U_0 - U_0\frac{\mathrm{D}\ln u^0}{\mathrm{D}t} - \frac{1}{(Dh)_0 u^0}\left(\partial_j p_0 + \zeta\partial_j\left[\frac{\pi_0}{\zeta}\right]\right) + \Gamma_{\rho\nu t}U_0^2(e_r)^\rho(e_r)^\nu, \tag{C.10}$$

to get

$$\begin{aligned}\left[\frac{\zeta}{(Dh)_0 u^0}\partial_r\left(\frac{\pi_0}{\zeta}\right)\right]^{n,*} =& -\frac{U_0^{n+1/2,*} - U_0^{n-1/2}}{(\Delta t^n + \Delta t^{n-1})/2} - U_0^{n,*}\left[\partial_r U_0\right]^{n,*} \\ &- U_0^{n,*}\left[\frac{\mathrm{D}\ln u^0}{\mathrm{D}t}\right]^{n,*} - \frac{[\partial_r p_0]^{n,*}}{[(Dh)_0 u^0]^{n,*}} \\ &+ \left[\Gamma_{\rho\nu t}U_0^2(e_r)^\nu(e_r)^\rho\right]^{n,*},\end{aligned} \tag{C.11}$$

where we define

$$U_0^{n,*} = \frac{\Delta t^n U_0^{n-1/2} + \Delta t^{n-1}U_0^{n+1/2,*}}{\Delta t^n + \Delta t^{n-1}}, \tag{C.12}$$

$$\begin{aligned}&[\partial_r U_0]^{n,*} \\ &\quad= \frac{1}{\Delta t^n + \Delta t^{n-1}}\left(\Delta t^n\left[\partial_r U_0\right]^{n-1/2} + \Delta t^{n-1}\left[\partial_r U_0\right]^{n+1/2,*}\right),\end{aligned}$$

$$[\partial_r p_0]^{n,*} = \frac{1}{\Delta t^n + \Delta t^{n-1}}\left(\Delta t^n\left[\partial_r p_0\right]^{n-1/2} + \Delta t^{n-1}\left[\partial_r p_0\right]^{n+1/2,*}\right).$$

If $n = 0$, then set $\Delta t^{-1} = \Delta t^0$.

Steps 3 & 7 *Construct the provisional time-centred advective velocity on the edges,* $\tilde{U}^{i,ADV,*}$.

Use the equation for the velocity perturbation,

$$\partial_t\tilde{U}_j = -\tilde{U}^i\partial_i U_j - U^i\partial_i\tilde{U}_j - \tilde{U}_j\frac{\mathrm{D}\ln u^0}{\mathrm{D}t} - \frac{1}{(Dh)'u^0}\left(\partial_j p_0 + \zeta\partial_j\left[\frac{\pi}{\zeta}\right]\right) + 2\Gamma_{\rho\nu j}\tilde{U}^\nu U^\rho, \tag{C.13}$$

to compute the time-centred edge velocities, $\tilde{U}^{i,\mathrm{ADV},\dagger,*}$, using $U^i = \tilde{U}^{i,n} + U_0^{n+1/2,*}$, and where † refers to the fact that the predicted velocity field does not satisfy the divergence constraint. $\tilde{U}^{i,\mathrm{ADV},*}$ is then constructed from $\tilde{U}^{i,\mathrm{ADV},\dagger,*}$ using a MAC projection.

Begin by upwinding $\tilde{U}^{i,n}$ to get $\tilde{U}^{i,\mathrm{ADV},\dagger,*}$ at each edge. Then enforce the divergence constraint by solving

$$\mathrm{D}^{\mathrm{MAC}}\left(\frac{(\zeta^n)^2}{Dh^n(u^0)^n}G^{i,\mathrm{MAC}}\phi^{\mathrm{MAC},*}\right) = \mathrm{D}^{\mathrm{MAC}}\left(\zeta^n\tilde{U}^{i,\mathrm{ADV},\dagger,*}\right) - \zeta^n\left(S^{n+1/2,*} - \bar{S}^{n+1/2,*}\right) \tag{C.14}$$

for ϕ, where $\mathrm{D}^{\mathrm{MAC}}$ represents a centred approximation to a cell-based divergence from edge-based velocities and $G^{i,\mathrm{MAC}}$ represents a centred approximation to edge-based gradients from cell-centred data. This is solved using multigrids. The provisional advective velocity is then

$$\tilde{U}^{i,\mathrm{ADV},*} = \tilde{U}^{i,\mathrm{ADV},\dagger,*} - \frac{1}{Dh^n(u^0)^n}G^{i,\mathrm{MAC}}\phi^{\mathrm{MAC},*}, \tag{C.15}$$

which satisfies

$$\mathrm{D}^{\mathrm{MAC}}\left(\zeta^n\tilde{U}^{i,\mathrm{ADV},*}\right) = \zeta^n\left(S^{n+1/2,*} - \bar{S}^{n+1/2,*}\right). \tag{C.16}$$

Steps 4H & 8H

Update the enthalpy using a discretised version of (B.43), omitting the source terms to get

$$\partial_t(Dh) = -\partial_i\left(U^i Dh\right) + u^0\psi + u^0\tilde{U}^r\partial_r p_0. \tag{C.17}$$

i Compute the time-centred enthalpy edge state, $(Dh)^{n+1/2,*,\mathrm{pred}}$, for the conservative update of $(Dh)^{(1)}$. Using (B.47) and ignoring the source terms,

$$\partial_t(Dh)' = -\partial_i\left(U^i(Dh)' + \tilde{U}^i(Dh)_0\right) + u^0\tilde{U}^r\partial_r p_0 \tag{C.18}$$

to predict $(Dh)' = (Dh)^{(1)} - (Dh)_0^n$ to the time-centred edges, using $U^i = \tilde{U}^{i,\mathrm{ADV},*} + U_0^{n+1/2,*}$ to obtain $(Dh)'^{n+1/2,*,\mathrm{pred}}$. We convert this to the full state using

$$(Dh)^{n+1/2,*,\mathrm{pred}} = (Dh)'^{n+1/2,*,\mathrm{pred}} + \frac{(Dh)_0^n + (Dh)_0^{n+1,*}}{2}. \tag{C.19}$$

ii Evolve $(Dh)^{(1)} \rightarrow (Dh)^{(2),*}$.

$$\begin{aligned}(Dh)^{(2),*} =& (Dh)^{(1)} \\ &- \Delta t\left\{\partial_i\left[\left(\tilde{U}^{i,\mathrm{ADV},*} + U_0^{n+1/2,*}(e_r)^i\right)(Dh)^{n+1/2,*,\mathrm{pred}}\right]\right\} \\ &+ \Delta t u^0\tilde{U}^{r,\mathrm{ADV},*}\partial_r p_0^n + \Delta t u^0\psi^{n+1/2,*}\end{aligned} \tag{C.20}$$

For each Cartesian cell where $D^{(2)} < D_{\text{cutoff}}$, recompute enthalpy using

$$(Dh)^{(2),*} = D^{(2),*} h\left(D^{(2),*}, p_0^{n+1,*}\right). \tag{C.21}$$

Step 11 *Update velocity.*

First compute the time-centred edge velocities, $\tilde{U}^{i,n+1/2,\text{pred}}$, then define

$$D^{n+1/2} = \frac{1}{2}(D^n + D^{n+1}), \qquad D_0^{n+1/2} = \frac{1}{2}(D_0^n + D_0^{n+1}). \tag{C.22}$$

Update the velocity field $\tilde{U}^{i,n}$ to $\tilde{U}^{i,n+1,\dagger}$ by discretising (C.13) as

$$\begin{aligned}\tilde{U}^{i,n+1,\dagger} = \tilde{U}^{i,n} &- \Delta t\left[\left(\tilde{U}^{j,\text{ADV}} + U_0^j\right)\partial_j\left(\tilde{U}^{i,n+1/2,\text{pred}}\right)\right] \\ &- \Delta t \tilde{U}^{r,\text{ADV}}\left[\partial_r\left(U_0\right)\right]^{n+1/2} \\ &+ \Delta t\Bigg[-g^{ij}\frac{\zeta^{n+1/2}}{(Dh)^{n+1/2}u^0}G_j\left(\frac{\pi}{\zeta}\right)^{n-1/2} + \left(\frac{\zeta}{(Dh)_0 u^0}\partial_r\left[\frac{\pi}{\zeta}\right]\right)^n (e_r)^i \\ &+ \left\{2g^{ij}\tilde{U}^\nu U^\rho \Gamma_{\rho\nu j}\right\}^{n+1/2}\Bigg],\end{aligned} \tag{C.23}$$

where G_j approximates a cell-centred gradient from edge-centred data. Finally, use an approximate edge-centred projection to define $\tilde{U}^{i,n+1}$ from $\tilde{U}^{i,n+1,\dagger}$ such that $\tilde{U}^{i,n+1}$ satisfies the divergence constraint.

D Derivation of the analytic solution of the Newtonian dam break problem

In this section, we shall summarise the derivation of the analytic solution of the Newtonian dam break problem given by Stoker (1957). The solution at time $t > 0$ consists of four distinct regions, as illustrated in Figure D.1: undisturbed water (1), a rarefaction fan (3) and a constant wave region with constant water depth (2) which is separated by a shock from another region of undisturbed water (0).

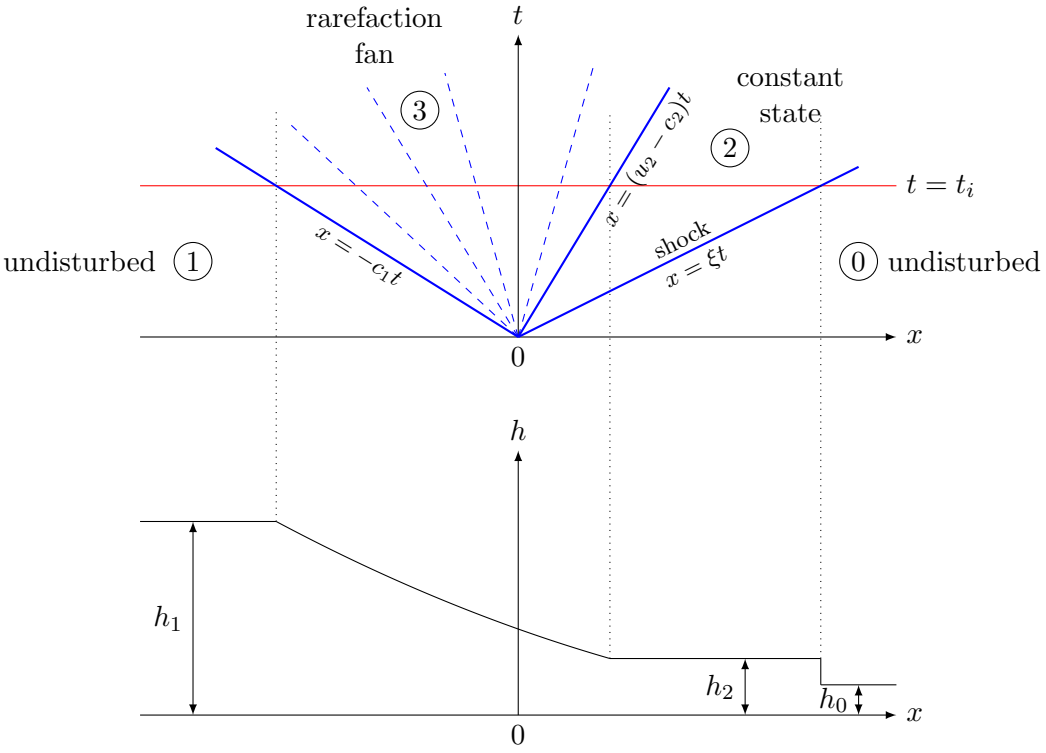

FIGURE D.1: Solution of the dam break problem, adapted from (Stoker, 1957). The top plot shows the characteristic wave structure of the solution of the Riemann problem, the lower plot shows the fluid height at time $t_i > 0$. We can see that the solution is much like that of the Sod shock tube test seen in Section 5.5.1.1.

The characteristics separating the regions are as follows:

- $x = -c_1 t$ between regions 1 and 3, where $c_1 = \sqrt{g h_1}$ is the sound speed in region 1,
- $x = (u_2 - c_2)t$ between regions 3 and 2,
- $x = \xi t$ between regions 2 and 0, where ξ is the shock speed.

To find the height profile across the various regions, we first need to find the height in the constant wave region, h_2. An expression for this depending only on h_0 and h_1 was found by Wu et al. (1999) to be

$$\left(\frac{h_2}{h_1}\right)^3 - 9\left(\frac{h_2}{h_1}\right)^2\left(\frac{h_0}{h_1}\right) + 16\left(\frac{h_2}{h_1}\right)^{3/2}\left(\frac{h_0}{h_1}\right) - \left(\frac{h_2}{h_1}\right)\left(\frac{h_0}{h_1}\right)\left(\frac{h_0}{h_1} + 8\right) + \left(\frac{h_0}{h_1}\right)^3 = 0. \tag{D.1}$$

From this, the sound speed in region 2 can be found using

$$\frac{c_2^2}{c_0^2} = \frac{h_2}{h_0}, \tag{D.2}$$

and the fluid speed calculated using the following relation found by considering the rarefaction region,

$$u_2 + 2c_2 = 2c_1. \tag{D.3}$$

This allows us to calculate the shock speed ξ as

$$\xi = c_0 \left\{ \frac{1}{8}\left[\left(2\left(\frac{c_2}{c_0}\right)^2 + 1 \right)^2 - 1 \right] \right\}^{1/2}. \tag{D.4}$$

Finally, we can calculate the fluid height in the rarefaction region, h_3. Using the relation

$$u_3 + 2c_3 = 2c_1, \tag{D.5}$$

and the formula for the characteristic

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \frac{x}{t} = u_3 - c_3 = 2c_1 - 3c_3 = \frac{3}{2}u_3 - c_1, \tag{D.6}$$

we find that

$$c_3 = \frac{1}{3}\left(2c_1 - \frac{x}{t}\right), \tag{D.7}$$

finally giving us the fluid height in the rarefaction region

$$h_3 = \frac{c_3^2}{g}. \tag{D.8}$$

References

D. Abarca and W. Kluźniak. Radial oscillations of a radiation-supported levitating shell in Eddington luminosity neutron stars. *Monthly Notices of the Royal Astronomical Society*, 461(3):3233–3238, Sep 2016. ISSN 0035-8711. DOI 10.1093/mnras/stw1432.

M. A. Abramowicz, W. Kluźniak, and J. P. Lasota. The centrifugal force reversal and X-ray bursts. *Astronomy & Astrophysics*, 374:L16, May 2001. DOI 10.1051/0004-6361:20010835.

F. Alcrudo and P. Garcia Navarro. A high resolution Godunov type scheme in finite volumes for the 2D shallow water equations. *International Journal For Numerical Methods in Fluids*, 16(June 1992):489–505, 1993. ISSN 0271-2091. DOI 10.1002/fld.1650160604.

M. Alcubierre. *Introduction to 3+1 Numerical Relativity*. Oxford University Press, Apr 2008. ISBN 9780199205677. DOI 10.1093/acprof:oso/9780199205677.001.0001.

A. S. Almgren, J. B. Bell, C. a. Rendleman, and M. Zingale. Low Mach Number Modeling of Type Ia Supernovae. I. Hydrodynamics. *The Astrophysical Journal*, 637 (2):922–936, Feb 2006a. ISSN 0004-637X. DOI 10.1086/498426.

A. S. Almgren, J. B. Bell, C. A. Rendleman, and M. Zingale. Low Mach Number Modeling of Type Ia Supernovae. II. Energy Evolution. *The Astrophysical Journal*, 637(2):927–938, Jun 2006b. ISSN 0004-637X. DOI 10.1086/507089.

A. S. Almgren, J. B. Bell, and M. Zingale. MAESTRO: A Low Mach Number Stellar Hydrodynamics Code. *Journal of Physics: Conference Series*, 78(1):012085, Jul 2007. ISSN 1742-6596. DOI 10.1088/1742-6596/78/1/012085.

A. S. Almgren, J. B. Bell, A. Nonaka, and M. Zingale. Low Mach Number Modeling of Type Ia Supernovae. III. Reactions. *The Astrophysical Journal*, 684(1):449–470, Sep 2008. ISSN 0004-637X. DOI 10.1086/590321.

A. S. Almgren, V. E. Beckner, J. B. Bell, M. S. Day, L. H. Howell, C. C. Joggerst, M. J. Lijewski, A. Nonaka, M. Singer, and M. Zingale. Castro: a New Compressible Astrophysical Solver. I. Hydrodynamics and Self-Gravity. *The Astrophysical Journal*, 715(2):1221–1238, 2010. ISSN 0004-637X. DOI 10.1088/0004-637X/715/2/1221.

N. Andersson and G. L. Comer. Relativistic Fluid Dynamics: Physics for Many Different Scales. *Living Reviews in Relativity*, 10, 2007. ISSN 1433-8351. DOI 10.12942/lrr-2007-1.

N. Andersson and C. S. Lopez-Monsalvo. A consistent first-order model for relativistic heat flow. *Classical and Quantum Gravity*, 28(19):195023, Oct 2011. ISSN 0264-9381. DOI 10.1088/0264-9381/28/19/195023.

M. Andrew and J. Sethian. The Derivation and Numerical Solution of the Equations for Zero Mach Number Combustion. *Combustion Science and Technology*, 42(3-4): 185–205, Jan 1985. ISSN 0010-2202. DOI 10.1080/00102208508960376.

J. Antoniadis, P. C. C. Freire, N. Wex, T. M. Tauris, R. S. Lynch, M. H. van Kerkwijk, M. Kramer, C. Bassa, V. S. Dhillon, T. Driebe, et al. A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340(6131):1233232–1233232, Apr 2013. ISSN 0036-8075. DOI 10.1126/science.1233232.

B. O. Arani, C. E. Frouzakis, J. Mantzaras, F. Lucci, and K. Boulouchos. Direct numerical simulation of turbulent channel-flow catalytic combustion: Effects of Reynolds number and catalytic reactivity. *Combustion and Flame*, 187:52–66, 2018. ISSN 15562921. DOI 10.1016/j.combustflame.2017.09.001.

W. D. Arnett, C. Meakin, M. Viallet, S. W. Campbell, J. C. Lattanzio, and M. Mocák. Beyond Mixing Length Theory: A Step Toward 321D. *The Astrophysical Journal*, 809(1):30, Aug 2015. ISSN 1538-4357. DOI 10.1088/0004-637X/809/1/30.

R. Arnowitt, S. Deser, and C. W. Misner. The dynamics of general relativity. In L. Witten, editor, *Gravitation: an introduction to current research*, chapter 7, pages 227–264. Wiley, New York, 1962. ISBN 0013-0133. DOI 10.1007/s10714-008-0661-1.

A. J. Aspden, J. B. Bell, M. S. Day, S. E. Woosley, and M. Zingale. Turbulence-Flame Interactions in Type Ia Supernovae. *The Astrophysical Journal*, 689:1173–1185, 2008. ISSN 0004-637X. DOI 10.1086/592726.

A. J. Aspden, J. B. Bell, and S. E. Woosley. Distributed Flames in Type Ia Supernovae. *The Astrophysical Journal*, 710:1654–1663, 2010. ISSN 0004-637X. DOI 10.1088/0004-637X/710/2/1654.

S. Ayasli and P. C. Joss. Thermonuclear processes on accreting neutron stars: a systematic study. *The Astrophysical Journal*, 256:637–665, 1981. ISSN 0004-637X. DOI 10.1086/159940.

L. Baiotti, I. Hawke, P. J. Montero, F. Löffler, L. Rezzolla, N. Stergioulas, J. A. Font, and E. Seidel. Three-dimensional relativistic simulations of rotating neutron-star collapse to a Kerr black hole. *Physical Review D*, 71(2), 2005. ISSN 15507998. DOI 10.1103/PhysRevD.71.024035.

F. Banyuls, J. a. Font, J. M. Ibanez, J. M. Marti, and J. a. Miralles. Numerical {3 + 1} General Relativistic Hydrodynamics: A Local Characteristic Approach. *The Astrophysical Journal*, 476:221–231, 1997. ISSN 0004-637X. DOI 10.1086/303604.

G. K. Batchelor. The conditions for dynamical similarity of motions of a frictionless perfect-gas atmosphere. *Quarterly Journal of the Royal Meteorological Society*, 29: 224–235, 1953. ISSN 00359009. DOI 10.1002/qj.49707934004.

G. K. Batchelor and A. A. Townsend. The Nature of Turbulent Motion at Large Wave-Numbers. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 199(1057):238–255, 1949. ISSN 1364-5021. DOI 10.1098/rspa.1949.0136.

M. Bauböck, D. Psaltis, and F. Özel. Narrow Atomic Features From Rapidly Spinning Neutron Stars. *The Astrophysical Journal*, 766(2):87, 2013. ISSN 0004-637X. DOI 10.1088/0004-637X/766/2/87.

G. Baym and C. Pethick. Physics of Neutron Stars. *Annual Review of Astronomy and Astrophysics*, 17(1):415–443, Sep 1979. ISSN 0066-4146. DOI 10.1146/annurev.aa.17.090179.002215.

A. Beccantini and E. Studer. The reactive Riemann problem for thermally perfect gases at all combustion regimes. *International Journal For Numerical Methods in Fluids*, 64(3):269–313, Sep 2010. ISSN 02712091. DOI 10.1002/fld.

B. Becker, B. Guo, K. Bandilla, M. A. Celia, B. Flemisch, and R. Helmig. An adaptive multiphysics model coupling vertical equilibrium and full multidimensions for multiphase flow in porous media. pages 1–20, Nov 2017. URL `http://arxiv.org/abs/1711.09614`.

J. Bell and M. Day. Adaptive Methods for Simulation of Turbulent Combustion. In T. Echekki and E. Mastorakos, editors, *Turbulent Combustion Modeling*, number d in Fluid Mechanics and Its Applications, pages 301–329. Springer, Dordrecht, 2011. ISBN 9789400704114. DOI 10.1007/978-94-007-0412-1.

J. Bell, M. Berger, J. Saltzman, and M. Welcome. Three-Dimensional Adaptive Mesh Refinement for Hyperbolic Conservation Laws. *SIAM Journal on Scientific Computing*, 15(1):127–138, Jan 1994. ISSN 1064-8275. DOI 10.1137/0915008.

J. Bell, M. Day, C. Rendleman, S. Woosley, and M. Zingale. Adaptive low Mach number simulations of nuclear flame microphysics. *Journal of Computational Physics*, 195(2): 677–694, Apr 2004. ISSN 00219991. DOI 10.1016/j.jcp.2003.10.035.

R. Belvedere, J. A. Rueda, and R. Ruffini. On the Magnetic Field of Pulsars with Realistic Neutron Star Configurations. *The Astrophysical Journal*, 799(1):23, Jan 2015. ISSN 1538-4357. DOI 10.1088/0004-637X/799/1/23.

M. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:64–84, 1989. ISSN 00219991. DOI 10.1016/0021-9991(89)90035-1.

M. J. Berger and J. Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53(3):484–512, 1984. ISSN 10902716. DOI 10.1016/0021-9991(84)90073-1.

R. G. Berkhout and Y. Levin. Evolution of the bursting-layer wave during a type I X-ray burst. *Monthly Notices of the Royal Astronomical Society*, 385(2):1029–1035, Apr 2008. ISSN 0035-8711. DOI 10.1111/j.1365-2966.2008.12897.x.

H. Bethe, G. Börner, and K. Sato. Nuclei in Neutron Matter. *Astronomy and Astrophysics*, 7:279, 1970. URL http://adsabs.harvard.edu/abs/1970A&A.....7..279B.

S. Bhattacharyya and T. E. Strohmayer. Evidence for Harmonic Content and Frequency Evolution of Oscillations during the Rising Phase of X-ray Bursts from 4U 1636-536. *The Astrophysical Journal*, 634:L157–L160, 2005. ISSN 0004-637X. DOI 10.1086/499100.

S. Bhattacharyya and T. E. Strohmayer. A Non-PRE Double-peaked Burst from 4U 1636-536: Evidence of Burning Front Propagation. *The Astrophysical Journal*, 636 (2):L121–L124, Jan 2006a. ISSN 0004-637X. DOI 10.1086/500199.

S. Bhattacharyya and T. E. Strohmayer. Signature of Temporary Burning Front Stalling from a Non-Photospheric Radius Expansion Double-peaked Burst. *The Astrophysical Journal*, 641(1):L53–L56, Apr 2006b. ISSN 0004-637X. DOI 10.1086/503768.

S. Bhattacharyya and T. E. Strohmayer. Thermonuclear Flame Spreading on Rapidly Spinning Neutron Stars: Indications of the Coriolis Force? *The Astrophysical Journal*, 666(2):L85–L88, Sep 2007. ISSN 0004-637X. DOI 10.1086/521790.

S. Bhattacharyya, T. E. Strohmayer, M. C. Miller, and C. B. Markwardt. Constraints on Neutron Star Parameters from Burst Oscillation Light Curves of the Accreting Millisecond Pulsar XTE J1814338. *The Astrophysical Journal*, 619(1):483–491, Jan 2005. ISSN 0004-637X. DOI 10.1086/426383.

S. Bhattacharyya, M. C. Miller, and F. K. Lamb. The Shapes of Atomic Lines from the Surfaces of Weakly Magnetic Rotating Neutron Stars and Their Implications. *The Astrophysical Journal*, 644(2):1085–1089, 2006. ISSN 0004-637X. DOI 10.1086/503860.

L. Bildsten and L. Hall. Thermonuclear Burning on Rapidly Accreting Neutron Stars. In R. Buccheri, J. van Paradijs, and A. Alpar, editors, *NATO Advanced Science Institutes (ASI) Series C*, volume 515, page 419. 1998. ISBN 0792351940. URL http://arxiv.org/abs/astro-ph/9709094.

J. M. Blondin, A. Mezzacappa, and C. DeMarino. Stability of Standing Accretion Shocks, with an Eye toward CoreCollapse Supernovae. *The Astrophysical Journal*, 584(2):971–980, 2003. ISSN 0004-637X. DOI 10.1086/345812.

M. Bocquet, S. Bonazzola, E. Gourgoulhon, and J. Novak. Rotating Neutron Star Models With Magnetic Field. *Astronomy and Astrophysics*, 301:757, 1995. URL http://arxiv.org/abs/gr-qc/9503044.

S. Bonazzola, L. Villain, and M. Bejger. Magnetohydrodynamics of rotating compact stars with spectral methods: description of the algorithm and tests. *Classical and Quantum Gravity*, 24(12):S221–S234, Jun 2007. ISSN 0264-9381. DOI 10.1088/0264-9381/24/12/S15.

F. Bouchot, V. Zeitlin, F. Bouchut, and V. Zeitlin. A robust well-balanced scheme for multi-layer shallow water equations. *Discrete and Continuous Dynamical Systems - Series B*, 13(4):739–758, Mar 2010. ISSN 1531-3492. DOI 10.3934/dcdsb.2010.13.739.

J. Boussinesq. *Théorie analytique de la chaleur*. Gauthier-Villars, Paris, 1903.

R. P. Brent. An Algorithm with Guaranteed Convergence for Finding a Zero of a Function. In *Algorithms for minimization without derivatives*, chapter 4. Prentice-Hall, Englewood Cliffs, NJ, 1973. ISBN 0-13-022335-2.

E. F. Brown. Superburst Ignition and Implications for Neutron Star Interiors. *The Astrophysical Journal*, 614(1):4, 2004. ISSN 0004-637X. DOI 10.1086/425562.

E. F. Brown, L. Bildsten, and P. Chang. Variability in the Thermal Emission from Accreting Neutron Star Transients. *The Astrophysical Journal*, 574(2):920–929, Aug 2002. ISSN 0004-637X. DOI 10.1086/341066.

G. L. Bryan, M. L. Norman, B. W. O'Shea, T. Abel, J. H. Wise, M. J. Turk, D. R. Reynolds, D. C. Collins, P. Wang, S. W. Skillman, et al. Enzo: An Adaptive Mesh Refinement Code for Astrophysics. *The Astrophysical Journal Supplement Series*, 211 (2):19, Mar 2014. ISSN 0067-0049. DOI 10.1088/0067-0049/211/2/19.

C. Cadeau, S. M. Morsink, D. Leahy, and S. S. Campbell. Light Curves for Rapidly Rotating Neutron Stars. *The Astrophysical Journal*, 654(1):458–469, Jan 2007. ISSN 0004-637X. DOI 10.1086/509103.

M. A. Calkins, K. Julien, and P. Marti. The breakdown of the anelastic approximation in rotating compressible convection: implications for astrophysical systems. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2175): 20140689–20140689, Feb 2015. ISSN 1364-5021. DOI 10.1098/rspa.2014.0689.

B. W. Carroll and D. A. Ostlie. *An Introduction to Modern Astrophysics*. Addison-Wesley, 1996. ISBN 0-201-54730-9.

J. Casares, P. G. Jonker, and G. Israelian. X-Ray Binaries. In *Handbook of Supernovae*, pages 1–28. Springer International Publishing, Cham, 2016. ISBN 9783319218465. DOI 10.1007/978-3-319-20794-0˙111-1.

Y. Cavecchi. *Accreting Neutron Stars : Strong Gravity and Type I Bursts*. PhD thesis, University of Amsterdam, 2013. URL http://hdl.handle.net/11245/1.378880.

Y. Cavecchi, A. Patruno, B. Haskell, A. L. Watts, Y. Levin, M. Linares, D. Altamirano, R. Wijnands, and M. Van Der Klis. Implications of burst oscillations from the slowly rotating accreting pulsar IGR J17480-2446 in the globular cluster terzan 5. *Astrophysical Journal Letters*, 740(1), 2011. ISSN 20418205. DOI 10.1088/2041-8205/740/1/L8.

Y. Cavecchi, A. L. Watts, Y. Levin, and J. Braithwaite. Rotational effects in thermonuclear type I bursts: equatorial crossing and directionality of flame spreading. *Monthly Notices of the Royal Astronomical Society*, 448(1):445–455, Mar 2015. ISSN 1365-2966. DOI 10.1093/mnras/stu2764.

Y. Cavecchi, Y. Levin, A. L. Watts, and J. Braithwaite. Fast and slow magnetic deflagration fronts in Type I X-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 459(2):1259–1275, Sep 2016. DOI 10.1093/mnras/stw728.

Y. Cavecchi, A. L. Watts, and D. K. Galloway. On the Dependence of the X-Ray Burst Rate on Accretion and Spin Rate. *The Astrophysical Journal*, 851(1):1, Dec 2017. ISSN 1538-4357. DOI 10.3847/1538-4357/aa9897.

J. Centrella and J. R. Wilson. Planar numerical cosmology. II - The difference equations and numerical tests. *The Astrophysical Journal Supplement Series*, 54:229, 1984. ISSN 0067-0049. DOI 10.1086/190927.

D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. *Nature*, 424(6944):42–44, Jul 2003. ISSN 00280836. DOI 10.1038/nature01732.

M. Chakraborty, Y. E. Bahar, and E. Göü. Time- and Energy-dependent Characteristics of Thermonuclear Burst Oscillations. *The Astrophysical Journal*, 851(2):79, Dec 2017. ISSN 1538-4357. DOI 10.3847/1538-4357/aa984e.

N. Chamel and P. Haensel. Physics of Neutron Star Crusts. *Living Reviews in Relativity*, 11(1):10, Dec 2008. ISSN 2367-3613. DOI 10.12942/lrr-2008-10.

S. Chandrasekhar. 101. The case of two uniform fluids in relative horizontal motion separated by a horizontal boundary. In *Hydrodynamic and Hydromagnetic Stability*, chapter XI, pages 483–486. Dover Publications, New York, third edition, 1961a.

S. Chandrasekhar. X. The Stability of Superposed Fluids: The Rayleigh-Taylor Instability. In *Hydrodynamic and Hydromagnetic Stability*, pages 428–480. Dover Publications, New York, third edition, 1961b.

S. Chandrasekhar and E. Fermi. Problems of Gravitational Stability in the Presence of a Magnetic Field. *The Astrophysical Journal*, 118:116, Jul 1953. ISSN 0004-637X. DOI 10.1086/145732.

J. G. Charney and Y. Ogura. A Numerical Model for Thermal Convection in the Atmosphere. *Journal of the Meteorological Society of Japan. Ser. II*, 38(6):19a, 1960. ISSN 0026-1165. URL https://www.jstage.jst.go.jp/article/jmsj1923/38/6/38_6_19a/_article.

G. W. Clark, J. G. Jernigan, H. Bradt, C. Canizares, W. H. G. Lewin, F. K. Li, W. Mayer, J. McClintock, and H. Schnopper. Recurrent brief X-ray bursts from the globular cluster NGC 6624. *The Astrophysical Journal*, 207:L105, Jul 1976. ISSN 0004-637X. DOI 10.1086/182190.

P. Colella, D. Graves, T. Ligocki, D. Martin, D. Modiano, D. Serafini, and B. Van Straalen. Chombo Software Package for AMR Applications (Design Document). Technical report, 2013. URL http://seesar.lbl.gov/anag/chombo/ChomboDesign-3.1.pdf.

R. L. Cooper and R. Narayan. The Latitude of Type I X-Ray Burst Ignition on Rapidly Rotating Neutron Stars. *The Astrophysical Journal*, 657(1):L29–L32, Mar 2007. ISSN 0004-637X. DOI 10.1086/513077.

R. Cornelisse, J. Heise, E. Kuulkers, F. Verbunt, and J. J. M. i. t. Zand. The longest thermonuclear X-ray burst ever observed? (A BeppoSAX Wide Field Camera observation of 4U 1735-44). *Astronomy & Astrophysics*, 357:L21–L24, Mar 2000. ISSN 00046361. URL http://arxiv.org/abs/astro-ph/0003454.

R. Cornelisse, J. J. M. in 't Zand, F. Verbunt, E. Kuulkers, J. Heise, P. R. den Hartog, M. Cocchi, L. Natalucci, A. Bazzano, P. Ubertini, et al. Six years of BeppoSAX Wide Field Cameras observations of nine galactic type I X-ray bursters. *Astronomy & Astrophysics*, 405(3):1033–1042, Jul 2003. ISSN 0004-6361. DOI 10.1051/0004-6361:20030629.

R. Courant, K. Friedrichs, and H. Lewy. On the Partial Difference Equations of Mathematical Physics. *IBM Journal of Research and Development*, 11(2):215–234, 1967. ISSN 0018-8646. DOI 10.1147/rd.112.0215.

T. G. Cowling. The Non-radial Oscillations of Polytropic Stars. *Monthly Notices of the Royal Astronomical Society*, 101(8):367–375, Dec 1941. ISSN 0035-8711. DOI 10.1093/mnras/101.8.367.

A. Cumming. *Non-radial oscillations of accreting neutron stars and rotational evolution during type I X-ray bursts*. PhD thesis, University of California, Berkeley, 2000. URL http://link.springer.com/chapter/10.1007/978-94-009-1263-2_2.

A. Cumming. Latitudinal Shear Instabilities during Type I X-Ray Bursts. *The Astrophysical Journal*, 630(1):441–453, Sep 2005. ISSN 0004-637X. DOI 10.1086/431731.

A. Cumming and L. Bildsten. Rotational Evolution during Type I XRay Bursts. *The Astrophysical Journal*, 544(1):453–474, Nov 2000. ISSN 0004-637X. DOI 10.1086/317191.

A. Cumming, E. G. Zweibel, and L. Bildsten. Magnetic screening in accreting neutron stars. *The Astrophysical Journal*, 557(2):958–966, 2001. ISSN 0004-637X. DOI 10.1086/321658.

A. Cumming, S. M. Morsink, L. Bildsten, J. L. Friedman, and D. E. Holz. Hydrostatic expansion and spin changes during type I X-Ray bursts. *The Astrophysical Journal*, 564:343–352, 2002. ISSN 0004-637X. DOI 10.1086/324157.

A. Cumming, J. Macbeth, J. J. M. i. T. Zand, and D. Page. Long Type I X-Ray Bursts and Neutron Star Interior Physics. *The Astrophysical Journal*, 646(1):429, 2006. ISSN 0004-637X. DOI 10.1086/504698.

M. S. Day and J. B. Bell. Numerical simulation of laminar reacting flows with complex chemistry. *Combustion Theory and Modelling*, 4(4):535–556, Dec 2000. ISSN 1364-7830. DOI 10.1088/1364-7830/4/4/309.

A. Deibel, Z. Meisel, H. Schatz, E. F. Brown, and A. Cumming. Urca Cooling Pairs in the Neutron Star Ocean and their Effect on Superbursts. *The Astrophysical Journal*, 831(1):13, Oct 2016. ISSN 1538-4357. DOI 10.3847/0004-637X/831/1/13.

A. Dubey, A. Almgren, J. Bell, M. Berzins, S. Brandt, G. Bryan, P. Colella, D. Graves, M. Lijewski, F. Löffler, et al. A survey of high level frameworks in block-structured adaptive mesh refinement packages. *Journal of Parallel and Distributed Computing*, 74(12):3217–3227, 2014. ISSN 07437315. DOI 10.1016/j.jpdc.2014.07.001.

D. R. Durran. Improving the Anelastic Approximation. *Journal of the Atmospheric Sciences*, 46(11):1453–1461, Jun 1989. ISSN 0022-4928. DOI 10.1175/1520-0469(1989)046¡1453:ITAA¿2.0.CO;2.

Einstein Toolkit. Einstein Toolkit, 2010. URL `http://www.einsteintoolkit.org`.

P. Esposito, N. Rea, and G. L. Israel. Magnetars: a short review and some sparse considerations. pages 1–49, 2018. DOI arXiv:1803.05716v1.

G. L. Eyink and T. D. Drivas. Cascades and Dissipative Anomalies in Relativistic Fluid Turbulence. *Physical Review X*, 8(1):011023, Feb 2018. ISSN 2160-3308. DOI 10.1103/PhysRevX.8.011023.

R. P. Fedkiw, T. Aslam, B. Merriman, and S. Osher. A Non-oscillatory Eulerian Approach to Interfaces in Multimaterial Flows (the Ghost Fluid Method). *Journal of Computational Physics*, 152(2):457–492, 1999. ISSN 00219991. DOI 10.1006/jcph.1999.6236.

J. L. Fisker, F.-K. Thielemann, and M. Wiescher. The nuclear reaction waiting points, Mg22, Si26, S30, and Ar34, and bolometrically double peaked type I X-ray bursts. *The Astrophysical Journal*, 608(1982):5, 2003. ISSN 15384357. DOI 10.1086/422215.

J. A. Font. Numerical Hydrodynamics and Magnetohydrodynamics in General Relativity. *Living Reviews in Relativity*, 11, 2008. ISSN 1433-8351. DOI 10.12942/lrr-2008-7.

B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo. FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *The Astrophysical Journal Supplement Series*, 131(1):273–334, 2000. ISSN 0067-0049. DOI 10.1086/317361.

M. Y. Fujimoto, T. Hanawa, and S. Miyaji. Shell flashes on accreting neutron stars and X-ray bursts, 1981. ISSN 0004-637X. DOI 10.1086/159034.

I. Fushiki and D. Q. Lamb. New insights from a global view of X-ray bursts. *The Astrophysical Journal*, 323:L55, Dec 1987. ISSN 0004-637X. DOI 10.1086/185056.

D. K. Galloway and L. Keek. Thermonuclear X-ray bursts. pages 1–57, Dec 2017. URL http://arxiv.org/abs/1712.06227.

D. K. Galloway, M. P. Muno, J. M. Hartman, D. Psaltis, and D. Chakrabarty. Thermonuclear (Type I) XRay Bursts Observed by the Rossi XRay Timing Explorer. *The Astrophysical Journal Supplement Series*, 179(2):360–422, Dec 2008. ISSN 0067-0049. DOI 10.1086/592044.

D. K. Galloway, Z. Johnston, A. Goodwin, and A. Heger. High-energy transients: thermonuclear (type-I) X-ray bursts. In *IAU Symposium 339, "Southern Horizons in Time Domain Astronomy"*, number 339, pages 1–6, Stellenbosch, South Africa, Mar 2018. URL http://arxiv.org/abs/1803.00223.

Y. Gao and C. K. Law. Rankine-Hugoniot relations in relativistic combustion waves. *The Astrophysical Journal*, 760(2):1–33, 2012. ISSN 15384357. DOI 10.1088/0004-637X/760/2/122.

A. L. Garcia, J. B. Bell, W. Y. Crutchfield, and B. J. Alder. Adaptive Mesh and Algorithm Refinement Using Direct Simulation Monte Carlo. *Journal of Computational Physics*, 154(1):134–155, Sep 1999. ISSN 00219991. DOI 10.1006/jcph.1999.6305.

N. Gehrels, G. Chincarini, P. Giommi, K. O. Mason, J. A. Nousek, A. A. Wells, N. E. White, S. D. Barthelmy, D. N. Burrows, L. R. Cominsky, et al. The Swift GammaRay Burst Mission. *The Astrophysical Journal*, 611(2):1005–1020, Aug 2004. ISSN 0004-637X. DOI 10.1086/422091.

K. C. Gendreau, Z. Arzoumanian, P. W. Adkins, C. L. Albert, J. F. Anders, A. T. Aylward, C. L. Baker, E. R. Balsamo, W. A. Bamford, S. S. Benegalrao, et al. The Neutron star Interior Composition Explorer (NICER): design and development. *Space*

Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray, (July 2016):99051H, Jul 2016. ISSN 1996756X. DOI 10.1117/12.2231304.

K. Glampedakis and P. D. Lasky. The freedom to choose neutron star magnetic field equilibria. *Monthly Notices of the Royal Astronomical Society*, 463(3):2542–2552, 2016. ISSN 13652966. DOI 10.1093/mnras/stw2115.

S. K. Godunov. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Matematicheskii Sbornik*, 47:271 – 306, 1959.

T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, and J. Shalf. The Cactus Framework and Toolkit: Design and Applications. In J. Palma, A. Sousa, J. Dongarra, and V. Hernández, editors, *High Performance Computing for Computational Science - VECPAR 2002*, volume 2565 of *Lecture Notes in Computer Science*, pages 197–227. Springer, Berlin, Heidelberg, 2003. DOI 10.1007/3-540-36569-9˙13.

S. Gottlieb, D. I. Ketcheson, and C.-W. W. Shu. High order strong stability preserving time discretizations. *Journal of Scientific Computing*, 38(3):251–289, Mar 2009. ISSN 08857474. DOI 10.1007/s10915-008-9239-z.

D. O. Gough. The Anelastic Approximation for Thermal Convection. *Journal of the Atmospheric Sciences*, 26(3):448–456, May 1969. ISSN 0022-4928. DOI 10.1175/1520-0469(1969)026¡0448:TAAFTC¿2.0.CO;2.

E. Gourgoulhon. An introduction to relativistic hydrodynamics. *EAS Publications Series*, 21(January 2006):43–79, Mar 2006. DOI 10.1051/eas:2006106.

E. Gourgoulhon. An introduction to the theory of rotating relativistic stars. In *Lectures at the COMPSTAR 2010 School GANIL, Caen, France*, number September, 2010. URL http://arxiv.org/abs/1003.5015.

J. M. Greenberg and A. Y. Leroux. A Well-Balanced Scheme for the Numerical Processing of Source Terms in Hyperbolic Equations. *SIAM Journal on Numerical Analysis*, 33(1):1–16, Feb 1996. ISSN 0036-1429. DOI 10.1137/0733001.

P. M. Gresho and S. T. Chan. On the theory of semiimplicit projection methods for viscous incompressible flow and its implementation via a finite element method that also introduces a nearly consistent mass matrix. Part 2: Implementation. *International Journal for Numerical Methods in Fluids*, 11(5):621–659, 1990. ISSN 10970363. DOI 10.1002/fld.1650110510.

J. Grindlay, H. Gursky, H. Schnopper, D. R. Parsignault, J. Heise, A. C. Brinkman, and J. Schrijver. Discovery of intense X-ray bursts from the globular cluster NGC 6624. *The Astrophysical Journal*, 205(2):L127, May 1976. ISSN 0004-637X. DOI 10.1086/182105.

Â. Guillard, H. Guillard, and C. Viozat. On the behaviour of upwind schemes in the low Mach number limit. *Computers & Fluids*, 28(1):63–86, Jan 1999. ISSN 00457930. DOI 10.1016/S0045-7930(98)00017-6.

H. Guillard and A. Murrone. On the behavior of upwind schemes in the low Mach number limit: II. Godunov type schemes. *Computers & Fluids*, 33(4):655–675, May 2004. ISSN 00457930. DOI 10.1016/j.compfluid.2003.07.001.

Y. Hadjimichael, C. B. Macdonald, D. I. Ketcheson, and J. H. Verner. Strong Stability Preserving Explicit Runge–Kutta Methods of Maximal Effective Order. *SIAM Journal on Numerical Analysis*, 51(4):2149–2165, Jan 2013. ISSN 0036-1429. DOI 10.1137/120884201.

P. Haensel, A. Y. Potekhin, and D. G. Yakovlev. *Neutron Stars 1: Equation of state and structure*. Springer Science & Business Media, 2007. ISBN 0387473017.

A. Harpole and I. Hawke. R3D2: Relativistic Reactive Riemann problem solver for Deflagrations and Detonations. *The Journal of Open Source Software*, 1(1):114501, May 2016. DOI 10.21105/joss.00016.

F. A. Harrison, S. Boggs, F. Christensen, W. Craig, C. Hailey, D. Stern, W. Zhang, L. Angelini, H. An, V. Bhalereo, et al. The Nuclear Spectroscopic Telescope Array (NuSTAR). pages 1–9, 2010. ISSN 0277786X. DOI 10.1117/12.857298.

A. Harten, P. D. Lax, and B. van Leer. On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Review*, 25(1):35–61, Jan 1983. ISSN 0036-1445. DOI 10.1137/1025002.

M.-a. Hashimoto, H. Seki, and M. Yamada. Shape of Nuclei in the Crust of Neutron Star. *Progress of Theoretical Physics*, 71(2):320–326, Feb 1984. ISSN 0033-068X. DOI 10.1143/PTP.71.320.

B. Haskell, L. Samuelsson, K. Glampedakis, and N. Andersson. Modelling magnetically deformed neutron stars. *Monthly Notices of the Royal Astronomical Society*, 385(1): 531–542, Feb 2008. ISSN 00358711. DOI 10.1111/j.1365-2966.2008.12861.x.

H. Helmholtz. On discontinuous fluid motions. *Monatl. Konigl. Preuss. Akad. Wiss*, 23 (1868):215–288, 1868. URL http://www.neo-classical-physics.info/uploads/3/0/6/5/3065888/helmholtz_-_discontinuous_fluid_motions.pdf.

C. Helzel, R. J. Leveque, and G. Warnecke. A Modified Fractional Step Method for the Accurate Approximation of Detonation Waves. *SIAM Journal on Scientific Computing*, 22(4):1489–1510, 2000. ISSN 1064-8275. DOI 10.1137/S1064827599357814.

J. S. Heyl. R-Modes on Rapidly Rotating, Relativistic Stars: I. Do Type-I Bursts Excite Modes in the Neutron-Star Ocean? *The Astrophysical Journal*, 600(939945):18, Aug 2004. ISSN 15384357. DOI 10.1086/379966.

W. Hillebrandt, M. Kromer, F. K. Röpke, and A. J. Ruiter. Towards an understanding of Type Ia supernovae from a synthesis of theory and observations. *Frontiers of Physics*, 8(2):116–143, Apr 2013. ISSN 2095-0462. DOI 10.1007/s11467-013-0303-2.

J. A. Hoffman, H. L. Marshall, and W. H. G. Lewin. Dual character of the rapid burster and a classification of X-ray bursts. *Nature*, 271(5646):630–633, 1978. ISSN 0028-0836. DOI 10.1038/271630a0.

J. D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science and Engineering*, 9(3):99–104, 2007. ISSN 15219615. DOI 10.1109/MCSE.2007.55.

N. A. Inogamov and R. A. Sunyaev. Spread of matter over a neutron-star surface during disk accretion. *Astronomy Letters*, 36(12):848–894, Dec 2010. ISSN 1063-7737. DOI 10.1134/S1063773710120029.

T. Inoue, K. Asano, and K. Ioka. Three-dimensional simulations of magnetohydrodynamic turbulence behind relativistic shock waves and their implications for gamma-ray bursts. *The Astrophysical Journal*, 734(2):77, Jun 2011. ISSN 0004-637X. DOI 10.1088/0004-637X/734/2/77.

A. M. Jacobs, M. Zingale, A. Nonaka, A. S. Almgren, and J. B. Bell. Low Mach Number Modeling of Convection in Helium Shells on Sub-Chandrasekhar White Dwarfs II: Bulk Properties of Simple Models. *The Astrophysical Journal*, 827(1):84, Aug 2016. ISSN 1538-4357. DOI 10.3847/0004-637X/827/1/84.

E. Jones, T. Oliphant, P. Peterson, and Others. SciPy: Open source scientific tools for Python. URL http://www.scipy.org/.

P. Joss. Helium-burning flashes on an accreting neutron star-A model for X-ray burst sources. *The Astrophysical Journal*, 225:L123–L127, 1978. URL http://adsabs.harvard.edu/full/1978ApJ...225L.123J.

J. J. E. Kajava, J. Nättilä, O.-M. Latvala, M. Pursiainen, J. Poutanen, V. F. Suleimanov, M. G. Revnivtsev, E. Kuulkers, and D. K. Galloway. The influence of accretion geometry on the spectral evolution during thermonuclear (type-I) X-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 445(4):4218–4234, 2014. ISSN 13652966. DOI 10.1093/mnras/stu2073.

R. Käppeli and S. Mishra. A well-balanced finite volume scheme for the Euler equations with gravitation. *Astronomy & Astrophysics*, 587:A94, 2016. ISSN 0004-6361. DOI 10.1051/0004-6361/201527815.

L. Keek, N. Langer, and J. J. M. in 't Zand. The effect of rotation on the stability of nuclear burning in accreting neutron stars. *Astronomy & Astrophysics*, 502(3): 871–881, Aug 2009. ISSN 0004-6361. DOI 10.1051/0004-6361/200911619.

L. Kelvin. Hydrokinetic solutions and observations. In *Mathematical and Physical Papers*, volume 4, pages 69–85. Cambridge, May 1910. URL http://www.ams.org/journal-getitem?pii=S0002-9904-1914-02527-3.

S. Klainerman and A. Majda. Compressible and incompressible fluids. *Communications on Pure and Applied Mathematics*, 35(5):629–651, Sep 1982. ISSN 00103640. DOI 10.1002/cpa.3160350503.

R. Klein. Numerics in Combustion. *Introduction to Turbulent Combustion*, 1998. URL http://star-www.st-and.ac.uk/~ch80/LITERATUR/KleinsLecture.ps.gz.

R. Klein and O. Pauluis. Thermodynamic Consistency of a Pseudoincompressible Approximation for General Equations of State. *Journal of the Atmospheric Sciences*, 69 (3):961–968, Mar 2012. ISSN 0022-4928. DOI 10.1175/JAS-D-11-0110.1.

T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, et al. Jupyter Notebooksa publishing format for reproducible computational workflows. In *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, pages 87–90, 2016. ISBN 9781614996491. DOI 10.3233/978-1-61499-649-1-87.

R. H. Kraichnan. Inertial Ranges in Two-Dimensional Turbulence. *Physics of Fluids*, 10(7):1417, 1967. ISSN 00319171. DOI 10.1063/1.1762301.

F. Kupka and H. J. Muthsam. *Modelling of stellar convection*, volume 3. Springer International Publishing, 2017. ISBN 4111501700019. DOI 10.1007/s41115-017-0001-9.

E. Kuulkers, J. J. M. in 't Zand, M. H. van Kerkwijk, R. Cornelisse, D. A. Smith, J. Heise, A. Bazzano, M. Cocchi, L. Natalucci, and P. Ubertini. A half-a-day long thermonuclear X-ray burst from KS 1731-260. *Astronomy & Astrophysics*, 382(2): 503–512, Feb 2002. ISSN 0004-6361. DOI 10.1051/0004-6361:20011654.

N. Lampe, A. Heger, and D. K. Galloway. The influence of accretion rate and metallicity on thermonuclear bursts: predictions from KEPLER models. *The Astrophysical Journal*, 819(1), Dec 2016. DOI 10.3847/0004-637X/819/1/46.

J. M. Lattimer and M. Prakash. The Physics of Neutron Stars. *Science*, 304(5670): 536–542, 2004. ISSN 0036-8075. DOI 10.1126/science.1090720.

J. M. Lattimer and M. Prakash. The equation of state of hot, dense matter and neutron stars. *Physics Reports*, 621:127–164, 2016. ISSN 03701573. DOI 10.1016/j.physrep.2015.12.005.

R. Lau, M. Beard, S. S. Gupta, H. Schatz, A. V. Afanasjev, E. F. Brown, A. Deibel, L. R. Gasques, G. W. Hitt, W. R. Hix, et al. Nuclear Reactions in the Crusts of Accreting Neutron Stars. pages 1–25, Mar 2018. URL http://arxiv.org/abs/1803.03818.

C. K. Law. *Combustion Physics*. Cambridge University Press, first edition, 2006. ISBN 9780521870528.

P. D. Lax and B. Wendroff. Systems of conservation laws. *Communications on Pure and Applied Mathematics*, 13(2):217–237, 1960. ISSN 00103640. DOI 10.1002/cpa.3160130205.

U. Lee and T. E. Strohmayer. Light curves of oscillating neutron stars. *Monthly Notices of the Royal Astronomical Society*, 361(2):659–672, Aug 2005. ISSN 0035-8711. DOI 10.1111/j.1365-2966.2005.09198.x.

R. J. LeVeque. *Numerical Methods for Conservation Laws*, volume 57. 1992. ISBN 3764327235. DOI 10.2307/2938728.

R. J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, third edition, 2002. ISBN 0 521 00924 3.

R. Liska and B. Wendroff. Comparison of Several Difference Schemes on 1D and 2D Test Problems for the Euler Equations. *SIAM Journal on Scientific Computing*, 25 (3):995–1017, 2003. ISSN 1064-8275. DOI 10.1137/S1064827502402120.

Q. Z. Liu, J. van Paradijs, and E. P. J. van den Heuvel. A catalogue of low-mass X-ray binaries in the Galaxy, LMC, and SMC (Fourth edition). *Astronomy & Astrophysics*, 469(2):807–810, Jul 2007. ISSN 0004-6361. DOI 10.1051/0004-6361:20077303.

F. Löffler, J. Faber, E. Bentivegna, T. Bode, P. Diener, R. Haas, I. Hinder, B. C. Mundim, C. D. Ott, E. Schnetter, et al. The Einstein Toolkit: a community computational infrastructure for relativistic astrophysics. *Classical and Quantum Gravity*, 29:115001, 2012. ISSN 0264-9381. DOI 10.1088/0264-9381/29/11/115001.

P. MacNeice, K. M. Olson, C. Mobarry, R. De Fainchtein, and C. Packer. PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126(3):330–354, 2000. ISSN 00104655. DOI 10.1016/S0010-4655(99)00501-9.

S. Mahmoodifar and T. Strohmayer. X-ray Burst Oscillations: From Flame Spreading to the Cooling Wake. *The Astrophysical Journal*, 818(1), Oct 2016. DOI 10.3847/0004-637X/818/1/93.

C. M. Malone, A. Nonaka, A. S. Almgren, J. B. Bell, and M. Zingale. Multidimensional Modeling of Type I X-ray Bursts. I. Two-Dimensional Convection Prior to the Outburst of a Pure He Accretor. *The Astrophysical Journal*, 728(2):118, Feb 2011. ISSN 0004-637X. DOI 10.1088/0004-637X/728/2/118.

C. M. Malone, M. Zingale, A. Nonaka, A. S. Almgren, and J. B. Bell. Multidimensional Modeling of Type I X-ray Bursts. II. Two-Dimensional Convection in a Mixed H/He Accretor. *The Astrophysical Journal*, 788(2):115, Jun 2014. ISSN 0004-637X. DOI 10.1088/0004-637X/788/2/115.

A. Maniopoulou and N. Andersson. The traditional approximation in general relativity. *Monthly Notices of the Royal Astronomical Society*, 351(4):1349–1358, Jul 2004. ISSN 00358711. DOI 10.1111/j.1365-2966.2004.07872.x.

J. M. Martí and E. Müller. The analytical solution of the Riemann problem in relativistic hydrodynamics. *Journal of Fluid Mechanics*, 258:317–333, 1994. ISSN 0022-1120. DOI 10.1017/S0022112094003344.

J. M. Martí and E. Müller. Numerical Hydrodynamics in Special Relativity. *Living Reviews in Relativity*, 6(7):1–100, 2003. ISSN 1433-8351. DOI 10.12942/lrr-2003-7.

J. M. Martí and E. Müller. Grid-based Methods in Relativistic Hydrodynamics and Magnetohydrodynamics. *Living Reviews in Computational Astrophysics*, 1(1):3, Dec 2015. ISSN 2367-3621. DOI 10.1007/lrca-2015-3.

J. M. Martí, J. M. Ibáñez, and J. A. Miralles. Numerical relativistic hydrodynamics: Local characteristic approach. *Physical Review D*, 43(12):3794–3801, Jun 1991. ISSN 0556-2821. DOI 10.1103/PhysRevD.43.3794.

M. Matsuoka, K. Kawasaki, S. Ueno, H. Tomida, M. Kohama, M. Suzuki, Y. Adachi, M. Ishikawa, T. Mihara, M. Sugizaki, et al. The MAXI mission on the ISS: Science and instruments for monitoring All-sky X-ray Images. *Publications of the Astronomical Society of Japan*, 61(5):999–1010, 2009. ISSN 00046264. DOI 10.1093/pasj/61.5.999.

I. Maurer and A. L. Watts. Ignition latitude and the shape of Type I X-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 383(1):387–398, Dec 2008. ISSN 00358711. DOI 10.1111/j.1365-2966.2007.12558.x.

C. P. McNally, W. Lyra, and J.-C. Passy. A well-posed Kelvin-Helmholtz test and comparison. *The Astrophysical Journal Supplement Series*, 201(2):18, Aug 2012. ISSN 0067-0049. DOI 10.1088/0067-0049/201/2/18.

C. A. Meakin and D. Arnett. Turbulent Convection in Stellar Interiors. I. Hydrodynamic Simulation. *The Astrophysical Journal*, 667(1):448–475, 2007. ISSN 0004-637X. DOI 10.1086/520318.

Z. Medin and A. Cumming. Compositionally-driven convection in the oceans of accreting neutron stars. *The Astrophysical Journal*, 730:97, 2011. ISSN 0004-637X. DOI 10.1088/0004-637X/730/2/97.

Z. Meisel and A. Deibel. Constraints on Bygone Nucleosynthesis of Accreting Neutron Stars. *The Astrophysical Journal*, 837(1):1–5, 2017. ISSN 1538-4357. DOI 10.3847/1538-4357/aa618d.

K. Menou. Differential rotation in weakly magnetized neutron star atmospheres. *Monthly Notices of the Royal Astronomical Society*, 352(4):1381–1389, Aug 2004. ISSN 00358711. DOI 10.1111/j.1365-2966.2004.08029.x.

F. Miczek, F. K. Röpke, and P. V. F. Edelmann. New numerical solver for flows at various Mach numbers. *Astronomy and Astrophysics*, 576:15, Sep 2014. DOI 10.1051/0004-6361/201425059.

A. Mignone and G. Bodo. An HLLC Riemann solver for relativistic flows – I. Hydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 364(1):126–136, Nov 2005. ISSN 0035-8711. DOI 10.1111/j.1365-2966.2005.09546.x.

M. C. Miller. A Characterization of the Brightness Oscillations During Thermonuclear Bursts From 4U 1636-536. *The Astrophysical Journal*, 531:458–466, 2000. DOI 10.1086/308438.

M. C. Miller. Astrophysical Constraints on Dense Matter in Neutron Stars. pages 1–51, Nov 2013. URL http://arxiv.org/abs/1312.0029.

S. T. Millmore. *Interfaces in Numerical Relativistic Hydrodynamics*. PhD thesis, University of Southampton, 2010. URL http://eprints.soton.ac.uk/170233/.

S. Mineo, M. Gilfanov, and R. Sunyaev. X-ray emission from star-forming galaxies - I. High-mass X-ray binaries. *Monthly Notices of the Royal Astronomical Society*, 419 (3):2095–2115, 2012. ISSN 00358711. DOI 10.1111/j.1365-2966.2011.19862.x.

S. M. Morsink, D. A. Leahy, C. Cadeau, and J. Braga. The Oblate Schwarzschild Approximation for Light Curves of Rapidly Rotating Neutron Stars. *The Astrophysical Journal*, 663(2):1244–1251, 2007. ISSN 0004-637X. DOI 10.1086/518648.

P. Mösta, C. D. Ott, D. Radice, L. F. Roberts, E. Schnetter, and R. Haas. A large-scale dynamo and magnetoturbulence in rapidly rotating core-collapse supernovae. *Nature*, 528(7582):1–24, Nov 2015. ISSN 1476-4687. DOI 10.1038/nature15755.

E. Motheau, M. Duarte, A. Almgren, and J. B. Bell. A hybrid adaptive low-mach-number/compressible method. Part I: Euler equations. *Journal of Computational Physics*, 2018. ISSN 00219991. DOI 10.1016/j.jcp.2018.01.036.

O. Møyner and H. M. Nilsen. Multiresolution Coupled Vertical Equilibrium Model for Fast Flexible Simulation of CO$_2$ Storage. pages 1–26, Oct 2017. URL http://arxiv.org/abs/1710.08735.

J. Muddle. *Advanced numerical methods for neutron star interfaces*. PhD thesis, University of Southampton, 2015. URL http://eprints.soton.ac.uk/375551/.

D. Mukherjee, P. Bult, M. van der Klis, and D. Bhattacharya. The magnetic-field strengths of accreting millisecond pulsars. *Monthly Notices of the Royal Astronomical Society*, 452(4):3994–4012, Oct 2015. ISSN 0035-8711. DOI 10.1093/mnras/stv1542.

B. Müller. The Status of Multi-Dimensional Core-Collapse Supernova Models. *Publications of the Astronomical Society of Australia*, 33(1990):e048, Sep 2016. ISSN 1323-3580. DOI 10.1017/pasa.2016.40.

Multi-INstrument Burst ARchive (MINBAR). X-ray Burst Source Table. URL https://burst.sci.monash.edu/wiki/index.php?n=MINBAR.SourceTable.

M. P. Muno, D. Chakrabarty, D. K. Galloway, and D. Psaltis. The Frequency Stability of Millisecond Oscillations in Thermonuclear X-Ray Bursts. *The Astrophysical Journal*, 580:1048, 2002. ISSN 0004-637X. DOI 10.1086/343793.

M. P. Muno, D. K. Galloway, D. Chakrabarty, and ;. The Effect of Neutron Star Rotation on the Properties of Thermonuclear X-ray Bursts. *The Astrophysical Journal*, 608: 930–934, 2004. ISSN 0004-637X. DOI 10.1086/420812.

R. Narayan and J. S. Heyl. Thermonuclear Stability of Material Accreting onto a Neutron Star. *The Astrophysical Journal*, 599(1):419–449, Dec 2003. ISSN 0004-637X. DOI 10.1086/379211.

J. C. Niemeyer and S. E. Woosley. The Thermonuclear Explosion Of Chandrasekhar Mass White Dwarfs. *The Astrophysical Journal*, 475(1):740–753, 1997. ISSN 0004-637X. DOI 10.1086/303544.

J. C. Niemeyer, W. K. Bushe, and G. R. Ruetsch. Small-scale Interaction of Turbulence with Thermonuclear Flames in Type Ia Supernovae. *The Astrophysical Journal*, 524 (1):290, 1999. ISSN 0004-637X. DOI 10.1086/307813.

A. Nonaka, A. S. Almgren, J. B. Bell, M. J. Lijewski, C. M. Malone, and M. Zingale. Maestro: an Adaptive Low Mach Number Hydrodynamics Algorithm for Stellar Flows. *The Astrophysical Journal Supplement Series*, 188(2):358–383, Jun 2010. ISSN 0067-0049. DOI 10.1088/0067-0049/188/2/358.

A. Nonaka, A. S. Almgren, J. B. Bell, M. J. Lijewski, C. Malone, and M. Zingale. MAESTRO User's Guide, 2014. URL http://bender.astro.sunysb.edu/Maestro/staging/MAESTRO/Docs/MaestroUsersGuide.pdf.

M. Oertel, M. Hempel, T. Klähn, and S. Typel. Equations of state for supernovae and compact stars. *Reviews of Modern Physics*, 89(1), 2017. ISSN 15390756. DOI 10.1103/RevModPhys.89.015007.

Y. Ogura and N. A. Phillips. Scale Analysis of Deep and Shallow Convection in the Atmosphere, 1962. ISSN 0022-4928.

T. E. Oliphant. *A guide to NumPy*. USA: Trelgol Publishing, 2006.

L. S. Ootes, A. L. Watts, D. K. Galloway, and R. Wijnands. The Accretion Rate Dependence of Burst Oscillation Amplitude. *The Astrophysical Journal*, 834(1):21, Dec 2016. ISSN 1538-4357. DOI 10.3847/1538-4357/834/1/21.

K. Ooyama. Numerical Simulation of the Life Cycle of Tropical Cyclones, 1969. ISSN 0022-4928. DOI 10.1175/1520-0469(1969)026¡0003:NSOTLC¿2.0.CO;2.

B. W. O'Shea, G. Bryan, J. Bordner, M. L. Norman, T. Abel, R. Harkness, and A. Kritsuk. Introducing Enzo, an AMR Cosmology Application. 2004. DOI 10.1007/3-540-27039-6˙24.

K. Oyamatsu, M.-a. Hashimoto, and M. Yamada. Further Study of the Nuclear Shape in High-Density Matter. *Progress of Theoretical Physics*, 72(2):373–375, Aug 1984. ISSN 0033-068X. DOI 10.1143/PTP.72.373.

F. Özel and P. Freire. Masses, Radii, and the Equation of State of Neutron Stars. *Annual Review of Astronomy and Astrophysics*, 54(1):401–440, 2016. ISSN 0066-4146. DOI 10.1146/annurev-astro-081915-023322.

J. Palaniappan, R. B. Haber, and R. L. Jerrard. A spacetime discontinuous Galerkin method for scalar conservation laws. *Computer Methods in Applied Mechanics and Engineering*, 193(33-35):3607–3631, Aug 2004. ISSN 00457825. DOI 10.1016/j.cma.2004.01.028.

V. Paschalidis, N. Stergioulas, and N. Stergioulas. Rotating stars in relativity. *Living Reviews in Relativity*, 20(1):7, Dec 2017. ISSN 2367-3613. DOI 10.1007/s41114-017-0008-x.

F. Peng, E. F. Brown, and J. W. Truran. Sedimentation and Type I XRay Bursts at Low Accretion Rates. *The Astrophysical Journal*, 654(2):1022–1035, Jan 2007. ISSN 0004-637X. DOI 10.1086/509628.

R. Perna, P. J. Armitage, and B. Zhang. Flares in Long and Short Gamma-Ray Bursts: A Common Origin in a Hyperaccreting Accretion Disk. *The Astrophysical Journal*, 636(1):L29–L32, 2006. ISSN 0004-637X. DOI 10.1086/499775.

M. E. Pessah and J. Goodman. On the Saturation of the Magnetorotational Instability via Parasitic Modes. *The Astrophysical Journal*, 698(1):L72–L76, Jun 2009. ISSN 0004-637X. DOI 10.1088/0004-637X/698/1/L72.

N. Peters. *Turbulent Combustion*. Cambridge University Press, third edition, 2000. ISBN 0 521 66082 3.

A. Piro. *Oscillations and Surface Physics of Neutron Stars*. PhD thesis, University of California, Santa Barbara, 2006. URL `http://adsabs.harvard.edu/abs/2006PhDT.....14P`.

A. L. Piro and L. Bildsten. Turbulent Mixing in the Surface Layers of Accreting Neutron Stars. *The Astrophysical Journal*, 663:1252–1268, 2007. ISSN 0004-637X. DOI 10.1086/518687.

J. A. Pons, J. A. Font, J. M. Ibanez, J. M. Marti, and J. A. Miralles. General Relativistic Hydrodynamics with Special Relativistic Riemann Solvers. *Astronomy &*

Astrophysics, 339:638–642, Jul 1998. ISSN 00046361. URL http://arxiv.org/abs/astro-ph/9807215.

J. A. Pons, J. M. Martí, and E. Müller. The exact solution of the Riemann problem with non-zero tangential velocities in relativistic hydrodynamics. *Journal of Fluid Mechanics*, 422:125–139, 2000. ISSN 0022-1120. DOI 10.1017/S0022112000001439.

A. Y. Potekhin. The physics of neutron stars. *Physics-Uspekhi*, 53(12):1235–1256, Dec 2010. ISSN 1063-7869. DOI 10.3367/UFNe.0180.201012c.1279.

D. Radice. *Advanced Numerical Approaches in the Dynamics of Relativistic Flows*. PhD thesis, Gottfried Wilhelm Leibniz Universität Hannover, 2013. URL http://inspirehep.net/record/1346306/.

D. Radice and L. Rezzolla. Universality and Intermittency in Relativistic Turbulent Flows of a Hot Plasma. *The Astrophysical Journal*, 766(1):L10, 2013. ISSN 2041-8205. DOI 10.1088/2041-8205/766/1/L10.

L. Rayleigh. Investigation of the Character of the Equilibrium of an Incompressible Heavy Fluid of Variable Density. *Scientific Papers*, 2:200–207, 1900.

L. Rayleigh. LIX. On convection currents in a horizontal layer of fluid, when the higher temperature is on the under side. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 32(192):529–546, 1916.

H. S. Reall. Part III General Relativity, 2011. URL http://www.damtp.cam.ac.uk/user/hsr1000/lecturenotes_2012.pdf.

R. G. R. Rehm, H. R. Baumt, N. Bureau, and H. Baum. The equations of motion for thermally driven, buoyant flows. *Journal of Research of the National Bureau of Standards*, 83(3):297, May 1978. ISSN 0160-1741. DOI 10.6028/jres.083.019.

M. Reinecke, W. Hillebrandt, and J. C. Niemeyer. Thermonuclear explosions of Chandrasekhar-mass C+O white dwarfs. *Astronomy and Astrophysics*, 347:739–747, 1999. ISSN 0004-6361. URL http://arxiv.org/abs/astro-ph/9812120.

L. Rezzolla and O. Zanotti. New Relativistic Effects in the Dynamics of Nonlinear Hydrodynamical Waves. *Physical Review Letters*, 89(11):114501, 2002. ISSN 0031-9007. DOI 10.1103/PhysRevLett.89.114501.

L. Rezzolla and O. Zanotti. *Relativistic Hydrodynamics*. OUP Oxford, 2013. ISBN 0198528906.

B. E. Robertson, A. V. Kravtsov, N. Y. Gnedin, T. Abel, and D. H. Rudd. Computational Eulerian hydrodynamics and Galilean invariance. *Monthly Notices of the Royal Astronomical Society*, 401(4):2463–2476, 2010. ISSN 00358711. DOI 10.1111/j.1365-2966.2009.15823.x.

M. Ruderman, T. Zhu, and K. Chen. Neutron Star Magnetic Field Evolution, Crust Movement, and Glitches. *The Astrophysical Journal*, 492(1):267–280, 1998. ISSN 0004-637X. DOI 10.1086/305026.

H. Sawai and S. Yamada. The Evolution and Impacts of Magnetorotational Instability in Magnetized Core-Collapse Supernovae. *The Astrophysical Journal*, 817(2):153, Jan 2016. ISSN 1538-4357. DOI 10.3847/0004-637X/817/2/153.

H. Sawai, S. Yamada, and H. Suzuki. Global simulations of magnetorotational instability in the collapsed core of a massive star. *Astrophysical Journal Letters*, 770(2):2–7, 2013. ISSN 20418205. DOI 10.1088/2041-8205/770/2/L19.

H. Schatz. Rare isotopes in thermonuclear explosions on neutron stars. *Progress in Particle and Nuclear Physics*, 66(2):277–282, 2011. ISSN 01466410. DOI 10.1016/j.ppnp.2011.01.020.

H. Schatz and K. E. Rehm. X-ray binaries. *Nuclear Physics A*, 777:601–622, Oct 2006. ISSN 03759474. DOI 10.1016/j.nuclphysa.2005.05.200.

H. Schatz, L. Bildsten, A. Cumming, and M. Wiescher. The Rapid Proton Process Ashes from Stable Nuclear Burning on an Accreting Neutron Star. *The Astrophysical Journal*, 524(2):1014–1029, Oct 1999. ISSN 0004-637X. DOI 10.1086/307837.

V. Schneider, U. Katscher, D. H. Rischke, B. Waldhauser, J. A. Maruhn, and C. D. Munz. New Algorithms for Ultra-relativistic Numerical Hydrodynamics. *Journal of Computational Physics*, 105(1):92–107, Mar 1993. ISSN 00219991. DOI 10.1006/jcph.1993.1056.

E. Schnetter, S. H. Hawley, and I. Hawke. Evolutions in 3D numerical relativity using fixed mesh refinement. *Classical and Quantum Gravity*, 21(6):1465–1488, Mar 2004. ISSN 0264-9381. DOI 10.1088/0264-9381/21/6/014.

M. Schwarzschild and R. Härm. Thermal Instability in Non-Degenerate Stars. *The Astrophysical Journal*, 142:855, Oct 1965. ISSN 0004-637X. DOI 10.1086/148358.

M. Serino, W. Iwakiri, T. Tamagawa, T. Sakamoto, S. Nakahira, M. Matsuoka, K. Yamaoka, and H. Negoro. MAXI observations of long X-ray bursts. *Publications of the Astronomical Society of Japan*, 68(6):95, Dec 2016. ISSN 0004-6264. DOI 10.1093/pasj/psw086.

V. Sguera, E. J. Barlow, A. J. Bird, D. J. Clark, A. J. Dean, A. B. Hill, L. Moran, S. E. Shaw, D. R. Willis, A. Bazzano, et al. INTEGRAL observations of recurrent fast X-ray transient sources. *Astronomy & Astrophysics*, 444(1):221–231, 2005. ISSN 0004-6361. DOI 10.1051/0004-6361:20053103.

M. M. Shara. Localized thermonuclear runaways and volcanoes on degenerate dwarf stars. *The Astrophysical Journal*, 261:649, Oct 1982. ISSN 0004-637X. DOI 10.1086/160376.

P. S. Shternin, D. G. Yakovlev, C. O. Heinke, W. C. Ho, and D. J. Patnaude. Cooling neutron star in the Cassiopeia A supernova remnant: Evidence for superfluidity in the core. *Monthly Notices of the Royal Astronomical Society: Letters*, 412(1):108–112, 2011. ISSN 17453933. DOI 10.1111/j.1745-3933.2011.01015.x.

K. P. Singh, S. N. Tandon, P. C. Agrawal, H. M. Antia, R. K. Manchanda, J. S. Yadav, S. Seetha, M. C. Ramadevi, A. R. Rao, D. Bhattacharya, et al. ASTROSAT mission. In T. Takahashi, J.-W. A. den Herder, and M. Bautz, editors, *Proc. SPIE*, volume 9144, page 91441S, Jul 2014. ISBN 9780819496126. DOI 10.1117/12.2062667.

V. Smiljanovski, V. Moser, and R. Klein. A capturing - tracking hybrid scheme for deflagration discontinuities. *Combustion Theory and Modelling*, 1(2):183–215, Apr 1997. ISSN 1364-7830. DOI 10.1088/1364-7830/1/2/004.

G. A. Sod. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *Journal of Computational Physics*, 27(1):1–31, 1978. ISSN 10902716. DOI 10.1016/0021-9991(78)90023-2.

E. A. Spiegel and G. Veronis. On the Boussinesq Approximation for a Compressible Fluid. *The Astrophysical Journal*, 131(1):442, Mar 1960. ISSN 0004-637X. DOI 10.1086/146849.

A. Spitkovsky, Y. Levin, and G. Ushomirsky. Propagation of Thermonuclear Flames on Rapidly Rotating Neutron Stars: Extreme Weather during Type I XRay Bursts. *The Astrophysical Journal*, 566(2):1018–1038, Feb 2002. ISSN 0004-637X. DOI 10.1086/338040.

A. St-cyr, C. Jablonowski, J. M. Dennis, H. Tufo, and S. J. Thomas. A Comparison of Two Shallow Water Models with Non-Conforming Adaptive Grids: classical tests. *Monthly Weather Review*, 136(6):1898–1922, 2008. ISSN 0027-0644. DOI 10.1175/2007MWR2108.1.

J. J. Stoker. *Water Waves. The Mathematical Theory with Applications*. Interscience Publ., New York, 1957. ISBN 9780471570349.

G. Strang. On the Construction and Comparison of Difference Schemes. *SIAM Journal on Numerical Analysis*, 5(3):506–517, Sep 1968. ISSN 0036-1429. DOI 10.1137/0705041.

T. E. Strohmayer and L. Bildsten. New views of thermonuclear bursts. In W. H. G. Lewin and M. van der Klis, editors, *Compact Stellar X-ray Sources*. Cambridge University Press, Jan 2003. URL http://arxiv.org/abs/astro-ph/0301544.

T. E. Strohmayer and C. B. Markwardt. Evidence for a Millisecond Pulsar in 4U 163653 during a Superburst. *The Astrophysical Journal*, 577(1):337–345, Sep 2002. ISSN 0004-637X. DOI 10.1086/342152.

T. E. Strohmayer, W. Zhang, J. H. Swank, A. Smale, L. Titarchuk, C. Day, and U. Lee. Millisecond X-Ray Variability from an Accreting Neutron Star System. *The Astrophysical Journal*, 469(1):L9–L12, Sep 1996. ISSN 0004637X. DOI 10.1086/310261.

T. E. Strohmayer, K. Jahoda, A. B. Giles, and U. Lee. Millisecond Pulsations from a LowMass XRay Binary in the Galactic Center Region. *The Astrophysical Journal*, 486(1):355–362, Sep 1997a. ISSN 0004-637X. DOI 10.1086/304522.

T. E. Strohmayer, W. Zhang, and J. H. Swank. 363 Hz Oscillations during the Rising Phase of Bursts from 4U 172834: Evidence for Rotational Modulation. *The Astrophysical Journal*, 487(1):L77–L80, Sep 1997b. ISSN 0004637X. DOI 10.1086/310880.

T. E. Strohmayer, W. Zhang, J. H. Swank, and I. Lapidus. The Long-Term Stability of Oscillations during Thermonuclear X-Ray Bursts: Constraining the Binary X-Ray Mass Function. *The Astrophysical Journal*, 503(2):L147–L150, 1998. ISSN 0004637X. DOI 10.1086/311545.

R. E. Taam, S. E. Woosley, and D. Q. Lamb. The Effect of Deep Hydrogen Burning in the Accreted Envelope of a Neutron Star on the Properties of X-Ray Bursts. *The Astrophysical Journal*, 459:271, Mar 1996. ISSN 0004-637X. DOI 10.1086/176890.

T. M. Tauris and E. van den Heuvel. Formation and Evolution of Compact Stellar X-ray Sources. In W. Lewin and M. van der Klis, editors, *Compact stellar X-ray sources*, volume 39 of *Cambridge Astrophysics Series*, pages 623 – 665. Cambridge University Press, Mar 2006. ISBN 978-0-521-82659-4. DOI 10.2277/0521826594.

G. Taylor. The Instability of Liquid Surfaces when Accelerated in a Direction Perpendicular to their Planes. I. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 202(1):192–196, 1950. ISSN 1364-5021. DOI 10.1098/rspa.1950.0086.

K. S. Thorne. Relativistic Shocks: the Taub Adiabat. *The Astrophysical Journal*, 179: 897, Feb 1973. ISSN 0004-637X. DOI 10.1086/151927.

K. S. Thorne. The relativistic equations of stellar structure and evolution. *The Astrophysical Journal*, (212):825–831, 1977. URL http://adsabs.harvard.edu/full/1977ApJ...212..825T.

E. Toro. *Shock-capturing methods for free-surface shallow flows*. John Wiley, 2001. ISBN 0471987662. DOI 10.1080/00221680309499935.

E. F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2nd edition, 1999. ISBN 978-3-540-25202-3. DOI 10.1007/b79761.

E. F. Toro, M. Spruce, and W. Speares. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves*, 4(1):25–34, Jul 1994. ISSN 0938-1287. DOI 10.1007/BF01414629.

D. M. Townsley, B. J. Miles, F. X. Timmes, A. C. Calder, and E. F. Brown. A Tracer Method for Computing Type IA Supernova Yields: Burning Model Calibration, Reconstruction of Thickended Flames, and Verification for Planar Detonations. *The Astrophysical Journal Supplement Series*, 225(1):3, Jul 2016. ISSN 1538-4365. DOI 10.3847/0067-0049/225/1/3.

M. J. Turk, B. D. Smith, J. S. Oishi, S. Skory, S. W. Skillman, T. Abel, and M. L. Norman. Yt: A multi-code analysis toolkit for astrophysical simulation data. *Astrophysical Journal, Supplement Series*, 192(1), 2011. ISSN 00670049. DOI 10.1088/0067-0049/192/1/9.

M. van der Klis. Millisecond Oscillations in X-Ray Binaries. *Annual Review of Astronomy and Astrophysics*, 38(September):717–760, Jan 2000. ISSN 0066-4146. DOI 10.1146/annurev.astro.38.1.717.

B. van Leer. Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method. *Journal of Computational Physics*, 32(1):101–136, Jul 1979. ISSN 00219991. DOI 10.1016/0021-9991(79)90145-1.

G. M. Vasil, D. Lecoanet, B. P. Brown, T. S. Wood, and E. G. Zweibel. Energy Conservation and Gravity Waves in Sound-Proof Treatments of Stellar Interiors. II. Lagrangian Constrained Analysis. *The Astrophysical Journal*, 773(2):169, Aug 2013. ISSN 0004-637X. DOI 10.1088/0004-637X/773/2/169.

M. Viallet, I. Baraffe, and R. Walder. Towards a new generation of multi-dimensional stellar evolution models: development of an implicit hydrodynamic code. *Astronomy & Astrophysics*, 531:A86, 2011. ISSN 0004-6361. DOI 10.1051/0004-6361/201016374.

R. K. Wallace and S. E. Woosley. Explosive hydrogen burning. *The Astrophysical Journal Supplement Series*, 45:389, Feb 1981. ISSN 0067-0049. DOI 10.1086/190717.

R. Walter, A. A. Lutovinov, E. Bozzo, and S. S. Tsygankov. High-mass X-ray binaries in the Milky Way: A closer look with INTEGRAL. *Astronomy and Astrophysics Review*, 23(1):1–98, 2015. ISSN 09354956. DOI 10.1007/s00159-015-0082-6.

J. Wang, C. Liang, and M. S. Miesch. A compressible high-order unstructured spectral difference code for stratified convection in rotating spherical shells. *Journal of Computational Physics*, 290:90–111, 2015. ISSN 10902716. DOI 10.1016/j.jcp.2015.02.047.

G. Watanabe. Understanding nuclear "pasta": Current status and future prospects. *AIP Conference Proceedings*, 891:373–381, 2007. ISSN 0094243X. DOI 10.1063/1.2713539.

A. L. Watts. Thermonuclear Burst Oscillations. *Annual Review of Astronomy and Astrophysics*, 50(1):609–640, Sep 2012. ISSN 0066-4146. DOI 10.1146/annurev-astro-040312-132617.

A. L. Watts. Burst oscillation library, 2015. URL http://staff.fnwi.uva.nl/a.l.watts/bosc/bosc.html.

A. L. Watts, T. E. Strohmayer, and C. B. Markwardt. Analysis of Variability in the Burst Oscillations of the Accreting Millisecond Pulsar XTE J1814338. *The Astrophysical Journal*, 634(1):547–564, Nov 2005. ISSN 0004-637X. DOI 10.1086/496953.

A. L. Watts, N. Andersson, D. Chakrabarty, M. Feroci, K. Hebeler, G. Israel, F. K. Lamb, M. C. Miller, S. Morsink, F. Özel, et al. Colloquium : Measuring the neutron star equation of state using x-ray timing. *Reviews of Modern Physics*, 88(2):021001, Apr 2016. ISSN 0034-6861. DOI 10.1103/RevModPhys.88.021001.

E. Weinan. *Principles of Multiscale Modeling*. Cambridge University Press, 2011. ISBN 9781107096547. DOI 10.1063/PT.3.1609.

N. N. Weinberg, P. Arras, and J. Burkart. An instability due to the nonlinear coupling of p-modes to g-modes: implications for coalescing neutron star binaries. *The Astrophysical Journal*, 769(2):121, May 2013. ISSN 0004-637X. DOI 10.1088/0004-637X/769/2/121.

M. Wielgus, W. Kluźniak, A. Sądowski, R. Narayan, and M. Abramowicz. Stable, levitating, optically thin atmospheres of Eddington-luminosity neutron stars. *Monthly Notices of the Royal Astronomical Society*, 454(4):3766–3770, May 2015a. ISSN 13652966. DOI 10.1093/mnras/stv2191.

M. Wielgus, A. Sądowski, W. Kluźniak, M. Abramowicz, and R. Narayan. Levitating atmospheres of Eddington-luminosity neutron stars. *Monthly Notices of the Royal Astronomical Society*, 458:3420–3428, May 2015b. DOI 10.1093/mnras/stv2191.

J. R. Wilson. Numerical Study of Fluid Flow in a Kerr Space. *The Astrophysical Journal*, 173:431, Apr 1972. ISSN 0004-637X. DOI 10.1086/151434.

K.-H. A. Winkler and M. L. Norman. Astrophysical Radiation Hydrodynamics. In *NATO Advanced Science Institutes (ASI) Series C*, volume 188 of *NATO Advanced Science Institutes (ASI) Series C*. 1986. URL http://adsabs.harvard.edu/abs/1986ASIC..188.....W.

C. Wu, G. Huang, and Y. Zheng. Theoretical solution of dam-break shock wave. *Journal of Hydraulic Engineering*, 125:1210, 1999. ISSN 07339429. DOI 10.1061/(ASCE)0733-9429(1999)125:11(1210).

F. Xu, M. Gunzburger, J. Burkardt, and Q. Du. A Multiscale Implementation Based on Adaptive Mesh Refinement for the Nonlocal Peridynamics Model in One Dimension. *Multiscale Modeling & Simulation*, 14(1):398–429, Jan 2016. ISSN 1540-3459. DOI 10.1137/15M1010300.

S.-C. Yoon, N. Langer, M. van der Sluys, and M. V. D. Sluys. On the stability of thermonuclear shell sources in stars. *Astronomy & Astrophysics*, 425(1):207–216, 2004. ISSN 0004-6361. DOI 10.1051/0004-6361:20040231.

J. i. t. Zand. Understanding superbursts. Feb 2017. URL http://arxiv.org/abs/1702.04899.

J. i. t. Zand. Galactic X-ray bursters list, 2018. URL https://personal.sron.nl/~jeanz/bursterlist.html.

G. Zhang, M. Méndez, D. Altamirano, T. M. Belloni, and J. Homan. A very rare triple-peaked type-I X-ray burst in the low-mass X-ray binary 4U 163653. *Monthly Notices of the Royal Astronomical Society*, 398(1):368–374, Sep 2009. ISSN 00358711. DOI 10.1111/j.1365-2966.2009.15148.x.

G. Zhang, M. Mariano, T. M. Belloni, J. Homan, M. Mendez, T. M. Belloni, and J. Homan. Coherent oscillations and the evolution of the apparent emission area in the decaying phase of radius-expansion bursts from 4U 1636-53. *Monthly Notices of the Royal Astronomical Society*, 436(3):2276–2285, Oct 2013. ISSN 0035-8711. DOI 10.1093/mnras/stt1728.

G. Zhang, M. Mendez, M. Zamfir, and A. Cumming. The link between coherent burst oscillations, burst spectral evolution and accretion state in 4U 1728-34. *Monthly Notices of the Royal Astronomical Society*, (October):2004–2017, 2015. ISSN 0035-8711. DOI 10.1093/mnras/stv2482.

T. Zhang and Y. Zheng. Riemann problem for gasdynamic combustion. *Journal of Differential Equations*, 77(2):203–230, Feb 1989. ISSN 00220396. DOI 10.1016/0022-0396(89)90142-3.

W. Zhang, A. Almgren, M. Day, T. Nguyen, J. Shalf, and D. Unat. BoxLib with Tiling: An AMR Software Framework. *ArXiv e-prints*, 1604.03570, Apr 2016. URL http://arxiv.org/abs/1604.03570.

M. Zingale. pyro: A teaching code for computational astrophysical hydrodynamics. *Astronomy and Computing*, 6:52–62, Oct 2014. ISSN 22131337. DOI 10.1016/j.ascom.2014.07.003.

M. Zingale, A. S. Almgren, M. G. B. Sazo, V. E. Beckner, J. B. Bell, B. Friesen, A. M. Jacobs, M. P. Katz, C. M. Malone, A. J. Nonaka, et al. Meeting the Challenges of Modeling Astrophysical Thermonuclear Explosions: Castro, Maestro, and the AMReX Astrophysics Suite. Nov 2017. URL http://arxiv.org/abs/1711.06203.

J. Zrake and A. MacFadyen. Numerical simulations of driven supersonic relativistic MHD turbulence. *AIP Conference Proceedings*, 1358:102–105, 2011. ISSN 0094243X. DOI 10.1063/1.3621748.