University of Southampton Research Repository

University of Southampton
Faculty of Social Sciences
School of Mathematical Sciences

Radiation in Neutron Star Merger Remnant Simulations

by

Grant Curtis Schomberg
ORCiD: 0009-0003-3407-2854

A thesis submitted for the degree of
Doctor of Philosophy

October 2024

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL SCIENCES
SCHOOL OF MATHEMATICAL SCIENCES

Thesis for the degree of Doctor of Philosophy

RADIATION IN NEUTRON STAR MERGER REMNANT SIMULATIONS

by Grant Curtis Schomberg

Binary neutron star mergers are extreme astrophysical events. They involve mass densities multiple times greater than that of an atomic nucleus, velocities nearing half the speed of light, temperatures soaring up to 100 billion Kelvin, and gravitational acceleration reaching up to 100 billion times that at the Earth's surface. The binary neutron star merger known as GW170817 marked a pivotal moment as the first observation of a gravitational wave with an electromagnetic counterpart. The complete evolution and dynamics of neutron star merger remnants remain not fully understood, underscoring the need for further research. Investigation of these events holds the potential to refine both our theories of gravity and our understanding of the properties of matter at extreme densities and temperatures.

Numerical simulations are an appropriate tool to approach an understanding of these complex processes. By employing such simulations, we can explore which physical approximations and numerical techniques are sufficient to capture the range of behaviours involved. However, binary neutron star merger remnants present a formidable computational challenge due to the vast range of length scales involved and the system of highly non-linear coupled partial differential equations governing mass, momentum, and energy evolution.

Addressing this challenge, this project has involved the creation of a novel three-dimensional physics code capable of simulating hydrodynamic evolution, radiative transfer, and the general relativistic effects induced by fluid motion and spacetime curvature. This code has been implemented into the AMReX software framework to make use of block-structured adaptive mesh refinement for enhanced accuracy and efficiency.

In this thesis, I present state-of-the-art simulations of binary neutron star merger remnants. Making use of a pre-calculated merger simulation generated using the Einstein Toolkit, and employing a tabulated physical equation of state. We find that the introduction of a radiation field can disperse material surrounding a remnant object, and that coordinate transforms can help reduce numerical advective flux errors. Additionally, I utilize this code to explore uncertainty quantification by modelling turbulent simulations and analysing the convergence of statistical solutions.

Table of Contents

List of Figures

List of Tables

Declaration of Authorship

I, Grant Curtis Schomberg, declare that this thesis entitled *Radiation in Neutron Star Merger Remnant Simulations* and the work presented in it are my own and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

Signed:

Date:

Acknowledgements

As I look back on the past seven years dedicated to crafting this thesis, I am humbled by the significant changes that have shaped my life. Throughout this journey, I have had the privilege of crossing paths with an array of remarkable individuals who have greatly influenced and supported me.

First and foremost, I am deeply grateful to my family, whose unwavering support has been the bedrock of my journey. To my parents, thank you for your enduring patience and belief in me. Your encouragement has been the driving force behind this endeavour.

I'm also grateful to the many individuals who have played pivotal roles in shaping my experience over these years. From the rotating cast of housemates I have had, to the fellow doctoral students with whom I have shared the second floor research student offices at the School of Mathematical Sciences, Building 54, Highfield Campus.

To my supervisor, Ian Hawke, I owe an immeasurable debt of gratitude. Ian's boundless supply of knowledge and guidance have been instrumental in both my academic pursuits and personal growth. Without his patience and expert mentorship this milestone would have remained out of reach. While my name may be printed on the cover of this thesis, I recognise that its contents are a product of the collaborative efforts and support of many. I am deeply appreciative of the sponsorship of this work provided by AWE. The guidance and insights offered by my industry supervisors and managers have been indispensable.

Chapter 1

Introduction

1.1 Motivation

A number of processes in astrophysics involve extreme conditions such as intense mass-densities, temperatures, velocities, gravitational forces, and radiation energies. The astrophysical events known as neutron star mergers are indeed extreme.

A binary neutron star is a pair of neutron stars orbiting about a common centre. Over time the binary loses energy due to gravitational wave emission (see for example [1]). Eventually the neutron stars will collide and merge. During a neutron star merger, densities up to multiple times that of an atomic nucleus, velocities up to half of the speed of light, and temperatures up to 100 billion Kelvin can be reached. This process is described further in section 1.2.

The full evolution and dynamics of neutron star merger remnants are not yet fully understood. Understanding these events has the potential to constrain both theories of gravity, and our understanding of the properties of matter at extremes of densities and temperatures. Numerical simulations are an appropriate tool to approach an understanding of these processes. Numerical simulations allow us to model the evolution of the dynamic system and investigate behaviours such as the stability of the remnant object, observable signal generation, and the partitioning of energy between thermal, rotational, and radiation forms. With numerical simulations we can explore which physical approximations and numerical efficiencies are sufficient to represent the full complex physics involved. A numerical model of a neutron star merger remnant would at least require accurate modelling of hydrodynamics, radiative transfer, and general relativistic effects. Further modelling capabilities may include magnetic fields and chemical processes.

Codes such as the Einstein toolkit [2] and GRaM-X [3] are designed to be general purpose relativistic hydrodynamics codes, they make use of high order numerical methods and are capable of evolving the spacetime. However they do not make full use of flexible block-structured adaptive mesh refinement approaches. There exists a gap for an efficient code which is specifically designed for approximately stationary spacetimes, and is capable of taking advantage of the flexibility gained when using many small refined patches. This project has involved the creation of a novel three dimensional physics code. This code simulates hydrodynamic evolution, radiative transfer, and the general relativistic effects due to fluid motion and a curved background spacetime. This code has been implemented in a particularly efficient computational framework and has been used to investigate neutron star merger remnants.

1.2 The Origins of Neutron Star Merger Remnants

1.2.1 Progenitor Stars and Core-Collapse

Neutron stars are compact objects which form from the gravitational collapse of the core of a main sequence progenitor star of mass $8M_\odot \lesssim M \lesssim 30M_\odot$, resulting in an astrophysical explosion known as a type-II core-collapse supernova [4].

Type-II supernovae are the core-collapse and resulting explosion of massive ($M \gtrsim 8M_\odot$) stars [5]. They are identified by the presence of hydrogen, which can be identified by particular lines in the emission spectra of observed optical signals.

Throughout the lifetime of such a star, lighter elements are fused into heavier elements, starting with the fusion of hydrogen into helium. These nuclear fusion reactions become slower and release less energy as the elements involved become heavier. As nuclear fuel is depleted the rate of energy released into the star's core drops below the threshold required to support against gravitational collapse, the core of the star begins to compress resulting in increases in temperature and density. This increase allows for the possibility of fusion reactions involving increasingly heavier elements. This fusion process can continue up until the production of iron and nickel for sufficiently massive stars ($M \gtrsim 10M_\odot$). At this point no further energy is deposited into the core from fusion reactions, due to the products reaching maximal binding energy.

With no more energy being released in the core of the star it begins to gravitationally collapse in on itself. Once the mass of the core reaches the Chandrasekhar limit ($\approx 1.4 M_\odot$ [6] [7]) the electron degeneracy pressure (arising from the Pauli exclusion principle which forbids identical fermions from occupying the same energy states) is no longer sufficient to support the core. It becomes energetically favourable for the protons and electrons in the core to combine via electron capture, producing neutrons and releasing neutrinos, forming a proto-neutron star. The density of this core compresses beyond that of nuclear density and becomes supported by repulsive neutron interactions and neutron degeneracy pressure, thus halting the collapse. The infalling matter rebounds off of this core, creating an outgoing shock wave which accelerates stellar material outwards. The neutrinos generated in the electron capture process typically react very weakly with matter, but under such extreme conditions about 1% of their energy is deposited into the shock wave. This small fraction is believed to be sufficient to result in a supernova explosion. However, effects due to magnetic fields and angular rotation may also play a role in this explosion mechanism, this process is not yet fully understood. If the core mass reaches beyond the Tolman-Oppenheimer-Volkoff (TOV) limit ($\approx 2M_\odot$), then the neutron degeneracy pressure is still insufficient to support against collapse, and the core collapses to a black hole. If the core mass does not reach the Chandrasekhar limit, then the progenitor star collapses to a white dwarf, and no supernova explosion occurs as there is no neutrino emission.

1.2.2 Neutron Star Properties

Neutron stars typically have a mass in the range $1.4M_\odot \lesssim M \lesssim 3M_\odot$, and a radius in the range $10 \text{ km} \lesssim R \lesssim 15 \text{ km}$ [6]. This is illustrated in figure 1.1. A neutron star's compactness ($(M/M_\odot)/(R/\text{km})$) is approximately 105 times greater than that of the Sun. They are extremely dense objects, with very strong gravitational acceleration at their surface, over 10^{12} m s^{-2}. As a comparison the gravitational acceleration at the surface of the Earth is around 10 m s^{-2}, and at the surface of the Sun it is around 274 m s^{-2}. There is significant spacetime curvature around neutron stars, requiring general relativistic effects to be included for accurate modelling. Thus our hydrodynamic and radiation evolution descriptions (introduced in section 2) must include general relativistic effects.

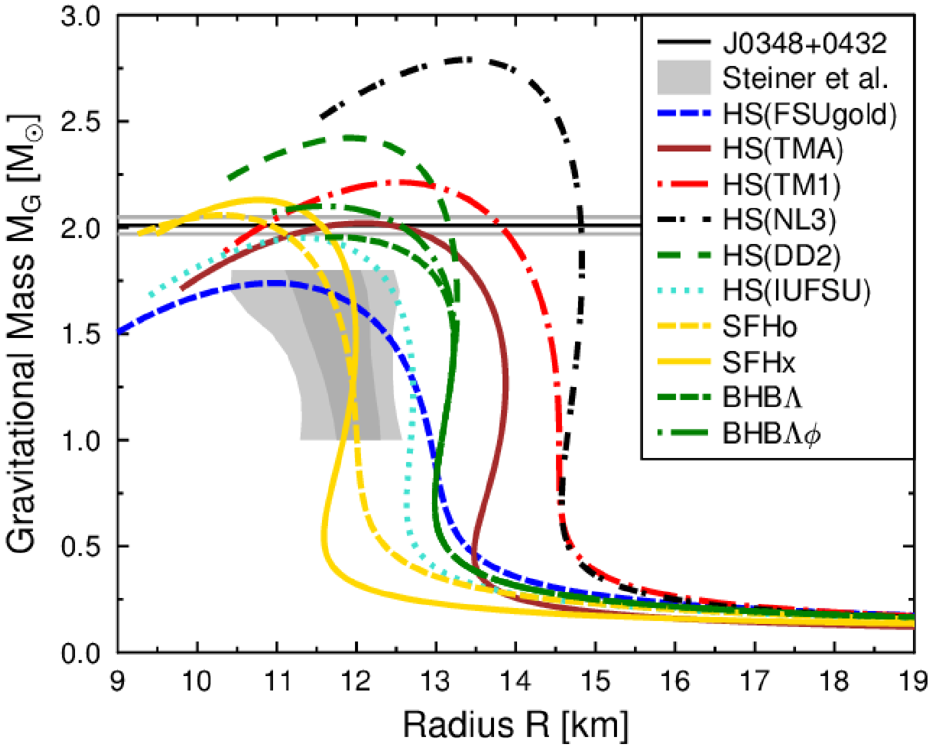

Figure 1.1: Neutron star mass-radius curves for various supernova equations of state. The grey region covers the expected range of mass and radius for cold neutron stars which have reached equilibrium states. Image from [8], **Copyright** © T. Fischer, M. Hempel, I. Sagert, Y Suwa, and J Schaffner-Bielich.

1.2.3 Neutron Star Binaries

For neutron star binaries, most often a pair of progenitor stars are ejected from a star forming region, eventually leading to a pair of neutron stars [9, 10]. If the two neutron stars are in a tight enough orbit to merge within the age of the universe ($P_{orbit} \lesssim 1$ day) the system will lose energy in the form of gravitational waves. This results in the system losing angular momentum, reducing the radius of the orbit at an increasing rate until the two compact objects eventually merge.

There are thought to be up to a billion neutron stars throughout the Milky Way. There are around $3,200$ known neutron stars in the Milky Way and the Magellanic Clouds, and around 5% of them exist in a binary system (either with a main sequence star, another neutron star, or a black hole). Ten double neutron star binaries are currently known [11].

1.2.4 Neutron Star Mergers

Neutron star mergers feature some of the most extreme physical conditions to exist in the entire universe, including densities greater than that of atomic nuclei, velocities approaching half the speed of light, and temperatures exceeding 10^{11}K. Neutron star mergers are the source of some gravitational wave detections. In recent cases a neutron star merger has been detected via multiple independent means. Specifically, a gravitational wave detection can have an electromagnetic counterpart [12]. The detection of multiple signals from astrophysical events is known as multi-messenger astronomy. The gravitational waves originating from neutron star mergers include the characteristic inspiral, merger, and post-merger (sometimes known as ringdown) phases. The electromagnetic counterpart signals span a range of frequencies including the optical and ultra-violet band [13], as well as radio, X-ray, and γ-ray emissions [14].

The gravitational wave detection GW170817 on the 17th of August 2017 was the first observation of a gravitational wave with an electromagnetic counterpart. This optical kilonova signal is thought to be caused by the radioactive decay of r-process synthesised nuclei in the ejected matter.

1.2.5 Neutron Star Merger Remnants

The compact remnant of a neutron star merger can promptly collapse to a black hole due to the mass of the remnant object significantly exceeding the TOV limit. Alternatively, if the remnant object has sufficient total support from thermal and rotational effects, as the remnant cools, or as its rotation slows over time this support can be reduced, leading to a delayed collapse to a black hole. The gravitational wave observation GW170817 collapsed to a black hole within a couple of seconds. The observations of objects in a mass range close to the TOV limit can help us to constrain our models of relativistic equations of state.

This work focuses specifically on the case where a pair of (light) neutron stars merge, and do not quickly collapse into a black hole, instead forming a (heavy) neutron star. I aim to study the effects that photon transport has on the evolution of a binary neutron star merger remnant. Radiative transfer can increase the energy loss and cooling rate of the remnant providing a channel for the loss of thermal support. As energy is exchanged between the fluid and the radiation field, energy-momentum can be redistributed around the system affecting its angular momentum. A transition from differential to uniform rotation can reduce the remnant's rotational support. These effects can sufficiently reduce the remnant's support over time, leading to a delayed collapse.

This effect would be noticeable in the observable gravitational wave and electromagnetic wave signals. As the angular momentum is affected so too are the observable gravitational waves emitted by the rotating remnant. Radiative transfer plays a role in creating the observable electromagnetic signals, the luminosity, spectrum, and timings of the signals depend directly on the photon transport.

1.2.6 Merger Simulation

To study this I will be making use of a precalculated binary neutron star merger simulation. A pair of neutron stars have been hydrodynamically evolved through inspiral to merger. The solution is reconstructed in a new code, and the evolution has been continued in an efficient computational framework, which takes into account radiative transfer effects.

The remnant consists of a hot, rotating, and oscillating (heavy) neutron star, with a surrounding torus of ejected material. This is a complicated system. I build up towards its description in chapter 6.

1.3 Hydrodynamics

In astrophysical events such as neutron star mergers, the huge amount of material involved renders a particle-based description inefficient and impractical. A more effective approach involves adopting a fluid approximation. This method treats the material as a continuous medium, using bulk properties such as mass density, velocity, and pressure. Astrophysical phenomena often involve the motion of compressible fluids. My goal is to model this fluid's motion using hydrodynamic techniques, enabling us to understand how these quantities evolve across the relevant domain.

A perfect fluid, often referred to as an ideal fluid, is characterized solely by its mass density, energy density, and pressure, without accounting for any thermal conductivity or viscous effects. It is expected that the main factors driving evolution in the fluid will be bulk fluid motion, including shock waves. Any thermal or viscous effects are expected to have minimal influence in comparison. Hence employing an ideal fluid approximation is deemed sufficient at this stage.

The bulk properties of a compressible ideal fluid are governed by a set of balance laws for mass, momentum, and energy. These quantities evolve due to their fluxes, and are sometimes influenced by external factors such as gravitational forces which can induce fluid acceleration. To close this system of conservation laws, an equation of state is required. This equation allows for the calculation of the fluid pressure based on known parameters such as the mass density and internal energy density. With this framework in place, the evolution of bulk fluid motion and shock waves can be accurately described.

1.4 Radiative Transfer

In neutron star mergers fluid temperatures can reach levels where effects due to radiative transfer significantly influence the fluid's physical properties. The fluid behaves as a black-body emitter, emitting electromagnetic radiation according to Planck's Law. The wide range of temperatures and densities encountered in neutron star mergers leads to varying opacities. The mean-free-paths of emitted radiation determine whether energy is absorbed back into the system or radiated outward. Radiation emission can escape as observable signals, providing valuable insights into the processes occurring during the merger, as demonstrated in the observed signal from GW170817 [14].

The radiative transfer equation describes the motion, emission, absorption, and scattering of a radiation intensity field. By integrating angular moments of this equation, we can describe the evolution of radiation energy and radiation momentum, determined by their respective fluxes and source terms. These evolution equations are linked to the hydrodynamic evolution equations through absorption and emission mechanisms, allowing for the exchange of energy and momentum between radiation and fluid. To fully describe the evolution of the radiation field, a closure relation is required. This approximates the radiation pressure tensor from known quantities such as the radiation energy density and the radiation flux. This tensor transitions from a fully isotropic form in diffusive scenarios to a fully anisotropic form in free-streaming scenarios. This transition depends on the magnitude of radiation flux relative to radiation energy density. For minimal flux magnitudes, the radiation evolves diffusively, whilst for flux magnitudes comparable to radiation energy density, the radiation streams freely.

1.5 Relativistic Effects

The aftermath of a binary neutron star merger results in the formation of a compact object. This object has significant gravitational forces at its surface. The resulting neutron star or black hole significantly curves the surrounding spacetime, leading to relativistic phenomena such as gravitational time dilation. Furthermore, the rotation of compact objects induces frame dragging effects. The curvature of spacetime around these compact objects can be described with a spacetime metric.

During a neutron star merger, the fluid velocities can reach a significant fraction of the speed of light. This high-speed fluid motion causes special relativistic effects such as Lorentz contraction, which can significantly affect the fluid's evolution.

In order to accurately incorporate these relativistic effects, I measure quantities in both the observer reference frame and the fluid reference frame. The conservation equations governing mass, momentum, and energy, whilst accounting for general relativistic effects, can be mathematically expressed through the vanishing covariant derivative of the mass density current and the energy-momentum tensor. This tensor has contributions from both the fluid and the radiation field and can be split into its distinct fluid and radiation components. Consequently, this framework results in a system of balance laws which govern the evolution of mass, fluid momentum, fluid energy, radiation momentum, and radiation energy. These equations involve fluxes, and sources which arise from gravitational forces and radiation-hydrodynamic coupling.

1.6 Numerical Methods

The hydrodynamic evolution of the fluid within a neutron star merger remnant is complex, largely due to the nonlinear behaviour of propagating shock fronts. Consequently, it is appropriate to employ numerical techniques to compute a solution. Closed-form solutions and perturbation approaches, whilst valuable in certain contexts, lack the generality required to accurately capture the diverse range of behaviors expected in this scenario.

The set of balance laws form a system of partial differential equations which involve temporal and spatial derivatives. These equations are designed to evolve quantities which may contain sharp spatial gradients, reflecting the nature of the astrophysical phenomena I aim to model. such as the propagation of shock waves.

I make use of high-resolution shock-capturing methods which are capable of handling discontinuities and traveling shock fronts. Employing the method of lines, I discretise the system of balance laws through a finite volume scheme. This approach involves calculating the fluxes at cell interfaces and evaluating the source terms at cell centres, allowing us to update the cell-centred conserved quantities.

In order to calculate the fluxes at cell interfaces, the known cell-centred quantities must be reconstructed to either side of the cell interfaces. Polynomial reconstruction approaches determine the spatial order of accuracy of the numerical scheme. The reconstruction must be limited so as to avoid introducing any new local extremal values. Once quantities have been reconstructed to either side of a cell interface, it is treated as a Riemann problem. The signal speeds either side of the interface dictate how the interface flux will be calculated, and whether the wave is a shock, rarefaction, or a contact wave. With the flux and source terms calculated, the evolved quantities can be updated by a single time-step. For this system to be solved explicitly the time-step must be restricted by the time taken for the fastest wave-speed to traverse a single cell-width, this Courant–Friedrichs–Lewy (CFL) condition is a requisite for numerical stability. The evolved quantities can be updated with a time integration technique such as a Runge-Kutta method, this determines the temporal order of accuracy for the numerical scheme. Once the evolved quantities have been updated, the set of primitive quantities must be recovered. From these the fluxes and source terms can be calculated for the next update. In this case the primitive quantity recovery requires iterative approaches.

1.7 Adaptive Mesh Refinement

The spatial scales involved in binary neutron star mergers, on which these hydrodynamic quantities vary, range from the vast diameter of the debris torus ($\approx 10^6$ m), all the way down to the widths of the shock fronts propagating through the fluid, and the resulting turbulent flows which form behind them ($\approx 10^{-3}$ m). An efficient numerical approach to a problem involving such a vast range of spatial scales is the use of adaptive mesh refinement (AMR), in which certain regions of the spatial domain are adaptively discretised at a higher spatial resolution than they otherwise would be. The adaptivity can be implemented to track features of interest throughout the simulation. Regions of interest can be tagged for refinement based on the values of physical quantities within them. In numerical simulations, unphysical approximations at the domain boundaries can cause errors to propagate into the regions of interest. AMR allows for the domain boundaries to be placed significantly far away, such that the effects of these errors are sufficiently small, take a long time to affect results, or are able to dissipate.

1.8 Statistical Results

In astrophysical contexts, encounters between dense, cool fluids and hot, rarefied fluids are frequent occurrences. Examples include various phenomena such as the interstellar medium, circumgalactic medium, intracluster medium, as well as supernova remnants, superbubbles, cosmic filaments, galactic winds, protoplanetary disks, protostellar jets, and active galactic nuclei jets. Burning and energy release often takes place within the turbulent media in stellar interiors and during supernova events [15]. These fluids typically maintain an approximate state of hydrostatic and local thermodynamic equilibrium. At the interface where these fluids shear against each other, a mixing layer forms, leading to the onset of a Kelvin-Helmholtz instability (KHI).

In scenarios where KHI occurs, it is often challenging to precisely determine the initial conditions with the required accuracy to replicate the resulting turbulence using physics models. However, by conducting numerous simulations with initial conditions covering a range of likely scenarios, we can obtain a distribution of results. This distribution allows us to predict the expected behavior with a quantified level of uncertainty.

In the Newtonian limit of hydrodynamics, it has been observed that in numerical simulations of KHI, a conserved quantity does not converge locally as spatial resolution increases. However, conducting numerous simulations with a range of initial condition perturbations, leading to varied turbulent motion, reveals that both the mean and variance of the quantity taken across multiple samples does indeed converge with increasing spatial resolution. This convergence also extends to other ensemble statistics derived from the conserved quantities.

I investigate whether the same statistical results hold for relativistic hydrodynamics, and how introducing radiative transfer affects the solutions.

1.9 Other Relevant Physical Processes Outside the Scope of this Project

Other physical processes which are expected to affect the evolution of astrophysical phenomena are magnetic fields, electrically charged fluids, neutrino radiative transfer, non-ideal fluid effects, and chemical processes. These mechanisms will not be the focus of this project. However neutrino radiative transfer shares many similarities with electro-magnetic radiative transfer processes.

1.10 Notation and Units

Throughout this report Greek indices indicate all four spacetime dimensions (eg. $\lambda, \mu, \nu \in \{0, 1, 2, 3\}$), whereas Latin indices indicate the three spatial dimensions (eg. $i, j, k \in \{1, 2, 3\}$). The mostly-positive spacetime metric signature $(-, +, +, +)$ is used. The geometric unit system is employed, such that the speed of light in a vacuum, the gravitational constant, and the solar mass are all set to unity: $c = G = M_\odot = 1$. A summary of the conventions and unit conversions is given in table 1.1.

Quantity	cgs units	Geometric units
Length	1 cm	6.771×10^{-6}
Mass	1 g	5.028×10^{-34}
Time	1 s	2.030×10^{5}
Density	1 g/cm^3	1.619×10^{-18}
Energy	1 g·cm^2/s^2	5.593×10^{-55}
Speed	1 cm/s	3.336×10^{-11}
Opacity	1 cm^2/g	9.117×10^{22}
Attenuation	1 /cm	1.477×10^{5}
c	2.998×10^{10} cm/s	1
G	6.674×10^{-8} cm^3/g/s^2	1
$M_\odot$	1.989×10^{33} g	1
m_p	1.661×10^{-24} g	2.338×10^{-56}

Table 1.1: The key units used in this thesis, and how to convert between them. For the fluid temperature we multiply by the Boltzmann constant $k_B = 8.617 \times 10^{-5}$ eV / K such that 1 K corresponds to 8.617×10^{-11} MeV, and for the radiation constant $a_{rad} = 4\sigma/c$, where σ is the Stefan-Boltzmann constant, $a_{rad} = 7.5657 \times 10^{-15} \mathrm{g\,cm^{-1}\,s^{-2}\,K^{-4}} = 2.472721 \times 10^{-13} \mathrm{MeV}^{-4}$.

1.11 Summary

In chapter 2 of this report I outline the background physics necessary for understanding extreme astrophysical phenomena such as neutron star mergers. I will describe compressible hydrodynamics, radiative transfer, general relativistic effects, and appropriate closure schemes. This description results in a system of balance laws for evolving radiation hydrodynamic quantities in a curved spacetime.

In chapter 3 the system of evolution equations is discretised for numerical evaluation, employing a finite volume scheme, reconstruction methods, approximate Riemann solvers, time integration techniques, and primitive recovery algorithms.

In this thesis a new general relativistic radiation hydrodynamics code has been created. In chapter 4, a series of test problems are described to verify and validate that the numerical methods employed are capable of accurately recovering the desired physics, and that the physics implemented is appropriate to explain the astrophysical phenomena of interest. These tests include relativistic hydrodynamic evolution in both smooth and discontinuous regimes, radiative transfer in both diffusive and streaming limits, radiation hydrodynamic coupling, evolution in curved spacetimes, and adaptive mesh refinement performance.

In chapter 5 I create neutron star models, and investigate how they evolve within our numerical simulation. This includes novel treatments of the coordinates to address rotational effects.

In chapter 6 I take remnant data from a precalculated binary neutron star merger simulation and use my numerical code to model the continued evolution, investigating its properties, behaviours, and characteristics.

In chapter 7 I conduct uncertainty quantification using the numerical evolution code to model Kelvin-Helmholtz instabilities. These involve discontinuous and turbulent features similar to the shock-turbulence interactions occurring in extreme astrophysical phenomena. These experiments explore how statistical results in Newtonian hydrodynamics extend to the relativistic hydrodynamic and relativistic radiation hydrodynamic regimes.

Finally in chapter 8, I will conclude and briefly discuss extensions of this research.

The novel work done in this thesis is primarily contained in chapters 5, 6, and 7. The key results which rely on the efficiency of the new code implemented and discussed here are presented in chapters 5-6. The uncertainty quantification results discussed in chapter 7 are aided by the efficiency of the new code and the insights gained from the remnant evolution, but are conceptually separate.

Chapter 2

Background Physics

2.1 Introduction

In this chapter I will outline the background physics required to describe the evolution of extreme astrophysical phenomena. This includes compressible fluid dynamics, radiative transfer, and general relativistic effects. In this thesis we ignore the impact of magnetic fields, although in general these will be crucial for describing such astrophysical behaviour. We will focus on neutron stars as the most extreme objects requiring matter in their description.

Neutron stars have extreme gravitational forces at their surfaces. I have used general relativity to describe how these gravitational forces affect the fluid and radiation fields we are interested in modelling. Compact objects such as white dwarfs, neutron stars, and black holes significantly curve their surrounding spacetime, resulting in general relativistic effects such as gravitational time dilation, and in cases of rotating compact objects, frame dragging effects. The spacetime curvature can be described with a spacetime metric.

Due to the extremely large amount of material involved in astrophysical phenomena such as neutron star mergers, a particle description would not be efficient or practical, a fluid approximation would be more appropriate as outlined in [16–18]. A viable approach is to treat the material as a continuum, and work with its bulk properties such as mass density, velocity, and pressure. The astrophysical phenomena we are interested in involve the motion of a compressible fluid as discussed in [19]. I have described the motion of this fluid with hydrodynamic processes to capture the evolution of these quantities throughout the region of interest.

The properties of a fluid can be described with an energy-momentum tensor. An ideal fluid (or perfect fluid) is fully characterised by its mass density and its pressure tensor. It assumes negligible thermal conductivity and viscous effects within the fluid. When considering supernova explosions and neutron star mergers it is expected that shock waves and bulk fluid motion will be the dominant source of evolution and any thermal or viscous behaviour would have comparatively little effect, thus an ideal fluid approximation will suffice at this stage. The accuracy of the ideal fluid approximation depends on the size of the viscosity coefficients, the isotropic shear terms, and the timescale of interest. For neutron star merger remnants we consider sub-second timescales, and on average small viscosity coefficients (as discussed in [19]). As the purpose of this thesis is to make remnant simulations efficient, in order to run to longer times, as a first approximation the simpler ideal fluid approximation is used.

The bulk properties of a compressible ideal fluid obey balance laws for mass, momentum, and energy. Their evolution is determined by their fluxes, and in some cases by source terms describing a gravitational field, within which the fluid is evolving, resulting in acceleration of the fluid. Balance laws are capable of describing smooth motion as well as discontinuities and shocks. An equation of state is required to close the system of balance laws. This equation of state allows the calculation of the fluid pressure from known quantities such as the internal energy density.

The temperatures of the fluid involved in neutron star mergers can reach sufficiently high values that effects arising from the transport of emitted radiation can have a significant impact on how the physical properties of the fluid evolve. For example, [20] sees temperatures above 50 MeV $\approx 6 \times 10^{11}$ K in the post-merger phase. The high temperature fluid emits electromagnetic radiation as a black-body emitter, with a spectrum described by Planck's Law [21]. The vast range of temperatures and densities involved in a neutron star merger result in a correspondingly vast range of material opacities. Thus the mean-free-paths of the emitted spectral radiation dictate where energy may be deposited within the system via re-absorption, or if the energy can radiate out of the system. Radiation can escape as an observable signal revealing information about the processes involved, as was observed in the signal from supernova (SN)1987A [22], the electromagnetic counterpart of the gravitational wave signal from binary neutron star merger GW 170817, GRB 170817A and AT 2017gfo [12, 23, 24].

The radiative transfer equation describes the motion, emission, absorption, and scattering of a radiation intensity field [25]. Taking the angularly integrated moments of this transfer equation describes the evolution of radiation energy density and radiation momentum, determined by their fluxes and source terms. This is analogous to the conservation laws describing the evolution of a fluid's energy and momentum. These evolution equations can couple to the hydrodynamic evolution equations via radiation absorption and emission source terms, allowing energy and momentum to be exchanged between the radiation and fluid. This system of evolution equations for the radiation field also requires a closure relation, similar to the equation of state for the fluid equations. This closure relation approximates the radiation pressure tensor. The radiation pressure tensor transitions from a fully isotropic form in the case of diffusion radiation, to a fully anisotropic form in the case of streaming radiation. This transition can be approximated using the radiation flux magnitude relative to the radiation energy density. For very small flux magnitudes the radiation evolves diffusively, conversely, for flux magnitudes comparable to the radiation energy density the radiation streams freely along its direction of travel.

The fluid velocities reached during a neutron star merger can approach a significant fraction of the speed of light (see, for example, [26]). Fluid motion at relativistic velocities results in special relativistic effects, which can affect the evolution of the fluid. In order to track any relativistic effects, quantities can be calculated in both the observer reference frame and the co-moving, fluid reference frame. The conservation of mass, momentum, and energy accounting for general relativistic effects can be described by the vanishing of the 4-divergence (covariant derivative) of the mass density current and the energy-momentum tensor. The energy-momentum tensor has contributions from both the fluid and the radiation field and can be split into fluid and radiation parts. This results in a system of balance laws for evolving mass, fluid momentum, fluid energy, radiation momentum, and radiation energy, involving fluxes, and source terms from gravitational forces and radiation-hydrodynamic coupling.

I build up towards the complex system of coupled radiation and hydrodynamic evolution equations on a curved spacetime background, fully detailed in section 2.7. This system of equations form the basis for a numerical simulation code outlined in chapter 3.

2.2 Spacetime in General Relativity

The extreme mass density of a neutron star (exceeding that of an atomic nucleus) curves its surrounding spacetime, resulting in strong gravitational forces near its surface. The evolution of the fluid and radiation around this compact object is affected by this curvature in spacetime. The spacetime curvature is dominated by the mass contained within the compact neutron star, allowing us to approximate the spacetime as stationary, and treat the evolving surrounding material as non-self-gravitating (a test fluid). This stationary spacetime approximation (the Cowling approximation [27]) allows us to avoid the expensive calculations required for spacetime evolution.

2.2.1 3+1 Split

The general relativistic spacetime interval (line-element) describes the differential distance in spacetime,

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu, \tag{2.1}$$

where $g_{\mu\nu}$ is the spacetime 4-metric. Following the standard $3+1$ decomposition of spacetime [28, 29], the four dimensional spacetime is foliated into a set of non-intersecting space-like hypersurfaces Σ_t [30], where t is a coordinate time used to label each hypersurface. This 3+1 decomposition is useful because we wish to arrive at a system of hyperbolic partial differential equations, in the form of an initial value problem, with which we will calculate the time evolution of an initial condition on a set of coordinates within a spatial domain. The lapse function α is used to measure the proper time between adjacent hypersurfaces, and the shift vector β^i measures the change of coordinates between hypersurfaces. The spatial 3-metric γ_{ij} is used to describe the spatial curvature within a hypersurface [30]. These quantities are illustrated in Figure 2.1.

The spacetime line-element in equation (2.1) can be written in terms of this $3+1$ decomposition,

$$ds^2 = \left(-\alpha^2 + \beta_k \beta^k\right) dt^2 + 2\beta_k dx^k dt + \gamma_{ij} dx^i dx^j. \tag{2.2}$$

The spacetime 4-metric and its inverse can be used to switch between the covariant and contravariant forms of tensors, $V_\mu = g_{\mu\nu} V^\nu$, and $V^\mu = g^{\mu\nu} V_\nu$. The spacetime 4-metric can be written in terms of the $3+1$ decomposition,

$$g_{\mu\nu} = \begin{pmatrix} -\alpha^2 + \beta_k \beta^k & \beta_i \\ \beta_j & \gamma_{ij} \end{pmatrix}, \tag{2.3}$$

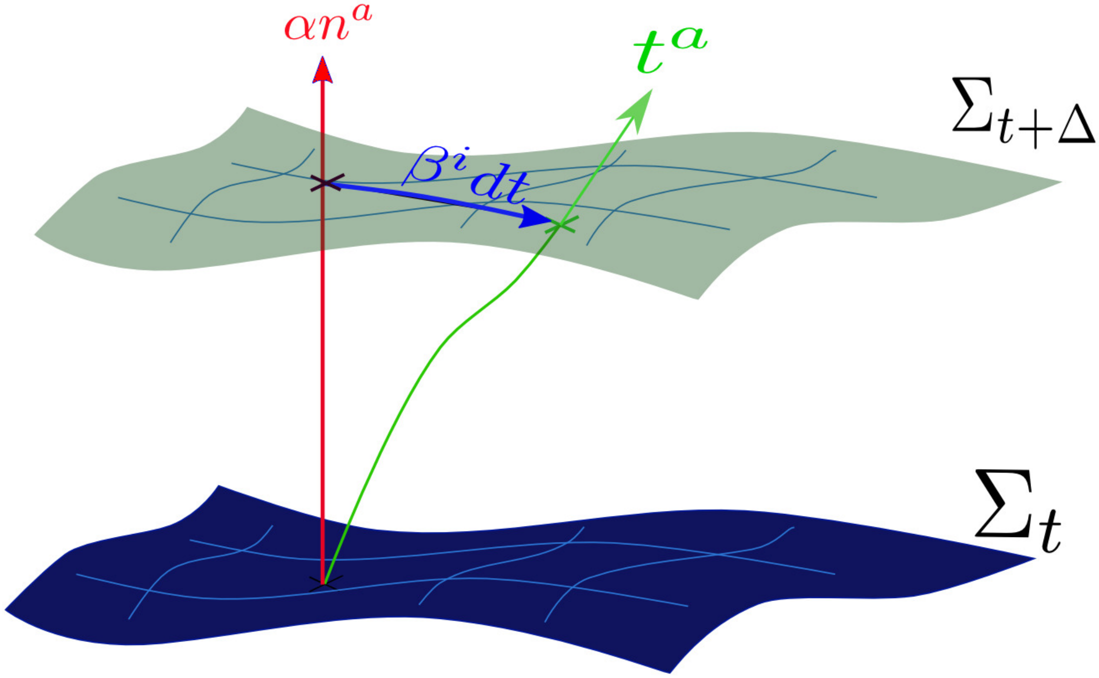

Figure 2.1: A pair of adjacent space-like hypersurfaces showing how the lapse function and shift vector describe changes in coordinates between them. Image from [31], **Copyright** © 2020 Palenzuela.

$$g^{\mu\nu} = \begin{pmatrix} -\alpha^{-2} & \beta^i\alpha^{-2} \\ \beta^j\alpha^{-2} & \gamma^{ij} - \beta^i\beta^j\alpha^{-2} \end{pmatrix}. \tag{2.4}$$

The future pointing, time-like unit vector orthogonal to the space-like hypersurface is

$$n^\mu = \left(\alpha^{-1}, -\beta^i\alpha^{-1}\right), \tag{2.5}$$

$$n_\mu = (-\alpha, 0_i). \tag{2.6}$$

This can be used to project tensors along the direction orthogonal to the space-like hypersurface.

2.2.2 Example Spacetime Metrics

Here some standard spacetime metrics are introduced, they illustrate the convenience of the 3+1 split, and will be useful for validation tests with simple spacetimes in section 4.

Minkowski Spacetime

The Minkowski spacetime metric describes the flat, static spacetime used in special relativity [32]. It uses the simple Euclidean space and allows for time-dilation, Lorentz-contraction, and non-rest-mass special relativistic effects. Its $3+1$ decomposed description in the Cartesian coordinate system gives the spacetime interval

$$ds^2 = -dt^2 + dx^2 + dy^2 + dz^2, \tag{2.7}$$

and thus the spacetime metric components are,

$$\alpha = 1, \tag{2.8}$$

$$\beta^i = 0^i, \tag{2.9}$$

$$\gamma_{ij} = \delta_{ij}. \tag{2.10}$$

For spherically symmetric situations of interest, spherical polar coordinates can provide a convenient description. In spherical polar coordinates the Minkowski spacetime interval is

$$ds^2 = -dt^2 + dr^2 + r^2 d\Omega^2, \tag{2.11}$$

where $d\Omega^2 = d\theta^2 + \sin^2\theta d\phi^2$, and r, θ, ϕ are the radial, polar and equatorial coordinates respectively. The corresponding spacetime metric has the non-zero components

$$\alpha = 1, \tag{2.12}$$

$$\gamma_{rr} = 1, \tag{2.13}$$

$$\gamma_{\theta\theta} = r^2, \tag{2.14}$$

$$\gamma_{\phi\phi} = r^2 \sin^2\theta. \tag{2.15}$$

Schwarzschild Spacetime

The exterior Schwarzschild spacetime metric describes the spherically symmetric static spacetime surrounding a non-rotating and uncharged compact object, in a vacuum. It is the unique solution to the Einstein equations for the exterior spacetime for any compact object such as a black hole or neutron star under these conditions. The $3+1$ decomposed description in Schwarzschild coordinates of the spacetime interval is

$$ds^2 = -\left(1 - \frac{2GM}{c^2 r}\right) dt^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} dr^2 + r^2 d\Omega^2, \tag{2.16}$$

and thus its non-zero metric components are

$$\alpha = \sqrt{1 - \frac{2GM}{c^2 r}}, \tag{2.17}$$

$$\gamma_{rr} = \left(1 - \frac{2GM}{c^2 r}\right)^{-1}, \tag{2.18}$$

$$\gamma_{\theta\theta} = r^2, \tag{2.19}$$

$$\gamma_{\phi\phi} = r^2 \sin^2\theta. \tag{2.20}$$

Here M is the mass of the compact object resulting in the spacetime curvature, G is the gravitational constant, and c is the speed of light in a vacuum, included explicitly. The Schwarzschild spacetime contains a physical singularity at $r = 0$. There are other coordinate systems that could be used to, for example, show that the singularity at $r = 2GM/c^2$ is purely a coordinate effect. However, we will not need them further in this work.

2.3 Relativistic Fluid and Radiation Fields

In order to evaluate the equations for general relativistic radiation hydrodynamic evolution, they must first be written in terms of calculable quantities. The rest-mass density as measured in the co-moving fluid reference frame is $\rho = \sum_x N_x m_x$, where summation is over baryon species, m_x is the baryon mass and N_x is the number of baryons of that type per unit volume. The matter current density (J^μ) is the product of this rest-mass density and the fluid's 4-velocity (u^μ). The 4-velocity is the rate of change of 4-position along an object's world line, with respect to its proper time (αdt), which is the time as measured by the object. From this we have

$$J^\mu = \rho u^\mu. \tag{2.21}$$

The fluid's 4-velocity can be related to the fluid's 3-velocity (v^i) via

$$u^\mu = W\left(\alpha^{-1}, v^i - \beta^i \alpha^{-1}\right), \tag{2.22}$$

$$u_\mu = W\left(-\alpha + v_k \beta^k, v_i\right). \tag{2.23}$$

Here we have used the Lorentz factor

$$W = \frac{1}{\sqrt{1 - v_k v^k}}, \tag{2.24}$$

to measure special relativistic effects. It is necessary for accurate calculation when the fluid velocity is a significant fraction of the speed of light in a vacuum [33, 34]. We remember that in this work we choose units where the speed of light is set to $c = 1$.

The energy-momentum tensor for an ideal fluid can be written as

$$T_M^{\mu\nu} = \rho\left(1 + \epsilon\right) u^\mu u^\nu + P(g^{\mu\nu} + u^\mu u^\nu), \tag{2.25}$$

where P is the isotropic fluid pressure and ϵ is the specific internal energy as measured in the co-moving fluid frame [33–35]. These are related to the specific enthalpy of the fluid $h = 1 + \epsilon + P/\rho$.

The ideal fluid approximation assumes negligible heat conduction and viscosity terms. The contraction of this energy-momentum tensor with the fluid's 4-velocity and any rank-1 tensor orthogonal to the 4-velocity vanishes. This shows that the ideal fluid pressure tensor has no anisotropic components as measured in the co-moving fluid frame. For a non-ideal fluid we would not be able to write the rightmost term as

$P(g^{\mu\nu} + u^\mu u^\nu)$ with a scalar fluid pressure, due to anisotropic components in the fluid frame pressure tensor.

The radiation field can be considered a fluid composed of photons, as opposed to baryons. It too can be described with an energy scalar, momentum vector, and pressure tensor. The energy-momentum tensor for a radiation field can be written as

$$T_R^{\mu\nu} = Eu^\mu u^\nu + F^\mu u^\nu + F^\nu u^\mu + P^{\mu\nu}, \tag{2.26}$$

where E, F^μ, $P^{\mu\nu}$ are the radiation energy density, momentum density, and pressure tensor as measured in the fluid frame. Their temporal components can be calculated from orthogonality conditions, $F^\mu u_\mu = 0$, $P^{\mu\nu} u_\mu = 0^\nu$. Thus $F^0 = v_i F^i/(\alpha - v_k \beta^k)$, $P^{i0} = v_j P^{ij}/(\alpha - v_k \beta^k)$, and $P^{00} = v_i v_j P^{ij}/(\alpha - v_k \beta^k)^2$. In general we would have non-zero anisotropic components arising from the $P^{\mu\nu}$ and F^μ tensors.

2.3.1 Observer Frame Quantities

The co-moving fluid frame quantities $(\rho, \epsilon, P, E, F^i, P^{ij})$ can be used to define the fluid and radiation tensors in terms of the fluid 4-velocity (u^μ). We now require these energy-momentum tensors in terms of the time-like normal vector (n^μ), to show how these quantities appear in the coordinates used in our numerical simulation. This gives (as shown in [36, 37]),

$$T_M^{\mu\nu} = Un^\mu n^\nu + S^\mu n^\nu + S^\nu n^\mu + W^{\mu\nu}, \tag{2.27}$$

$$T_R^{\mu\nu} = \overline{U} n^\mu n^\nu + \overline{S}^\mu n^\nu + \overline{S}^\nu n^\mu + \overline{W}^{\mu\nu}. \tag{2.28}$$

Here U, S^μ, and $W^{\mu\nu}$ are the energy, momentum, and pressure tensors as measured in the observer-frame (here an overline is used to denote the radiation quantities). These are obtained by projecting the known fluid-frame energy-momentum tensors along the time-like normal, and into the space-like hypersurfaces. They are related to the fluid-frame quantities via

$$\begin{array}{lclcl}
U & = & T_M^{\mu\nu} n_\mu n_\nu & = & \rho h W^2 - P, \\
S^i & = & -T_M^{\mu\nu} n_\mu \gamma^i_\nu & = & \rho h W^2 v^i, \\
W^{ij} & = & T_M^{\mu\nu} \gamma^i_\mu \gamma^j_\nu & = & \rho h W^2 v^i v^j + P\gamma^{ij}, \\
\overline{U} & = & T_R^{\mu\nu} n_\mu n_\nu & = & E W^2 + 2W\alpha F^0 + \alpha^2 P^{00}, \\
\overline{S}^i & = & -T_R^{\mu\nu} n_\mu \gamma^i_\nu & = & E W^2 v^i + W\left[F^0(\alpha v^i + \beta^i) + F^i\right] + \alpha P^{i0}, \\
\overline{W}^{ij} & = & T_R^{\mu\nu} \gamma^i_\mu \gamma^j_\nu & = & E W^2 v^i v^j + W\left[(F^i + \beta^i F^0) v^j + (F^j + \beta^j F^0) v^i\right] + P^{ij},
\end{array} \tag{2.29}$$

with $S^0 = \overline{S}^0 = 0$, and $W^{\mu 0} = \overline{W}^{\mu 0} = 0^\mu$.

The numerous terms involved illustrate the apparent anisotropic stresses resulting from the use of coordinates defined by the time-like normal, which are not aligned with the fluid's 4-velocity.

2.3.2 Slow-Motion Limit

When the fluid velocity magnitude is sufficiently small the slow-motion approximation can be used to simplify the equations involved. This is useful to compare with the Newtonian limit, in which much of our intuition is derived. Taking the limit as $v^i \to 0^i$, so that $W \to 1$,

$$D = \rho \tag{2.30}$$

$$U = \rho + \rho\epsilon \tag{2.31}$$

$$S^i = 0^i \tag{2.32}$$

$$W^{ij} = P\gamma^{ij} \tag{2.33}$$

$$\overline{U} = E \tag{2.34}$$

$$\overline{S}^i = F^i \tag{2.35}$$

$$\overline{W}^{ij} = P^{ij}. \tag{2.36}$$

We will discuss the Newtonian equations in more detail in section 2.9.

2.4 Field Evolution

The covariant derivative (also known as the 4-divergence) of the rest-mass density current J^μ vanishes, $\nabla_\mu (J^\mu) = 0$, and gives the relativistic continuity equation. This equation describes the conservation of the number of baryons within a given control volume. The covariant derivative of a rank-1 tensor is

$$\nabla_\mu\left(J^\mu\right) = \frac{1}{\alpha\sqrt{\gamma}}\partial_\mu\left(\alpha\sqrt{\gamma}J^\mu\right) \tag{2.37}$$

$$= 0. \tag{2.38}$$

This gives the rest-mass density evolution equation,

$$\partial_\mu\left(\alpha\sqrt{\gamma}J^\mu\right) = 0. \tag{2.39}$$

This is the continuity equation written in terms of partial derivatives in the coordinate system defined by our $3+1$ decomposition. We can re-write this as

$$\partial_t\left(\alpha\sqrt{\gamma}J^0\right) + \partial_i\left(\alpha\sqrt{\gamma}J^i\right) = 0. \tag{2.40}$$

Here the differential operator has been split into its temporal and spatial parts. The 4-dimensional volume element is given as $\alpha\sqrt{\gamma} = \sqrt{-g}$, where $\gamma = \det\left(\gamma_{ij}\right)$ and $g = \det(g_{\mu\nu})$ are the determinants of the spatial 3-metric and spacetime 4-metric respectively.

From the Einstein equation and the Bianchi identities, the covariant derivative of the energy-momentum tensor $T^{\mu\nu}$ also vanishes, $\nabla_\mu\left(T^{\mu\nu}\right) = 0^\nu$. This leads to the relativistic balance laws for energy and momentum.

The covariant derivative contracted with a rank-2 tensor such as the energy-momentum tensor gives [29, 35]

$$\nabla_\mu\left(T^{\mu\nu}\right) = \frac{1}{\alpha\sqrt{\gamma}}\partial_\mu\left(\alpha\sqrt{\gamma}T^{\mu\nu}\right) - T^{\mu\lambda}\Gamma^\nu_{\mu\lambda} \tag{2.41}$$

$$= 0^\nu. \tag{2.42}$$

Here geometric source terms are introduced, unlike in the relativistic continuity equation, they involve Christoffel symbols $\Gamma^\nu_{\mu\lambda}$ which contain derivatives of the spacetime metric,

$$\Gamma^\nu_{\mu\lambda} = \frac{1}{2}g^{\nu\kappa}\left(\partial_\mu g_{\kappa\lambda} + \partial_\lambda g_{\mu\kappa} - \partial_\kappa g_{\mu\lambda}\right). \tag{2.43}$$

These geometric source terms describe the effect the spacetime curvature has on the energy-momentum tensor, such as gravitational acceleration and frame dragging effects.

Contracting equation (2.41) with n_ν gives

$$\nabla_\mu \left(T^{\mu\nu} n_\nu\right) = n_\nu \nabla_\mu (T^{\mu\nu}) + T^{\mu\nu} \nabla_\mu (n_\nu) = T^{\mu\nu} \nabla_\mu (n_\nu). \tag{2.44}$$

Recalling that $n_\nu = (-\alpha, 0_i)$, so that

$$\nabla_\mu \left(-\alpha T^{\mu 0}\right) = T^{\mu\nu} \nabla_\mu (n_\nu), \tag{2.45}$$

this gives the relativistic conservation of energy equation

$$\partial_\mu \left(\alpha^2 \sqrt{\gamma} T^{\mu 0}\right) = -\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\mu (n_\nu). \tag{2.46}$$

This source term can be expanded as [35]

$$-\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\mu (n_\nu) = \sqrt{\gamma} \left(\alpha T^{ik} K_{ik} - \alpha T^{j0} \partial_j \alpha\right), \tag{2.47}$$

where $K_{\mu\nu} = -(\delta^\lambda_\mu + n^\lambda n_\mu) \nabla_\lambda (n_\nu)$ is the extrinsic curvature, and is orthogonal to the time-like normal ($K_{\mu\nu} n^\mu = 0_\nu$). Thus its temporal components are $K_{i0} = \beta^i K_{ij}$ and $K_{00} = \beta^i \beta^j K_{ij}$.

When the spacetime is stationary this source term can be written as [35]

$$-\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\mu (n_\nu) = \sqrt{\gamma} \left(\frac{1}{2} T^{ik} \beta^j \partial_j \gamma_{ik} + T^j_i \partial_j \beta^i - \alpha T^{j0} \partial_j \alpha\right). \tag{2.48}$$

The energy-momentum equation (2.41) can be expanded as

$$\begin{aligned}
\nabla_\mu \left(T^{\mu\nu}\right) &= \frac{1}{\alpha\sqrt{\gamma}} \partial_\mu \left(\alpha \sqrt{\gamma} T^{\mu\nu}\right) - T^{\mu\lambda} \Gamma^\nu_{\mu\lambda} && (2.49)\\
&= g^{\nu\lambda} \left[\frac{1}{\alpha\sqrt{\gamma}} \partial_\mu \left(\alpha \sqrt{\gamma} T^\mu_\lambda\right) - \frac{1}{2} T^{\alpha\beta} \partial_\lambda g_{\alpha\beta}\right] && (2.50)\\
&= 0^\nu. && (2.51)
\end{aligned}$$

Restricting the free index ν to the spatial index j gives the relativistic conservation of momentum equations [35]

$$\partial_\mu \left(\alpha \sqrt{\gamma} T^\mu_j\right) = \frac{1}{2} \alpha \sqrt{\gamma} T^{\mu\lambda} \partial_j g_{\mu\lambda}. \tag{2.52}$$

These source terms can be expanded as [35]

$$\frac{1}{2}\alpha\sqrt{\gamma}T^{\mu\lambda}\partial_j g_{\mu\lambda} = \sqrt{\gamma}\left(\frac{1}{2}\alpha T^{ik}\partial_j\gamma_{ik} + \alpha T^0_i\partial_j\beta^i - \alpha^2 T^{00}\partial_j\alpha\right). \tag{2.53}$$

The equations for conservation of energy (2.46) and momentum (2.52) contain sufficiently in-depth information about the energy-momentum tensor that in conjunction with the continuity equation (2.39), we can use them to model the fluid's evolution.

2.5 Splitting the Energy-Momentum Tensor

For radiation-hydrodynamics calculations the total energy-momentum tensor can be split into its matter part and its radiation part,

$$T^{\mu\nu} = T^{\mu\nu}_M + T^{\mu\nu}_R, \tag{2.54}$$

where $T^{\mu\nu}_M$ is the matter contribution to the energy-momentum tensor, and $T^{\mu\nu}_R$ is the radiation contribution to the energy-momentum tensor. Introducing this split to the energy-momentum conservation equations returns a system of balance laws involving coupling source terms,

$$\begin{aligned}
\nabla_\mu\left(T^{\mu\nu}\right) &= \nabla_\mu\left(T^{\mu\nu}_M + T^{\mu\nu}_R\right) && (2.55)\\
&= \nabla_\mu\left(T^{\mu\nu}_M\right) + \nabla_\mu\left(T^{\mu\nu}_R\right) && (2.56)\\
&= 0^\nu, && (2.57)\\
\nabla_\mu\left(T^{\mu\nu}_M\right) &= G^\nu, && (2.58)\\
\nabla_\mu\left(T^{\mu\nu}_R\right) &= -G^\nu, && (2.59)
\end{aligned}$$

where G^ν is the radiation 4-force, which allows energy and momentum to be transferred between fluid and radiation forms. By decomposing G^ν with respect to the fluid 4-velocity, and comparing with the Newtonian limit, we can write its most general form using both thermal and scattering opacities χ^t and χ^s, and a radiation emission term Θ,

$$G^\nu = \chi^t\left(E - \Theta\right)u^\nu + \left(\chi^t + \chi^s\right)F^\nu. \tag{2.60}$$

2.6 Species Conservation

When the fluid represents a mixture of different particles, the different constituent species can be tracked. For example, consider the populations of electrons, protons, and neutrons within a fluid volume. We aim to construct equations of motion for each separate species.

Let N_x be the number density for electrons, protons, or neutrons, where $x \in \{e, p, n\}$. Here the only baryons taken into consideration are protons and neutrons. Therefore we define the number density of baryons to be $N_b = N_p + N_n$. We assume that the masses of the baryons are approximately equal ($m_b = m_p = m_n$), and that the electron mass is negligible ($m_e = 0$). Thus the total number of baryons in the fluid volume can be determined by $N_b = \rho/m_p$. In addition, in an electrically neutral fluid the number of electrons and protons is equal, $N_e = N_p$. We define the species per baryon fractions to be $Y_x = N_x/N_b \in [0, 1]$. We therefore have that $Y_e = Y_p$ and, $Y_p + Y_n = 1$.

From these definitions we can modify the continuity equation to obtain a species number conservation equation. The species number density can be advected around the spatial domain along with the fluid flux, giving

$$\nabla_\mu \left(Y_x J^\mu\right) = 0. \tag{2.61}$$

This encodes that there are no reactions to change the local particle number, which would not be the case were neutrinos considered, as discussed in [26].

2.7 Balance Laws

Expanding the relativistic balance laws for continuity, matter, radiation, and species conservation (equations (2.39), (2.58), (2.59), and (2.61)) into temporal and spatial components, including any geometric and coupling source terms results in a system of partial differential equations to be solved for the evolution of our general relativistic radiation hydrodynamics calculation. These are

$$\begin{array}{rcrcl}
\partial_t\left(\alpha\sqrt{\gamma}J^0\right) & + & \partial_i\left(\alpha\sqrt{\gamma}J^i\right) & = & 0, \\
\partial_t\left(\alpha\sqrt{\gamma}Y_x J^0\right) & + & \partial_i\left(\alpha\sqrt{\gamma}Y_x J^i\right) & = & 0, \\
\partial_t\left(\alpha^2\sqrt{\gamma}T^{00}_M\right) & + & \partial_i\left(\alpha^2\sqrt{\gamma}T^{i0}_M\right) & = & \sqrt{\gamma}\left(\alpha T^{ik}_M K_{ik} - \alpha T^{j0}_M \partial_j\alpha + \alpha^2 G^0\right), \\
\partial_t\left(\alpha\sqrt{\gamma}T^0_{Mj}\right) & + & \partial_i\left(\alpha\sqrt{\gamma}T^i_{Mj}\right) & = & \sqrt{\gamma}\left(\frac{1}{2}\alpha T^{ik}_M\partial_j\gamma_{ik} + \alpha T^0_{Mi}\partial_j\beta^i - \alpha^2 T^{00}_M\partial_j\alpha + \alpha G_j\right), \\
\partial_t\left(\alpha^2\sqrt{\gamma}T^{00}_R\right) & + & \partial_i\left(\alpha^2\sqrt{\gamma}T^{i0}_R\right) & = & \sqrt{\gamma}\left(\alpha T^{ik}_R K_{ik} - \alpha T^{j0}_R \partial_j\alpha - \alpha^2 G^0\right), \\
\partial_t\left(\alpha\sqrt{\gamma}T^0_{Rj}\right) & + & \partial_i\left(\alpha\sqrt{\gamma}T^i_{Rj}\right) & = & \sqrt{\gamma}\left(\frac{1}{2}\alpha T^{ik}_R\partial_j\gamma_{ik} + \alpha T^0_{Ri}\partial_j\beta^i - \alpha^2 T^{00}_R\partial_j\alpha - \alpha G_j\right).
\end{array} \tag{2.62}$$

These equations are in a form convenient for numerical evolution [38, 39]. The source terms only contain spatial derivatives of the spacetime metric terms, not the matter or radiation quantities. Hence they are well behaved around discontinuities in the matter or fluid fields.

The system of coupled balance laws for general relativistic radiation hydrodynamics can be written in the Valencia formulation [33] as (following [18])

$$\frac{\partial}{\partial t}\left(\underline{U}\right) + \frac{\partial}{\partial x^i}\left(\underline{F}^i\right) = \underline{S}. \tag{2.63}$$

Here $\underline{U}$ is the vector of conserved quantities,

$$\underline{U} = \sqrt{\gamma}\begin{pmatrix} D \\ Y_x D \\ S_j \\ \tau \\ \overline{S}_j \\ \overline{U} \end{pmatrix}. \tag{2.64}$$

The vector $\underline{F}^i$ contains their corresponding fluxes,

$$\underline{F}^i = \sqrt{\gamma}\begin{pmatrix} \alpha D v^i - \beta^i D \\ \alpha Y_x D v^i - \beta^i Y_x D \\ \alpha W^i_j - \beta^i S_j \\ \alpha(\tau + P)v^i - \beta^i\tau \\ \alpha\overline{W}^i_j - \beta^i\overline{S}_j \\ \alpha\overline{S}^i - \beta^i\overline{U} \end{pmatrix}. \tag{2.65}$$

The geometric and coupling source terms are contained in the vector $\underline{S}$,

$$\underline{S} = \sqrt{\gamma}\begin{pmatrix} 0 \\ 0 \\ \frac{1}{2}\alpha W^{ik}\partial_j\gamma_{ik} + S_i\partial_j\beta^i - (\tau + D)\,\partial_j\alpha \\ \alpha W^{ik}K_{ik} - S^j\partial_j\alpha \\ \frac{1}{2}\alpha\overline{W}^{ik}\partial_j\gamma_{ik} + \overline{S}_i\partial_j\beta^i - \overline{U}\partial_j\alpha \\ \alpha\overline{W}^{ik}K_{ik} - \overline{S}^j\partial_j\alpha \end{pmatrix} + \sqrt{\gamma}\begin{pmatrix} 0 \\ 0 \\ \alpha G_j \\ \alpha^2 G^t \\ -\alpha G_j \\ -\alpha^2 G^t \end{pmatrix}. \tag{2.66}$$

To simplify notation we have defined the non-rest mass density $D = \rho W = J^0$. In addition we have defined $\tau = U - D = \rho h W^2 - P - \rho W$, which is a useful way to reformulate the fluid energy equation (2.46) because in non-relativistic cases, the mass density is the dominant source of energy density, meaning the fluid energy U tends to the mass D.

A set of primitive quantities from which each of the fluid and radiation terms in these balance laws can be calculated is $\{\rho, v^i, P, F^i, E\}$, the rest-mass density, fluid velocity, fluid pressure, radiation flux and radiation energy density respectively. These primitive quantities provide an efficient approach to computing each of the terms required for evolution. This will be discussed further in section 3.12.

2.8 Closure Schemes

The set of balance laws requires further information to close the system of equations. This information loosely encapsulates the microscopic behaviour of the material which has been lost when moving to the continuum approximation.

2.8.1 Fluid Equation of State

The isotropic fluid pressure can be calculated from known quantities, such as the fluid temperature and mass density, with an equation of state.

A simple example is the Gamma-law equation of state, requiring a material dependent adiabatic index Γ,

$$P = (\Gamma - 1)\rho\epsilon. \tag{2.67}$$

As a more general example, the fluid pressure can be determined from rest mass density, fluid temperature, and electron fraction,

$$P = P(\rho, T, Y_e). \tag{2.68}$$

The isotropic pressure depends on the small scale forces between particles, their temperature dependent fluctuations, and the types of particles involved. Hence the fluid's composition, which is encoded in the species fraction can be a useful quantity to keep track of. For our purposes the effect of the electron fraction dominates this.

We have access to the CompOSE library of tabulated equations of state [6, 40]. A particularly relevant equation of state is the three parameter SFHx (with electrons) table, which is useful for neutron stars from core collapse supernova calculations and accounts for the electron fraction.

2.8.2 Radiation Pressure Tensor

The radiation pressure tensor can be fully isotropic in the co-moving fluid frame, which recovers diffusive radiation behaviour in the optically thick case. This simple approximation is known as the Eddington closure, and takes the form $P^{\alpha\beta} = \frac{1}{3}E(g^{\alpha\beta} + u^\alpha u^\beta)$. This is known as the $M0$ closure, as it only depends on the zeroth moment of the radiation field (E).

Alternatively the radiation pressure tensor can be fully anisotropic, which recovers streaming behaviour in the optically thin case. Our approach is to approximate the physical radiation pressure tensor with an interpolation between these two limiting cases. This interpolation relies on the fluid-frame radiation quantities E and F^i. This closure scheme is known as the $M1$ method, as it requires up to the first moment of the radiation field. In this case

$$P^{\alpha\beta} = \frac{3\chi(\xi) - 1}{2} P^{\alpha\beta}_{thin} + \frac{3[1 - \chi(\xi)]}{2} P^{\alpha\beta}_{thick}, \tag{2.69}$$

$$P^{\alpha\beta}_{thin} = \frac{F^\alpha F^\beta}{E}, \tag{2.70}$$

$$P^{\alpha\beta}_{thick} = \frac{1}{3}E(g^{\alpha\beta} + u^\alpha u^\beta). \tag{2.71}$$

The interpolation factor $\chi(\xi)$ can take numerous forms. Multiple closure options are explored in [41]. We will focus on the simple interpolation given by the Kershaw closure,

$$\chi(\xi) = \frac{1+2\xi^2}{3}, \tag{2.72}$$

where the anisotropy measure $\xi = \sqrt{F_\mu F^\mu}/E \in [0,1]$, and $\chi(\xi) \in [1/3, 1]$. As required, where the radiation flux is small compared to the radiation energy, the diffusive, fully isotropic radiation pressure tensor is recovered, and when the radiation flux is of similar magnitude to the radiation energy, the streaming, fully anisotropic radiation pressure tensor is recovered.

2.8.3 Opacity

The opacity values are required to evaluate the radiation 4-force terms from equation (2.60). In the Newtonian limit they determine the rate of energy and momentum exchange between matter and radiation forms. For a detailed discussion of the calculation of the opacities in some situations, which I have used for the basis of the radiation hydrodynamics simulations here, see [42]. The thermal and scattering attenuation coefficients are given the forms

$$\chi^t = 1.928 \times 10^{17}\ Y_e^2\ \rho^2\ T^{-7/2}, \tag{2.73}$$

$$\chi^s = 4.387 \times 10^{-24}\ Y_e\ \rho, \tag{2.74}$$

here the mass density is in geometric units, the fluid temperature is in MeV, and attenuation coefficients are the inverse mean-free-path in geometric units. The thermal opacity varies as $\rho^2\ T^{-7/2}$; this is referred to as a Kramers opacity. For our applications we expect the scattering opacity to have a negligible contribution compared to the thermal opacity, this is due to the difference in magnitude of their constant coefficients. These opacities take into account the mass density, electron fraction, and fluid temperature, which are quantities we have access to in our radiation hydrodynamic evolution model.

2.9 Newtonian Limit

In the Newtonian limit, we recover simplified formulations for hydrodynamics and radiation evolution. With fluid velocities significantly less than the speed of light in a vacuum, and negligible spacetime curvature, we obtain equations for radiation and hydrodynamic evolution appropriate for terrestrial applications. From these Newtonian evolution equations a huge amount of numerical techniques have been developed for efficient computational simulation. Many of these techniques are also applicable to the general relativistic regime.

2.9.1 Newtonian Hydrodynamics

The Newtonian limit of hydrodynamic motion introduces a system of equations in conservation form. These equations describe how fluid momentum and pressure can affect the evolution of the fluid quantities [43]. Newtonian gravitational effects can be included via source terms, forming a set of balance laws [44]. This set of equations requires a closure relation which can be included with a material dependent equation of state.

The system of conservation laws describing Newtonian hydrodynamics is

$$\frac{\partial}{\partial t}\left(\underline{U}\right) = \frac{\partial}{\partial x_i}\left(\underline{F}^i\right) + \underline{S}, \tag{2.75}$$

where $\underline{U}$ is the vector of conserved quantities: mass density, momentum densities and total energy density E,

$$\underline{U} = \begin{pmatrix} \rho \\ \rho v^j \\ E \end{pmatrix}. \tag{2.76}$$

The vector $\underline{F}^j$ contains the fluxes corresponding to these conserved quantities,

$$\underline{F}^i = \begin{pmatrix} \rho v^i \\ \rho v^i v^j + P\delta^{ij} \\ \left(E + P\right) v^i \end{pmatrix}. \tag{2.77}$$

The vector $\underline{S}$ contains any source terms. For example gravitational source terms would be described by,

$$\underline{S} = \begin{pmatrix} 0 \\ -\rho \partial_i \phi \\ -\rho v_j \partial_j \phi \end{pmatrix}. \tag{2.78}$$

where ϕ is the Newtonian gravitational potential. In simple cases $\phi = gz$, when using uniform gravitational acceleration towards the negative z-direction.

This Newtonian balance law is analogous to the fluid parts of the general relativistic balance law in equation (2.63).

The equation for conservation of mass density states that mass is advected and there are no source (or sink) terms, meaning that mass is never created or destroyed. The equations for conservation of momentum density state that momenta are advected, there is an acceleration from spatial gradients of fluid pressure, and there can be acceleration effects due to gravitational source terms. The equation for conservation of total energy density states that total energy is advected, there is work done by the pressure term and there can be changes in kinetic energy due to a gravitational source term.

2.9.2 Newtonian Radiative Transfer

The quasi-Newtonian radiative transfer equations introduce how the evolution of the radiation field couples to the fluid properties. Radiation absorption and emission allows the transfer of energy between the radiation and fluid fields.

The time dependent radiative transfer equation (to first order v/c) is a Boltzmann transport equation describing how the intensity field I for electro-magnetic radiation changes over time due to absorption, emission and scattering,

$$\frac{1}{c}\frac{\partial}{\partial t}I + \hat{n}\cdot\nabla I = -\left(\chi^t + \chi^s\right)I + \eta + \chi^s\frac{1}{4\pi}\int_{4\pi} I d\Omega. \tag{2.79}$$

Here, χ^t, χ^s are the absorption and scattering opacities respectively, η is the radiation emission term, and Ω is a solid-angle variable (note: the speed of light in a vacuum is written explicitly as c). Accounting for thermal radiation emission the emission term $\eta = \chi^t B$, where B is the Planckian distribution for black-body emission. In Local Thermodynamic Equilibrium (LTE) the radiation field (I) matches the thermal emission.

The time dependent radiative transfer equation is a seven dimensional equation ($\underline{x}, t, \hat{n}, \nu$) (three spatial, three momentum, and one temporal, we choose to link the three momentum coordinates to two angles and one frequency), which would be incredibly expensive and complex to solve. To reduce the cost of calculating a numerical solution, approximations are required to reduce the dimensionality of the problem.

Diffusion Approximation

An approach to removing the two angular dimensions from the time dependent radiative transfer equation involves taking the zeroth moment of the radiative transfer equation by integrating it over its angular dependence,

$$\frac{\partial E}{\partial t} + \frac{\partial F^i}{\partial x^i} = -c\chi_a \left(E - \frac{4\pi}{c} B \right). \tag{2.80}$$

This Newtonian approach is analogous to the $M0$ closure we introduced for general relativistic radiation evolution in section 2.8.2. Here the radiation energy density E and the radiation fluxes F^i are related to the radiation intensity via,

$$E = \frac{1}{c} \int_{4\pi} I d\Omega, \tag{2.81}$$

$$\mathbf{F} = \int_{4\pi} \hat{n} I d\Omega, \tag{2.82}$$

where $\hat{n}$ are directional unit vectors. This zeroth moment equation can be viewed as a balance law for radiation energy density, with advective transport and a source term coupling it to radiation emission and absorption processes.

In the diffusion approximation the radiation flux is approximated as proportional to the spatial gradient of the radiation energy density. This proportionality factor D is known as the diffusion coefficient,

$$F^i = D \frac{\partial E}{\partial x_i}, \tag{2.83}$$

$$D = \frac{-c}{3\left(\chi^t + \chi^s\right)}. \tag{2.84}$$

This approximation is appropriate for the case of large opacities, and thus short mean free paths. This approach assumes an isotropic radiation intensity and results in the radiation diffusion equation,

$$\frac{\partial E}{\partial t} = \frac{\partial}{\partial x^i} \left(\frac{c}{3\left(\chi^t + \chi^s\right)} \frac{\partial E}{\partial x^i} \right) - c\chi^t \left(E - \frac{4\pi}{c} B \right). \tag{2.85}$$

The diffusion approximation to the radiative transport equation can be solved with implicit numerical techniques, allowing for significantly larger time steps.

Two Moment Approximations

The isotropic restriction on the radiation intensity can be relaxed by taking the first moment of the radiative transfer equation, along with the zeroth moment. Again this Newtonian approach is analogous to the $M1$ method we introduced for general relativistic radiation evolution in equation (2.69). This requires calculating the radiation fluxes using information from the radiation pressure tensor P^{ij},

$$\frac{\partial F^j}{\partial t} + c^2 \frac{\partial P^{ij}}{\partial x^i} = -c \left(\chi^t + \chi^s \right) F^j. \tag{2.86}$$

The radiation pressure tensor is related to the radiation intensity via

$$\mathrm{P} = \frac{1}{c} \int_{4\pi} \hat{n}\hat{n} I d\Omega. \tag{2.87}$$

The system of zeroth and first moment equations is in balance law form, and numerical solutions can be found using the same approach as for the hydrodynamic conservation laws. These first moment equations can be viewed as balance laws for radiation momentum density, with advective transport and flux dependent source terms.

The radiation pressure must be determined via a closure relation, relating it to known or previously calculated quantities, just as was required in the general relativistic case described in section 2.8.2.

Fluid-Radiation Coupling

For radiation coupled fluid dynamics calculations, the energy density of the fluid can be coupled to the radiation field via the material balance equation, which accounts for radiation emission and absorption by the fluid. This is given by

$$\rho \frac{\partial \epsilon}{\partial t} = c \int_0^\infty \chi_a \left(E_\nu - \frac{4\pi}{c} B \right) d\nu. \tag{2.88}$$

When using a constant (but fluid dependent) specific heat capacity at constant volume (C_V) the fluid temperature (T) can be obtained with $T = \epsilon / C_V$. The material balance equation involves integrating radiation quantities over the entire frequency spectrum. When using a multi-group approach this can be approximated by summation over the frequency groups used.

Multi-group Radiation

The radiative transfer equations can be applied to monochromatic (spectral), narrow band, wide band, or full-spectrum frequencies of electro-magnetic radiation. One approach to reducing the expense of working with the time dependent radiative transfer equation is to discretise the continuous frequency (energy) variable. The entire frequency domain can be integrated over to give a 'grey' radiation approximation, or the frequency range of interest can be split into a set of 'frequency groups'. The grey radiation approach can be thought of as using a single, very large frequency group. Discretising the frequency domain can drastically reduce the computational cost of radiative transfer equations compared to highly resolved monochromatic methods. The radiation calculations can be completed for each frequency group separately. The radiation energy density contained within a particular frequency group g, can be calculated by integrating radiation energy density over the corresponding frequency interval,

$$E_g = \int_{\nu_{g-1/2}}^{\nu_{g+1/2}} E d\nu. \tag{2.89}$$

Radiation quantities such as radiation flux and radiation pressure for a given frequency group can be calculated in terms of the multi-group radiation energy density.

The opacity of a material for a given frequency group can be found via a weighted average of the spectral radiation opacities. Two common averaging methods are the Planck mean and the Rosseland mean.

The Planck mean opacity χ_P is optimised for use in the radiation emission in the low opacity regime,

$$\chi_P = \frac{\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} \chi_\nu B d\nu}{\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} B d\nu}, \tag{2.90}$$

where the weighting function B is a function of both frequency and temperature. This weighting function is the Planckian distribution for black-body emission,

$$B(T, \nu) = \frac{h\nu^3}{2c} \left(e^{h\nu/kT} - 1 \right)^{-1}. \tag{2.91}$$

The Rosseland mean opacity χ_R is optimised for use in the radiation flux in the high opacity regime,

$$\chi_R^{-1} = \frac{\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} \chi_\nu^{-1} \frac{\partial B}{\partial T} d\nu}{\int_{\nu_{g-1/2}}^{\nu_{g+1/2}} \frac{\partial B}{\partial T} d\nu}, \tag{2.92}$$

where a harmonic average has been employed.

Some advantages of using a multi-group radiation scheme are that it captures behaviours at different frequencies, the regions of the frequency spectrum to concentrate on can be chosen freely, there are less computations involved than with a monochromatic approach, and it introduces a convenient format for storage of radiation quantities and data such as the multi-group opacities. However some disadvantages of using a multi-group scheme are that it may become computationally expensive as more frequency groups are used, and the black-body emission term B requires numerical integration over each frequency group. The grey black-body emission integrated over all frequencies has a very simple form,

$$\int_0^\infty B(T,\nu) d\nu = \frac{\sigma}{\pi} T^4, \tag{2.93}$$

where σ is the Stefan-Boltzmann constant.

2.10 Summary

In this chapter I have introduced a system of balance laws which describe the evolution of a compressible fluid and a radiation field within a curved spacetime (equations (2.63)). These physical processes have previously been described by [18, 36].

The fluid takes the ideal approximation, and its isotropic fluid-frame pressure can be dependent on numerous fluid quantities. The radiation field is capable of both diffusive and streaming behaviour. Its pressure tensor can utilise up to the first moment of the radiation field.

A more complete model could include further complexities such as multi-group radiation [45], magnetodynamics [46], and non-ideal fluid properties such as thermal conduction and viscosity [26].

This system of balance laws will form the basis for a numerical simulation code, employing computational techniques which were initially developed for terrestrial radiation hydrodynamic evolution in the Newtonian limit, but can be extended to also find application in numerical relativity, as outlined in chapter 3.

Chapter 3

Numerical Methods

3.1 Introduction

The system of balance laws for radiation hydrodynamic evolution within a curved spacetime, as outlined in chapter 2 equation (2.63), in combination with appropriate initial conditions, boundary conditions, and closure schemes, can form the basis of a numerical simulation code.

In this chapter I have detailed the discretisation methods, numerical approximations, and computational efficiencies which will allow the creation of a general relativistic radiation hydrodynamic evolution code. The outlined numerical methods are relevant for modelling astrophysical phenomena of interest such as post-bounce core collapse supernova explosions and neutron star merger remnants. These events involve smooth spacetimes which evolve on slow timescales compared to the fluid motion, a regime in which our stationary spacetime approximation is valid.

The hydrodynamic evolution of the fluid involved in a neutron star merger remnant is sufficiently complex due to the non-linear properties of propagating shock fronts (especially when not assuming any symmetries), that it is appropriate to employ numerical techniques to calculate a solution [47]. Closed-form analytic solutions and perturbation approaches are not general enough to capture the range of behaviours expected.

As our focus is on modelling the evolution of neutron star merger remnants, we will explore a collection of numerical techniques optimised to the range of lengthscales and timescales we expect to encounter for these scenarios. A particular challenge involves capturing an appropriate level of detail in the resulting turbulent flows within the spiralling material thrown off of the compact remnant during the merger phase, whilst still maintaining sufficient accuracy of the larger problem as a whole. A code which uses 1$^{\text{st}}$ or 2$^{\text{nd}}$ order accuracy everywhere will not be sufficient, as it would accumulate significant errors over time, as we will discuss in section 4. The radiation field we consider is coupled to both this fast moving, detailed fluid, as well as the comparatively slow, smooth spacetime.

The numerical techniques employed for binary black hole merger simulations will not necessarily be applicable for our uses. These focus on capturing the gravitational wave signal, accounting for the signal's wavelength, and concentrate on maintaining phase accuracy. Wave phases are not a priority for us, thus we will not be using a spectral method like those used in the SpEC code [48]. Furthermore the numerical approaches employed for terrestrial applications such as vehicle-drag and ventilation calculations are also unlikely to be directly applicable to our purposes. They can be very detailed, but not scalable to the spatial distances we require.

The balance laws consist of a system of first order hyperbolic partial differential equations involving both temporal and spatial derivatives. They are intended to evolve quantities which may possess steep spatial gradients, as the astrophysical phenomena we are interested in modelling can involve propagating shock waves.

Results from computational fluid dynamics suggest a number of efficient techniques to assist in achieving this objective. In order to discretise the system of balance laws for numerical evaluation [33], I introduce a computational grid to represent the spatial domain of interest. The resolution of this grid is linked to the accuracy of the numerical evaluation, quantified by the order of accuracy p, where the error is proportional to the grid spacing to the power p. The efficiency of a numerical simulation is given by its accuracy and the runtime taken to complete the calculation. Increasing the complexity of the numerical schemes used can increase the order of accuracy with respect to the grid resolution. The spatial order of accuracy is given by the spatial distribution of quantities reconstructed from cell-centred volume averaged data, such as piecewise constant, linear, or parabolic reconstructions.

Difficulties arise when attempting to represent shock fronts with high order spatial reconstructions as they can oscillate and smear out the steep gradient around the discontinuity [43, 49]. To avoid this, reconstructions can be restricted to employ low order, piecewise constant reconstruction around shocks. Reconstructions are also restricted to avoid introducing new local extremal values, which qualifies the scheme as total variation diminishing (TVD). Numerical schemes which recover high order reconstruction in regions of smooth data and low order reconstruction around discontinuities whilst conserving appropriate quantities are known as 'high resolution shock capturing schemes'. It is appropriate to employ high resolution shock capturing methods which can cope with shocks and travelling shock fronts when modelling supernova explosions and neutron star mergers. Since the problems we are interested in modelling involve the motion and interaction of shock waves we shall be employing high resolution shock capturing schemes in our numerical simulation code.

To reduce the cost of computation, adaptive mesh refinement (AMR) may be used. When using AMR, regions of the spatial domain can be refined to a sufficiently high resolution to capture any detailed structure or behaviour of interest. Simultaneously the large bulk of the domain away from these regions can be kept at a significantly coarser resolution [50]. This is advantageous over a uniform mesh approach which would require high resolution throughout the domain, and thus much more computation time to capture the same complex behaviour.

3.2 The Method of Lines

The conservation laws for relativistic radiation hydrodynamics form a system of first order hyperbolic partial differential equations (PDEs) involving both temporal and spatial derivatives. A numerical approach to solving this system aims to progress the vector of conserved quantities from one computational time-step to the next. This progression begins with the initial conditions describing the physical problem to be simulated.

We aim to evolve this initial value problem by means of the method of lines, similar to that used by live spacetime codes such as the Einstein Toolkit [2]. This is opposed to a direct explicit update seen in the IAMR, Maestro, and Castro codes [51–53] which make use of corner-transport-upwind methods. The method of lines approach is more convenient when we need to account for the coupling between the fluid, radiation, and the spacetime.

The methods which I use have up to third order temporal convergence. Higher order methods like those employed by spectral methods will not be necessary as we are not concerned about gravitational wave phase accuracy accumulating global error.

The system of balance laws (2.63) can be re-arranged to involve purely temporal gradients on the left hand side and purely spatial gradients and source terms on the right hand side,

$$\frac{\partial}{\partial t}\left(\underline{U}\right) = -\frac{\partial}{\partial x^i}\left(\underline{F}^i\right) + \underline{S}. \tag{3.1}$$

By applying a spatial discretisation, the spatial domain can be split into a finite number of locations which can be described with a three dimensional grid. Each discretised location represents a zone (cell) of the spatial domain within which the physical quantities can be described with an approximate distribution. Using the values of the physical quantities at only the discretised locations allows the spatial derivatives to be represented with algebraic expressions (as will be discussed in section 3.4) and allows the system of PDEs to be written as a system of time-dependent initial value problems (IVPs) [54],

$$\frac{\partial}{\partial t}\left(\underline{U}_j\right) = \underline{f}\left(x_j, t\right). \tag{3.2}$$

Here the equations have been restricted to a single spatial dimension, and the subscript index j denotes discretised spatial location. This system of IVPs is to be solved for each discretised spatial location, thus evolving the entire discretised spatial domain forwards in time [43, 49, 55].

3.3 Initial Value Problem Methods

The right hand side of our system of IVPs can be collected into a single term which is dependent on the conserved quantities we aim to evolve,

$$\frac{\partial}{\partial t}\underline{U} = \underline{f}\left(\underline{U}\right). \tag{3.3}$$

The vector of conserved quantities can be progressed in time by applying a temporal discretisation and employing an explicit numerical technique to evolve the quantities from one time-step to the next.

The solution at an earlier time (U^n) can be used to obtain the solution at a later time (U^{n+1}). For a multi-step method we can make use of a number of intermediate solutions ($U^{(m)}$). Below we introduce a number of minimum storage explicit

Runge-Kutta methods, with the particular Shu-Osher representation [56]:

$$U^{(m+1)} = a_m U^n + (1 - a_m) U^{(m)} + b_m \Delta t f\left(U^{(m)}\right). \tag{3.4}$$

For an s-step method we have $m = 0, \cdots, s-1$. The method is started using $U^{(0)} = U^n$ and completed with $U^{(s)} = U^{n+1}$.

3.3.1 The Forward Euler Method

The simplest approach to the numerical evolution of the simulation is the forward Euler method. The forward Euler method is the explicit, one-step, first-order in time method (the first-order explicit Runge-Kutta method). It involves evaluating the right-hand side term at the beginning of the time-step to approximate a temporal gradient for the conserved quantities, and uses that temporal gradient to extrapolate a solution for the conserved quantities at the end of the time-step,

$$U^{n+1} = U^n + \Delta t f\left(U^n\right), \qquad (a_0, b_0) = (1, 1). \tag{3.5}$$

3.3.2 The Mid-Point Method

The mid-point method is an explicit, two-step, second-order in time method in the family of second-order Runge-Kutta methods. It involves calculating half of a forward Euler step to obtain a solution at the mid-point of the time-step, hence the method's name. This mid-point solution is then used to calculate a new temporal gradient for the conserved quantities, which in turn is used to extrapolate the conserved quantities from the beginning of the time-step to the end,

$$U^{(1)} = U^n + \frac{1}{2}\Delta t f\left(U^n\right), \qquad (a_0, b_0) = \left(1, \frac{1}{2}\right), \tag{3.6}$$

$$U^{n+1} = U^n + \Delta t f\left(U^{(1)}\right), \qquad (a_1, b_1) = (1, 1). \tag{3.7}$$

3.3.3 Heun's Method

Heun's method (also known as the explicit trapezoidal rule) is another explicit, two-step, second-order in time method, in the family of second-order Runge-Kutta methods.

$$U^{(1)} = U^n + \Delta t f\left(U^n\right), \qquad (a_0, b_0) = (1, 1), \tag{3.8}$$

$$U^{n+1} = \frac{1}{2}U^n + \frac{1}{2}U^{(1)} + \frac{1}{2}\Delta t f\left(U^{(1)}\right), \qquad (a_1, b_1) = \left(\frac{1}{2}, \frac{1}{2}\right). \tag{3.9}$$

This method is strong stability preserving (SSP, sometimes referred to as total variation diminishing (TVD)). This gives some mathematical guarantees about the growth rate of solutions [57], and avoids oscillatory evolution when using a stable time-step. SSP methods satisfy the condition that the total variation of the solution does not increase as it is evolved, such that

$$\sum_j \left|U_{j+1}^{n+1} - U_j^{n+1}\right| \leq \sum_j \left|U_{j+1}^n - U_j^n\right|. \tag{3.10}$$

3.3.4 The Shu-Osher Method

A minimum storage, third-order in time, explicit, SSP Runge-Kutta method is given by [58]

$$U^{(1)} = U^n + \Delta t f\left(U^n\right), \qquad (a_0, b_0) = (1, 1), \tag{3.11}$$

$$U^{(2)} = \frac{3}{4}U^n + \frac{1}{4}U^{(1)} + \frac{1}{4}\Delta t f\left(U^{(1)}\right), \qquad (a_1, b_1) = \left(\frac{3}{4}, \frac{1}{4}\right), \tag{3.12}$$

$$U^{n+1} = \frac{1}{3}U^n + \frac{2}{3}U^{(2)} + \frac{2}{3}\Delta t f\left(U^{(2)}\right), \qquad (a_2, b_2) = \left(\frac{1}{3}, \frac{2}{3}\right). \tag{3.13}$$

3.4 The Finite Volume Method

The discretised zones of our spatial domain form a collection of finite volumes. Here we work towards obtaining a finite volume method (FVM) which evolves the volume integrated averages of our conserved quantities.

The spatial derivative terms in our system of balance laws (2.63) represent a divergence of fluxes. Integrating these balance laws over the finite volume allows us to apply the divergence theorem, replacing the divergence term with surface integrals. These represent the sum of fluxes though the surfaces of the finite volume. We obtain an exact expression for the volume averaged conserved quantities (shown here restricted to a single spatial dimension)

$$\frac{\partial}{\partial t}\left(\underline{U}_j\right) = -\frac{1}{\Delta x}\left(\underline{F}_{j+1/2} - \underline{F}_{j-1/2}\right) + \underline{S}_j. \tag{3.14}$$

Keeping the surface flux between neighbouring finite volumes consistent allows for the creation of conservative methods.

At the spatial domain boundaries obtaining the cell-interface flux may require data from outside of the spatial domain. We can achieve this whilst maintaining our cell-centred data structure by introducing ghost cells. We grow our numerical grid to add external computational zones, and populate them with data to satisfy any physical boundary conditions we require, such as reflective, transmissive, or in/outflow. The number of zones we grow our computational grid by is determined by the stencil size we decide to use for the interface flux calculation.

3.5 Quantity Reconstruction at Interfaces

Here examples of techniques known as high resolution shock capturing methods are described [59]. High resolution shock capturing methods are beneficial as the discontinuities between the discretised physical quantities used in the balance laws can evolve as shock fronts.

The finite-volume discretised evolution equations (3.14) require flux values at cell interfaces. We calculate these fluxes from the adjacent cell-centred primitive quantities as in the PIZZA, GRHydro, and Whisky codes [60–62]. We choose to reconstruct the primitive quantities over the conserved quantities, as it is much more straightforward to calculate the flux terms from them, otherwise a root-find would be required at every cell interface, a very computationally expensive process. We are unable to avoid evaluating this root-find entirely, but we aim to reduce the number of times it must be calculated, as will be discussed in section 3.12.

The approximation to the distribution of the primitive quantities within a computational cell can be chosen. For numerical stability the chosen distributions should avoid introducing artificial oscillations or new extremal points, and recover the cell integrated values required for the finite volume method. The simplest distribution approximation to use is that of a uniform profile. The disadvantage to using a constant profile is that it leads to substantial discontinuity jumps at cell boundaries, even for smooth solutions, and results in only first order convergence with spatial resolution. A larger difference in physical quantity either side of a cell boundary results in a more dissipative evolution of the numerical simulation. The magnitude of the discontinuity jumps can be influenced by employing higher order intra-cell distributions.

An illustration of a piecewise constant reconstruction of primitive quantities is shown in figure 3.1. Within each computational zone the quantity is uniform, and at zone

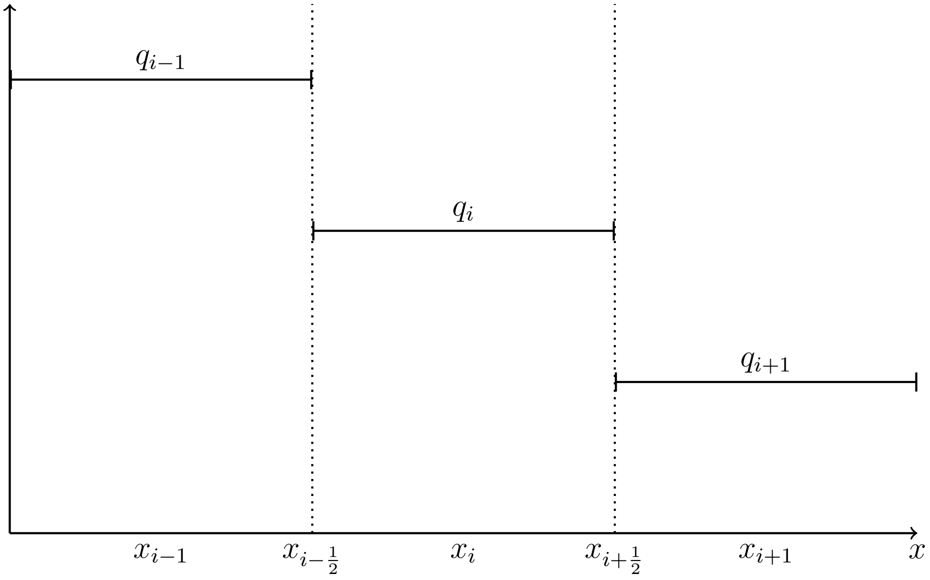

Figure 3.1: A one-dimensional illustration of a piecewise constant reconstruction of a cell-centred primitive quantity. These interface values can be used by a Riemann solver to obtain interface fluxes.

interfaces there can be a discontinuity. These values can be used by a Riemann solver to obtain interface flux values.

A danger of employing high order reconstruction methods is that they can introduce large, unwanted oscillations around discontinuities, similar to Runge's and Gibbs phenomena [63, 64]. To avoid this we aim to use lower order reconstructions around discontinuities or limit the magnitude of the reconstruction.

3.5.1 Linear Profiles

From cell-centred primitive quantities, we can linearly reconstruct values to either side of each cell interface [49, 65–67]. Let q represent some primitive quantity, and

$$R_i = \frac{q_{i+1} - q_i}{q_i - q_{i-1}}, \tag{3.15}$$

then the interface values can be reconstructed with some slope limiting function ϕ, as

$$q^L_{i+1/2} = q_i + \frac{1}{2}\phi(R_i)(q_i - q_{i-1}), \tag{3.16}$$

$$q^R_{i-1/2} = q_i - \frac{1}{2}\phi(1/R_i)(q_{i+1} - q_i). \tag{3.17}$$

Here the L and R superscripts indicate projection on to a cell interface from the left and right sides respectively, and the ϕ function calculates appropriate slopes from surrounding values [68]. These linear reconstructions require a stencil of size three, specifically $\{x_{i-1}, x_i, x_{i+1}\}$, corresponding to at least two required ghost zones at each boundary.

Possible choices for such a slope calculating function include

$$\phi_{MM}(R) = \min(1, R), \tag{3.18}$$

$$\phi_{VL}(R) = \frac{2R}{1+R}, \tag{3.19}$$

the MinMod and VanLeer methods [43]. Here $\phi(R) = 0$ if $(q_{i+1} - q_i)(q_i - q_{i-1}) \leq 0$ for all such methods. The slope functions are only non-zero with strictly monotonic data, ensuring that new local extrema are not introduced, and the methods are total variation diminishing.

Using piecewise linear profiles can result in the numerical method having second-order convergence with spatial resolution in regions of smoothly varying physical quantities. In regions of non-smooth physical quantities such as where shock-fronts are propagating through the fluid, a first-order method is maintained, preserving the robust and stable behaviour which is desired.

3.5.2 Parabolic Reconstruction

Approximating the spatial distribution of the primitive quantities as piecewise parabolic can result in a third order reconstruction in smoothly varying regions. Using the volume averaged values (in 1D) of a computational cell and its adjacent neighbours we can find a system of linear equations which can be solved for the coefficients of a quadratic polynomial distribution. Evaluating this polynomial at the cell interfaces gives a third order reconstruction for the interface values.

These interface values can be written as projections from the central value based on the differences in the primitive quantity between the neighbouring cell centres. In order to avoid introducing any new extremal values and to not overshoot the neighbouring values, and thus be total variation diminishing (TVD), this projection is limited to only be non-zero for locally strictly monotonic spatial primitive distributions. This is the Koren slope limiter as defined in [69]

$$\phi_K(R) = \min\left(2, 2R, \frac{1+2R}{3}\right). \tag{3.20}$$

This parabolic reconstruction also requires the three point stencil $\{x_{i-1}, x_i, x_{i+1}\}$ and at least two ghost zones at each boundary.

3.5.3 WENO

Weighted essentially non-oscillatory (WENO) schemes are a class of high order reconstruction methods. They can be used to reconstruct the value of primitive quantities at cell-interfaces from surrounding discrete cell-centred data points (q_i). They employ a convex linear combination of stencils to obtain high order reconstruction, whilst avoiding the undesirable oscillations resulting from high order reconstructions around discontinuities (see Runge's and Gibb's phenomena [63, 64]). This is achieved by giving any discontinuity containing stencils a lower weighting in the linear combination, as explained further in [70–72].

If used in conjunction with patch-based adaptive mesh refinement (AMR), each patch would require at least the minimum number of ghost cells demanded by the WENO stencil size (conversely the maximum stencil size available is restricted by the number of ghost cells at the domain boundaries). As AMR requires patch communication, this becomes more expensive with more ghost cells, limiting stencil size, and tying it to computational cost. AMR is discussed further in section 3.13.

In our case we construct a fifth order scheme from a set of third order accurate quadratic fits $\{p_1(x), p_2(x), p_3(x)\}$. Focusing on approaching the cell interface at $x_{i+1/2}$ from the left. The fit $p_j(x)$ is centred on x_{i-2+j}, and utilises the three-point stencil $\mathbf{S}_j = \{x_{i-3+j}, x_{i-2+j}, x_{i-1+j}\}$. The set of third order fits reconstruct the interface value to be

$$q^{i-1}_{i+1/2} = p_1(x_{i+1/2}) = \frac{1}{3}q_{i-2} - \frac{7}{6}q_{i-1} + \frac{11}{6}q_i, \tag{3.21}$$

$$q^{i}_{i+1/2} = p_2(x_{i+1/2}) = -\frac{1}{6}q_{i-1} + \frac{5}{6}q_i + \frac{1}{3}q_{i+1}, \tag{3.22}$$

$$q^{i+1}_{i+1/2} = p_3(x_{i+1/2}) = \frac{1}{3}q_i + \frac{5}{6}q_{i+1} - \frac{1}{6}q_{i+2}. \tag{3.23}$$

The fifth order reconstruction can be found from a convex linear combination of the third order reconstructions as

$$
\begin{aligned}
q^{5th}_{i+1/2} &= \frac{1}{30}q_{i-2} - \frac{13}{60}q_{i-1} + \frac{47}{60}q_i + \frac{9}{20}q_{i+1} - \frac{1}{20}q_{i+1} && (3.24)\\
&= a_1 q^{i-1}_{i+1/2} + a_2 q^{i}_{i+1/2} + a_3 q^{i+1}_{i+1/2}, && (3.25)\\
a_j &= \left(\frac{1}{10}, \frac{3}{5}, \frac{3}{10}\right). && (3.26)
\end{aligned}
$$

To avoid issues such as undesired oscillations, due to discontinuous data, an alternative convex combination of these quadratic fits can be found using smoothness indicators (d_j). In smooth regions the full fifth order reconstruction is recovered, and around discontinuities an at least third order reconstruction is created, free from the excessive oscillations. The alternative combination used here is

$$
q^{WENO}_{i+1/2} = b_1 q^{i-1}_{i+1/2} + b_2 q^{i}_{i+1/2} + b_3 q^{i+1}_{i+1/2}, \tag{3.27}
$$

where

$$
\begin{aligned}
b_j &= \frac{c_j}{\sum_k c_k}, && (3.28)\\
c_j &= \frac{a_j}{(d_j + 10^{-6})^2}, && (3.29)\\
d_j &= \sum_{l=1}^{2} \Delta x^{2l-1} \int_{x_{i-1/2}}^{x_{i+1/2}} \left(\frac{d^l}{dx^l} p_j(x)\right)^2 dx, && (3.30)\\
d_1 &= \frac{13}{12}(q_{i-2} - 2q_{i-1} + q_i)^2 + \frac{1}{4}(q_{i-2} - 4q_{i-1} + 3q_i)^2, && (3.31)\\
d_2 &= \frac{13}{12}(q_{i-1} - 2q_i + q_{i+1})^2 + \frac{1}{4}(q_{i-1} - q_{i+1})^2, && (3.32)\\
d_3 &= \frac{13}{12}(q_i - 2q_{i+1} + q_{i+2})^2 + \frac{1}{4}(3q_i - 4q_{i+1} + q_{i+2})^2. && (3.33)
\end{aligned}
$$

The surrounding cell-centred data points can be reverse-ordered to reconstruct data to the right of the interface $x_{i+1/2}$.

Where a reconstruction over or under-shoots the surrounding data-points, introducing new local extrema, physical restrictions can placed on the interpolation, such as limiting how low mass-density, temperature, or energy can become, to avoid vacuum or negative values.

The stencils used for this fifth order WENO scheme require three ghost zones at each boundary.

3.6 Spacetime Evaluation

The flux terms in our balance laws require the spacetime metric terms (lapse-function, shift-vector, and spatial-metric) to be reconstructed at the cell-interfaces, as seen in equation (2.65). We make the assumption that they are smoothly varying everywhere. For a numerical simulation with at least three ghost zones we are able to employ up to a sixth order Lagrange polynomial interpolation,

$$q_{i+1/2}^{2nd} = \frac{q_i + q_{i+1}}{2}, \tag{3.34}$$

$$q_{i+1/2}^{4th} = \frac{-q_{i-1} + 9q_i + 9q_{i+1} - q_{i+2}}{16}, \tag{3.35}$$

$$q_{i+1/2}^{6th} = \frac{3q_{i-2} - 25q_{i-1} + 150q_i + 150q_{i+1} - 25q_{i+2} + 3q_{i+3}}{256}. \tag{3.36}$$

The geometric source terms in our balance laws require the gradients of the spacetime metric terms at cell-centres, as seen in equation (2.66). Again with at least three ghost zones and smoothly varying quantities, we can employ up to a sixth order finite difference approximation to the spacetime gradients,

$$\partial_x q_i^{2nd} = \frac{-q_{i-1} + q_{i+1}}{2}, \tag{3.37}$$

$$\partial_x q_i^{4th} = \frac{q_{i-2} - 8q_{i-1} + 8q_{i+1} - q_{i+2}}{12}, \tag{3.38}$$

$$\partial_x q_i^{6th} = \frac{-q_{i-3} + 9q_{i-2} - 45q_{i-1} + 45q_{i+1} - 9q_{i+2} + q_{i+3}}{60}. \tag{3.39}$$

Provided that the quantities are smoothly varying in space, the higher order approximations are more accurate. In practice we use the highest order approximation we can, given the number of ghost cells available. Even the sixth order methods shown here are computationally inexpensive to evaluate.

3.7 Approximate Riemann Solvers

The system of IVPs in equation (3.14) require the evaluation of the flux terms at the cell interfaces. The interfaces separating the computational cells can have discontinuous physical quantities on either side. A conservation law with discontinuous initial conditions is known as a Riemann problem [43, 66, 73]. The fluxes through these cell interfaces can be calculated with the use of Riemann solvers. Treating each cell interface as a Riemann problem is known as a Godunov scheme.

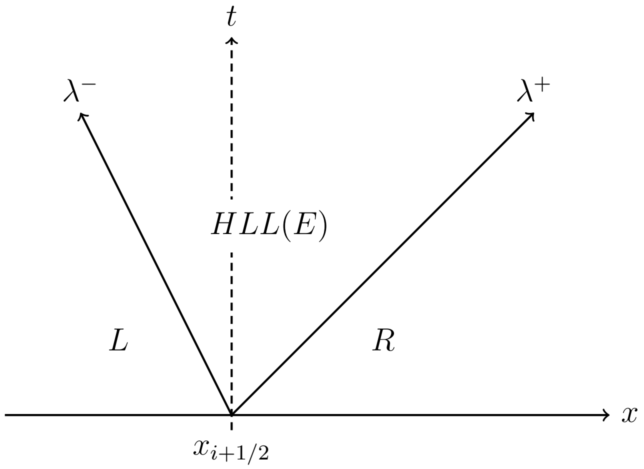

Figure 3.2: An illustration of a two-wave evolution of a Riemann problem, in this case the cell interface exists between the left and right-going waves.

Here we consider cell interfaces in the x-direction (the y and z-direction cases are equivalent),

$$\partial_t \underline{U} + \partial_x \underline{F}^x = \underline{S}. \tag{3.40}$$

We aim to capture the evolution of waves propagating away from the discontinuity, such as shock waves, rarefaction waves, and if applicable, contact waves. We will require their wave-speeds in order to determine the state of the flux at the cell interface, as in [43].

3.7.1 The HLL(E) Riemann Solver

The HLLE Riemann solver (often simply HLL) is an approximate Riemann solver which assumes a simple two-wave evolution of the discontinuity. The leftmost and rightmost waves have wave-speeds λ^- and λ^+ respectively, as in [34, 43, 74–76]. The location of the cell interface relative to the waves used by the approximate Riemann solver depends on the wave-speed approximations, this is illustrated in figure 3.2.

The interface fluxes are chosen depending on where in the wave-pattern the stationary cell interface sits,

$$\underline{F}_{i+1/2} = \begin{cases} \underline{F}_L & \text{if} \quad 0 \leq \lambda^-, \\ \underline{F}^{HLL} & \text{if} \quad \lambda^- < 0 < \lambda^+, \\ \underline{F}_R & \text{if} \quad \lambda^+ \leq 0. \end{cases} \tag{3.41}$$

Here the L subscript indicates that the physical quantities on the left side of the cell interface have been used in the calculation, and the R subscript indicates that the physical quantities on the right side of the cell interface have been used in the calculation. The upwind fluxes are used where appropriate. In the subsonic case, the conserved quantities $\underline{U}^{HLL}$ and their corresponding fluxes $\underline{F}^{HLL}$ in between the two waves in the HLL approximation are calculated as

$$\underline{U}^{HLL} = \frac{\lambda^{+}\underline{U}_R - \lambda^{-}\underline{U}_L + \underline{F}_L - \underline{F}_R}{\lambda^{+} - \lambda^{-}}, \tag{3.42}$$

$$\underline{F}^{HLL} = \frac{\lambda^{+}\underline{F}_L - \lambda^{-}\underline{F}_R + \lambda^{-}\lambda^{+}(\underline{U}_R - \underline{U}_L)}{\lambda^{+} - \lambda^{-}}. \tag{3.43}$$

For relativistic hydrodynamic simulations the wave-speeds are dependent on the fluid velocity and the local sound-speed c_s of the fluid. As in [77] the fastest left and right hydrodynamic wave-speeds propagating from a discontinuity at a cell interface are

$$\lambda^{\pm} = \frac{\alpha}{1 - v^2 c_s^2}\left[v^x\left(1 - c_s^2\right) \pm c_s\sqrt{(1 - v^2)\left[\gamma^{xx}(1 - v^2 c_s^2) - v^x v^x (1 - c_s^2)\right]}\right] - \beta^x. \tag{3.44}$$

When fluid motion is perpendicular to the cell interface (i.e. $v^y = v^z = 0$), this reduces to

$$\lambda^{\pm} = \alpha\left(\frac{v^x \pm c_s}{1 \pm v^x c_s}\right) - \beta^x. \tag{3.45}$$

When using a Gamma-law equation of state (as introduced in section 2.8.1) the local hydrodynamic sound speed is given as $c_s = \sqrt{\Gamma P / h\rho}$. These wave-speeds are calculated at the interface using values reconstructed from either side, the minimum of the two left wave-speeds and maximum of the two right wave-speeds are taken.

For simulations involving radiation the radiation wave-speeds can be much larger than the hydrodynamic wave-speeds. The radiation wave-speeds can be over-approximated with the speed of light in a vacuum. An over-approximation of wave-speed results in a more dissipative approximation to the evolution of the simulation. However this is safer than under-approximating the wave-speeds which can lead to instabilities in the simulation.

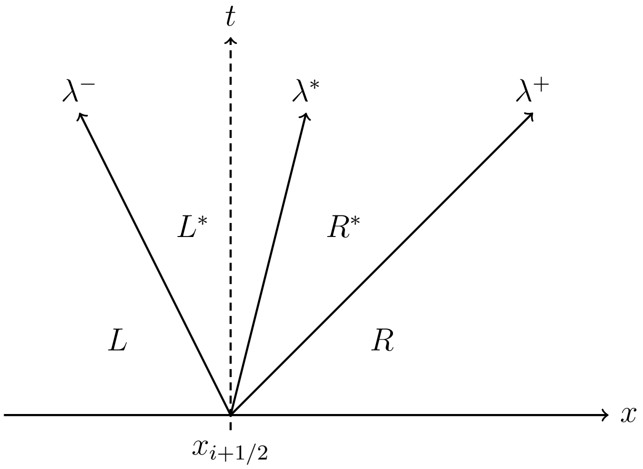

Figure 3.3: An illustration of a three-wave evolution of a Riemann problem.

3.7.2 The HLLC Riemann Solver

The HLL(E) Riemann solver is robust and reliable, however in astrophysical simulations such as for supernova and relativistic jets, a more physical description of the wave-pattern is required. Another approximate Riemann solver, more physical than the HLL(E) solver, is the HLLC Riemann solver. The HLLC solver assumes a 3-wave evolution of the discontinuity by restoring the contact discontinuity λ^* in the resulting wave pattern [73, 78], this is illustrated in figure 3.3.

As with the HLL(E) Riemann solver, the fluxes are chosen depending on where in the wave pattern the stationary cell interface exists,

$$\underline{F}_{i+1/2} = \begin{cases} \underline{F}_L & \text{if} \quad \lambda^- \geq 0, \\ \underline{F}_L^* & \text{if} \quad \lambda^- < 0 \leq \lambda^*, \\ \underline{F}_R^* & \text{if} \quad \lambda^* \leq 0 < \lambda^+, \\ \underline{F}_R & \text{if} \quad \lambda^+ \leq 0. \end{cases} \tag{3.46}$$

The $*$ indicates quantities in a region which has been affected by the propagating wave-fronts, with $^*_L$ indicating quantities to the left of the contact discontinuity, and $^*_R$ indicating quantities to the right of the contact discontinuity. The conserved quantities and their corresponding fluxes in the affected region are related via

$$\underline{F}_K^* = \underline{F}_K + \lambda_K \left(\underline{U}_K^* - \underline{U}_K \right), \tag{3.47}$$

where $K \in \{L, R\}$. The location of the cell interface relative to the three waves of the HLLC approximation depends on the set of wave-speeds $\{\lambda^-, \lambda^*, \lambda^+\}$.

For one-dimensional relativistic hydrodynamics, using the Valencia Formulation as in (2.63) [33],

$$\underline{U} = \sqrt{\gamma}\begin{pmatrix} D \\ S_x \\ \tau \end{pmatrix}, \quad \underline{F}^x = \sqrt{\gamma}\begin{pmatrix} D(\alpha v^x - \beta^x) \\ S_x(\alpha v^x - \beta^x) + \alpha P \\ \tau(\alpha v^x - \beta^x) + \alpha P v^x \end{pmatrix}. \tag{3.48}$$

The relativistic hydrodynamic flux is dependent on the conserved quantities, from equation (3.47) we have that

$$\underline{F}^*_K\left(\underline{U}^*_K\right) = \underline{F}_K + \lambda_K\left(\underline{U}^*_K - \underline{U}_K\right). \tag{3.49}$$

Rearranging this for the components of $\underline{U}^*_K$, we have that

$$\begin{aligned} D^*_K &= \frac{D_K(\alpha v^x_K - \beta^x - \lambda_K)}{(\alpha v^{x*} - \beta^x - \lambda_K)}, \\ S^*_{xK} &= \frac{S_{xK}(\alpha v^x_K - \beta^x - \lambda_K) + \alpha(P_K - P^*)}{(\alpha v^{x*} - \beta^x - \lambda_K)}, \\ \tau^*_K &= \frac{\tau_K(\alpha v^x_K - \beta^x - \lambda_K) + \alpha(P_K v^x_K - P^* v^{x*})}{(\alpha v^{x*} - \beta^x - \lambda_K)}. \end{aligned} \tag{3.50}$$

Once these conserved quantities are obtained we can calculate the HLLC interface fluxes from equation (3.47).

Calculating v^{x*} and P^*

We use that v^{x*} and P^* are constant across the contact discontinuity, which lies between λ_L and λ_R. Only ρ is discontinuous in this region. We use the HLL(E) Riemann solver to calculate U^{HLL}_D, $U^{HLL}_{S_i}$, U^{HLL}_τ, and $F^{HLL}_{S_x}$. Similar to [78], let

$$U^{HLL}_{S^x} = (U^{HLL}_\tau + U^{HLL}_D + \sqrt{\gamma}P^*)v^{x*}. \tag{3.51}$$

The term on the left can be calculated with the known HLL(E) quantities and the spatial metric via

$$U^{HLL}_{S^x} = U^{HLL}_{S_i}\gamma^{ix}. \tag{3.52}$$

Additionally we let

$$F^{HLL}_{S_x} = U^{HLL}_{S_x}(\alpha v^{x*} - \beta^x) + \alpha\sqrt{\gamma}P^*. \tag{3.53}$$

From the two relations (3.51) and (3.53), eliminating P^* leads to a quadratic equation for v^{x*}

$$v^{x*}F_{S_x}^{HLL} - \alpha U_{S^x}^{HLL} = v^{x*}U_{S_x}^{HLL}(\alpha v^{x*} - \beta^x) - \alpha(U_\tau^{HLL} + U_D^{HLL})v^{x*}, \tag{3.54}$$

where

$$v^{x*} = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}, \tag{3.55}$$

$$a = \alpha U_{S_x}^{HLL}, \tag{3.56}$$

$$b = -F_{S_x}^{HLL} - \beta^x U_{S_x}^{HLL} - \alpha(U_\tau^{HLL} + U_D^{HLL}), \tag{3.57}$$

$$c = \alpha U_{S^x}^{HLL}. \tag{3.58}$$

Requiring that $\lambda_L \leq \lambda^* = \alpha v^{x*} - \beta^x \leq \lambda_R$, we take the appropriate of the two roots of the quadratic equation. In the case $a = 0$, then we have that $v^{x*} = -c/b$. Now P^* can be calculated from rearranging (3.53) as

$$P^* = \frac{F_{S_x}^{HLL} - U_{S_x}^{HLL}(\alpha v^{x*} - \beta^x)}{\alpha\sqrt{\gamma}}. \tag{3.59}$$

Given $\lambda^* = \alpha v^{x*} - \beta^x$ and P^* we can now evaluate equations (3.50) and (3.47) in order to obtain the HLLC interface fluxes.

3.8 Time-Step Calculation

For numerical stability the time-step must be limited by the time it takes for the fastest wave-speed to propagate from the discontinuity on one side of a computational cell to the opposite side, as described in [43, 49]. This limit must hold throughout the spatial domain, so the minimum allowed time-step should be taken over all cells involved in the simulation. This limit can be ensured by including a CFL factor $c_{CFL} \in (0, 1]$ [79],

$$\Delta t = c_{CFL} \ \min\left(\frac{\Delta x}{|\lambda_K|}\right). \tag{3.60}$$

In practice we have found it safe to use $c_{CFL} = 1/10$, if excessively so. Since our wave-speeds are bound by the speed of light in a vacuum ($\max|\lambda_K| \leq 1$), we can safely use $\Delta t = c_{CFL}\Delta x$.

3.9 Operator Splitting

Our situation of evolving the general relativistic radiation hydrodynamics equations in the form of a system of balance laws contain contributions from directional fluxes, geometric sources, and energy-momentum exchange terms, as seen in equation (2.63). These contributions can be considered a series of non-linear sub-problems, and may benefit from being solved in isolation.

Consider an initial value problem, composed of linear operators L_1 and L_2

$$\partial_t y = L_1 y + L_2 y. \tag{3.61}$$

We can find the analytic evolution for this problem. Let $y_n = y(n\Delta t)$, then

$$\begin{aligned} y_{n+1} &= e^{(L_1+L_2)\Delta t} y_n \qquad (3.62)\\ &= e^{L_1\Delta t} e^{L_2\Delta t} y_n. \qquad (3.63) \end{aligned}$$

We could introduce an intermediate solution $\widetilde{y}_{n+1}$ where

$$\begin{aligned} \widetilde{y}_{n+1} &= e^{L_1\Delta t} y_n, \qquad (3.64)\\ y_{n+1} &= e^{L_2\Delta t} \widetilde{y}_{n+1}. \qquad (3.65) \end{aligned}$$

Noticing how each step solves a sub-problem, we can split our initial value problem into a series of these sub-problems

$$\begin{aligned} \partial_t y &= L_1 y, \qquad (3.66)\\ \partial_t y &= L_2 y. \qquad (3.67) \end{aligned}$$

This approach proves to still be useful when generalised to consider non-linear differential operators.

We have strict limits on the effect that radiation-fluid energy exchange can have, but not on what the effect the combined flux and exchange contributions can have. In order to enforce these strict limits, our series of non-linear differential operators benefit from operator splitting.

3.10 Energy-Momentum Exchange

The energy-momentum exchange operator in the Newtonain limit, and with constant opacity, describes exponential decay towards an equilibrium state. From this linear approximation we expect that the energy-momentum exchange between the fluid and radiation approaches the local thermodynamic equilibrium state (LTE), but does not overshoot it. In cases of large opacities or large time-steps an explicit update is likely to overshoot this LTE state, it could even return negative energies.

3.10.1 Stiff Source Terms

In cases where opacity is large the energy and momentum exchange can evolve on a much faster timescale than the advective fluxes. We operator split our system of balance laws in order to isolate the energy and momentum exchange source term (S_{RH}),

$$\frac{\partial}{\partial t}U = -\frac{\partial}{\partial x^i}F^i + S_{geo}, \tag{3.68}$$

$$\frac{\partial}{\partial t}U = S_{RH}. \tag{3.69}$$

The sub-problem of energy-momentum exchange can now be updated separately from the flux and geometric source update. This allows the change in conserved quantities due to the fluxes and spacetime to not affect the physical restrictions placed on the exchange update: $\widetilde{U}^{n+1} = U^n + \Delta t S_{RH}$.

3.10.2 Local Thermodynamic Equilibrium

The fluid and radiation exchange energy and momentum via radiation absorption and emission. Considering only the sub-problem of the energy-momentum exchange source terms in our system of balance laws we have from equations (2.63) and (2.60) that

$$\frac{\partial}{\partial t}\sqrt{\gamma}\begin{pmatrix} S_j \\ \tau \\ \overline{S}_j \\ \overline{U} \end{pmatrix} = \sqrt{\gamma}\begin{pmatrix} \alpha G_j \\ \alpha^2 G^0 \\ -\alpha G_j \\ -\alpha^2 G^0 \end{pmatrix}, \tag{3.70}$$

$$G^\nu = \chi^t \left(E - \Theta\right) u^\nu + \left(\chi^t + \chi^s\right) F^\nu. \tag{3.71}$$

Here the source term is the radiation 4-force introduced in section 2.5. The timescale of the exchange is determined by the opacities. The larger the opacity, the faster the exchange occurs, and the faster the local thermodynamic equilibrium state is reached. The stationary equilibrium solution is attained when $G^\mu = 0^\mu$. When the radiation is co-moving with the fluid (such that $F^i = 0^i$) this equilibrium occurs when $E := \Theta(T_R) = \Theta(T_M)$, where T_M and T_R are the fluid and radiation temperatures respectively. The radiation emission is related to the fluid temperature via $\Theta(T_M) = a_{rad}T_M^4$, and the radiation energy is related to the radiation temperature via $E = \Theta(T_R) = a_{rad}T_R^4$. An equilibrium is reached when the radiation flux vanishes in the fluid frame, and the radiation temperature is equal to the fluid temperature $(T_M = T_R = T_{LTE})$.

From equation (3.70) we can see that the total energy is conserved during the exchange,

$$\frac{\partial}{\partial t}\sqrt{\gamma}\left(\tau + \overline{U}\right) = 0, \tag{3.72}$$

$$\tau\left(T_{LTE}\right) + \overline{U}\left(T_{LTE}\right) = \tau^n + \overline{U}^n. \tag{3.73}$$

The temperature dependent energies are

$$\tau(T) = \left[\rho + \rho\epsilon(T) + P(T)\right]W^2 - P(T) - \rho W, \tag{3.74}$$

$$\overline{U}(T) = \frac{4}{3}E(T)W^2 + 2W\alpha F^0 - \frac{1}{3}E(T). \tag{3.75}$$

Here we are employing the Eddington closure, as near LTE the radiation flux should be small compared to radiation energy density.

A root finding algorithm can be used to converge on the equilibrium temperature T_{LTE}, which, from equation (3.73), is the root of the function

$$f(T) = \tau(T) + \overline{U}(T) - \tau^n - \overline{U}^n. \tag{3.76}$$

We can use initial guesses based on the fluid temperature (i.e., $T^{[0]} = T_M$, and $T^{[-1]} = 1.1T_M$). Root finding algorithms will be discussed in section 3.11.

Having evaluated the equilibrium temperature to within an acceptable tolerance, the fluid and radiation equilibrium energies are obtained, $\tau_{LTE} = \tau(T_{LTE})$, and $E_{LTE} = E(T_{LTE})$.

The continuous analytic solution would behave like a smooth monotonic decay from initial energy towards the equilibrium energy. Evolved energies due to emission or absorption should not overshoot their equilibrium values, and certainly should not become negative. When using an explicit update, this can occur when time-steps are large, or when the opacities are large. The change in fluid energy must be equal in magnitude and opposite in sign to the change in radiation energy. A limit on $|G^0|$ can be found which corresponds to the maximum permissible change in energy which does not overshoot the equilibrium solution. From a simple explicit temporal discretisation of the energy equations from (3.70), this limit is

$$|G^0_{lim}| = \frac{|\tau_{LTE} - \tau^n|}{\alpha^2 \Delta t} = \frac{\left|\overline{U}^{LTE} - \overline{U}^n\right|}{\alpha^2 \Delta t}. \tag{3.77}$$

We aim to find a restricted source term $\overline{G}^\nu$ which does not exceed this limit.

If G^0 is positive, then the material energy will increase, and correspondingly the radiation energy will decrease, we require that G^0 is not *too* positive,

$$\overline{G}^0 = \min\left(G^0, k\left|G^0_{lim}\right|\right). \tag{3.78}$$

If G^0 is negative, then the material energy will decrease, and correspondingly the radiation energy will increase, we require that G^0 is not *too* negative,

$$\overline{G}^0 = \max\left(G^0, -k\left|G^0_{lim}\right|\right). \tag{3.79}$$

Here $k \in (0, 1]$ is some safety factor to prevent overshooting equilibrium in a single time-step (in practice we use $k = 1/2$).

When the energy exchange source is restricted, the momentum exchange sources are scaled by the same amount,

$$\overline{G}^i = G^i k \left|\frac{G^0_{lim}}{G^0}\right|. \tag{3.80}$$

3.11 Root Finding Algorithm

As seen when dealing with stiff source terms in the previous section 3.10, root finding algorithms are a crucial part of solving non-linear equations. They will prove useful again when we approach the problem of recovering primitive quantities from updated conserved quantities in section 3.12. Here a one-dimensional root finding method is

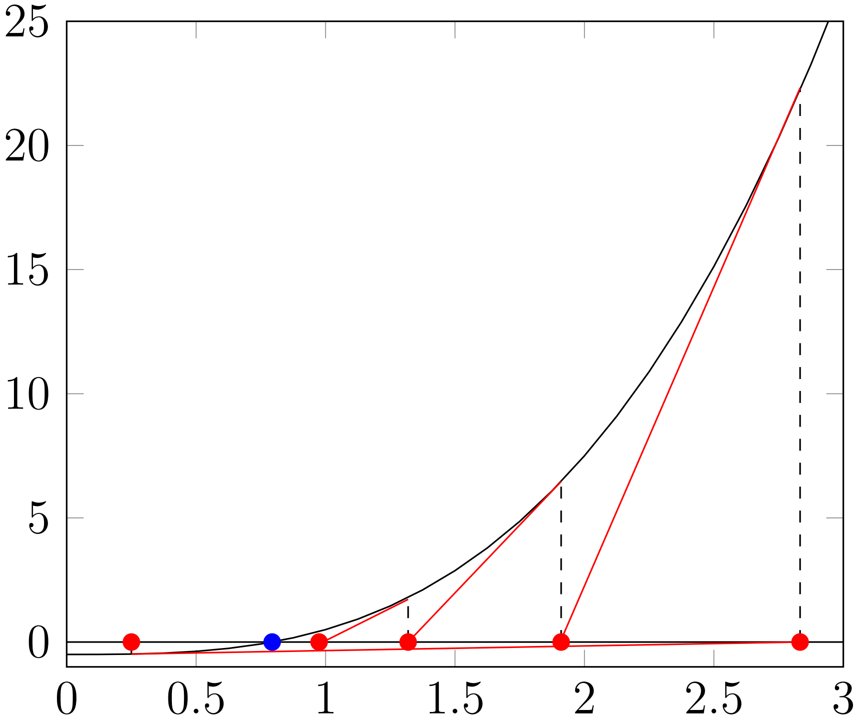

Figure 3.4: An illustration of the Newton-Raphson iteration converging on the root of $f(x) = x^3 - 1/2$, with an initial guess of $x = 1/4$.

introduced. However they can be generalised to arbitrary dimensions for use when dealing with systems of non-linear equations.

A convenient root finding approach is a Newton-Raphson iteration. We aim to find x^* such that $f(x^*) = 0$. We do this by constructing a sequence of guesses x_n using

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)} \xrightarrow{n\to\infty} x^*. \tag{3.81}$$

We stop when the relative change in successive guesses ($|x_{n+1} - x_n|/|x_n|$) shrinks to within some satisfactory tolerance (in practice this tolerance $= 10^{-6}$). This algorithm has second order convergence, assuming that the multiplicity of the root is one, and that the function is smooth. An illustration of this root-finding procedure is shown in figure 3.4.

When the derivative is inconvenient to evaluate, a finite difference approximation can be substituted into equation 3.81 in its place. This leads to the secant method

$$f'(x_n) \approx \frac{f(x_n) - f(x_{n-1})}{x_n - x_{n-1}}, \tag{3.82}$$

$$x_{n+1} = \frac{x_{n-1}f(x_n) - x_n f(x_{n-1})}{f(x_n) - f(x_{n-1})}. \tag{3.83}$$

This method requires two distinct initial guesses near the root, and has golden ratio $\left(\left[1+\sqrt{5}\right]/2 \approx 1.618\right)$ order convergence. The practical implementations we use always employ the secant method.

3.12 Primitive Recovery

Once the system of balance laws (equation (2.63)) has been updated, and the conserved quantities have been evolved to the next time-step, it is standard practice to recover the set of primitive variables. This is because it may not be possible to evaluate the updated flux terms and source terms in the system of equations for the next time-step purely from the updated conserved quantities themselves [74]. Due to relativistic terms the set of equations relating the conserved quantities to the primitive quantities is not algebraically closed. Fortunately an iterative approach can be used to converge on an accurate solution for the updated primitive quantities. We follow methodologies from the HARM and RePrimAnd codes [80, 81].

3.12.1 Hydrodynamic Primitive Quantities

The fluid primitive quantities can be recovered using only the fluid conserved quantities. We have the observer frame conserved quantities D, S^i, and τ, and we require the primitive quantities ρ, v^i, and P. The species fractions are recovered trivially from the conserved quantities by $Y_x = Y_x D/D$. In order to obtain an initial estimate, it is assumed that the change in fluid pressure P is small during the update. In the calculation for the Lorentz factor the value of the fluid pressure from the beginning of the time-step is used as an initial approximation $\widetilde{P}$ for the value of the fluid pressure at the end of the time-step, such that

$$\widetilde{W} = \left(1 - \frac{S_i S^i}{\left(\tau + D + \widetilde{P}\right)^2}\right)^{-1/2}. \tag{3.84}$$

This approximation to the updated Lorentz factor is then used to approximate the updated rest-mass density, specific enthalpy, and specific internal energy,

$$\widetilde{\rho} = \frac{D}{\widetilde{W}}, \tag{3.85}$$

$$\tilde{h} = \frac{\tau + D + \tilde{P}}{D\widetilde{W}}, \tag{3.86}$$

$$\tilde{\epsilon} = \tilde{h} - 1 - \frac{\tilde{P}}{\tilde{\rho}}. \tag{3.87}$$

The equation of state can now be used to calculate an improved approximation to the fluid pressure from these approximate hydrodynamic quantities,

$$P_* = P(\tilde{\rho}, \tilde{\epsilon}) \tag{3.88}$$

$$= (\Gamma - 1)\,\tilde{\rho}\tilde{\epsilon}, \text{ when using a Gamma-law EoS.} \tag{3.89}$$

If the initial fluid pressure approximation was accurate, then it will match the pressure value returned by the equation of state $\left(P_* \approx \tilde{P}\right)$. A sufficiently accurate value for the fluid pressure can be converged upon with a root finding algorithm (as described in section 3.11) for the function $f(P) = P_* - P$. When using a Gamma-law equation of state, the derivative of this function can be approximated with $f'(P) = v_i v^i c_s^2 - 1$ where,

$$v^i = \frac{S^i}{\tau + D + P_*}, \tag{3.90}$$

$$c_s^2 = \frac{\Gamma P_*}{\tilde{\rho}\tilde{h}}. \tag{3.91}$$

Otherwise a finite difference approximation can be used. This iterative root find can be a severe bottleneck for computation time, as it must be computed for every cell within the simulation at each time-step. Once an accurate fluid pressure has been obtained, the remaining primitive quantities (ρ and v^i) found during the process are sufficiently accurate to be carried forwards to the next time-step.

3.12.2 Radiation Primitive Quantities

Once the hydrodynamic primitive quantities have been recovered the fluid velocities can be utilised in the radiation primitive quantity recovery. We have the observer frame conserved quantities $\overline{S}^j$ and $\overline{U}$, and we require the fluid frame primitive quantities F^i and E. We have that

$$\begin{array}{rclcl} E & = & M^{\mu\nu} u_\mu u_\nu & = & W^2 \left(\overline{U} - 2\overline{S}^i v_i + \overline{W}^{ij} v_i v_j\right) \\ F^i & = & -M^{\mu\nu} h^i_\mu u_\nu & = & W \left[\left(\overline{U} - \overline{S}^j v_j\right) h^i_\mu n^\mu + \left(\overline{S}^j - \overline{W}^{jk} v_k\right) h^i_j\right] \end{array} \tag{3.92}$$

where $h^{\mu\nu} = g^{\mu\nu} + u^\mu u^\nu$.

A closure relation is required to determine $\overline{W}^{ij}$. It can be interpolated between fully isotropic (thick) and fully anisotropic (thin) limits,

$$\overline{W}^{ij} = \frac{3\chi(\xi) - 1}{2}\overline{W}^{ij}_{thin} + \frac{3\left[1 - \chi(\xi)\right]}{2}\overline{W}^{ij}_{thick}, \tag{3.93}$$

$$\overline{W}^{ij}_{thin} = \frac{\overline{S}^i\overline{S}^j}{\overline{U}}, \tag{3.94}$$

$$\overline{W}^{ij}_{thick} = E_{thick}W^2 v^i v^j + W(F^i_{thick}v^j + F^j_{thick}v^i) + P^{ij}_{thick}. \tag{3.95}$$

Here $\chi(\xi)$ is some chosen closure scheme, with the anisotropy measure $\xi = \sqrt{F_\mu F^\mu}/E \in [0, 1]$, as described in section 2.8.2 and [41].

The values E_{thick} and F^i_{thick} can be determined assuming the Eddington closure scheme. We have from equation (2.29) that

$$\overline{U} = E_{thick}W^2 + 2W\alpha F^0_{thick} + \alpha^2 P^{00}_{thick}, \tag{3.96}$$

$$\overline{S}^i = E_{thick}W^2 v^i + W\left(\alpha F^0_{thick}v^i + F^i_{thick}\right) + \alpha P^{i0}_{thick}, \tag{3.97}$$

with $F^0 = F^k v_k/(\alpha - v_i\beta^i)$ and $P^{\mu\nu}_{thick} = \frac{1}{3}E_{thick}h^{\mu\nu}$. Constructing $\overline{S}^i v_i$ allows us to invert these coupled linear equations, returning

$$E_{thick} = \frac{3}{2W^2 + 1}[(2W^2 - 1)\overline{U} - 2W^2\overline{S}^i v_i], \tag{3.98}$$

$$\alpha F^0_{thick} = W\frac{\overline{U}(4 - 4W^2) + \overline{S}^i v_i(4W^2 - 1)}{2W^2 + 1}. \tag{3.99}$$

From this we can rearrange (3.97) and obtain

$$F^i_{thick} = \frac{\overline{S}^i}{W} - E_{thick}\frac{4Wv^i}{3} - \alpha F^0_{thick}\left(v^i + \frac{\beta^i}{\alpha}\right). \tag{3.100}$$

Now that we can evaluate both $\overline{W}^{ij}_{thin}$ and $\overline{W}^{ij}_{thick}$, all we require to obtain $\overline{W}^{ij}$ is the anisotropy measure ξ. In order to construct a reasonable initial guess, the assumption is made that the change in ξ is small over one time-step. An initial approximation $\widetilde{\xi}$ can be found using the values for E and F^μ from the beginning of the time-step. Then $\overline{W}^{ij}$ is approximated and used to update E and F^i (as in equation (3.92)) [82]. If the initial anisotropy approximation was accurate then it will match the anisotropy value returned by the newly calculated energy and fluxes (that is $\widetilde{\xi} = \sqrt{F_\mu F^\mu}/E$).

A root finding algorithm (as described in section 3.11) can be used to converge upon an accurate anisotropy measure by approximating the zero of the scaled function

$$f(\xi) = \frac{E^2\xi^2 - F^\alpha F_\alpha}{\overline{U}^2}. \tag{3.101}$$

The derivative $f'(\xi)$ can be approximated by finite differencing,

$$f'(\xi) \approx \frac{f(\xi^+) - f(\xi^-)}{\xi^+ - \xi^-}, \tag{3.102}$$

$$\xi^+ = \min(1, \xi + 10^{-6}), \tag{3.103}$$

$$\xi^- = \max(0, \xi - 10^{-6}), \tag{3.104}$$

ensuring that all approximations to ξ are in the valid range $[0, 1]$.

If the Eddington closure is being used, this radiation primitive recovery simplifies as $\xi \equiv 0$. This results in the thick limits being taken (E_{thick}, F^i_{thick}), and no iterative root find is required.

3.13 Adaptive Mesh Refinement

To reduce the cost of computation, adaptive mesh refinement (AMR) may be used. When using AMR, regions of the spatial domain can be refined to a higher resolution to increase numerical accuracy, and capture any detailed structure or behaviour of interest. Simultaneously the large bulk of the domain away from these regions can be kept at a significantly coarser resolution [50]. This is advantageous over a uniform mesh approach which would require high resolution throughout the domain, and thus much more computation time to capture the same complex behaviour.

An alternative approach to increasing numerical accuracy would be to use higher order numerical methods. However low order methods have the advantages of being relatively robust, simple, and fast. The AMR approach allows for accurate simulations whilst using relatively low order methods.

Regions of interest where high resolution patches may be concentrated include shock-fronts, defined as where the spatial gradients of physical variables become steeper than some given critical threshold. As these regions of interest appear, disappear, and move around the mesh, the higher resolution regions can be adaptively created, destroyed, or moved in order to track them. This focuses high numbers of computational cells where they are useful, and maintains low numbers of computational cells where the resolution requirements can be relaxed.

I am using the AMReX library [83] developed at LBNL to implement adaptive mesh refinement into my simulation code. The AMReX library is a software framework containing all the functionality required to write massively parallel, block-structured adaptive mesh refinement applications.

3.13.1 Block-Structured AMR

A type of structured AMR is patch-based AMR, which is useful when using Cartesian meshes [84]. With patch-based AMR the cells on the coarsest level (referred to as the base level, or level-0, see the black grid in figure 3.5 [85]) which meet some refinement criteria are tagged for refinement, such as large pressure gradients or high velocities, which may result in larger errors or slower convergence in calculations. A load balancing algorithm then constructs rectangular patches of cells which cover these tagged locations [86], whilst aiming to optimise the balance between the memory cost of increasing the number of cells, and keeping the total number of patches small, as they will require a exchange of data at their boundaries. These patches are at a higher resolution than the base level, according to some refinement ratio (often twice the resolution of the underlying cells). This collection of refined patches are referred to as level-1 (see the blue grid in figure 3.5). This process can be repeated, nesting more levels until the desired resolution (or maximum number of levels) is reached. Typically level-0 is the coarsest level and covers the entire domain, with higher resolution levels being concentrated around regions of interest. Our system of balance laws can be solved on an individual patch. In order to use our spatial reconstruction schemes (such as the linear and WENO methods introduced in section 3.5) each AMR patch requires an appropriate number of ghost cells. At the interface between patches on the same level, these ghost cells can be filled by copying the data from the neighbouring patch. At a level boundary overlaying a coarser grid, the ghost cells can be filled by copying data up from the lower level, and up-scaling its resolution. At domain boundaries the AMR patch ghost cells can be filled according to the physical boundary conditions describing the problem to be modelled.

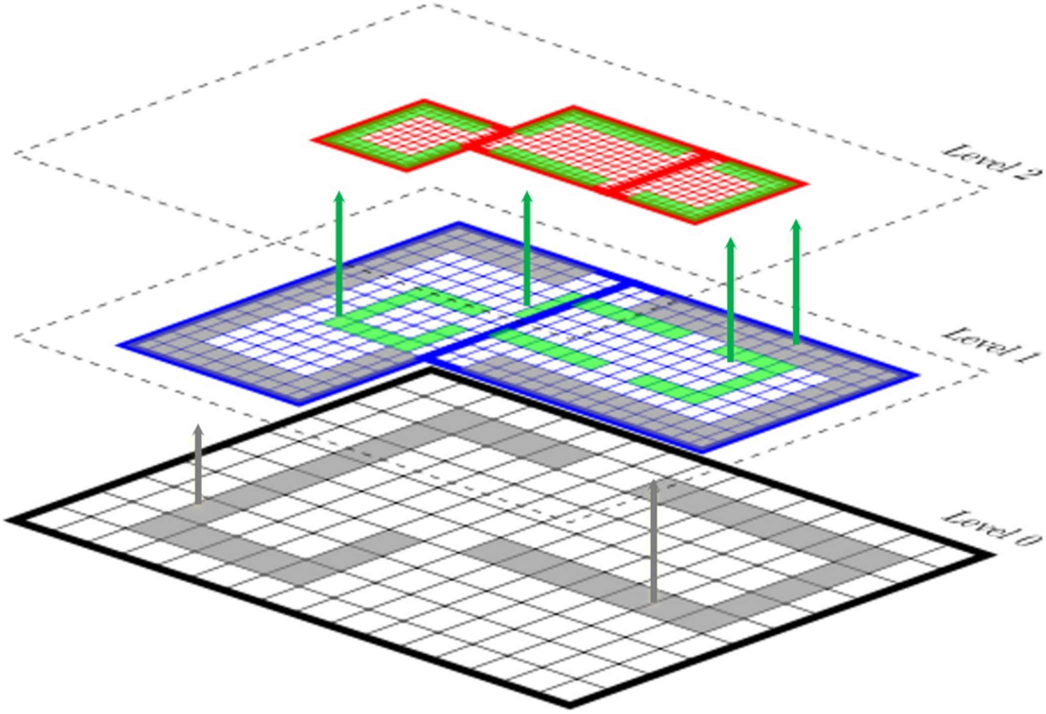

Figure 3.5: Block-structured adaptive mesh refinement illustrating patches and levels. In this example the black grid is level-0, the blue grid is level-1, and the red grid is level-2. © Copyright 2023, Rob Farber.

3.13.2 Sub-Cycling

In order to obtain a stable evolution, and satisfy time-step constraints which are linked to the spatial resolution (as in section 3.8), refined regions of the domain can be sub-cycled in time. This involves evolving a coarse level by a full time-step, and evolving a fine level by half of a time-step. The coarse solution is then interpolated in time to obtain a coarse solution at the half time-step, which is up-scaled and used to fill the ghost cells in the patch boundaries of the fine solution (any ghost cells at domain boundaries can be populated using the physical boundary conditions at the half time-step). Now the fine solution has the information it requires to evolve the remaining half time-step, re-synchronising with the coarse solution. The data on the fine grid is likely to be more accurate than that on the coarse grid. To take advantage of this the fine solution can be down-scaled and copied into the coarse grid, aiming to reduce its error. This process is illustrated in figure 3.6. This sub-cycling procedure is implemented recursively to account for multiple levels of refinement. Temporal sub-cycling allows the evolution to be stable, whilst improving computational efficiency by avoiding updating the entire domain with the smallest time-step

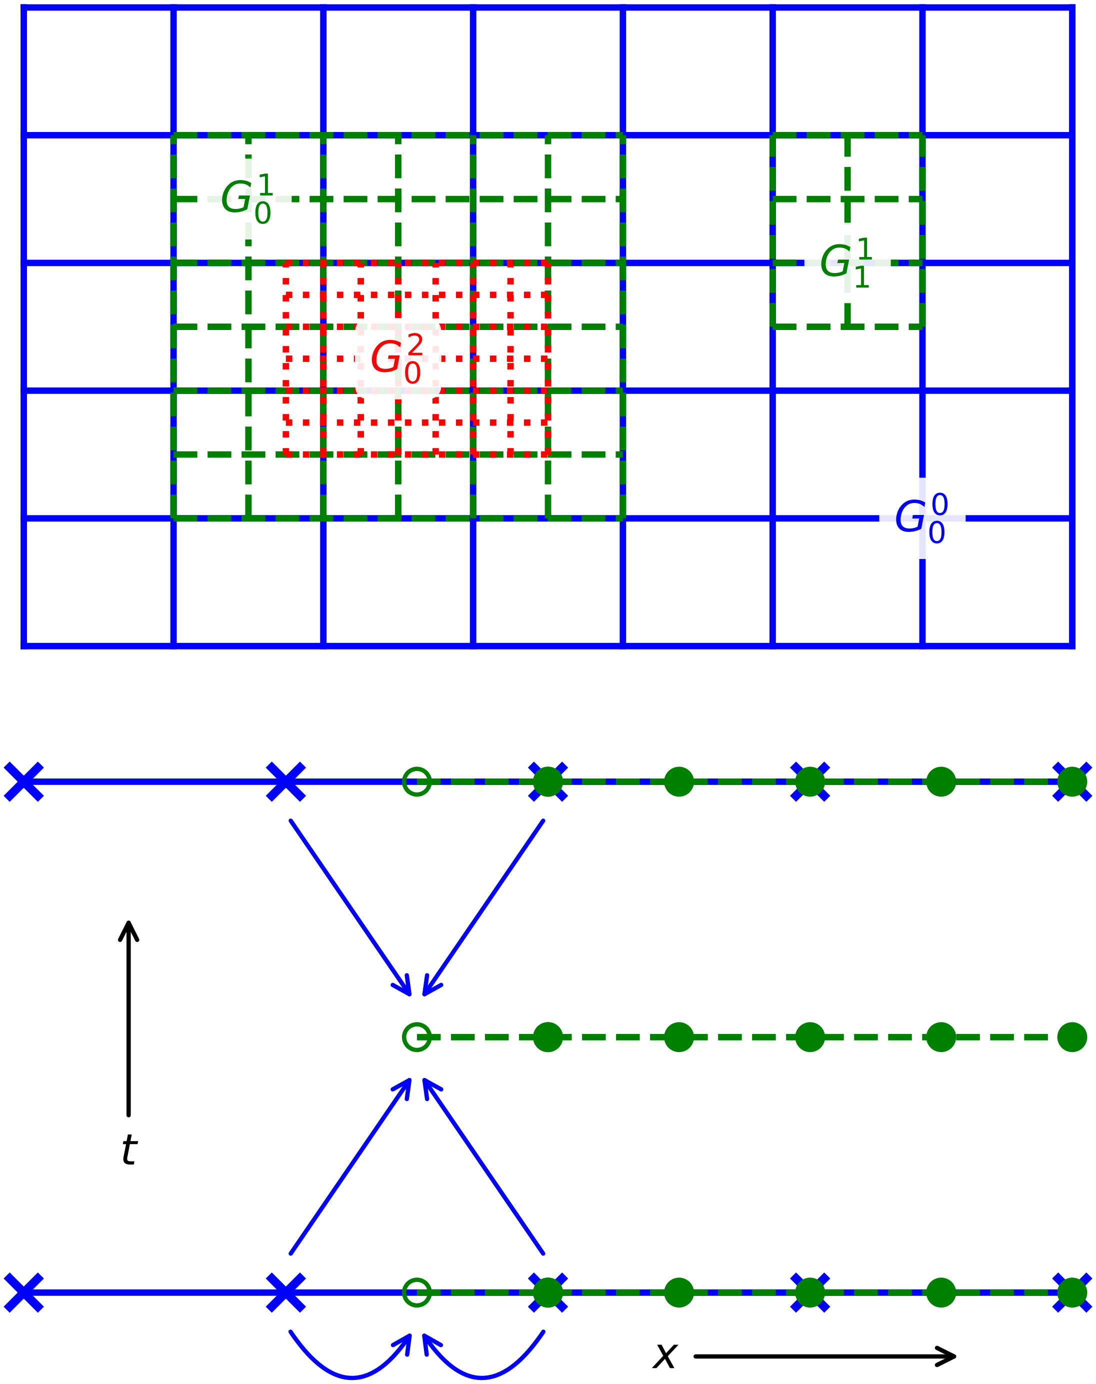

Figure 3.6: An illustration of a fine grid (in green) being sub-cycled, using information from a time interpolated coarse grid (in blue). Image from github.com/IanHawke [87].

determined by the highest resolution region.

3.13.3 Refluxing

A potential issue with patch-based AMR is that a fine level is updated with fluxes and source terms computed from high resolution data, whilst a coarse level is updated with fluxes and source terms computed from low resolution data. At the level interface we can have inconsistent fluxes, due to the differing data sources. This can lead to required quantities not being conserved.

This issue of inconsistent fluxes can be fixed with a process known as refluxing, discussed in [83]. Refluxing is a process which keeps track of level-boundary fluxes, and corrects the coarse solution with the difference in coarse and fine fluxes.

However this does not take into account that the coarse and fine solutions have been updated with different source terms. In Newtonian simulations the differences from gravitational sources can be re-synchronised as in the gravitational synchronisation step used in the Castro code [88].

Our source terms involve the coupling between radiation and fluid fields, where the energies and momenta evolve towards the local thermodynamic equilibrium (LTE) state, as in section 3.10.2. The coupling source terms are sensitive to the radiation and fluid temperatures (proportional to T^4), and the refluxing process may push the solutions away from their LTE limits.

Another approach to avoiding the coarse-fine flux mismatch is to avoid having shock-waves pass through level boundaries, as it is around shock-fronts where this mismatch in fluxes can cause larger errors. Our AMR patches can track propagating shocks, preventing their interaction with level boundaries until they have dissipated and weakened, or left the region of study. In our implementation we do not employ a refluxing procedure, as the flux inconsistency may be preferable to introducing LTE inconsistencies.

3.13.4 Parallelisation and Scaling

With an AMR level spatially decomposed into a number of patches, this allows efficient parallel computing algorithms to distribute the collection of patches amongst multiple processors. The evolution equations can be solved on the AMR patches in parallel, with any information required at the patch boundaries already being accessible from the ghost-zones.

In order to evolve our system of balance laws on an AMR patch, the processor requires knowledge of the primitive quantities on the patch and in its ghost cells. It also requires knowledge of the appropriate closure information, such as any tabulated equations of state, radiation pressure tensor formulae, and opacity data.

With modern high performance computing (HPC) we are able to massively parallelise our domain decomposed simulation with the use of MPI [89]. However memory restrictions become a significant bottleneck, as the required closure information can reach multiple gigabytes of data. This is difficult for CPU parallelisation, and unfeasible for GPU parallelisation. In our case we have low memory requirements for our empirical formulae for opacity and radiation pressure tensor calculations. However our fluid equation of state table (the three parameter SFHx data discussed in section 2.8.1) is very large. To overcome this we have used the shared memory capability available with OpenMP [90] to fit the EoS table into the shared memory, which can be accessed by each of the OpenMP threads updating the AMR patches. I have made use of the AMReX libraries capability to implement hybrid MPI and OpenMP parallelisation.

When spatially decomposing the domain into AMR patches, consideration must be taken of the number of patches being created. A naive approach would be to create a tiny patch for each coarse cell tagged for refinement. This would have the benefit of completely avoiding introducing fine cells over regions which do not require refinement. However each of these many tiny patches would require sufficient ghost zones for our spatial reconstruction scheme. The copying and up-scaling of data from the coarse grid into each of these ghost zones would become incredibly computationally expensive. Processor communication can become more expensive than computation, harshly reducing the code's performance. At the other extreme a significant proportion of the spatial domain could be used to create a single huge refined patch. This would have the benefit of minimising the number of ghost zones at the patch boundaries, and thus minimising the processor communication required to copy and upscale the coarse data. This approach would not be very useful as many unnecessary refined cells would be created, maximising the number of cells which would require updating. This would be equivalent to simply constructing a uniform grid at a higher resolution, avoiding any benefits we could attain from using AMR. We aim to approach some optimal balance, to minimise the number of fine cells to update, and minimise the number of ghost zones to populate via processor communication. Such a load balancing algorithm is included within the AMReX library and appropriate methodologies are discussed in [91].

The AMReX software framework supports the development of block-structured AMR algorithms for solving systems of partial differential equations on machines from laptops to exascale supercomputer architectures [85]. Alternative mesh refinement capable codes include the Einstein Toolkit, which makes use of the Carpet environment, allowing for fixed mesh refinement (FMR, also referred to as box-in-box). The Einstein Toolkit code's uniform grid capability (using the PUGH driver) is highly scalable, up to 130,000 cores [92]. Its AMR capability (using Carpet driver) has demonstrated scalability to over 10,000 cores on a weak scaling test [2].

A simulation code's weak scaling is how its computation time varies with the number of processors, whilst keeping the problem size per processor constant. An alternative measure of a code's scalability is the strong scaling. This is how the computation time varies with the number of processors, whilst keeping the total problem size fixed. The strong scaling is limited by the parts of the calculation which cannot be parallelised.

The CarpetX and GRaM-X projects are working to implement AMReX block-structured AMR algorithms into the Einstein Toolkit code, described in [3, 93]. This highlights the expectation in the field that AMReX will be useful for numerical relativity, supporting our choice of approach.

3.14 Summary

The techniques used to discretise the system of balance laws described in section 2.7 for implementation into a computational code are described in this chapter. The balance laws are spatially discretised with a finite volume scheme, requiring the fluxes to be calculated at cell interfaces and the sources to be cell-centred allowing the update of cell centred quantities. In order to calculate the fluxes at cell interfaces, the known cell-centred quantities must be reconstructed to either side of the cell interfaces. Constant, linear, parabolic, and high order WENO reconstruction schemes allow the numerical scheme to be 1st, 2nd, 3rd, and higher order in space respectively. The reconstruction must be limited so as to avoid introducing any new local extremal points, such reconstructions are total variation diminishing (TVD). Once quantities have been reconstructed at either side of a cell interface, it is treated as a Riemann problem. The signal speeds either side of the interface dictate how the interface flux will be calculated, and whether the wave is a shock, rarefaction or a contact wave.

With the flux and source terms calculated, the evolved quantities can be updated by a single time-step. For this system to be solved explicitly the time-step must be restricted by the time taken for the fastest wave speed to traverse a single cell-width. This CFL condition is a requisite for numerical stability. The evolved quantities can be updated with explicit time integration schemes such as Runge-Kutta methods. Once the evolved quantities have been updated, the primitive quantities must be recovered, from which the fluxes and source terms can be calculated for the next update. In our case the primitive quantity recoveries require iterative approaches.

Software packages such as AMReX [94, 95] contain libraries of routines which can be used to manage AMR patches during a numerical simulation. These AMR routines make use of a patch-solver which can be written to solve the required system of ODEs on each individual patch.

I have implemented a general relativistic radiation coupled hydrodynamics patch solver into the AMReX framework. This patch solver solves the set of balance laws to conserve mass, momentum, and total energy on a curved spacetime as described in (2.63). It uses the HLLC Riemann solver for the hydrodynamic equations and the HLL Riemann solver for the radiation equations. The radiation wave-speeds are over-estimated with the speed of light in a vacuum. The code is up to fifth order in space due to the WENO reconstruction used for projecting cell-centred quantities to cell faces, and the up to third order in time from the three-step third-order SSP Runge-Kutta time integration scheme.

Test problems demonstrating the performance of the numerous aspects of the general relativistic radiation hydrodynamics code are described and their results shown in chapter 4.

Chapter 4

Test Problems and Results

4.1 Introduction

For the purposes of simulating a hot, dense, complex material in a slowly evolving spacetime, I have written a code which evolves fluid and radiation within a fixed yet general background spacetime. As outlined in chapters 2 and 3, this code has been designed for general relativistic radiation hydrodynamics simulations. This chapter provides a description and discussion of a series of test problems aimed at validating and verifying that the numerical methods employed can accurately recover the desired physical processes, and that the processes implemented are appropriate to describe astrophysical phenomena of interest.

I quantitatively test smooth relativistic fluid flow, and make measurements of the order of accuracy. The qualitative properties I test include the accurate evolution of discontinuities and shocks, radiative transfer in both diffusive and streaming limits, radiation and fluid coupling, and how well the expected behaviours are maintained for high fluid velocities and curved background spacetimes.

I confirm that the adaptive mesh refinement can track features of interest and can improve numerical accuracy without increasing the spatial resolution everywhere. In some situations this can allow for more practical simulation times, due to maintaining a reduced number of computational cells (this will be demonstrated with 3D simulations in the next chapter).

Many of the following test scenarios can be found in [96, 97] which focus on relativistic simulations for black hole accretion problems. While these scenarios are not our intended application, they involve fluid and radiation evolution at relativistic velocities and utilise a fixed background spacetime, making them appropriate tests for reproduction.

A method of quantitatively measuring the accuracy of a numerical simulation is to calculate the deviation from an exact solution. One simple way of measuring this deviation is with an L_p-norm

$$L_p\left(\widetilde{f}\right) = \left(\sum_{n=1}^{N} \left|f(x_n) - \widetilde{f}_n\right|^p \Delta x_n\right)^{1/p}. \tag{4.1}$$

The numerical error for a simulation should decrease as the resolution is increased. The rate of this decrease indicates the spatial order of accuracy of the simulation. I will use the L_2 error metric in the following quantitative validation tests.

In this chapter I continue to use the $c = G = M_\odot = 1$ unit system.

4.2 Special Relativistic Hydrodynamics

The tests in this section are used to validate the code's capability to model hydrodynamic evolution involving special relativistic effects. These include non-rest mass, Lorentz contraction, and special relativistic time dilation. These effects are expected to occur in the fluid involved in neutron star merger remnants, where fluid velocities are expected to exceed 40% of the speed of light.

4.2.1 Relativistic Isentropic Pulse

This 1D test simulates the smooth hydrodynamic flow of a pulse of fluid travelling at relativistic velocities [96]. It is used to validate the accurate reconstruction of primitive quantities at zone interfaces. For smooth fluid profiles I expect greater than first order accuracy from the WENO reconstruction scheme introduced in section 3.5. As this test involves velocities over 0.5 it is also a validation test of special relativistic effects.

The initial pulse-shaped density profile is defined as

$$\rho(x) = \begin{cases} 1 + \left[(x/L)^2 - 1\right]^2, & x \in (-L, L), \\ 1, & x \notin (-L, L). \end{cases} \tag{4.2}$$

A polytropic equation of state $\left(P = K\rho^{\Gamma} = (\Gamma - 1)\rho\epsilon\right)$ is used to obtain the initial pressure profile. The sound speed can be obtained with $c_s = \sqrt{\Gamma P/h\rho}$, where $h = 1 + \epsilon + P/\rho$. The Riemann invariants are given by

$$J_{\pm} = \tanh^{-1}(v) \pm \frac{2}{\sqrt{\Gamma - 1}} \tanh^{-1}\left(\frac{c_s}{\sqrt{\Gamma - 1}}\right). \tag{4.3}$$

One of the Riemann invariants (J_-) is kept constant, and can be evaluated from a reference state where $\rho = 1$ and $v = 0$. This can then be used to obtain the velocity profile which keeps the Riemann invariant constant

$$v = \tanh\left(J_- + \frac{2}{\sqrt{\Gamma - 1}} \tanh^{-1}\left(\frac{c_s}{\sqrt{\Gamma - 1}}\right)\right). \tag{4.4}$$

An analytic solution can be found from characteristic analysis. The other Riemann invariant (J_+) is not constant. Its initial profile can be found with equation (4.3). It is invariant along the characteristic paths

$$x = x_0 + t\frac{v(x_0, 0) + c_s(x_0, 0)}{1 + v(x_0, 0)c_s(x_0, 0)}, \tag{4.5}$$

leading to $J_+(x, t) = J_+(x_0, 0)$. For any position x and time t we can use a numerical root-finding algorithm to solve the equation

$$J_+(x_0, 0) = J_+\left(x - t\frac{v(x_0, 0) + c_s(x_0, 0)}{1 + v(x_0, 0)c_s(x_0, 0)}, 0\right), \tag{4.6}$$

obtaining x_0, allowing us to evaluate $J_+(x, t) = J_+(x_0, 0)$.

With both Riemann invariants we can recover the velocity profile from equation (4.3),

$$v(x, t) = \tanh\left(\frac{J_- + J_+(x, t)}{2}\right). \tag{4.7}$$

Then equation (4.3) can be used to recover the sound speed, which in turn can be used to recover the pressure and density profiles.

I expect to see the pulse travel towards to the right, maintaining its width and height, with the front of the pulse steepening and approaching a shock-front. Eventually a shock will form, this simulation is stopped before this time. Behind the travelling pulse the fluid returns to its reference state.

For this problem I use the parameters $L = 0.3$, $K = 100$, and $\Gamma = 5/3$, leading to the constant Riemann invariant $J_{-} \approx -8.463$. I use the domain $x \in [-0.35, 1]$, discretised into 80 zones. The fluid is evolved to a time of 0.8.

Numerical results for the density, pressure, and velocity profiles have been plotted in figure 4.1. The results are consistent with those shown in [96]. The plots show the initial condition with a dashed curve, highlighting how the pulse has clearly travelled to the right, maintaining its height and width. Solid curves show the analytic solutions found via characteristic analysis. It can be seen that the numerical results shown with circle markers agree very closely with the analytic solutions.

This close agreement can be quantified with the L_2 error norm (4.1). The numerical test has been repeated with a range of explicit Runge-Kutta initial value problem solvers introduced in section 3.3, the first order Forward Euler method (RK1), the second order Heun's method (RK2), and the third order Shu-Osher method (RK3). For each method the numerical test has been repeated to cover a range of spatial resolutions $N_x \in \{20, 40, 80, 160, 320, 640\}$. With this information I can find how the error changes with spatial resolution for each method. The results are tabulated in table 4.1, and plotted in figure 4.2 along with linear regressions (which exclude the outlier $N_x = 20$ points). The RK1 method shows close to first order convergence, and the RK2 and RK3 methods show close to second order convergence. For practical purposes the RK2 and RK3 methods show qualitatively indistinguishable errors.

4.2.2 With Adaptive Mesh Refinement

In order to verify that the code's AMR meshing works as intended, the test has been repeated using the RK2 method, with two levels of factor-2 refinement covering the region $\rho \geq 1.01$. The refined patches adaptively track the density pulse. I discretise the coarsest level with $N_x = 80$ zones, which means the effective resolution is 320 zones. The results are shown in figure 4.3. This AMR simulation has an L_2-error of 2.389×10^{-3}, improving on the uniform grid $N_x = 160$ case, but is not quite as good as the comparable $N_x = 320$ uniform grid, highlighting the benefit of AMR.

4.2.3 Relativistic Shock Tube

This test is used to validate the code's capability to accurately capture shock waves, rarefaction waves, and contact discontinuities. This is expected from the implementation of the HLLC Riemann solver introduced in section 3.7.2.

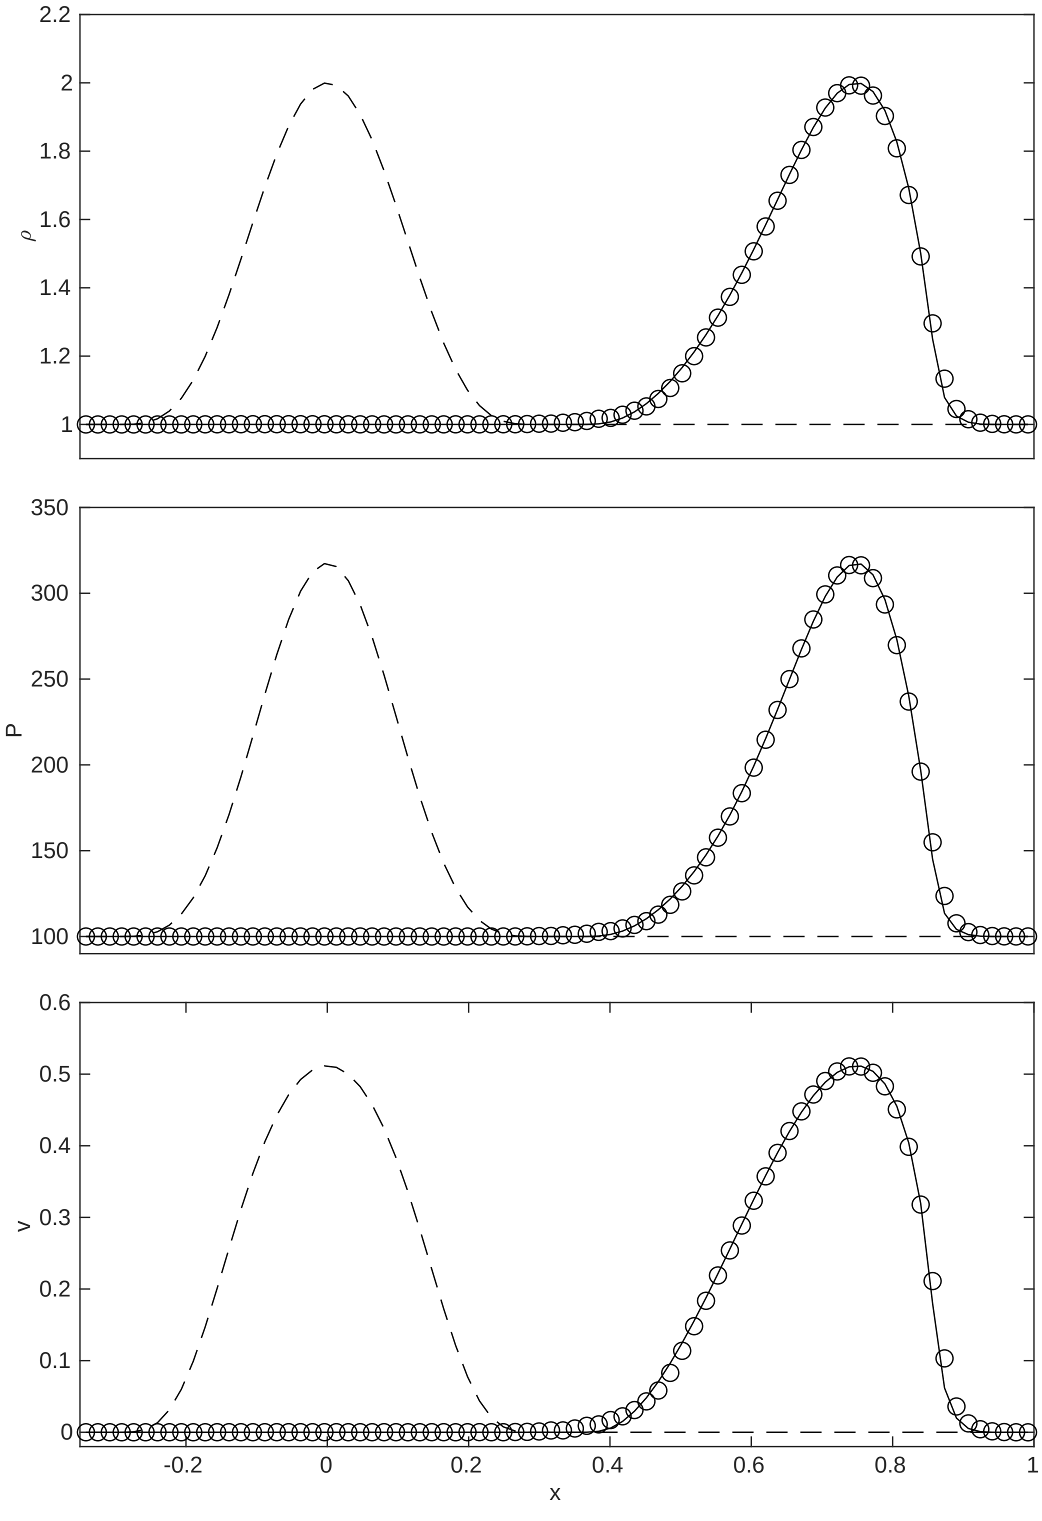

Figure 4.1: Solution profiles for the isentropic pulse test. Quantities are, top-to-bottom: rest-mass density, fluid pressure, and fluid 3-velocity. The initial profiles are shown with dashed curves and circle markers indicate the numerical results at time 0.8. We can see that the pulse width and height remains constant, with the leading front steepening over time reproducing the results in [96] and matching the exact solutions shown by the solid curves.

Scheme	N_x	L_2-error	Convergence Rate
RK1	20	5.857×10^{-2}	–
	40	5.265×10^{-2}	0.154
	80	3.011×10^{-2}	0.8059
	160	1.670×10^{-2}	0.851
	320	8.856×10^{-3}	0.915
	640	4.770×10^{-3}	0.893
RK2	20	7.151×10^{-2}	–
	40	3.139×10^{-2}	1.188
	80	1.061×10^{-2}	1.565
	160	3.141×10^{-3}	1.756
	320	7.939×10^{-4}	1.984
	640	2.250×10^{-4}	1.819
RK3	20	7.151×10^{-2}	–
	40	3.181×10^{-2}	1.168
	80	1.083×10^{-2}	1.555
	160	3.214×10^{-3}	1.753
	320	8.280×10^{-4}	1.957
	640	2.558×10^{-4}	1.695

Table 4.1: Isentropic pulse L_2 error norms for a range of IVP solvers and spatial resolutions. The final column shows the convergence rates. In each case the rate approaches the optimal value (1^{st} or 2^{nd} order convergence) expected.

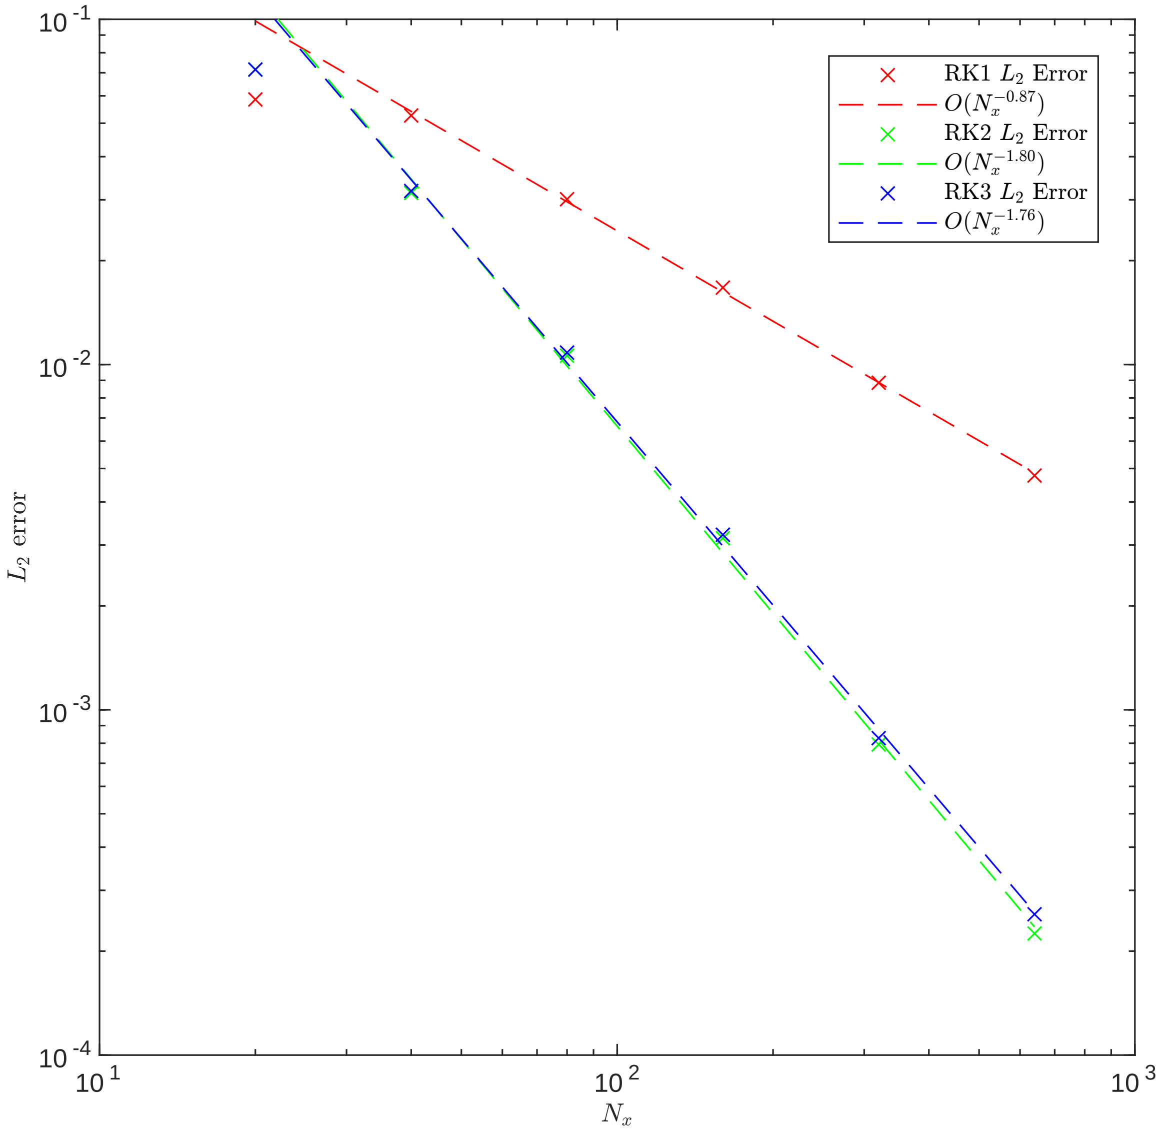

Figure 4.2: Isentropic pulse test L_2 error against spatial resolution for RK1, RK2, and RK3 methods. Convergence rates are 0.87, 1.80, and 1.76 respectively. The coarsest ($N_x = 20$) data is not included in the linear regressions. The RK2 and RK3 methods return qualitatively indistinguishable results.

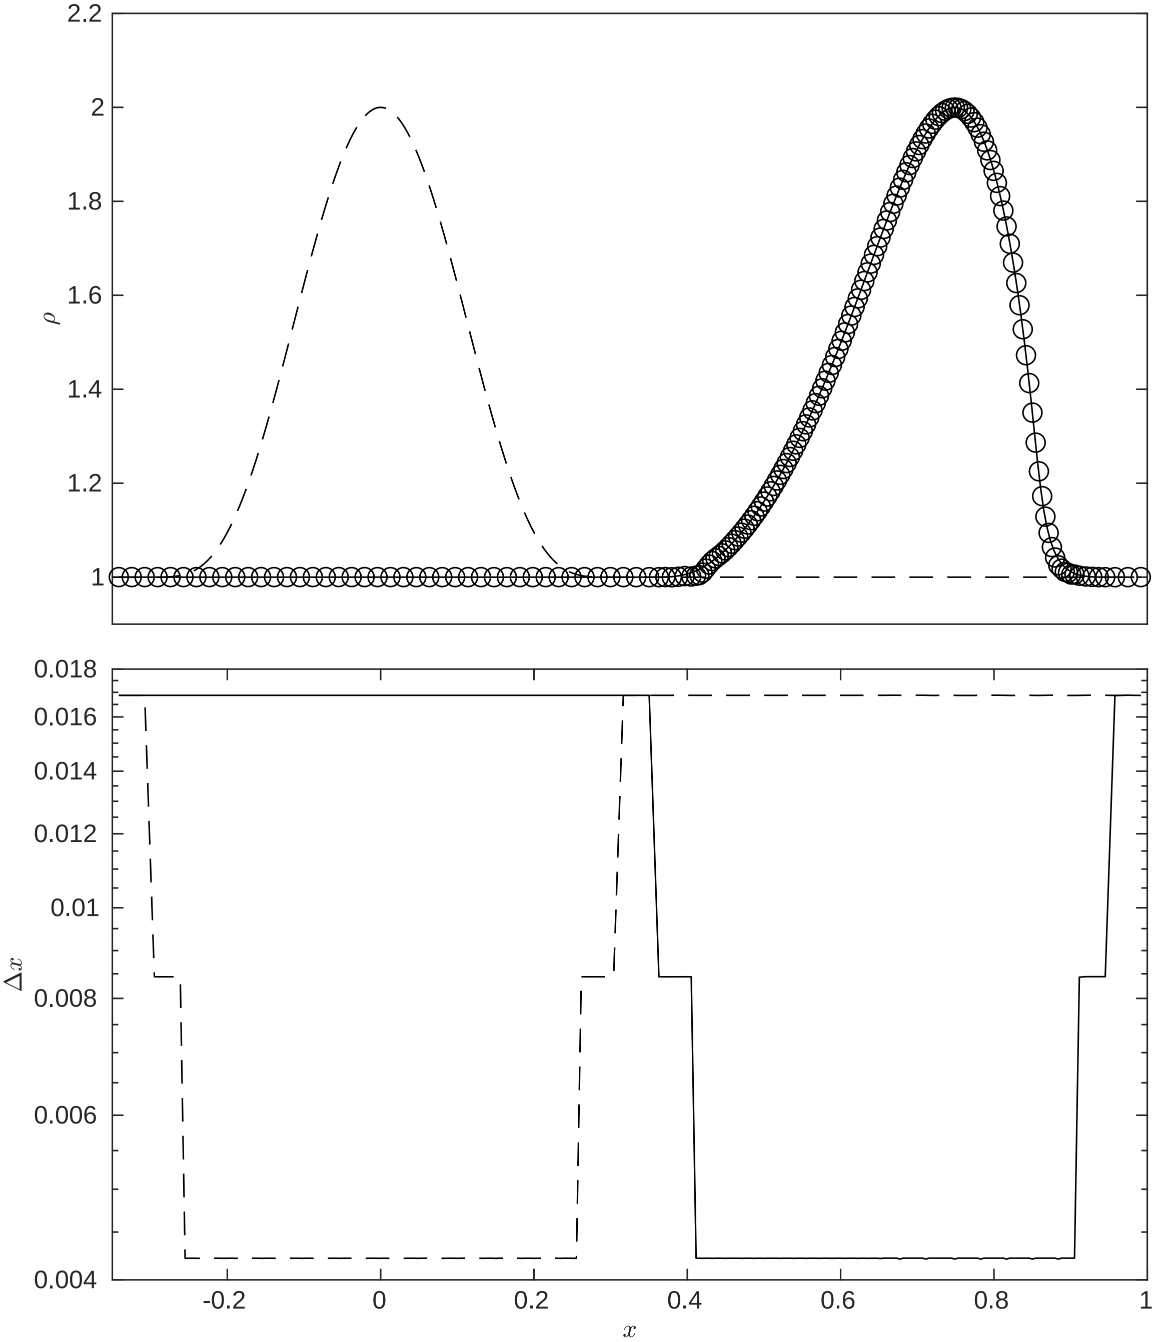

Figure 4.3: Isentropic pulse test, with two levels of factor-2 refinement. The base grid (AMR level-0) has a resolution of $N_x = 80$. The refined patches track the region where $\rho \geq 1.01$.

ρ_L	P_L	v_L	ρ_R	P_R	v_R
10	40/3	0	1	10^{-8}	0

Table 4.2: Relativistic shock tube initial conditions.

It is a 1D hydrodynamic test, modelling the evolution of a pressure and density discontinuity. The setup is as described in [96, 97]. The initial conditions are given in table 4.2, and describe two uniform physical states separated by a discontinuity. Here the ideal fluid equation of state with $\Gamma = 5/3$ has been used.

The initial discontinuity produces a shock wave which reaches relativistic velocities exceeding 0.7. It is a further verification test that special relativistic effects are captured accurately. I use the spatial domain $x \in [-1/2, 1/2]$, discretised into 200 zones, and evolve to time 0.4 using the second order Heun's method for temporal integration (as introduced in section 3.3), the fifth order WENO spatial reconstruction scheme (as introduced in section 3.5), and the three wave HLLC Riemann solver (as introduced in section 3.7.2).

I can obtain exact solutions from the Riemann solver code included in [34]. I expect to see the discontinuity evolve into a right-travelling shock wave, a left-travelling rarefaction wave, and a contact discontinuity in the density field.

Numerical results for the density, pressure, and velocity profiles are shown in figure 4.4. These match very well with results in both [96, 97]. The solid curves show the exact solutions generated by the code from [34]. It can be seen that the numerical results shown with circle markers agree very closely with the exact solutions.

4.2.4 2D Relativistic Shock

To validate that the code can accurately model special relativistic hydrodynamics with shock waves, rarefaction waves, and contact discontinuities in multiple dimensions, a 2D test has been constructed using the initial data and methods from section 4.2.3, given in table 4.2, with the initial discontinuity at a 45 degree angle to the Cartesian mesh. The density, pressure, and horizontal and vertical velocity surfaces at time 0.4 have been plotted in figure 4.5. The results are consistent with the 1D tests. It can be seen that effects from the unaligned boundaries are propagating into the domain from the top left and bottom right corners. This highlights the usefulness of AMR schemes to set the boundaries far away from features of interest.

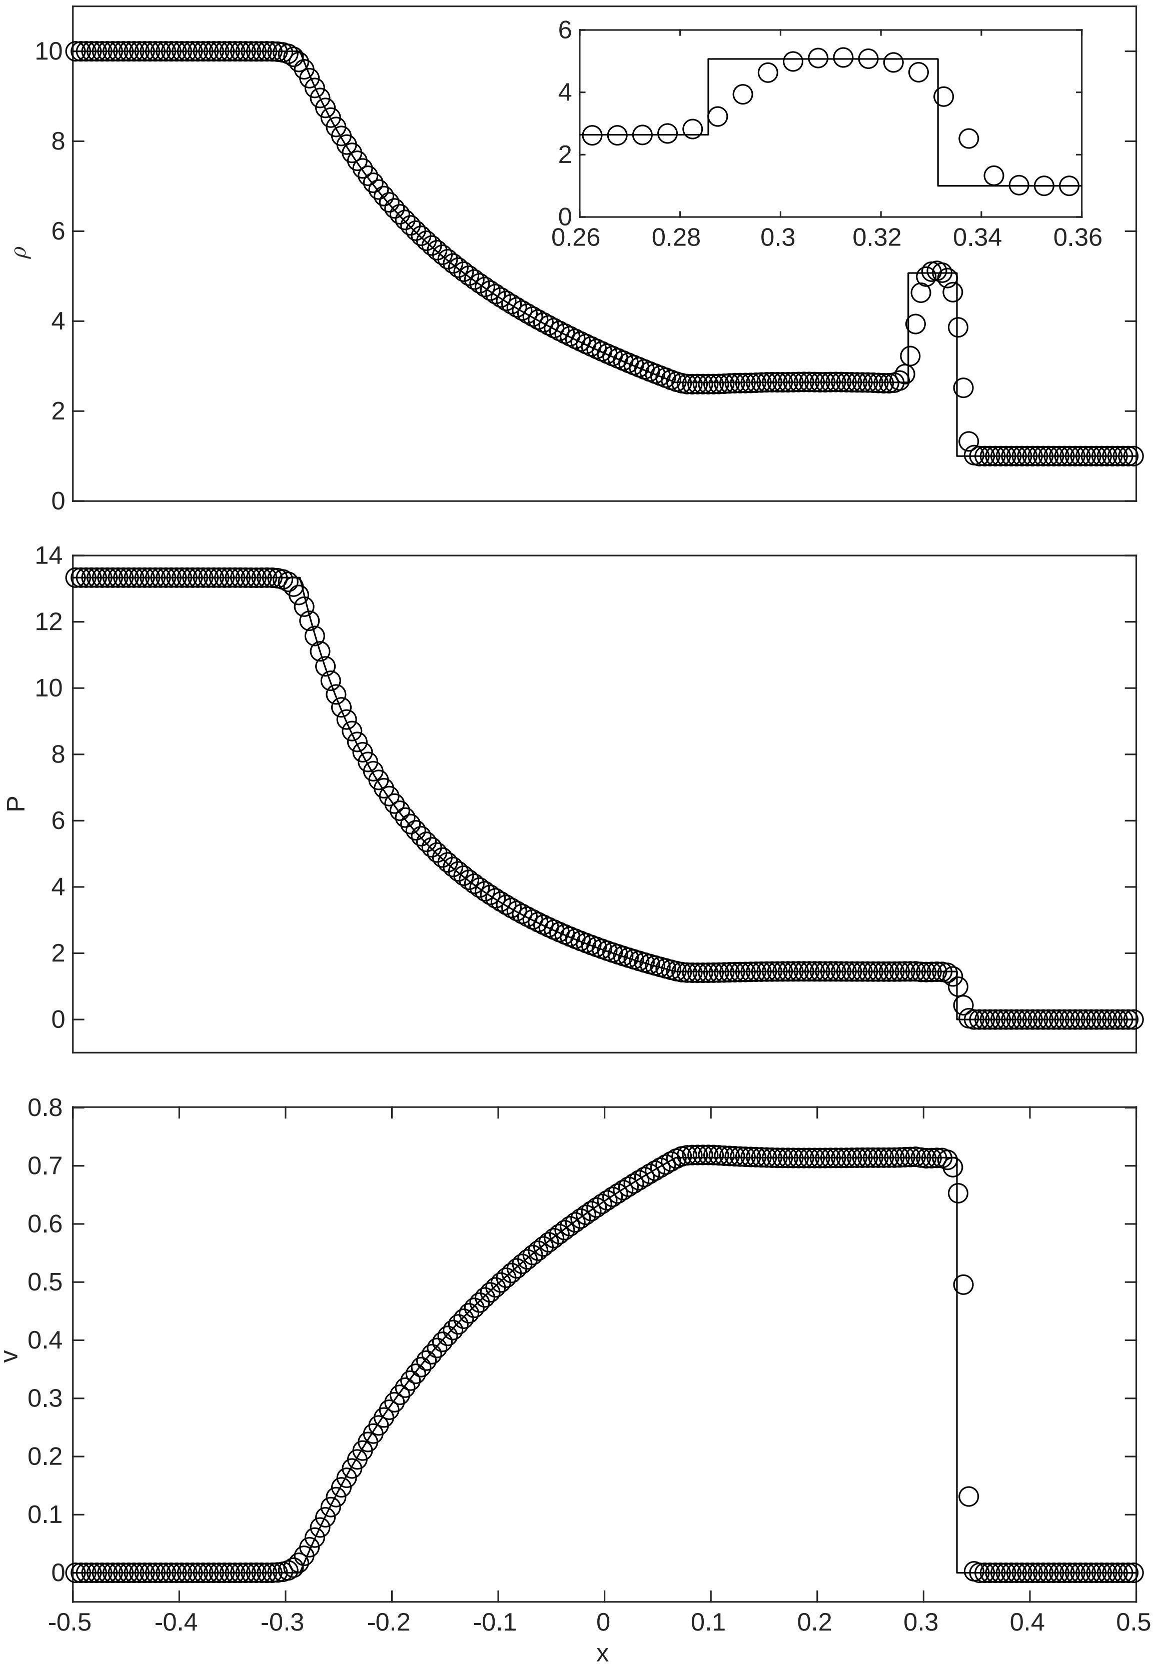

Figure 4.4: Profiles for the relativistic shock tube test. From top to bottom, are rest-mass density, fluid pressure, and fluid velocity. At a time of 0.4 we can see that a shock front, rarefaction wave, and a contact discontinuity have formed, and there are no significant oscillations around the discontinuities, matching the results in [97].

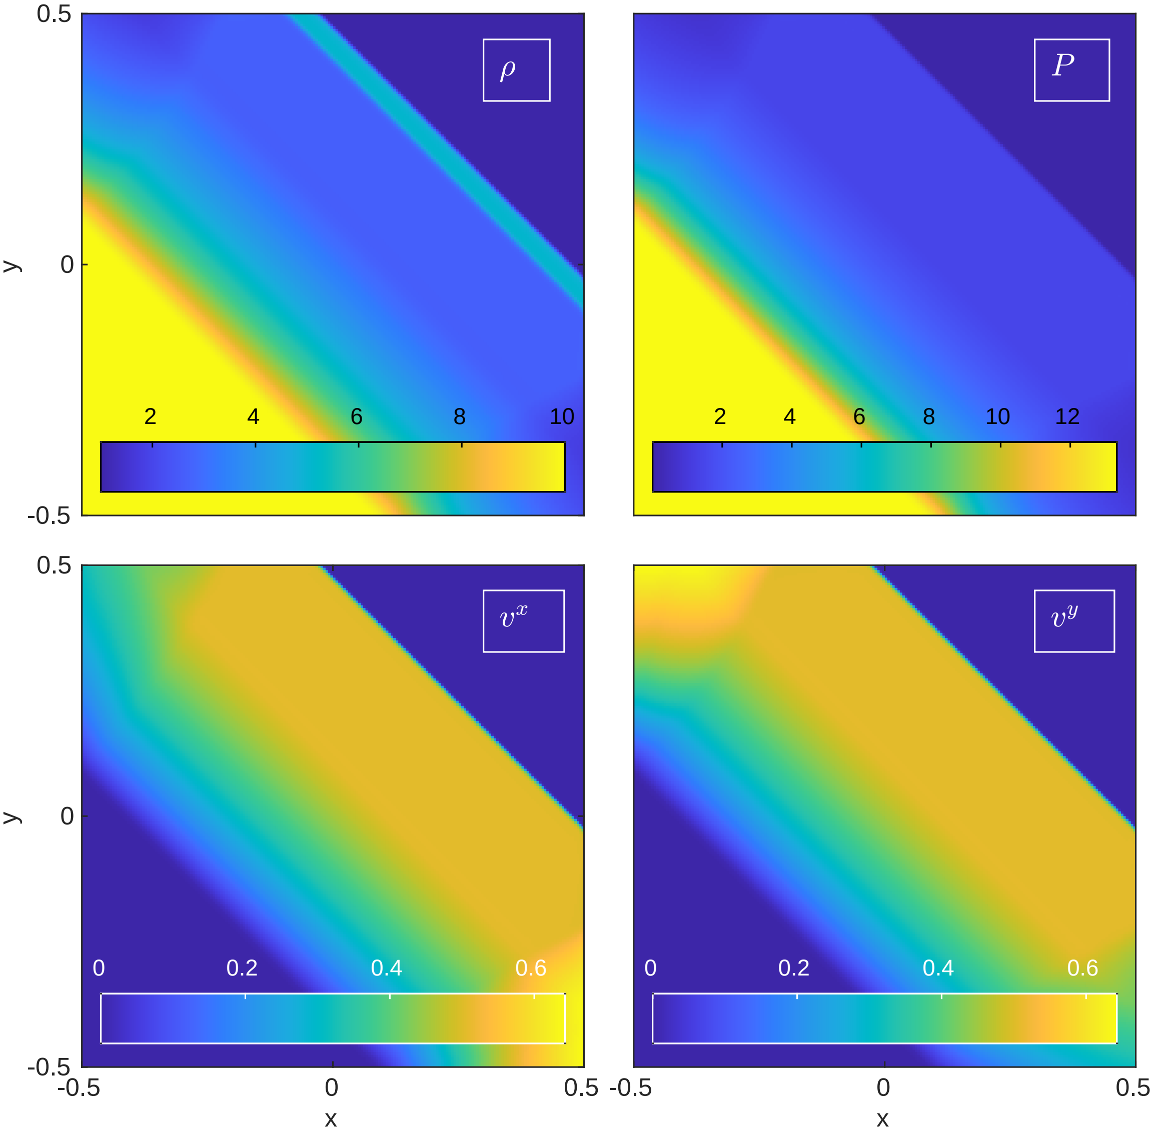

Figure 4.5: The 2D diagonal SR shock tube test results at time 0.4. The top left panel shows ρ, top right P, bottom left v^x, and bottom right v^y. Effects propagating in from boundaries can be seen in the top left and bottom right corners of each panel.

quadrant	ρ	P	v^x	v^y
$x < 0, y < 0$	0.5	1	0	0
$x > 0, y < 0$	0.1	1	0	0.87
$x < 0, y > 0$	0.1	1	0.87	0
$x > 0, y > 0$	0.1	0.01	0	0

Table 4.3: Special relativistic hydrodynamic 2-by-2 grid initial conditions.

4.2.5 Relativistic Shock Grid

This test is a further verification that the code can accurately model multidimensional special relativistic hydrodynamics involving shock waves. In more than 1D, effects due to tangential velocities can affect the solution structure. With larger velocities more extreme special relativistic effects are to be expected. This test involves a Lorentz factor greater than two.

This hydrodynamics test simulates the evolution of a 2-by-2 grid of uniform physical states separated by discontinuities. The initial conditions are described in table 4.3, similar to those in [96]. Here $\Gamma = 5/3$ is used, with a spatial domain $(x, y) \in [-0.5, 0.5]^2$ discretised into 100^2 zones, and evolved up to a time of 0.4.

The log-density, pressure, and horizontal and vertical velocity surfaces are shown in figure 4.6. The numerical results are qualitatively consistent with those shown in [96].

4.3 Radiation Tests

This set of tests verify the code's capability to model the evolution of radiation in a variety of optical thicknesses, and its coupling to a fluid via radiation absorption and emission.

4.3.1 Exchange Test

This test is used to validate the code's accurate exchange of energy between fluid and radiation fields via radiation emission and absorption. It evolves in time only, with no spatial dependencies. It is similar to the coupling test setup found in [82].

The exchange of energy is governed by the equations

$$\partial_t \tau = \chi_a (E - \Theta), \tag{4.8}$$

$$\partial_t E = \chi_a (\Theta - E). \tag{4.9}$$

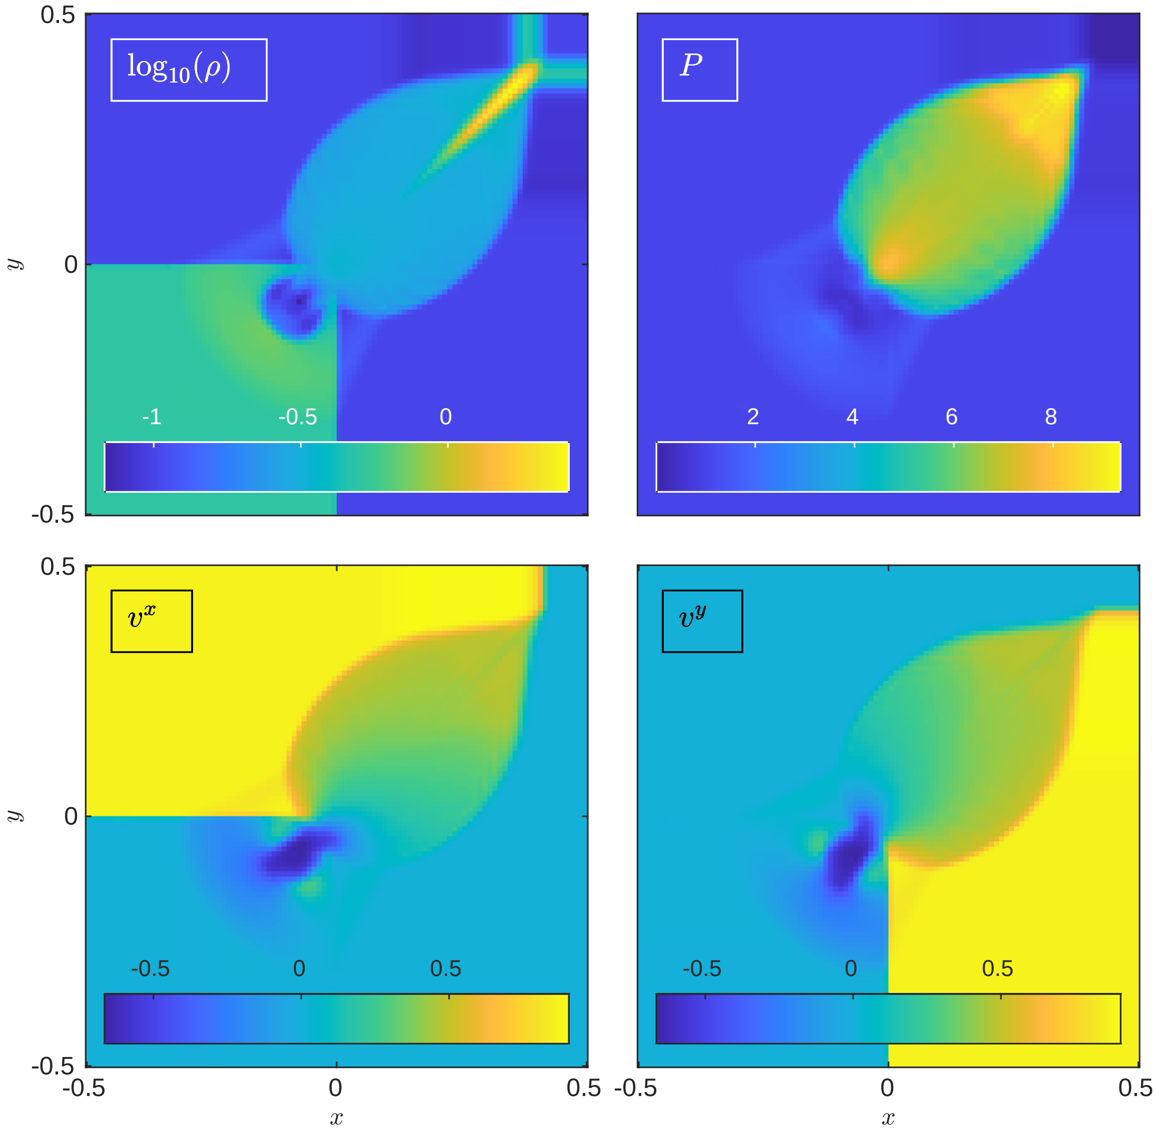

Figure 4.6: The 2-by-2 hydrodynamic shock grid evolved to time 0.4. Top left shows $\log_{10}(\rho)$, top right P, bottom left v^x, and bottom right v^y. This shows more complex multi-dimensional wave interactions. The results are consistent with [96].

Here τ and E are the fluid and radiation energy densities, χ_a is the absorption opacity, and Θ is the radiation emission term.

In order to find an analytic solution against which to validate the code, the equations have been linearised. A constant absorption opacity is used, and the energy densities and emission are related to the fluid and radiation temperatures (T_M and T_R) via

$$\tau = \rho C_V T_M^4, \tag{4.10}$$

$$E = a_{rad} T_R^4, \tag{4.11}$$

$$\Theta = a_{rad} T_M^4. \tag{4.12}$$

Here ρ is the mass density, C_V is the specific heat capacity, and a_{rad} is the radiation constant.

I expect the fluid and radiation to exchange energy, each approaching the local thermodynamic equilibrium temperature (T_{LTE}). The exchange equations conserve total energy such that

$$\tau_{LTE} + E_{LTE} = \tau(0) + E(0). \tag{4.13}$$

From this the equilibrium temperature can be obtained,

$$T_{LTE} = \left[\frac{\rho C_V T_M^4(0) + a_{rad} T_R^4(0)}{\rho C_V + a_{rad}}\right]^{1/4}. \tag{4.14}$$

The analytic evolution of the fluid and radiation energy densities are

$$\tau(t) = \rho C_V \left[T_{LTE}^4 + \left(T_M(0)^4 - T_{LTE}^4\right) e^{-t(\chi_a + a_{rad}\chi_a/\rho C_V)}\right], \tag{4.15}$$

$$E(t) = a_{rad} \left[T_{LTE}^4 + \left(T_R(0)^4 - T_{LTE}^4\right) e^{-t(\chi_a + a_{rad}\chi_a/\rho C_V)}\right]. \tag{4.16}$$

For this test the parameters are $\rho = C_V = \chi_a = 1$ and $a_{rad} = 10$, with initially distinct temperatures $T_M(0) = 10$ and $T_R(0) = 1$. The setup is evolved to a time of 1. It is repeated with a fine timestep ($\Delta t = 0.01$) without the local thermodynamic equilibrium overshoot prevention method introduced in section 3.10.2 (as the small timestep would not require it), and a coarse timestep ($\Delta t = 0.1$) both with and without the overshoot prevention method.

The fluid and radiation energy densities are shown in the top panel of figure 4.7, their corresponding temperatures are shown in the bottom panel. The fluid quantities are shown in red, and radiation quantities are shown in blue. The black dot-dashed line indicates the total energy, which can be seen to be kept constant, as is expected. The dotted lines indicate the LTE limits. The fine-timestep numerical results (dashed curves) agree very closely with the analytic evolution (solid cuves), and are qualitatively consistent with the behaviours seen in [82]. The 'X' markers show the coarse-timestep numerical results without the overshoot prevention method. We can see that the LTE limit is indeed overshot, resulting in unphysical oscillatory behaviour. This overshoot very nearly caused the fluid energy and temperature to become negative, due to more energy being absorbed out of the fluid within a single update than is available. This would have occurred if the coarse timestep were slightly larger, and if not caught could potentially crash the simulation. The circle markers show the coarse-timestep numerical results *with* the overshoot prevention method. We can see that the LTE overshoot does not occur, however the LTE limit is approached more slowly. We take this to be a worthwhile trade-off, as it prevents the troublesome unphysical oscillatory behaviour and negative energy and temperature issues.

4.3.2 Optically Thick Pulse

The purpose of this test is to validate the spatial evolution of a radiation field in the optically thick regime ($\chi_s \gg 0$). I use a highly scattering, but not absorbing or emitting fluid. I employ a similar optically thick pulse setup as found in [82, 97]. In the diffusion limit the governing 1D radiation equations reduce to

$$\partial_t E = \partial_x \left(D \partial_x E \right). \tag{4.17}$$

This is a diffusion equation where the diffusion coefficient D and the scattering opacity χ_s are related via $D = 1/3\chi_s$. The radiation pressure tensor approaches full isotropy.

An initial smooth 1D pulse is constructed

$$E(x,0) = A + Be^{-x^2/4D}, \tag{4.18}$$

with vanishing initial radiation flux $F(x,0) = 0$. When using a spatially uniform opacity the analytic evolution of this pulse follows

$$E(x,t) = A + Be^{-x^2/4D(1+t)} \frac{1}{\sqrt{1+t}}. \tag{4.19}$$

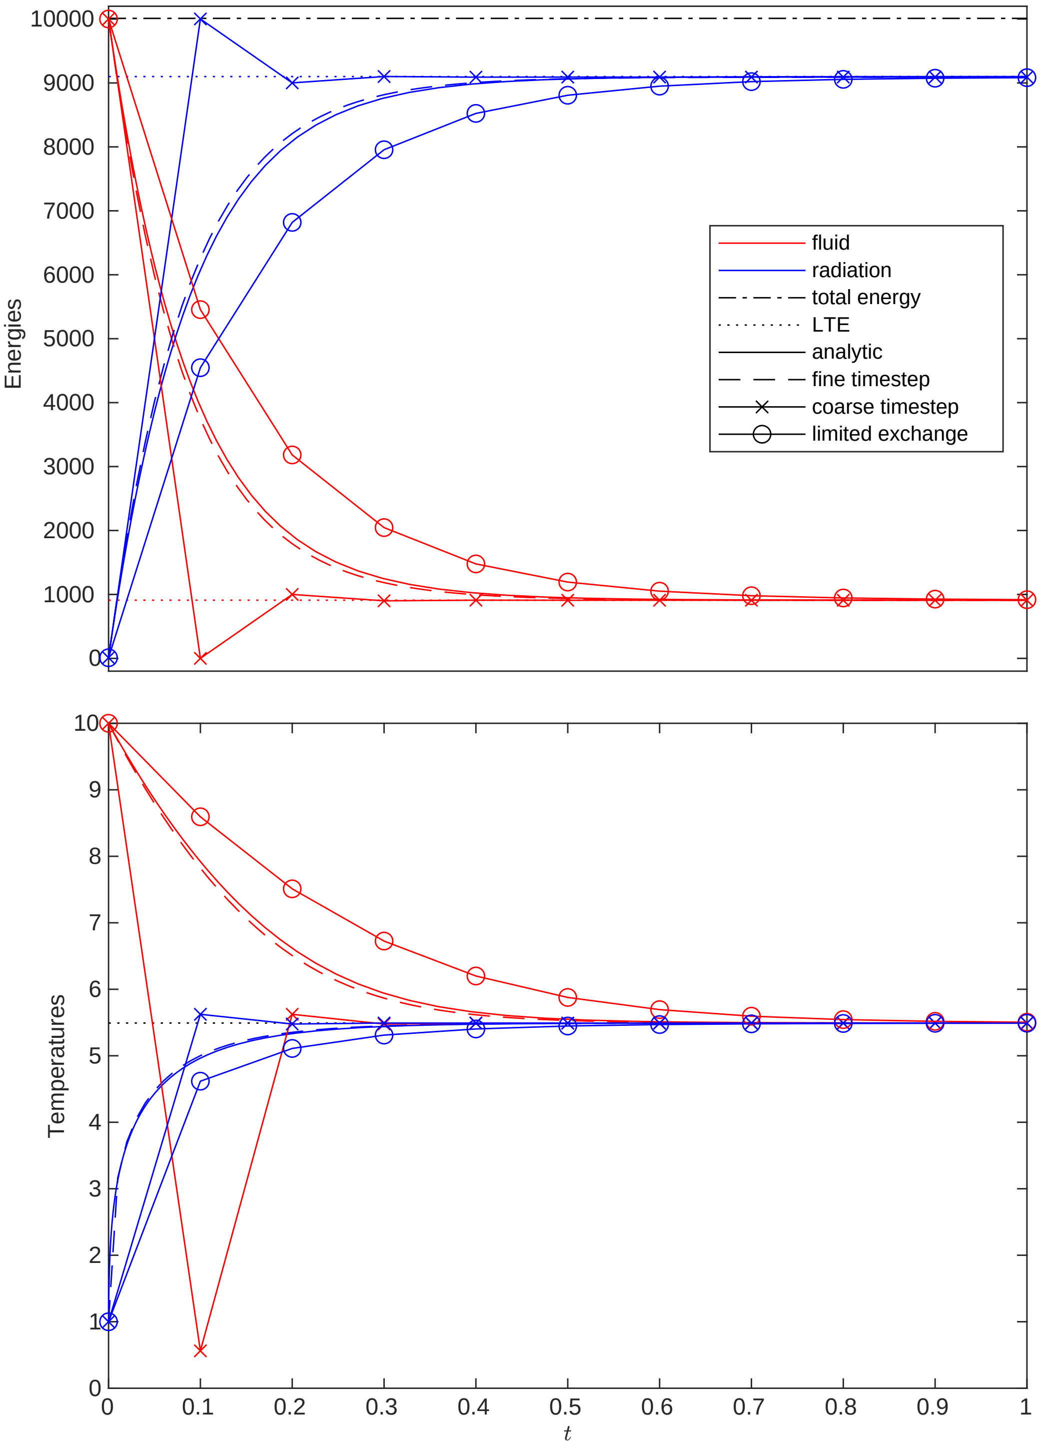

Figure 4.7: Illustration of the exchange of energy between fluid and radiation. The top panel shows the fluid and radiation energy density evolution. The bottom panel shows the fluid and radiation temperature evolution. Fluid quantities are indicated in red, and radiation quantities in blue. The dotted lines show the LTE limits and the black dot-dashed line shows total energy. The solid curves show analytic solutions, and the dashed curves show numerical results calculated with a fine timestep ($\Delta t = 0.01$). The 'X' markers indicate numerical results calculated with coarse timestep ($\Delta t = 0.1$) without the overshoot prevention method, and the circle markers indicate numerical results calculated with a coarse timestep *with* the overshoot prevention method.

For this test the parameters are $A = 1$, $B = 100 - A$, and $D = 1/30$. This uniform diffusion coefficient corresponds to a scattering opacity of $\chi = 10$, and a photon mean free path of 0.1. With a spatial domain two orders of magnitude larger, $x \in [-10, 10]$, this is sufficiently optically thick to lead to diffusive behaviour. The spatial domain has been discretised into 200 zones, and the pulse is allowed to evolve to a time of 100.

I expect to see the pulse of radiation energy density spread out, approaching a linear distribution, whilst conserving total energy within the spatial domain.

The radiation energy density profiles are shown in figure 4.8 for times 0, 10, 30, 60, and 100. The solid lines show the analytic solutions, and markers indicate the numerical results. They are qualitatively consistent with those shown in [97]. We can see that the numerical results agree very closely with the analytic solutions.

4.3.3 Optically Thin Case

The purpose of this test is to validate the spatial evolution of a radiation field in the optically thin regime ($\chi \approx 0$). I use a neither scattering, nor absorbing or emitting fluid. In the free streaming limit the governing 1D radiation equations reduce to

$$\partial_{tt} E = \partial_{xx} E. \tag{4.20}$$

This is a wave equation with speed of light wavespeed. The radiation pressure tensor becomes fully anisotropic.

I use the same initial smooth 1D pulse as in equation (4.18)

$$E(x, 0) = A + Be^{-x^2/4D}, \tag{4.21}$$

this time with maximal initial radiation flux in the positive direction $F^x(x, 0) = E(x, 0)$. This is a one dimensional analogue to the multidimensional optically thin pulse setup in [82, 97], differing in that I am initialising with non-vanishing radiation flux, to emphasise the free streaming wave behavior we wish to see. The analytic evolution of this pulse follows

$$E(x, t) = A + Be^{-(x-t)^2/4D}. \tag{4.22}$$

For this test the parameters are $A = 1$, $B = 100 - A$, $D = 1/30$. With a vanishing opacity the photon mean free path becomes unbounded, leading to free streaming

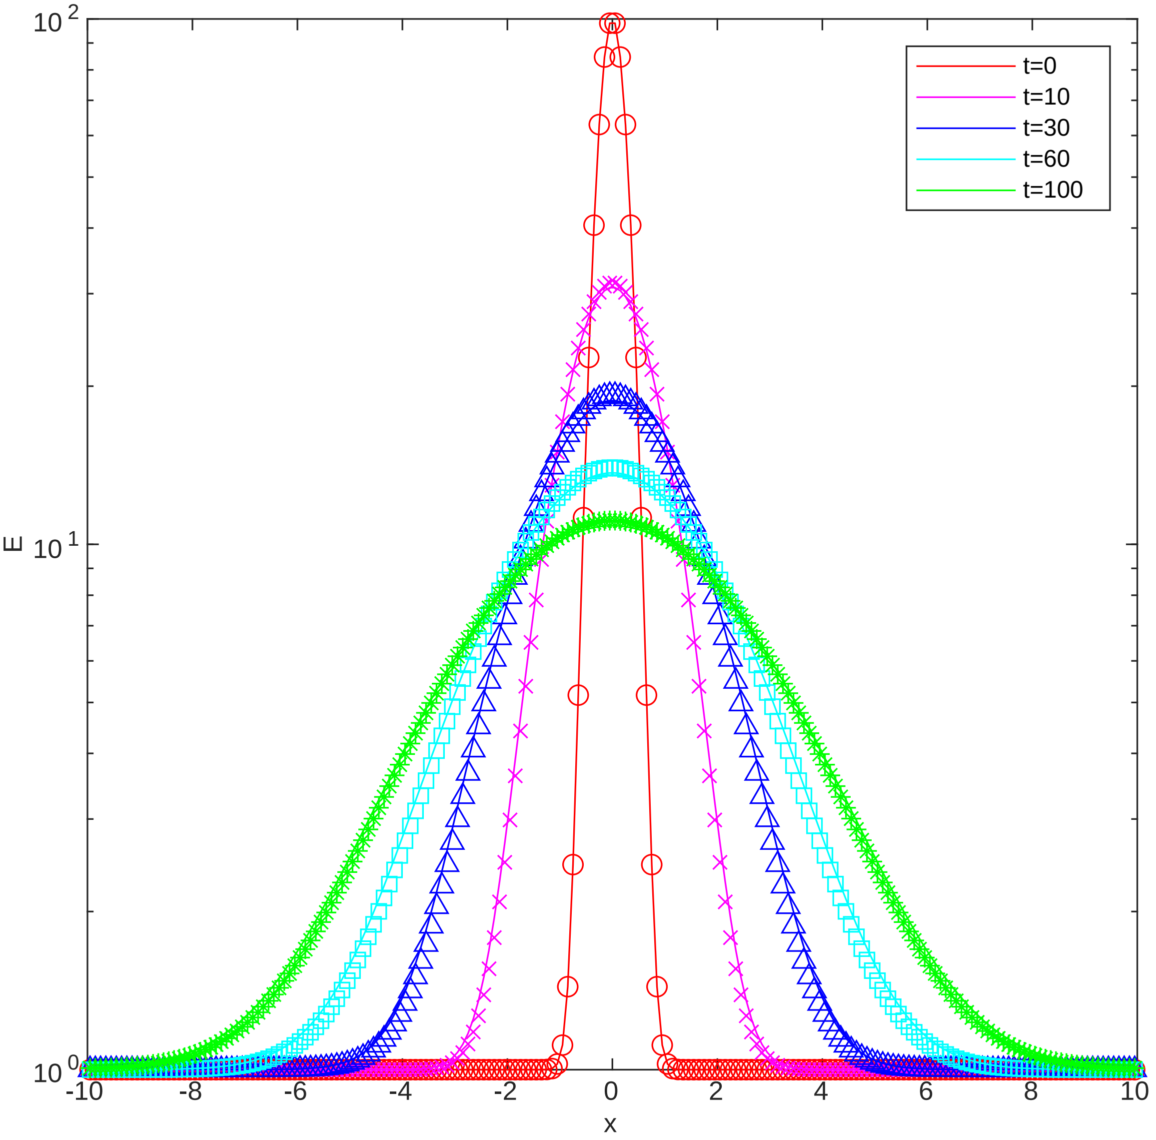

Figure 4.8: Diffusing radiation energy density profiles at times 0, 10, 30, 60, and 100. Analytic solutions are shown with solid lines and numerical results are indicated with markers. We see the initial pulse diffusing out approaching a uniform distribution as expected, consistent with the results found in [97].

behaviour. With a spatial domain of $x \in [-10, 10]$, discretised into 200 zones, the pulse is allowed to evolve to a time of 8.

I expect to see the pulse of radiation energy density travel in the direction of its flux at the speed of light, maintaining its shape and amplitude, thus conserving the total energy within the domain.

The radiation energy density profiles are shown in figure 4.9 for times 0, 2, 4, 6, and 8. The solid lines show the analytic solutions, and markers indicate the numerical results. They are qualitatively consistent with those shown in [97]. We can see that the numerical results agree very closely with the analytic solutions.

4.3.4 Shadow Test

The purpose of this test is to validate that the code can accurately model the multidimensional evolution of a radiation field within a domain containing both optically thick and optically thin regions. This shadow test simulates the 2D evolution of radiation around an opaque region illuminated from below. The setup is similar to those in [82, 97].

Within the domain $(x, y) \in [0, 2] \times [-2, 4]$, I set the density profile to transition (both smoothly and steeply) from a high density circular region to a low density background,

$$\rho(x, y) = \begin{cases} 10^6, & \sqrt{x^2 + y^2} \leq 0.9, \\ 10^{-6} + (10^6 - 10^{-6}) \, \dfrac{1 - \sin(5\pi(r-1))}{2}, & \sqrt{x^2 + y^2} \in (0.9, 1.1), \\ 10^{-6}, & \sqrt{x^2 + y^2} \geq 1.1. \end{cases} \tag{4.23}$$

The fluid is non-scattering, but has an absorption opacity of $\chi_a = \rho$. This results in a dense, high opacity circular region, embedded within a low opacity background. The domain is initially empty of radiation: $E(x, y, 0) = F^x(x, y, 0) = F^y(x, y, 0) = 0$. The radiation emission term is set to zero ($\Theta = 0$), so that radiation is not being emitted anywhere within the domain.

At the lower boundary ($y_L = -2$) the domain is illuminated with a high radiation energy density, with maximal radiation flux in the positive vertical direction, $E(x, y_L, t) = F^y(x, y_L, t) = 10^6$.

I run this simulation twice, once with the fully isotropic Eddington closure, and once with Kershaw closure introduced in equation (2.69), allowing the radiation pressure

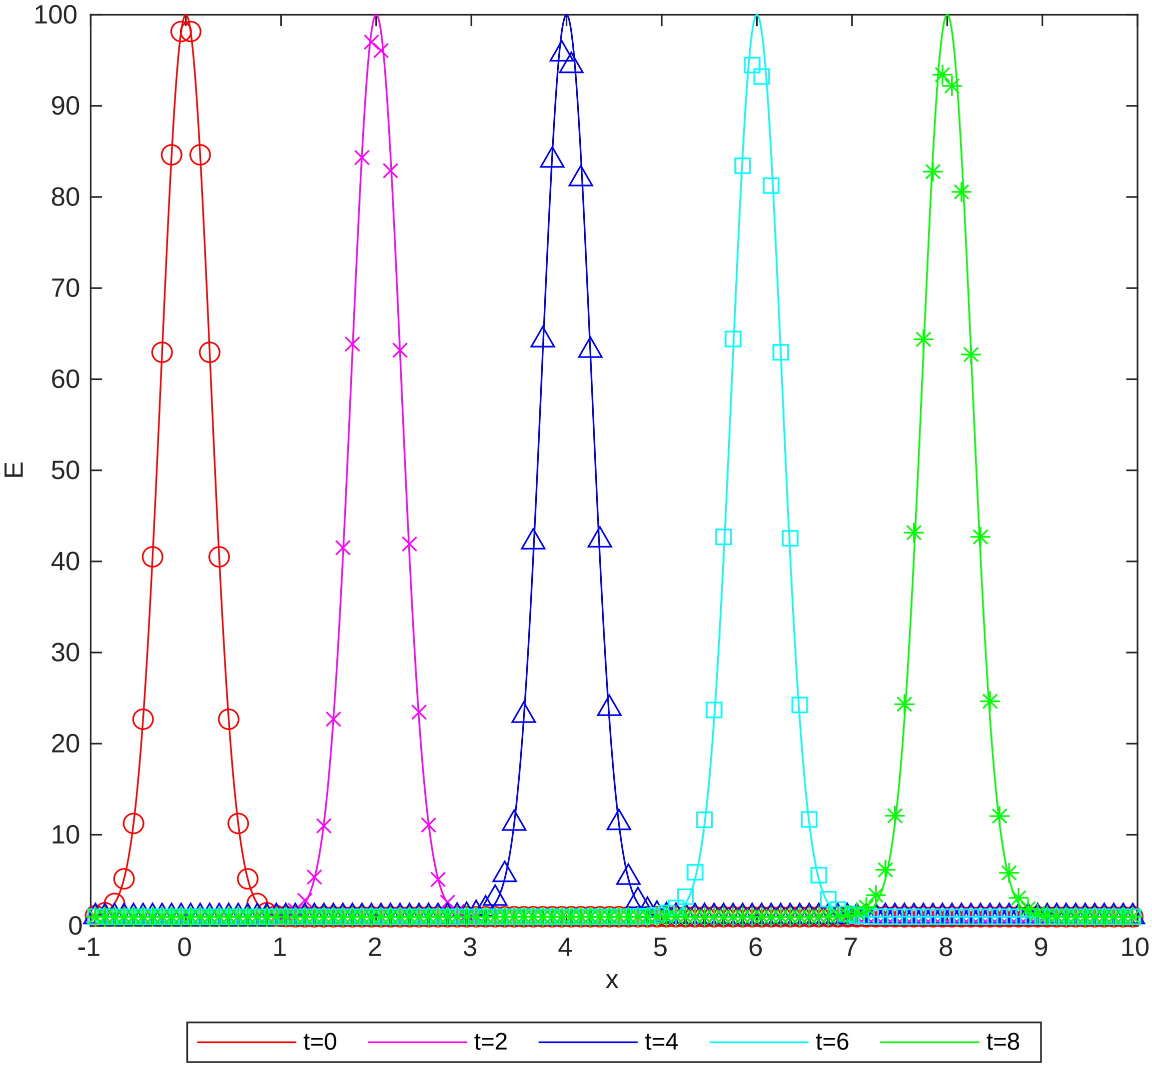

Figure 4.9: Streaming radiation energy density profiles at times 0, 2, 4, 6, and 8. Analytic solutions are shown with solid lines and numerical results are indicated with markers. We see the initial pulse travelling in the direction of its flux as expected, consistent with the results found in [97].

tensor to vary between fully isotropic and fully anisotropic forms.

When using the Eddington closure I expect to see the radiation field exhibit only diffusive behaviour, diffusing around the opaque region, and filling in the area behind it.

Conversely, with the Kershaw closure I expect to see a much more physical solution, with the radiation continuing to stream freely past the opaque region, and the opaque region blocking all incoming radiation, casting a shadow behind it.

The radiation energy density surfaces have been plotted for the early time evolution ($t = 3$) in figure 4.10. Density contours are shown in white, and the radiation flux vectors are indicated with black arrows. The solutions calculated with the Eddington closure are shown in the left panels, and the solutions calculated with the Kershaw closure are shown on the right. We can see the difference in radiation front propagation speeds. For the streaming behaviour the front is propagating at the speed of light ($c = 1$), and for the diffusive motion it is propagating at a speed of around 0.6. These are the expected wavespeeds for streaming and diffusing radiation respectively. The later time evolution ($t = 21$) in shown figure 4.11. At this time the solutions are approximately stationary. We can see that with the Eddington closure the energy density slowly diffuses around the opaque region and fills in the area behind it. However with the Kershaw closure a sharp shadow is indeed cast as radiation streams past the opaque region. These results are qualitatively consistent with those found in [82, 97].

4.4 Special Relativistic Radiation Hydrodynamics Tests

The purpose of these tests is to validate the code's accurate evolution of coupled fluid and radiation fields involving relativistic velocities.

This set of tests evolve 1D piecewise constant initial conditions separated by a discontinuity at $x = 0$. Tests 1, 2, and 3 approach stationarity, and can be compared to analytic solutions. Test 4 is an exception and does not have stationary solution, however can still be qualitatively validated. The fluid and radiation are initialised in local thermodynamic equilibrium with $E = \Theta = 4\sigma_B(P/\rho)^4$. For consistency with previous work by other authors the Eddington closure is used for these tests. The initial conditions and parameters for the four tests are given in table 4.4. These tests

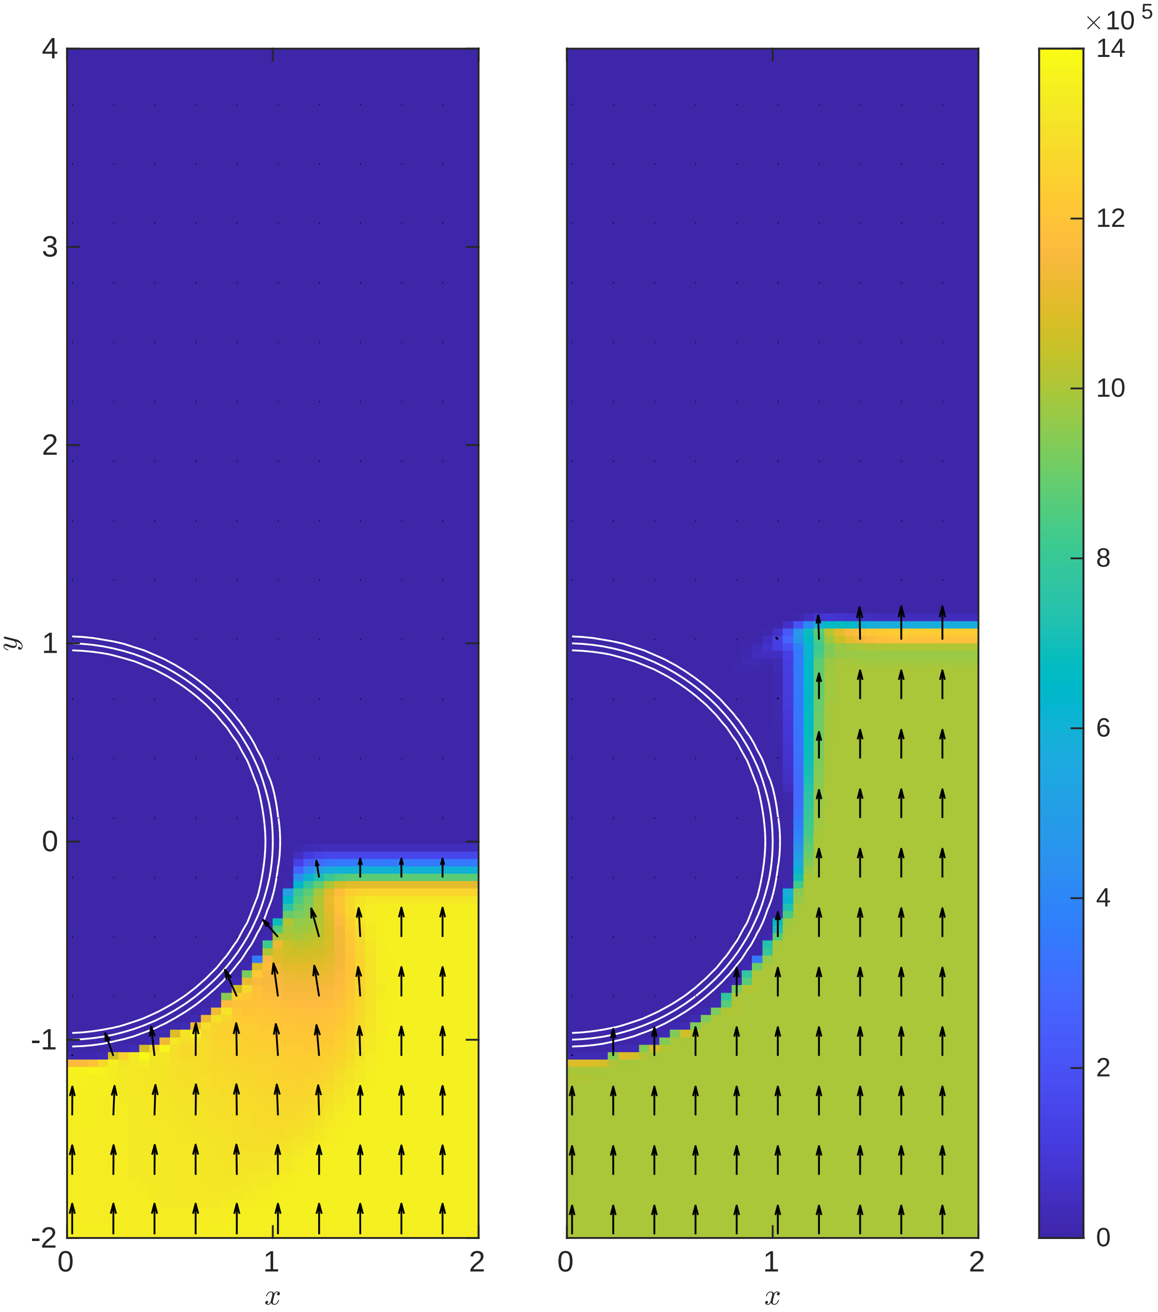

Figure 4.10: Shadow test Radiation energy density surfaces at time 3. Showing density contours in white, and with radiation flux vectors indicated by black arrows. The left panel shows a solution calculated using the Eddington closure, allowing for diffusive radiation behaviour only. The right panel shows a solution calculated using the Kershaw closure, allowing for both diffusive and streaming radiation behaviour. There is a clear difference in radiation front propagation speeds due to the closure schemes used.

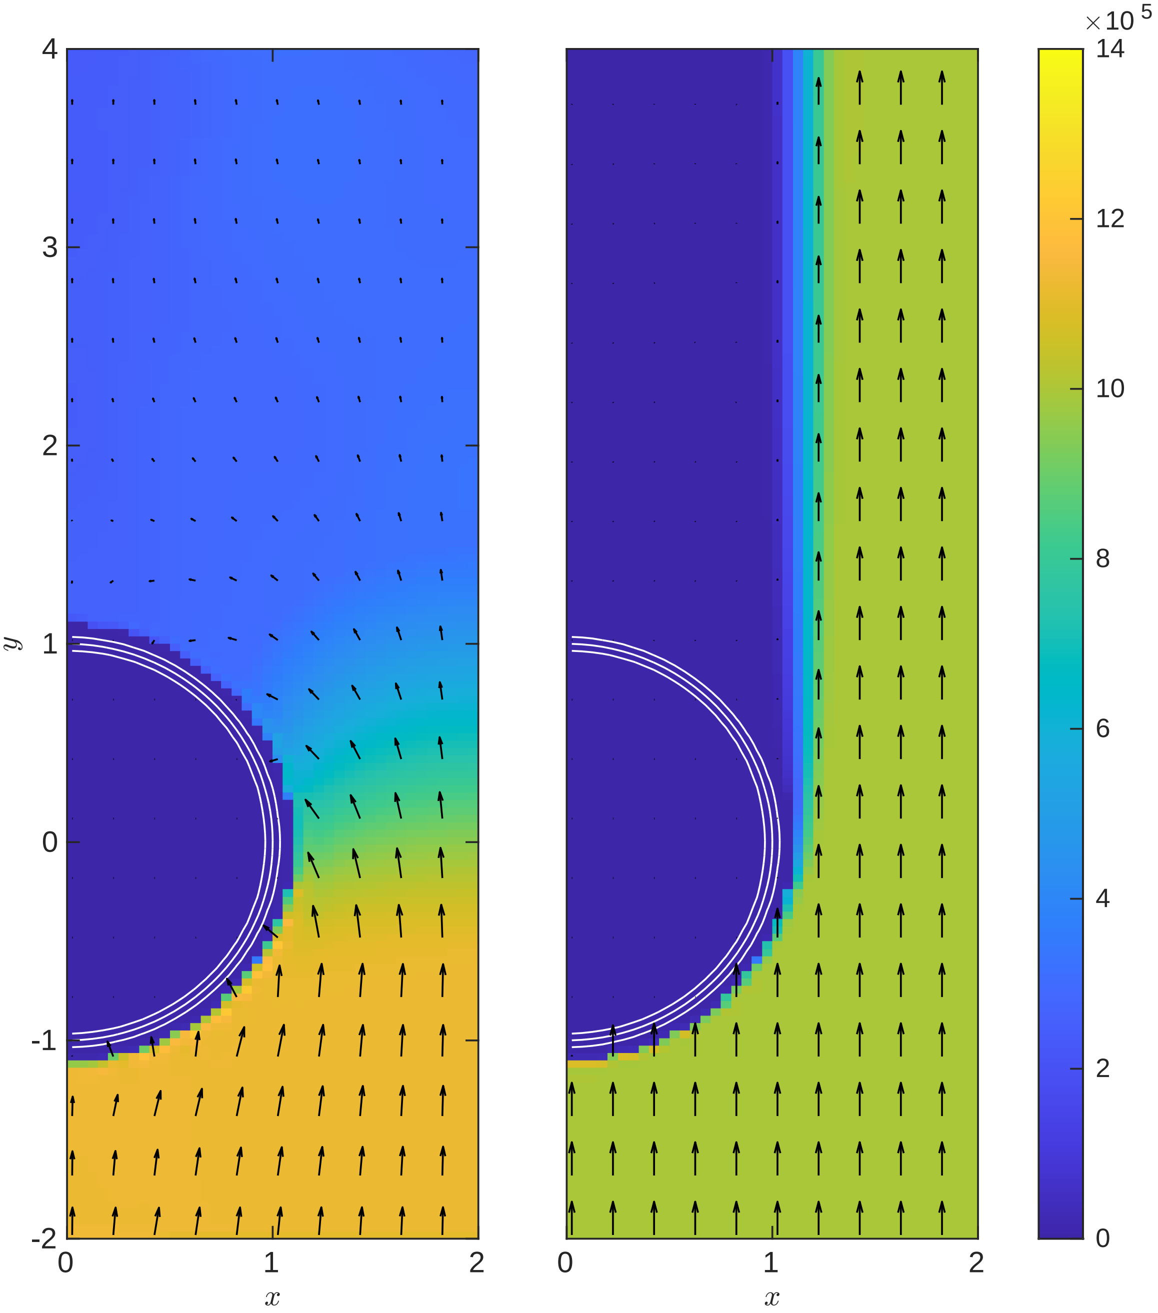

Figure 4.11: Shadow test radiation energy density surfaces at time 21. Showing density contours in white, and with radiation flux vectors indicated by black arrows. The left panel shows a solution calculated using the Eddington closure, allowing for diffusive radiation behaviour only. The right panel shows a solution calculated using the Kershaw closure, allowing for both diffusive and streaming radiation behaviour. The solutions have reached an approximately stationary state by this time.

Test	Left State	Right State	Γ	σ_B	χ_a/ρ	t_{end}
1	$\rho_L = 1$ $P_L = 3 \times 10^{-5}$ $u^x_L = 1.5 \times 10^{-2}$ $E_L = 10^{-8}$	$\rho_R = 2.4$ $P_R = 1.61 \times 10^{-4}$ $u^x_R = 6.25 \times 10^{-3}$ $E_R = 2.51 \times 10^{-7}$	5/3	3.085×10^9	0.4	10^4
2	$\rho_L = 1$ $P_L = 4 \times 10^{-3}$ $u^x_L = 0.25$ $E_L = 2 \times 10^{-5}$	$\rho_R = 3.11$ $P_R = 4.512 \times 10^{-2}$ $u^x_R = 8.04 \times 10^{-2}$ $E_R = 3.46 \times 10^{-3}$	5/3	1.953×10^4	0.2	500
3	$\rho_L = 1$ $P_L = 60$ $u^x_L = 10$ $E_L = 2$	$\rho_R = 8$ $P_R = 2.34 \times 10^3$ $u^x_R = 1.25$ $E_R = 1.14 \times 10^3$	2	3.858×10^{-8}	0.3	20
4	$\rho_L = 1$ $P_L = 60$ $u^x_L = 1.25$ $E_L = 2$	$\rho_R = 1$ $P_R = 60$ $u^x_R = 1.1$ $E_R = 2$	2	3.858×10^{-8}	10^3	15

Table 4.4: Special relativistic radiation hydrodynamics tube tests initial conditions. Here the fluid 3-velocity can be obtained with $v^x = u^x/\sqrt{1+(u^x)^2}$.

can be found amongst [18, 39, 97, 98]. The fluids involved are non-scattering, but are capable of both emitting and absorbing radiation.

I have used the spatial domain $x \in [-20, 20]$, discretised with 400 uniform zones.

Test 1 simulates the evolution of a weakly relativistic fluid pressure dominated shock, with a maximum velocity of around 0.015. Results are shown in figure 4.12 and are qualitatively consistent with those shown in [18, 39, 97].

Test 2 simulates the evolution of a mildly relativistic fluid pressure dominated shock, with a maximum velocity of around 0.243. Results are shown in figure 4.13 and are qualitatively consistent with those shown in [18, 39, 97], including the spike in radiation energy density near the shock front.

Test 3 simulates the evolution of a highly relativistic fluid pressure dominated wave, with a maximum velocity of around 0.995. The solution has reached a stationary state in the region $x \in [-5, 10]$, with small waves continuing to slowly propagate away from the initial discontinuity. Results are shown in figure 4.14 and are qualitatively consistent with those shown in [18, 39, 97].

Test 4 simulates the evolution of a strongly coupled optically thick flow with mildly relativistic fluid velocities. The maximum velocity is around 0.781. It results in a pair of separating shockwaves which are continuing to approach the edges of the domain. This test does not have an asymptotic stationary solution. Results are shown in figure 4.15 and are qualitatively consistent with those shown in [97, 98].

4.5 General Relativistic Test

The purpose of this test is to verify that the code can accurately evolve a radiation field within a curved spacetime. This test simulates a beam of light curving near a black hole.

Setting up an exterior Schwarzschild spacetime, the spacetime interval is given by

$$ds^2 = -\alpha^2 dt^2 + \gamma_{rr} dr^2 + r^2 d\Omega^2, \tag{4.24}$$

$$d\Omega^2 = d\theta^2 + \sin^2\theta d\phi^2. \tag{4.25}$$

Here the lapse function and radial spatial metric are

$$\alpha = \sqrt{1 - \frac{2M}{r}}, \tag{4.26}$$

$$\gamma_{rr} = \left(1 - \frac{2M}{r}\right)^{-1}. \tag{4.27}$$

In order to use this spacetime on a Cartesian grid, I transform this spacetime to isotropic coordinates. Introducing the isotropic radius ($\overline{r}$), we desire a spacetime interval of the form

$$ds^2 = -\alpha^2 dt^2 + \gamma_{\overline{r}\overline{r}} \left(d\overline{r}^2 + \overline{r}^2 d\Omega^2\right), \tag{4.28}$$

$$\gamma_{\overline{r}\overline{r}} = \gamma_{xx} = \gamma_{yy} = \gamma_{zz}, \tag{4.29}$$

$$d\overline{r}^2 + \overline{r}^2 d\Omega^2 = dx^2 + dy^2 + dz^2, \tag{4.30}$$

$$\overline{r} = \sqrt{x^2 + y^2 + z^2}. \tag{4.31}$$

We can transform between Schwarzschild and isotropic radii with the relations

$$\overline{r} = \frac{\sqrt{r^2 - 2Mr} + r - M}{2}, \tag{4.32}$$

$$r = \overline{r}\left(1 + \frac{M}{2\overline{r}}\right)^2. \tag{4.33}$$

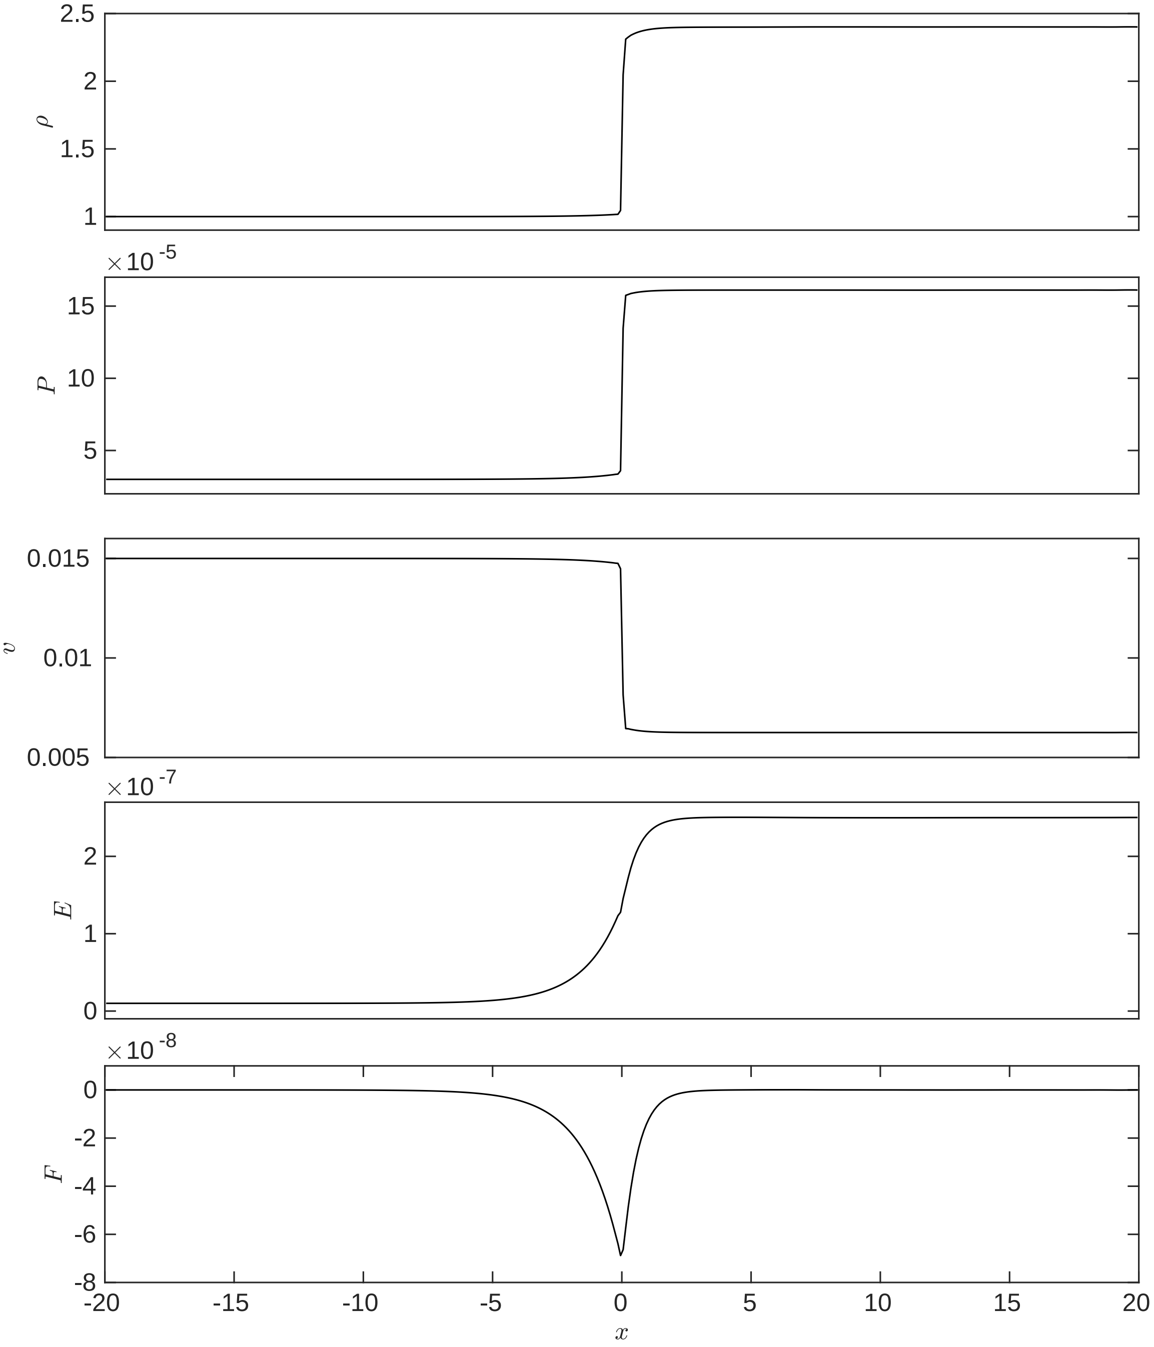

Figure 4.12: Test 1: a weakly relativistic fluid pressure dominated shock. The plots show rest-mass density, fluid pressure, fluid 3-velocity, radiation energy density, and radiation flux at time 10^4. The discontinuities have evolved correctly and the stationary solution matches that in [18, 39, 97].

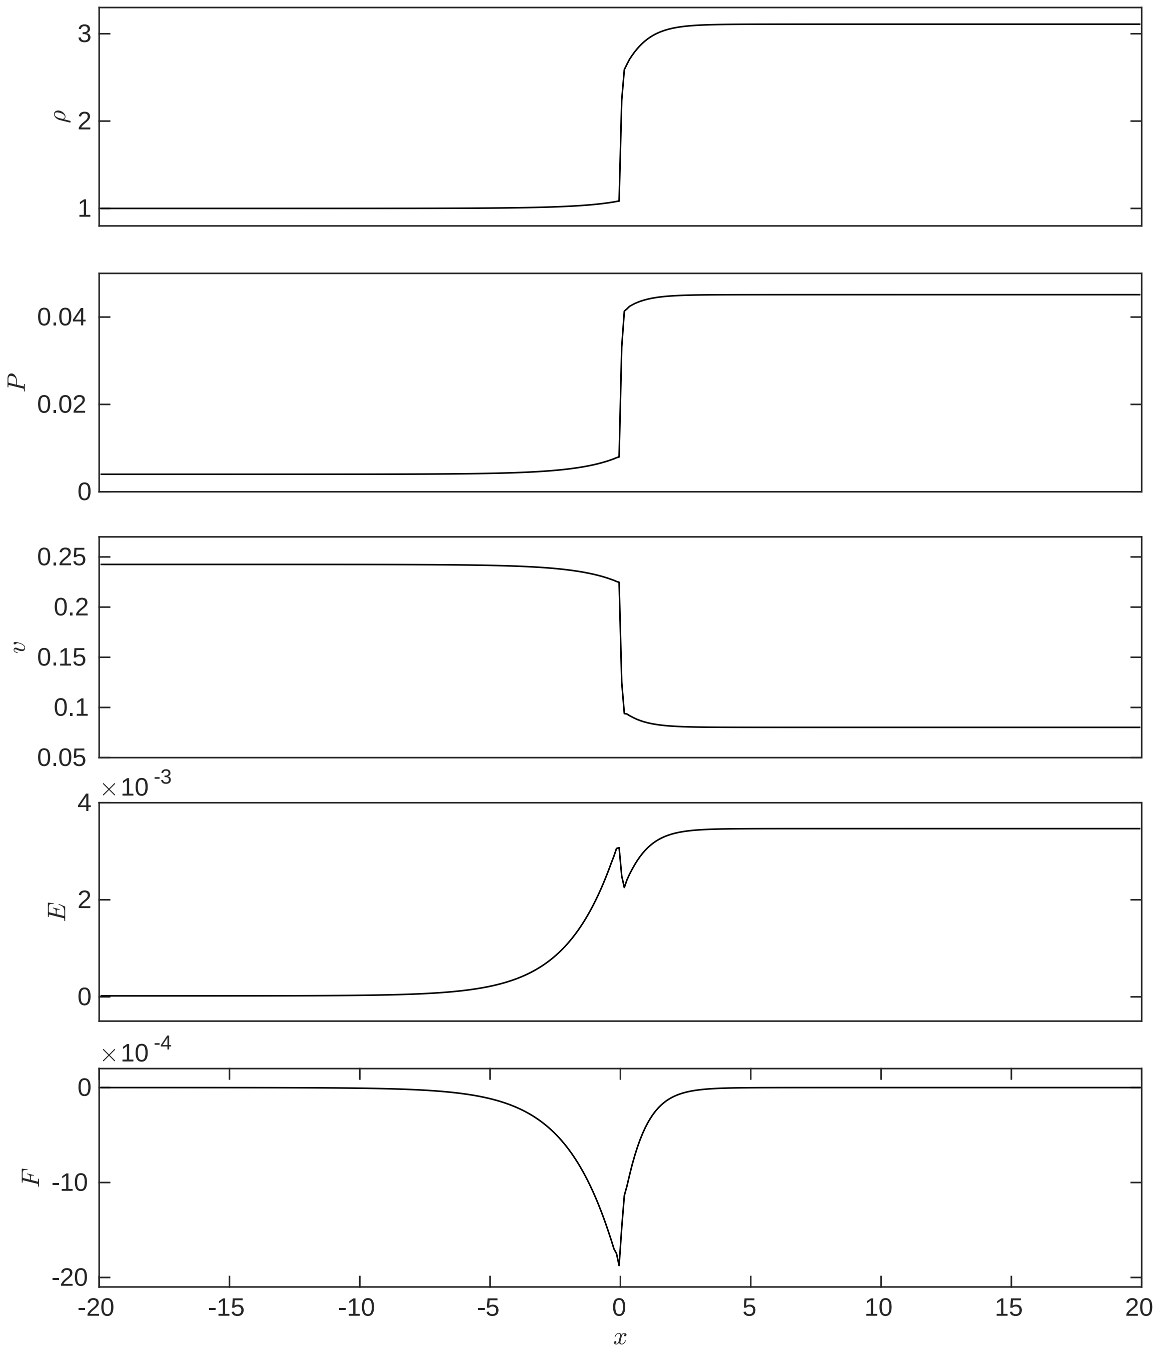

Figure 4.13: Test 2: a mildly relativistic fluid pressure dominated shock. The plots show rest-mass density, fluid pressure, fluid 3-velocity, radiation energy density, and radiation flux at time 500. The discontinuities have evolved correctly and the stationary solution matches that in [18, 39, 97].

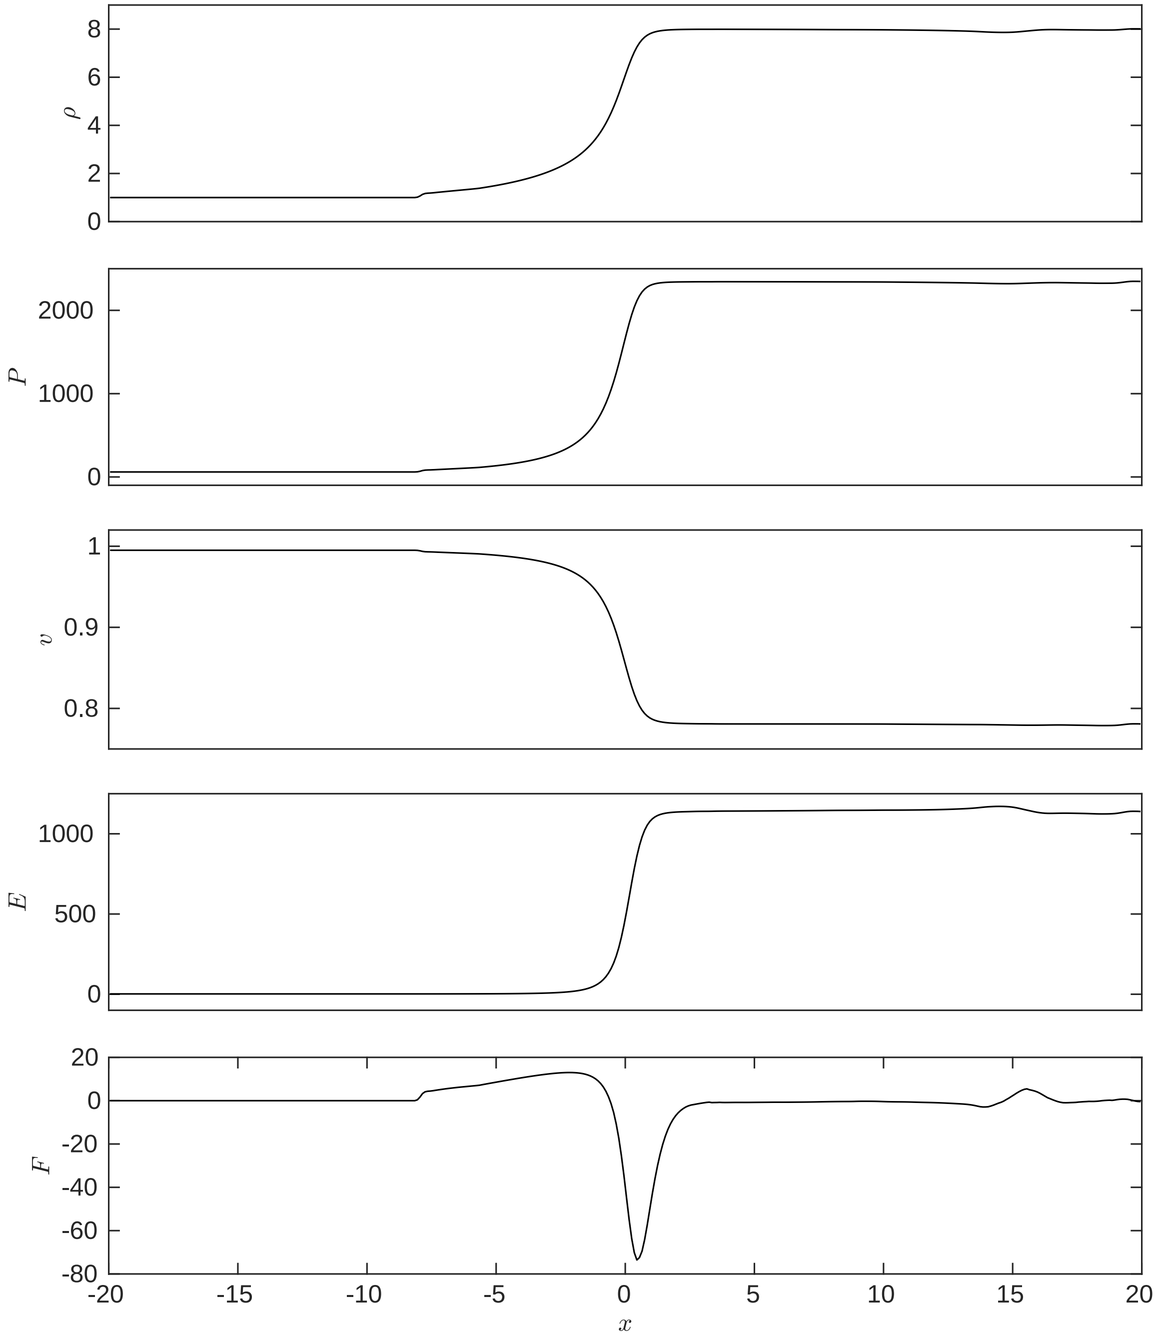

Figure 4.14: Test 3: a highly relativistic fluid pressure dominated wave. The plots show rest-mass density, fluid pressure, fluid 3-velocity, radiation energy density, and radiation flux at time 20. The discontinuities have evolved correctly and the solution in the stationary region matches that in [18, 39, 97].

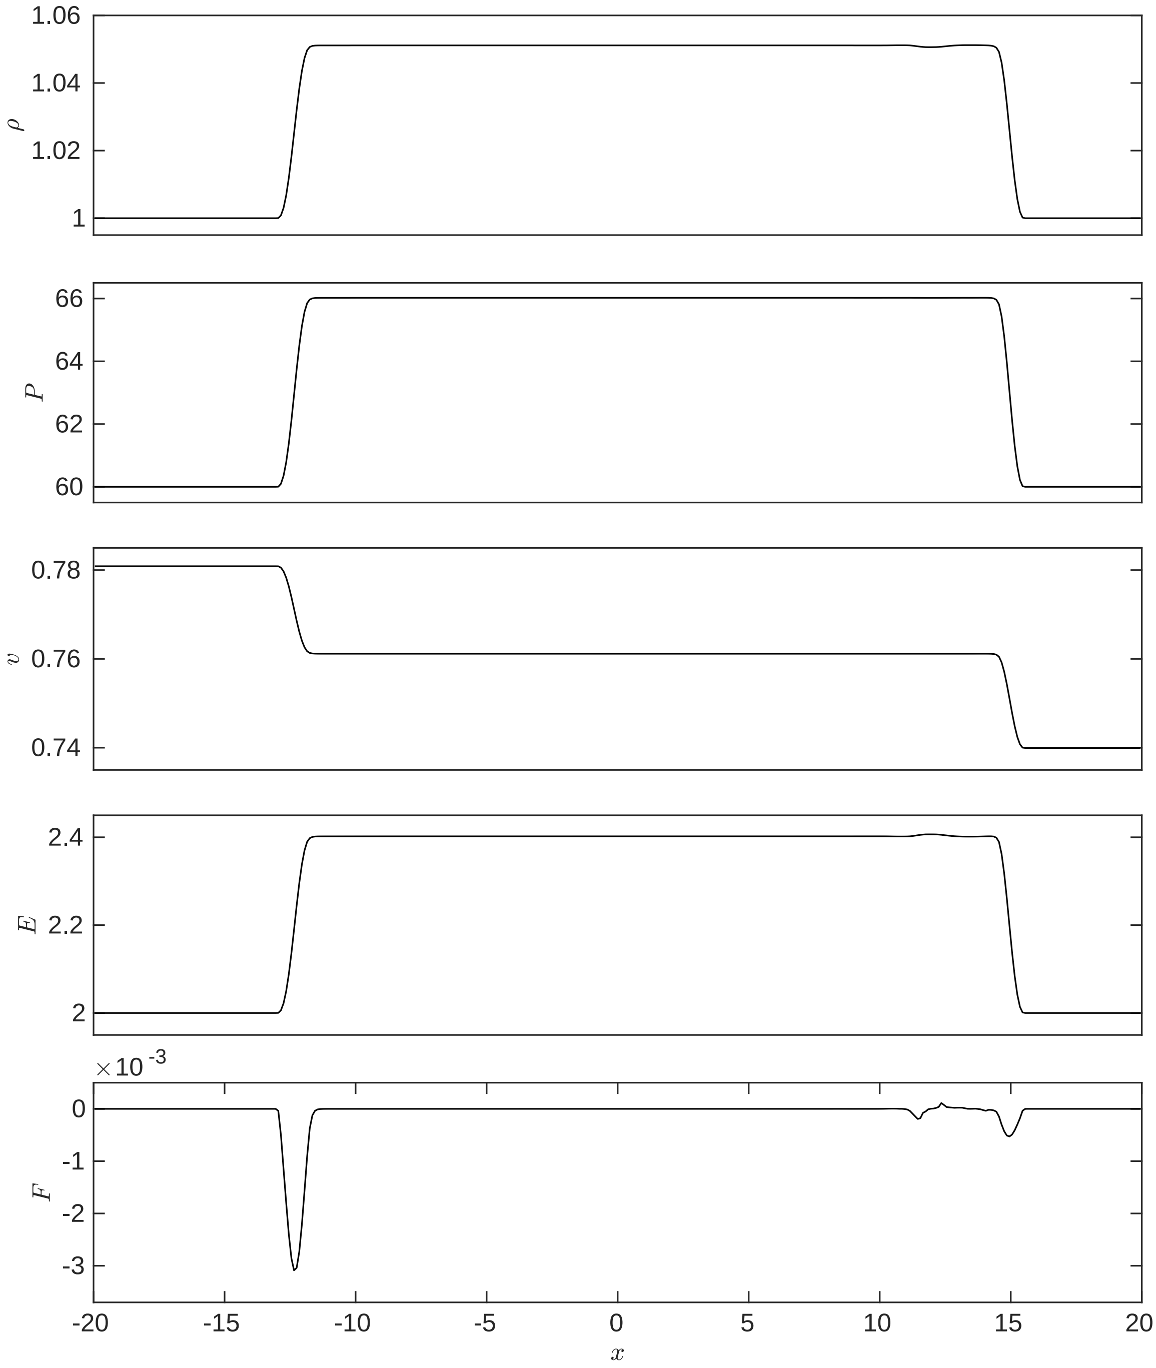

Figure 4.15: Test 4: a strongly coupled optically thick flow with mildly relativistic fluid velocities. The plots show rest-mass density, fluid pressure, fluid 3-velocity, radiation energy density, and radiation flux at time 15. The discontinuities have evolved correctly and the solution matches those in [97, 98].

The lapse function and the isotropic spatial metric in terms of the isotropic radius are given by

$$\alpha = \frac{M - 2\overline{r}}{M + 2\overline{r}}, \tag{4.34}$$

$$\gamma_{\overline{r}\overline{r}} = \left(1 + \frac{M}{2\overline{r}}\right)^4 . \tag{4.35}$$

With a solar mass black hole, $M = M_\odot = 1$. I use a 2D domain $(x, y) \in [0, 2.5]^2$ and $z = 0$, discretised with 100^2 zones. The spacetime is capped to $r \geq 2.5$ (that is $\overline{r} \gtrsim 1.309$), to avoid event horizon (at $r = 2$ and $\overline{r} = 0.5$) and black hole interior considerations. The domain is initially empty of radiation, $E(x, y, 0) = F^x(x, y, 0) = F^y(x, y, 0) = 0$. At the lower boundary ($y_L = 0$), in the interval $r \in [3, 3.1]$ (corresponding to the isotropic interval $x_{in} \in [1.866, 1.973]$), high radiation energy density is injected into the domain, with maximal radiation flux in the positive vertical direction $E(x_{in}, y_L, t) = F^y(x_{in}, y_L, t) = 10^6$.

With zero opacity ($\chi = 0$) no radiation scattering, or emission or absorption occurs. The radiation does not couple to any fluid allowing the radiation to stream freely with a highly anisotropic pressure tensor. This setup is similar to those from [82, 97].

Two problems have been simulated, one without any compact object centred at the origin resulting in a flat spacetime, allowing the light beam to travel along a straight line through the mesh. The other problem contains a solar mass black hole centred at the origin, resulting in a Schwarzschild spacetime which is expected to curve the radiation beam's path into a circular orbit.

The radiation energy density surfaces are shown for the flat spacetime in figure 4.16 and for the Schwarzschild spacetime in figure 4.17. Lapse function contours are shown in white, and the radiation flux vectors are indicated with black arrows. The numerical evolution of the radiation energy density field approximately follows that of a geodesic light path calculation and is qualitatively consistent with the results shown in [82, 97]. Due to numerical dissipation and diverging geodesics in the Schwarzschild spacetime case (as the radiation beam has a finite width), the radiation beam spreads out as it travels and decreases in intensity.

The expected path of the radiation beam has been refined with two levels of factor-2 refinement, this was set up at initialisation and not created with the dynamic tagging capability. The resulting mesh is shown in figure 4.18. As the path of the curved beam does not align with the Cartesian mesh, grid effects can be seen in the solution.

It is important to note that the refined grids are made up of many patches. This contrasts with other mesh refinement algorithms used in the field, such as the Carpet code used in the Einstein Toolkit. Carpet uses a box-in-box style approach where a few large grids as nested inside each other. This approach is very efficient for the bulk of a binary inspiral but does not effectively capture asymmetric features of the sort that appear in the remnant and the outflows. The ability of AMReX to efficiently work with these grids will be important in the problems tackled later.

4.6 Summary

The results from these test problems show that the numerical simulation performs as expected for both smooth and discontinuous fluid motion at relativistic velocities. The evolution of a radiation field transitions between diffusive behaviour and streaming behaviour appropriately. The radiation hydrodynamic coupling behaves as intended. The benefits of using adaptive mesh refinement have been demonstrated. The code can perform calculations with a curved spacetime. The results from the tests in this chapter show that the code is performing well enough to simulate problems of interest.

I have yet to demonstrate situations involving a more complex spacetime and more physical equations of state. These are both important for evolving neutron star merger remnants and will be tested in the next chapter.

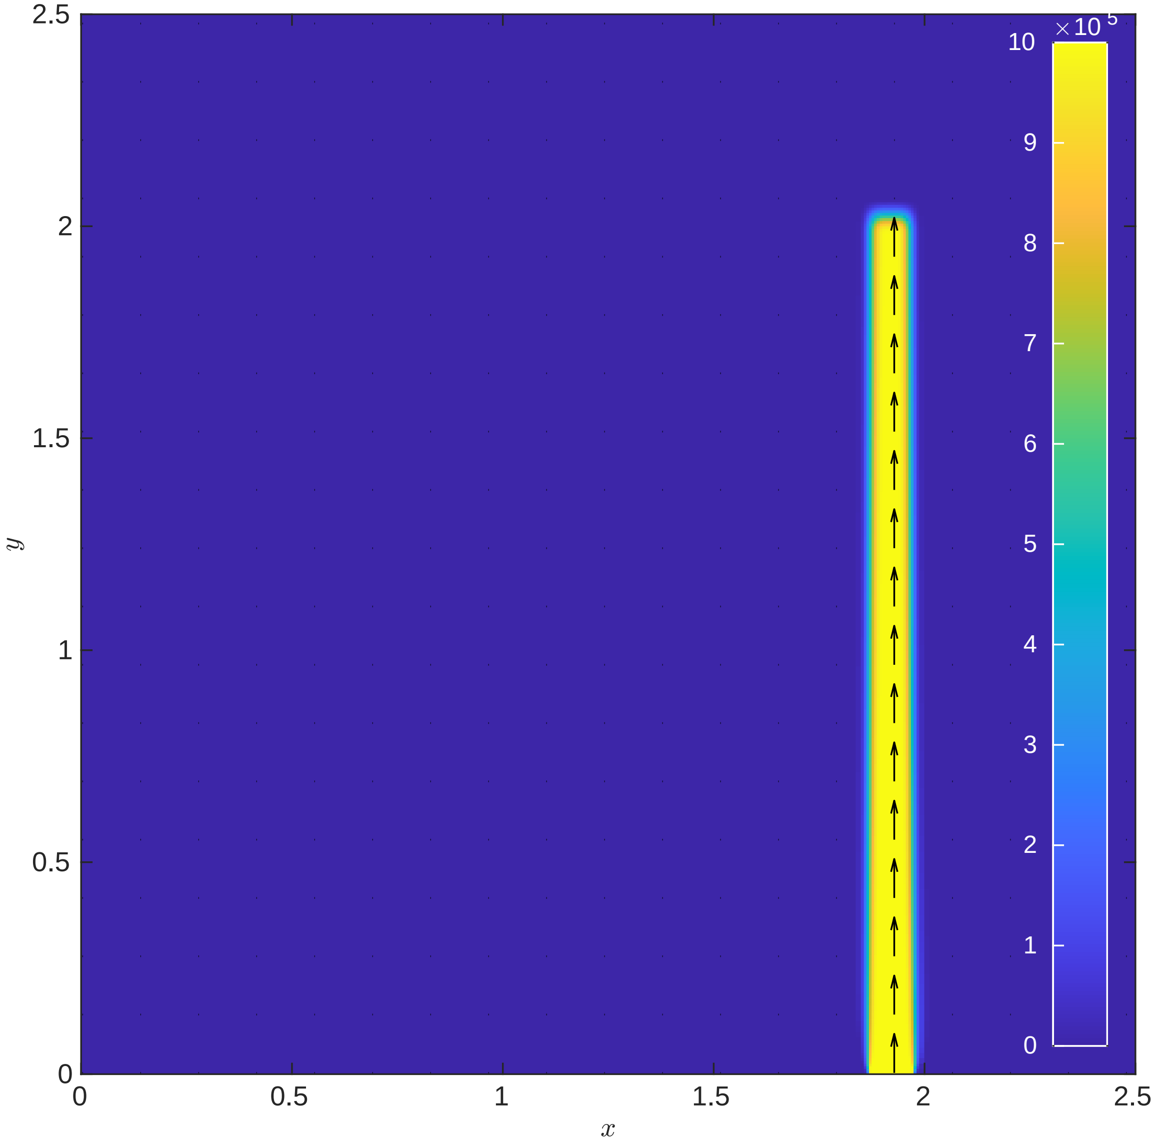

Figure 4.16: Evolution of a narrow beam of anisotropic radiation in a flat spacetime. We can see that the radiation travels along a straight path through the mesh.

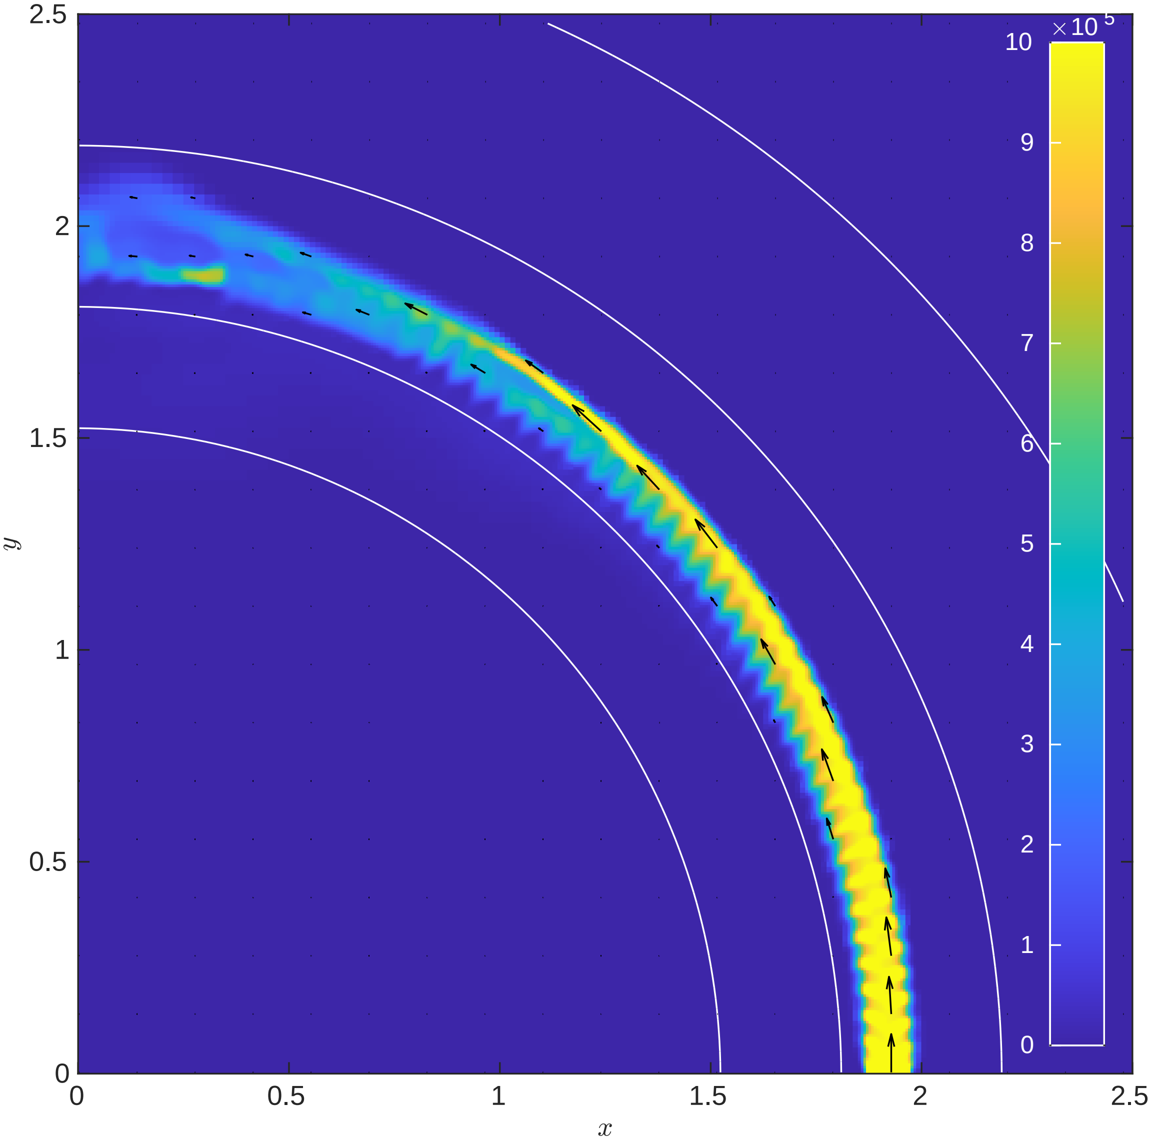

Figure 4.17: Evolution of a narrow beam of anisotropic radiation in a Schwarzschild spacetime. We can see that the path of the radiation is curved around the compact object, following along a circular geodesic curve at $\overline{r} \approx 1.866$. The beam spreads out and reduces in intensity at it travels due to numerical dissipation and from following diverging geodesics. These results are qualitatively consistent with those seen in [82, 97].

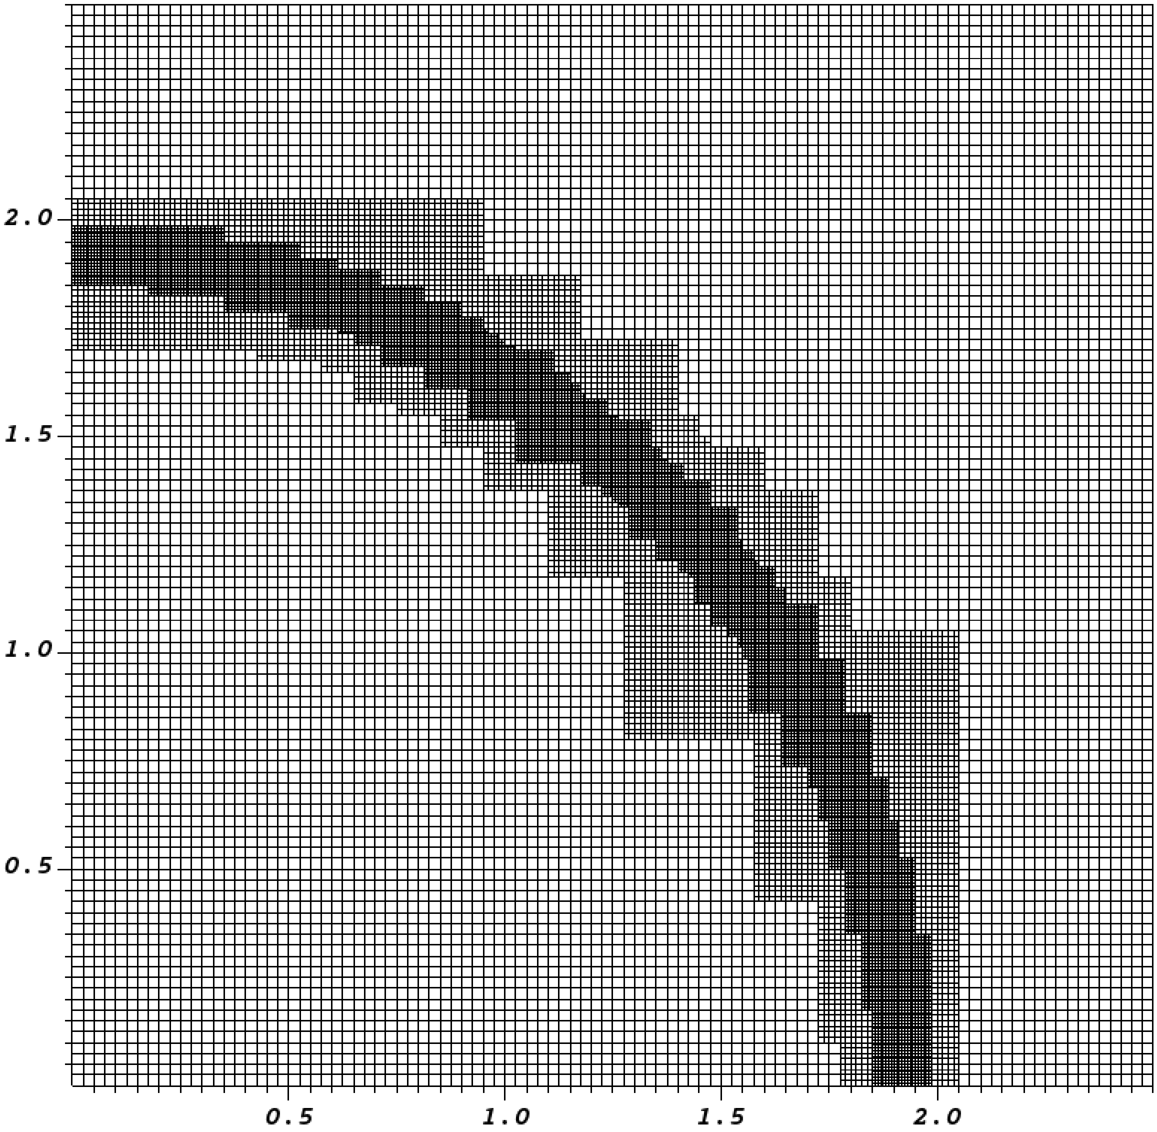

Figure 4.18: The expected path of this beam of radiation in a Schwarzschild spacetime. The ring $\bar{r} \in [1.866, 1.973]$, is covered with two levels of factor-2 mesh refinement.

Chapter 5

Neutron Star Models

5.1 Introduction

I want to validate that the code (as detailed in section 3) can accurately evolve a setup involving a physical fluid within a curved spacetime. I can obtain a model curved spacetime, induced by a horizonless compact object like a neutron star or neutron star merger remnant, with the Tolman-Oppenheimer-Volkoff (TOV) equation (which will be described in equation (5.1)). In this chapter I design setups to mimic a neutron star merger remnant, and ensure that the code is capable of evolving such setups as expected.

I make use of a physical three parameter equation of state (EoS). This is one which takes into account electrons and relativistic effects (such as relativistic mean field interactions). When setting up the initial data, in order to reduce the number of degrees of freedom available from this EoS (as with the initialisation of the Einstein Toolkit merger simulation which will be discussed in chapter 6) a uniform fluid temperature is chosen (the isothermal approximation), and β-equilibrium is enforced. This uses the assumption that the β-reactions via the Urca processes have had sufficient time to reach nuclear statistical equilibrium (NSE).

5.2 Physical Equation of State

To validate that the code can use a more complex and physical equation of state (and for compatibility with the Einstein Toolkit merger simulation discussed in chapter 6) I want to use an equation of state which takes into account temperature effects (such as from shock heating) so will be dependent on the fluid temperature (T), and out of β-equilibrium (nuclear statistical equilibrium) effects so will be dependent on the electron per baryon fraction (Y_e).

The β-reactions are assumed to occur sufficiently slowly that they do not have time to affect the electron fraction during the merger time (and hence they are not included in the merger evolution). In cold neutron star matter the β-reactions primarily involve the direct and modified Urca processes [99]. These account for neutron decay and electron capture reactions [20, 100], which affect the electron fraction of the fluid, and thus can affect its pressure.

A library of tabulated equations of state are available from the online service CompOSE. I have chosen to use the SFHx (with electrons) EoS [40] as I have access to a merger simulation which uses this EoS, and it covers the rest mass density, fluid temperature, and electron fraction ranges expected in remnant evolution. It returns consistent mass and radius ranges with observed data when used to calculate TOV solutions for spherically symmetric neutron star profiles.

This equation of state provides the fluid pressure (P), specific internal energy (ϵ), and sound speed (c_s) dependent on rest mass density (ρ), fluid temperature (T), and electron per baryon fraction (Y_e). The available quantities are interpolated from tabulated data points.

5.2.1 Nuclear Statistical Equilibrium

I aim to construct a neutron star model such that all of the fluid is in β-equilibrium. This is referred to as nuclear statistical equilibrium (NSE). This model would be an approximation to a neutron star which has had sufficient time for its β-reactions to reach equilibrium.

Following the discussions of β-equilibrium for neutron star initialisation found in chapters 8 and 9 of [26], there are numerous considerations (including temperature, neutrinos, and strong interactions) regarding when β-equilibrium is attained. For simplicity using the three parameter equation of state eg. $P = P(\rho, T, Y_e)$, I find the slice though the tabulated data such that the lepton chemical potential (available from the equation of state table) vanishes, i.e. $\mu_e(\rho, T, Y_e) = 0$. This is valid for 'cold' equilibrium ($T \lesssim 1$MeV), and any thermal effects on the equation of state become less pronounced for temperatures reaching $T \gtrsim 10$MeV. That is to say for each pair of tabulated rest mass density and fluid temperature values I interpolate the table to find the electron fraction which satisfies this chemical potential condition. This has the effect of reducing the dimensionality of our equation of state. I use this when finding a TOV solution for neutron star model initialisation, allowing us to use relations of the form $P = P(\rho, T)$ and $Y_e = Y_e(\rho, T)$. When evolving the data through time the full three parameter table is used, allowing for out of β-equilibrium fluid.

5.3 TOV Equations

An important demonstration that the code can perform as expected with a curved spacetime is the ability to maintain the hydrostatic equilibrium of a Tolman-Oppenheimer-Volkoff (TOV) solution. This is an isolated, spherically symmetric, and self-gravitating compressible fluid. A TOV solution can be constructed to approximate an isolated neutron star. A merger remnant is similar to the TOV solution in that they both have a central region of high density above that of atomic nuclei, which drops off with radius, and that the inwards and outwards forces are approximately balanced. They differ in that for a merger remnant the fluid and spacetime are rotating and still (slowly) evolving. Furthermore they do not have the spherical symmetry of the one dimensional TOV solution, and they have a surrounding cloud of ejected material.

The TOV equations (as described in the Einstein Toolkit documentation [101]) can be used to integrate out the central fluid pressure (P) and enclosed gravitational mass (m, including mass-energy contributions). The equations take the form

$$\frac{dP}{dr} = -\left[\rho + \rho\epsilon + P\right] \frac{m + 4\pi r^3 P}{r(r - 2m)}, \tag{5.1}$$

$$\frac{dm}{dr} = 4\pi r^2 \left[\rho + \rho\epsilon\right]. \tag{5.2}$$

Once the fluid pressure and enclosed mass profiles have been calculated, the resulting spacetime can be found with

$$\frac{d\phi}{dr} = \frac{m + 4\pi r^3 P}{r(r - 2m)}. \tag{5.3}$$

Here the gravitational potential ϕ is related to the lapse function α via $\phi = \ln(\alpha)$. The radial component of the spatial metric is given by the relation $\gamma_{rr} = (1 - 2m/r)^{-1}$.

5.4 TOV Solutions

I choose the central mass density (ρ_c) above nuclear density (which is $\approx 2.5 \times 10^{14}$ g/cm^3), and an isothermal fluid temperature (T_c) (common values I have used are $\rho_c = 7.913 \times 10^{14}$ g/cm$^3 = 1.28 \times 10^{-3}$ in $c = G = M_\odot = 1$ units, and $T_c = 10$ MeV). We can acquire the central electron fraction (Y_{ec}) required for β-equilibrium, and we can obtain the central fluid pressure (P_c), the central specific internal energy (ϵ_c), along with the central sound speed (c_{sc}) from the three parameter SFHx (with electrons) equation of state.

With the TOV equations (5.1) we can numerically advance the fluid pressure outwards from the centre with an IVP solver (as introduced in section 3.3). Given the fluid pressure (P_n) at some radius (r_n), I want to find the mass density, fluid temperature, and electron fraction at this same radius (ρ_n, T_n, and Y_{en}). If we assume an isothermal profile, then the temperature will be spatially uniform ($T_n = T_{n-1}$). At this point I want to satisfy the equation $P(\rho_n, T_n, Y_{en}) = P_n$, where $Y_{en} = Y_e(\rho_n, T_n)$. I use a root finding algorithm (as introduced in section 3.11) to solve the equation $P\left(\rho, T_n, Y_e(\rho, T_n)\right) = P_n$ for the mass density ρ_n. I use the initial guess ρ_{n-1}, as for this scenario I do not expect very large jumps across any discontinuities in the mass density profile. Once the mass density has been found the required specific internal energy (ϵ_n) can be obtained from the equation of state.

I now have the necessary quantities required to use the TOV equations (5.1) to advance the fluid pressure and enclosed mass outwards further, and begin the process again. This outwards integration is repeated until the mass density drops below some atmospheric threshold, at which point I define the neutron star's surface, obtaining its radius (R^*) and its total gravitational mass (M^*).

I am now in a position to match the lapse function (α) to the external Schwarzschild spacetime, which gives the exact spacetime representation outside of the TOV surface (R^*). I obtain that $\alpha\left(r \geq R^*\right) = \sqrt{1 - 2M^*/r}$. From here I can numerically integrate the gravitational potential (ϕ) back towards the centre of the TOV solution using equation (5.3). I thus obtain the internal TOV spacetime profile.

As in the curved spacetime test in section 4.5, in order to use this TOV spacetime on a Cartesian grid I transform to an isotropic coordinate system following (4.32). The coordinate transformation in the interior of the TOV solution (for $r < R^*$) requires modification and follows the procedure in [101].

The TOV solution for an isothermal temperature profile at $T = 10$ MeV, in β-equilibrium, with a central density of $\rho_c = 1.28 \times 10^{-3}$, using the SFHx (with electrons) equation of state [40] has been calculated. The fluid profiles are illustrated in figure 5.1, with the spacetime profiles shown in figure 5.2. As expected I obtain a high density central region which decreases with radius. It reaches an atmospheric threshold $\rho_{atm} = 10^{-10}$ at an isotropic radius of around $R^* = 10.2$. This gives a total mass of $M^* = 1.48\ M_\odot$. In the electron fraction profile we can see a small peak near the isotropic radius $r = 6$, at this same radius the mass density reaches its steepest decrease although the fluid pressure profile remains smooth throughout the profile (as is expected). This peak corresponds to a phase change in the equation of state. This TOV solution has a central lapse function of $\alpha_c = 0.634$.

5.4.1 Evolved Simulations

The TOV solution described in section 5.4 has been created on a Cartesian computational domain covering $(x, y, z) \in [-24, 24]^3$, with a uniform discretisation using 192^3 zones. This solution has been calculated using the tabulated equation of state, and has been transformed onto an isotropic coordinate system. The initial rest mass density and electron fraction profiles are illustrated in figure 5.3, which shows the $z = 0$ slice through the domain.

As this is a hydrostationary configuration, a perfect simulation should maintain its stationarity. However due to the finite spatial discretisation of the initial profiles and the numerical reconstruction and evolution methods, there is likely to exist a discrete numerical stationary configuration, qualitatively similar but quantitatively different to the continuous stationary configuration. Due to this I expect to see (ideally small) oscillations in the solution's evolution, as the numerical approximations break the precise continuous stationarity. These oscillations will be seen in section 5.8.

This initial state has been evolved to a time of 500 in geometric units (approximately 2.5 ms) with the HLLC Riemann solver, fifth order WENO reconstruction, and second order temporal integration. The evolved rest mass density and electron fraction profiles are shown in figure 5.4, along with their deviation from the initial states. We can see that the evolved profiles are qualitatively very close to the initial states, they maintain close to spherical symmetry, and the magnitudes and radial profiles values are well preserved. The four-lobed symmetry structure visible in these slices are a result of Cartesian grid imprinting effects. We still see a high density and neutron rich core region transitioning to a low density atmosphere of symmetric nuclear

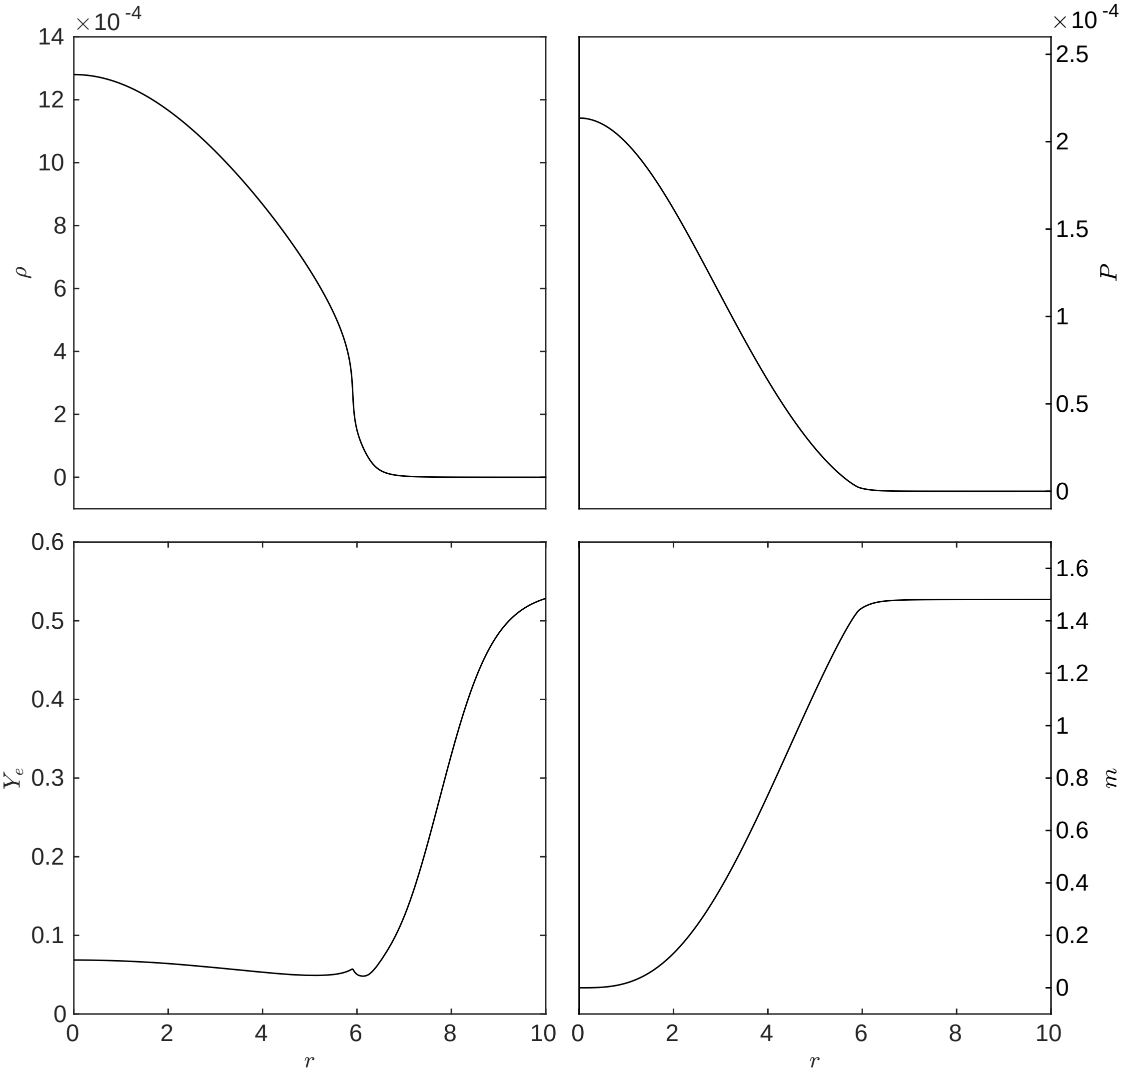

Figure 5.1: TOV solution fluid profiles against isotropic radius. The top left panel shows mass-density, the top right shows fluid pressure, the bottom left shows electron fraction, and the bottom right shows the enclosed gravitational mass. The peak in the electron fraction profile at a radius of around 6 corresponds to a phase transition in the equation of state. Here and throughout this chapter the radius is given in geometric units.

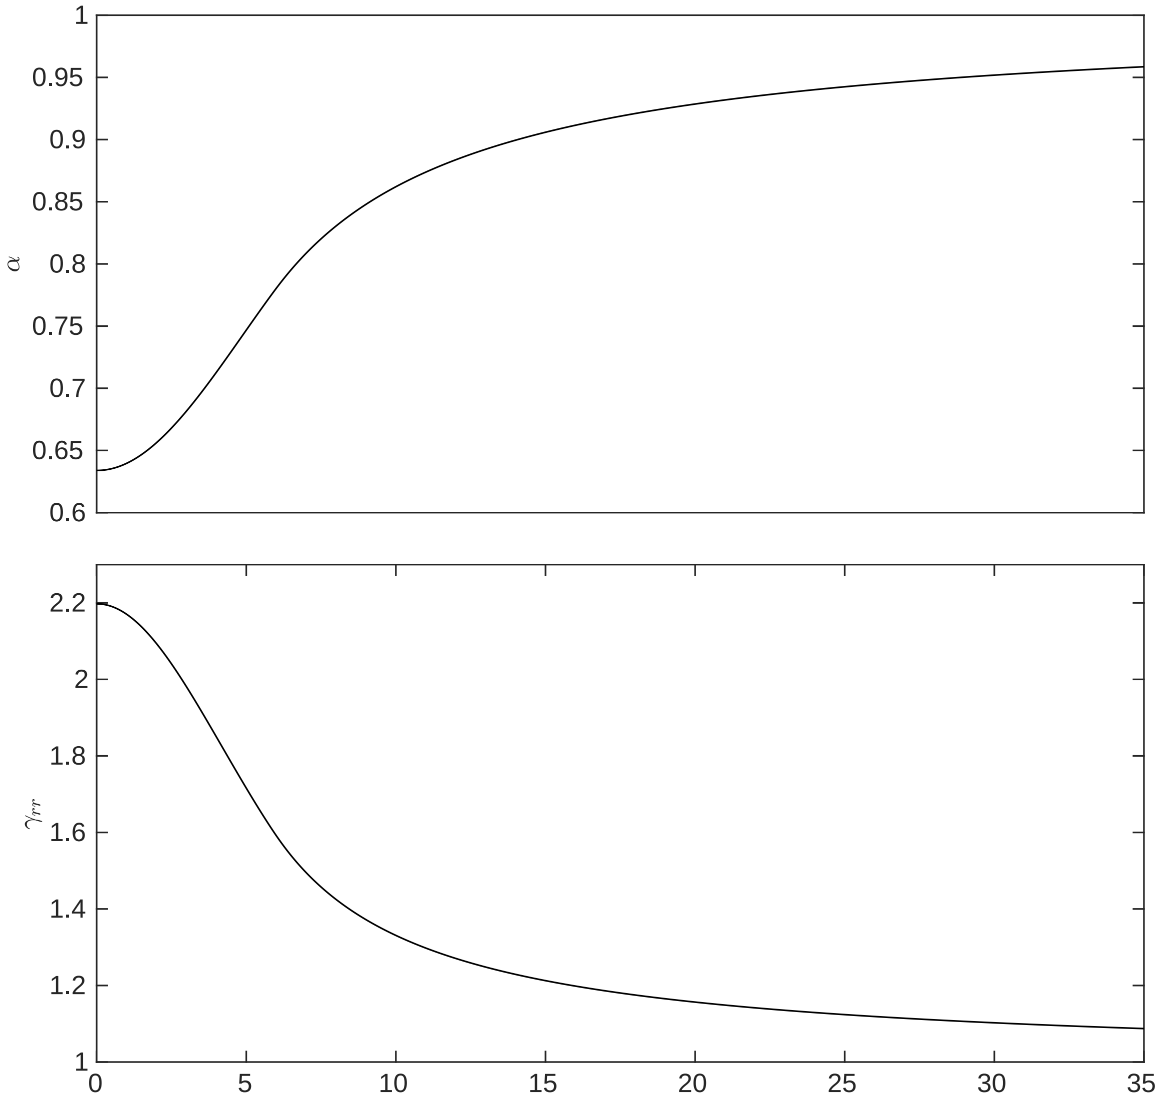

Figure 5.2: TOV solution spacetime profiles against isotropic radius. The top panel shows the lapse function, and the bottom shows the isotropic spatial metric.

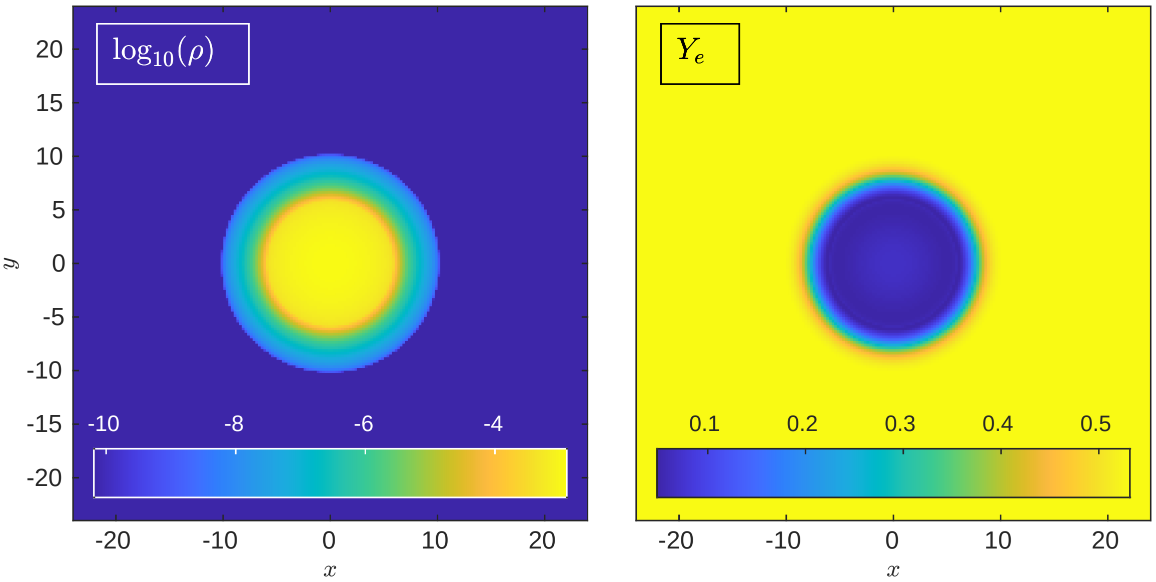

Figure 5.3: Initial neutron star model log-scale rest mass density, and electron fraction profiles.

matter. The deviation from initial state is at most approximately a 10% change.

The fluid temperature field is initially uniform, and the fluid velocity field is initially static. The evolved temperature and velocity fields are shown in figure 5.5. The fluid temperature field is showing thin shells of heating and cooling, approximately doubling in temperature in some regions, but away from these shells the temperature remains stable, the velocity magnitude has grown to reach a peak of nearly 12% of the speed of light. As with the rest mass density and electron fraction these deviations from the initial state are not problematic, indicating that the code can indeed maintain the stationary configuration of a physical fluid within a non-trivial curved spacetime. Waves passing through the phase transition at around $r = 6$ result in oscillations with cell-width wavelength. These oscillations are visible in the evolved mass density, electron fraction, fluid temperature, and fluid velocity fields. This is most likely due to the Riemann solver (described in section 3.7.2) lacking dissipative behaviour at low sound speeds (where the wave structure degenerates as the sound speed drops to zero). A naively more dissipative Riemann solver would have undesirable effects on the bulk density. I choose to accept these small scale oscillations to prioritise long term accuracy of evolution of bulk fluid and radiation properties.

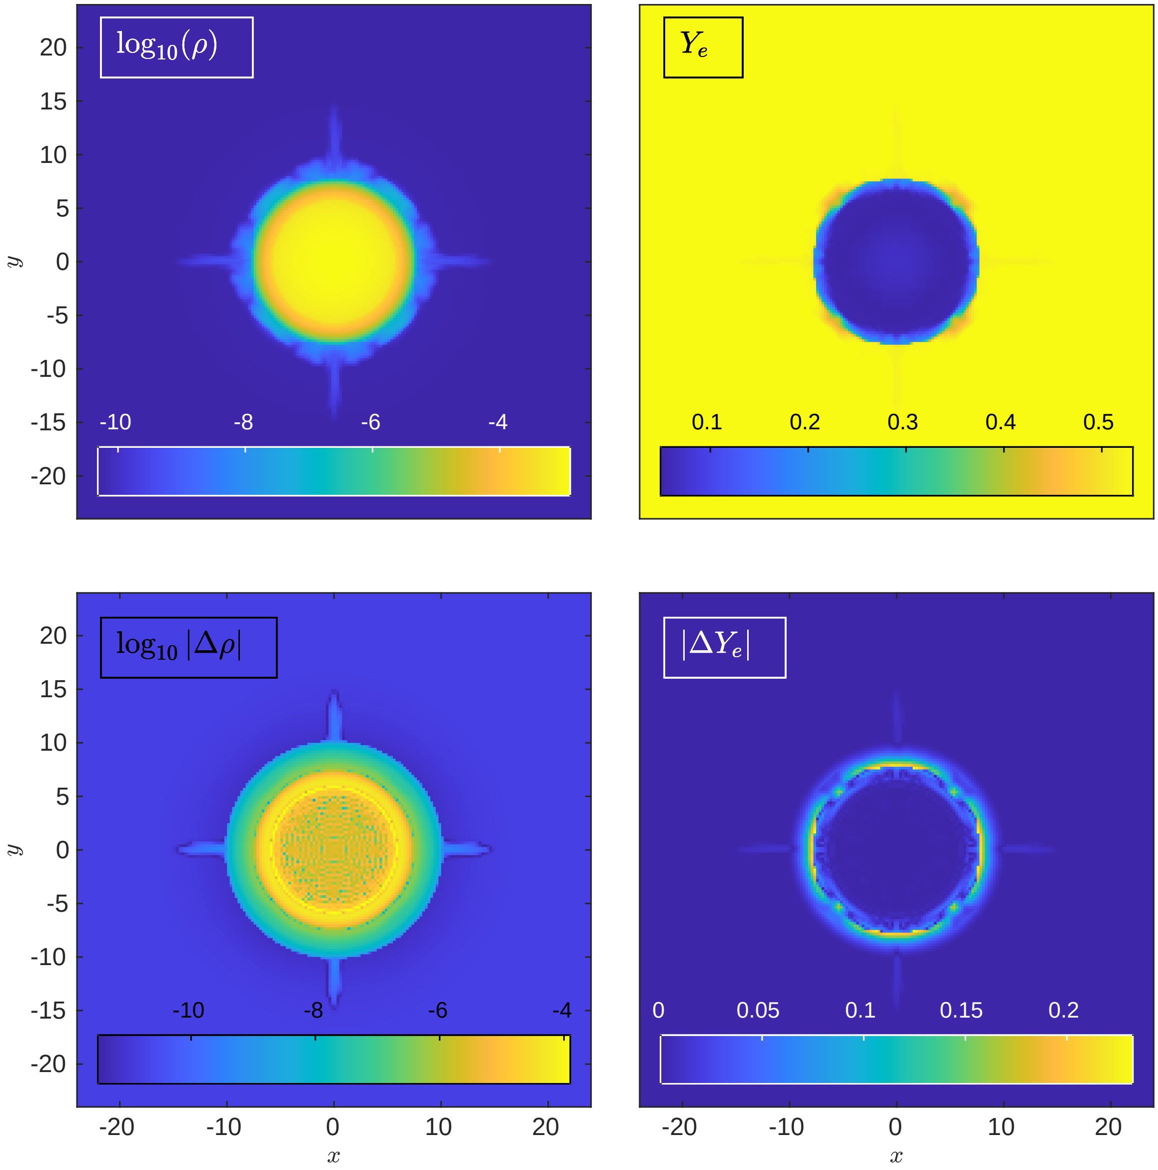

Figure 5.4: Evolved TOV model profiles. The top row shows the evolved neutron star model log-scale rest mass density, and electron fraction profiles. The bottom row shows the absolute deviations from the initial profiles. The quantities are shown at time $t = 500$ in geometric units.

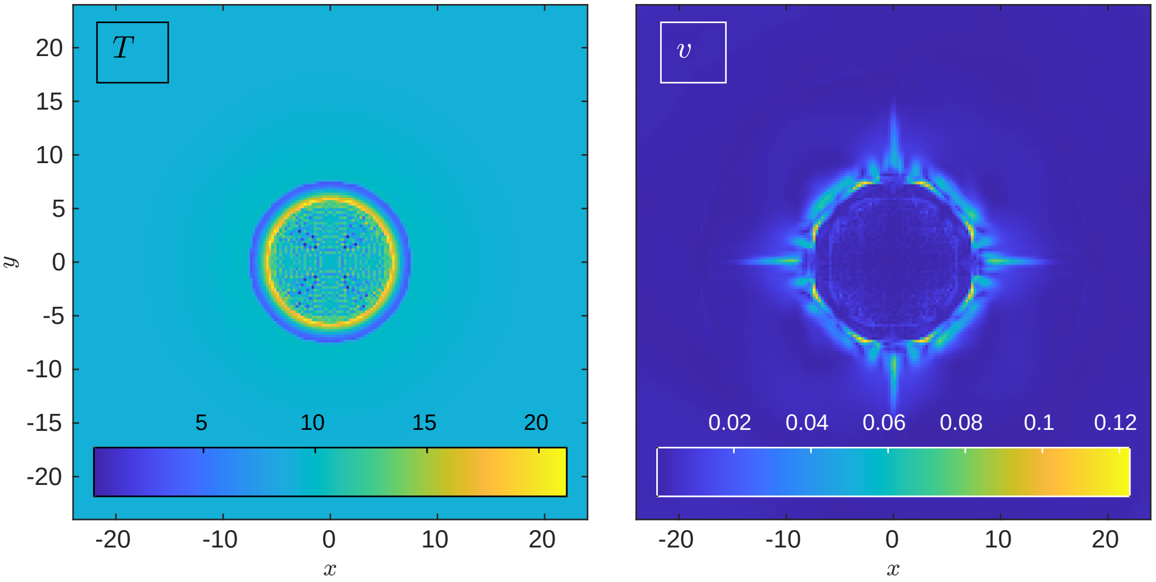

Figure 5.5: Evolved neutron star model fluid temperature, and fluid velocity magnitude profiles. The quantities are shown at time $t = 500$ in geometric units.

5.5 TOV With Radiation

If we wish to include radiation during the construction of the TOV fluid and spacetime profiles, a radiation field can be created to both contribute towards and balance against the spacetime curvature. The TOV equations can be split by considering both radiation and fluid components of the energy and pressure terms. Here the total energy is $E_{tot} = \rho + \rho\epsilon + E_R$, and the total pressure is $P_{tot} = P + P_R$. These split TOV equations take the form

$$\frac{\partial P}{\partial r} = -\left[\rho + \rho\epsilon + P\right] \frac{m + 4\pi r^3 (P + P_R)}{r(r - 2m)}, \tag{5.4}$$

$$\frac{\partial P_R}{\partial r} = -\left[E_R + P_R\right] \frac{m + 4\pi r^3 (P + P_R)}{r(r - 2m)}, \tag{5.5}$$

$$\frac{\partial m}{\partial r} = 4\pi r^2 \left[\rho + \rho\epsilon + E_R\right], \tag{5.6}$$

$$\frac{d\phi}{dr} = \frac{m + 4\pi r^3 (P + P_R)}{r(r - 2m)}. \tag{5.7}$$

Here the fluid pressure equation takes into account the fluid energy and fluid pressure, the radiation pressure equation accounts for the radiation energy and radiation pressure, the gravitational mass equation contains contributions from both the fluid and the radiation energies, and the gravitational potential equation is affected by both the fluid and the radiation pressures.

From a chosen central fluid temperature, I can obtain the central radiation energy density using local thermodynamic equilibrium (LTE): $E_{Rc} = a_{rad}T_c^4$, where $a_{rad} = 2.473 \times 10^{-13}$ MeV^{-4} with geometric units. With this the radiation pressure can be obtained from the Eddington closure (as introduced in section 2.8.2): $P_R = E_R/3$. As I am seeking a stationary solution, it is appropriate to assume that the radiation flux is small, and thus the radiation pressure tensor is near fully isotropic, a regime in which the Eddington closure is accurate. Using these LTE and Eddington relations I can integrate out the radiation pressure with the TOV radiation equation, obtaining radiation energy density and LTE temperature profiles. This approach to determining the temperature replaces the isothermal constraint used in section 5.4.

The TOV solution with a stationary and LTE radiation field, for a central temperature of $T_c = 10$ MeV, in β-equilibrium, with a central density of $\rho_c = 1.28 \times 10^{-3}$, using the SFHx (with electrons) equation of state [40] has been calculated. The radiation and temperature profiles are shown in figure 5.6. The mass density, fluid pressure, electron fraction, enclosed gravitational mass, and the spacetime profiles remain qualitatively unchanged by the introduction of this radiation field. This is a result of the radiation field having a much lower energy density than the fluid at the temperatures involved. The radiation field drops off with radius much slower than the mass density and fluid pressure does.

However, for a pre-calculated fluid and spacetime solution, I have a choice of approaches to including a radiation field on top of the given fluid and spacetime profiles. One approach is to construct a stationary radiation field balancing against (but not contributing towards) the pre-calculated spacetime

$$\frac{dP_R}{dr} = -\left(E_R + P_R\right)\frac{m + 4\pi r^3 P}{r(r-2m)}, \tag{5.8}$$

where

$$\frac{d\phi}{dr} = \frac{m + 4\pi r^3 P}{r(r-2m)} \tag{5.9}$$

as in equation (5.3). I use the Eddington closure as I expect the radiation flux to be small when in a stationary configuration, giving $E_R = 3P_R$. This leads to radiation energy density varying only with the lapse function, such that $E_R(r) = k\alpha(r)^{-4}$. The remaining degree of freedom allows us to fix the radiation field to be in local

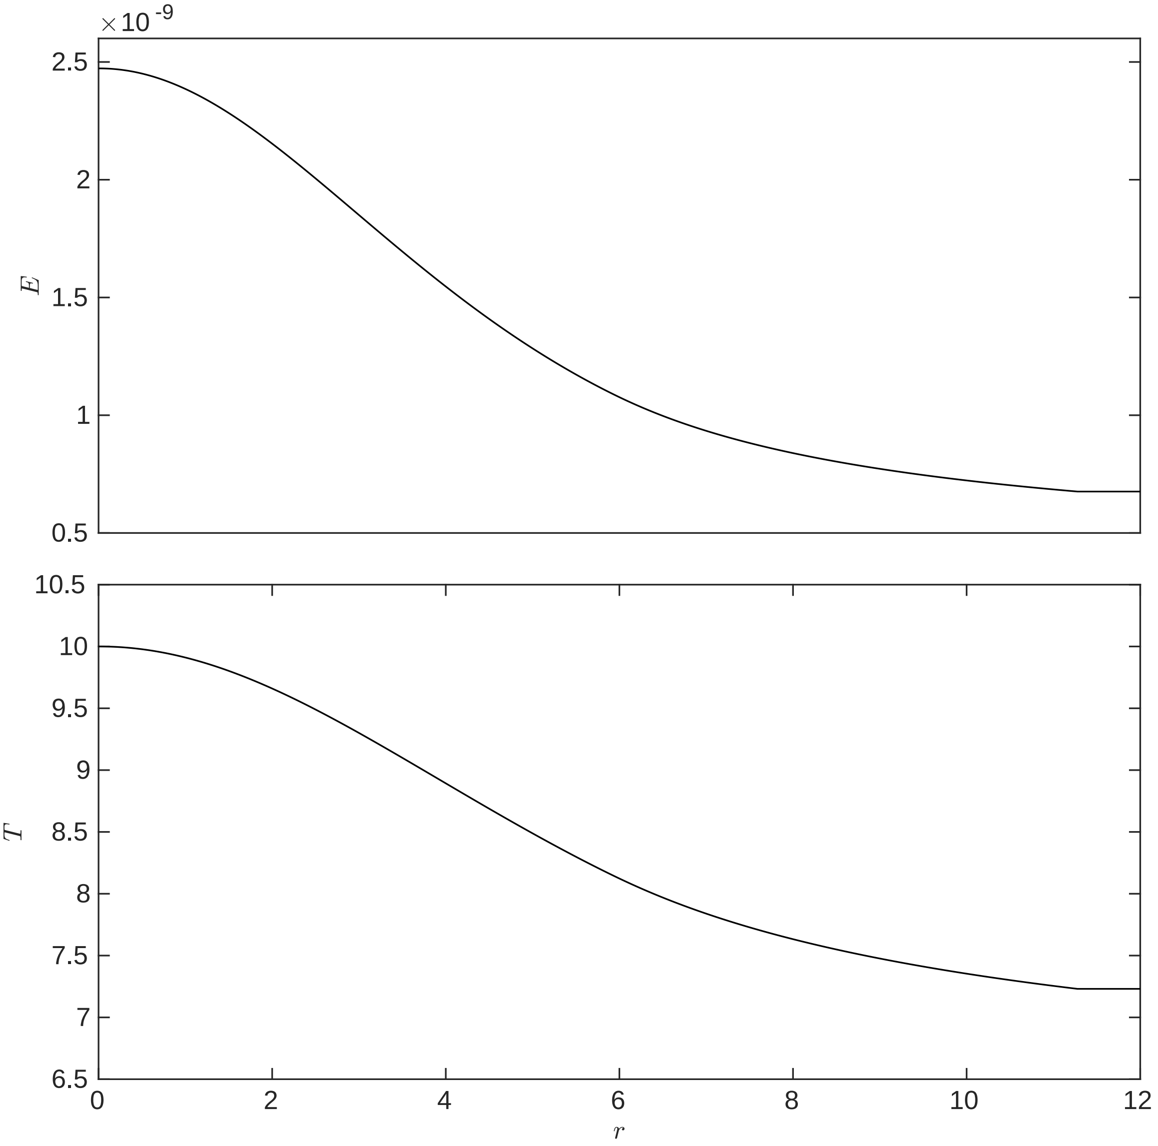

Figure 5.6: TOV solution radiation profiles against isotropic radius. The top panel shows the radiation energy density, and the bottom shows the local thermodynamic equilibrium temperature of the fluid and radiation.

thermodynamic equilibrium with the fluid at a single point, but not globally. For simplicity I can choose this point to be at the centre of the TOV solution, leading to

$$E_R(r) = a_{rad} T_c^4 \alpha_c^4 \alpha(r)^{-4}. \tag{5.10}$$

An alternative approach is to create a radiation field which is in local thermodynamic equilibrium with the fluid everywhere, leading to

$$E_R(r) = a_{rad} T(r)^4. \tag{5.11}$$

With an isothermal TOV solution this leads to a spatially uniform radiation field. This approach results in the radiation pressure gradient not balancing against the spacetime curvature (other than at the symmetry condition at the centre), nor contributing towards it, and is thus not a stationary configuration.

Each of the above approaches to introducing a radiation field to a TOV solution can lead to a configuration where the radiation energy density does not drop off quickly with radius. This is potentially problematic as I do not expect to see large radiation energy densities far away from a neutron star merger remnant. Another issue is that significant amounts of radiation energy-momentum could be lost through the boundaries of the numerical domain. This would require computationally expensive large spatial domains. Adaptive mesh refinement can be used to mitigate this issue by adding a large and coarse resolution grid around the simulation domain. An alternate way forward is to construct a radiation field in local thermodynamic equilibrium with the fluid only where the opacity is large (some threshold must be chosen). In this region I would expect LTE to be attained relatively quickly. These large opacities are found in the central, high mass density region of the TOV solution. In regions with low opacity I include no radiation field, as I would expect to see far away from a neutron star merger remnant. In these low opacity regions I expect the radiation to interact with the fluid relatively slowly, and have time to propagate away from the compact object. My numerical simulations have shown that having a radiation field out of balance with the spacetime curvature does not cause significant issues. The evolution equations allow the radiation field to rearrange towards a stable configuration in a smooth way.

I use the absorption opacity formula introduced in section 2.8.3. I do not include the scattering opacity as it has a negligible contribution to the evolution. The absorption

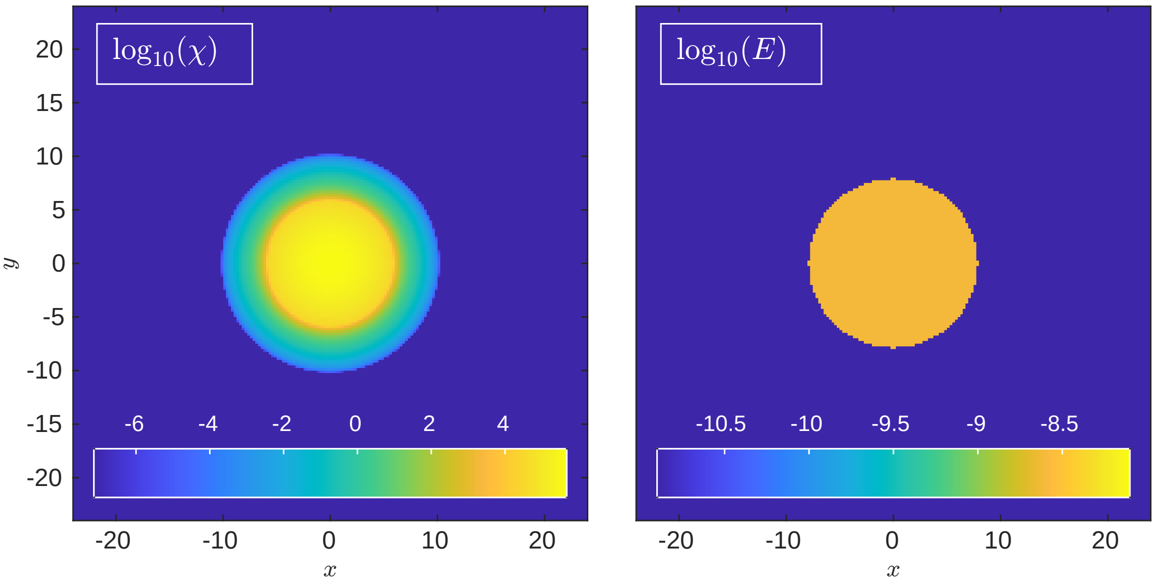

Figure 5.7: Initial neutron star model (with radiation) log-scale opacity, and radiation energy density profiles.

opacity varies between 10^{-7} in the background atmosphere, and around 10^{5} in the high density core.

5.5.1 Evolved Simulations with Radiation

As in section 5.4.1 the TOV solution has been created on a Cartesian computational domain covering $(x, y, z) \in [-24, 24]^3$, with a uniform discretisation using 192^3 zones. However here a radiation field has been introduced in local thermodynamic equilibrium with the isothermal fluid, in the region where the opacity is greater than unity in geometric units. The initial opacity and resulting radiation profiles are illustrated in figure 5.7.

This initial state has been evolved to a time of $t = 500$ in geometric units. The evolved rest mass density and radiation energy density profiles are shown in figure 5.8, along with their deviation from initial states. As was the case with the hydrodynamic setup in section 5.4.1, the evolved profiles are again qualitatively very close to the initial states. Spherical symmetry is still maintained, and the magnitudes and radial profiles are still well preserved. The deviation is similar to the hydrodynamic case, with the rest mass density changing by up to 10% around the surface of the TOV solution. This change was expected as the initial configuration was not stationary. The radiation field was not balanced against the spacetime curvature, and contained steep radiation pressure gradients. The radiation field has evolved to smooth out these steep pressure gradients, whilst simultaneously exchanging energy and momentum with the fluid. The effect of this evolution can be seen in the fluid temperature and radiation flux magnitude profiles shown in figure 5.9. The fluid has cooled and heated in shells just like in the purely hydrodynamic case. The opacity is around unity in this region, and the radiation field is initially far out of local thermodynamic equilibrium. The central short-wavelength oscillations are visible in fluid and radiation fields. These expected behaviours indicate that the code is capable of accurately evolving coupled radiation and fluid fields with a non-trivial curved spacetime.

5.6 TOV with Rotation

A more realistic neutron star merger remnant model requires a non-vanishing fluid velocity and shift vector. The TOV equations can provide a lapse function $\overline{\alpha}$, and diagonal spatial metric $\overline{\gamma}_{ab}$. From these I can construct a 4-velocity $\overline{u}_\alpha$ and diagonal spacetime 4-metric $\overline{g}_{\alpha\beta}$, which take the form

$$\overline{u}^\alpha = \left(1/\overline{\alpha}, 0^a\right), \tag{5.12}$$

$$\overline{g}^{\alpha\beta} = \begin{pmatrix} -1/\overline{\alpha}^2 & 0^a \\ 0^b & \overline{\gamma}^{ab} \end{pmatrix}. \tag{5.13}$$

Here the shift vector and fluid velocity are zero: $\overline{\beta}^a = \overline{v}^a = 0^a$. Overlined quantities denote the TOV coordinate system, with Greek spacetime indices $\{\alpha, \beta\}$ and Latin spatial indices $\{a, b\}$.

To investigate the effect of non-trivial velocities and spacetime backgrounds, I aim to perform a coordinate transformation of the static solution. In the full remnant simulation my goal is to use this freedom to minimise the impact of the physical rotation on the numerical solution.

I intend to transform this TOV solution to a new coordinate system, one in which I am free to choose the fluid 3-velocities v^i. In this new coordinate system we have a general 4-velocity u^μ and spacetime metric $g^{\mu\nu}$, given by

$$u^\mu = W\left(1/\alpha, v^i - \beta^i/\alpha\right), \tag{5.14}$$

$$g^{\mu\nu} = \begin{pmatrix} -1/\alpha^2 & \beta^i/\alpha^2 \\ \beta^j/\alpha^2 & \gamma^{ij} - \beta^i\beta^j/\alpha^2 \end{pmatrix}, \tag{5.15}$$

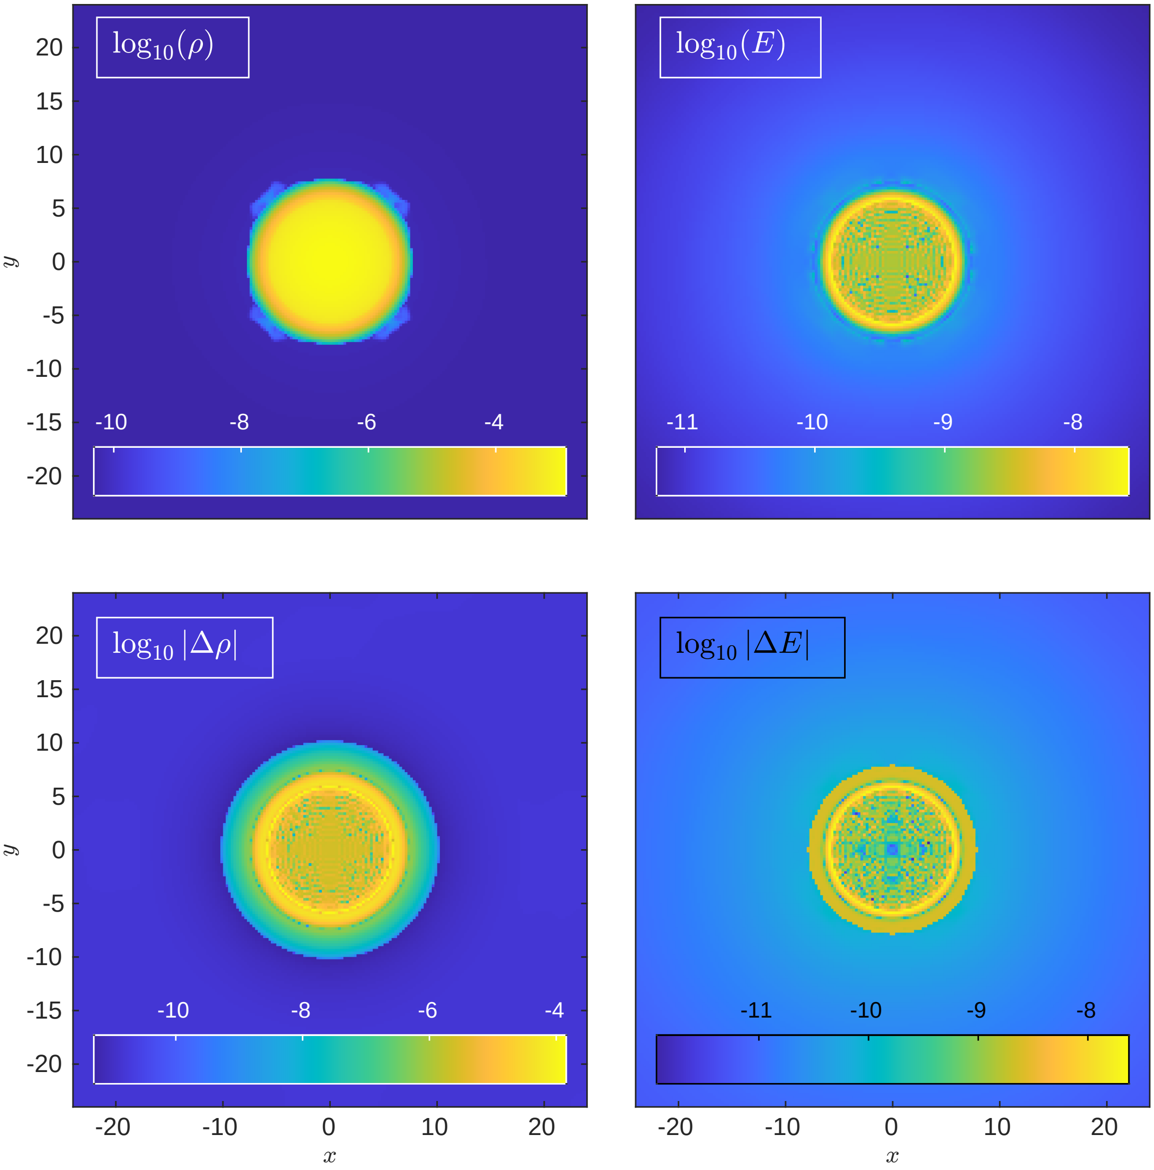

Figure 5.8: Evolved TOV model with radiation profiles. The top row shows evolved neutron star model (with radiation) log-scale rest mass density and radiation energy density profiles. The bottom row shows the absolute deviations from initial profiles. The quantities are shown at time $t = 500$ in geometric units.

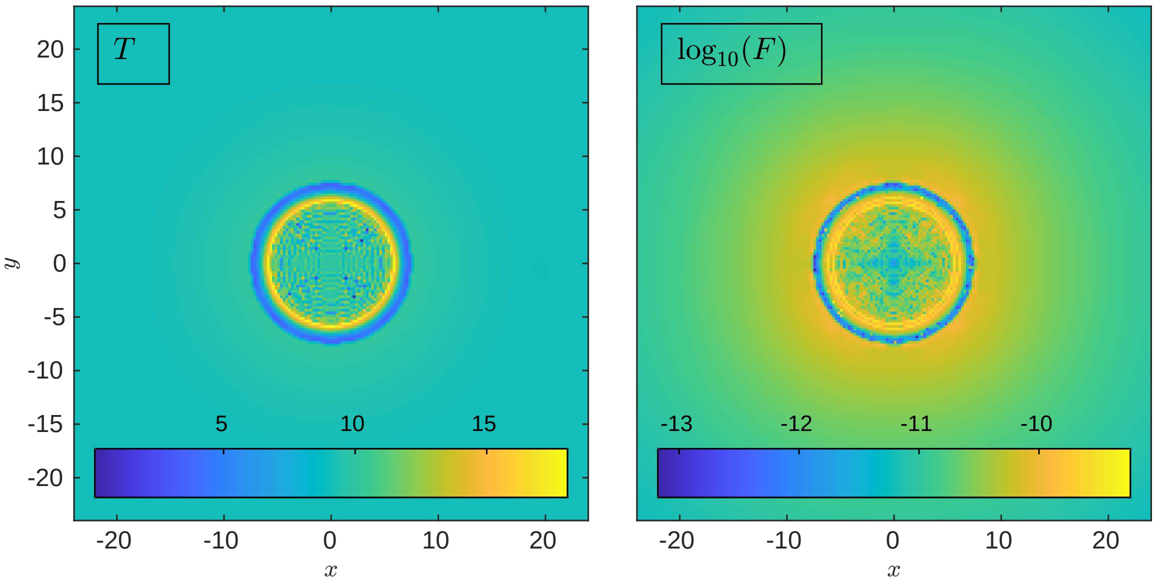

Figure 5.9: Evolved neutron star model (with radiation) fluid temperature and log-scale radiation flux magnitude profiles. The quantities are shown at time $t = 500$ in geometric units.

with Greek spacetime indices $\{\mu, \nu\}$ and Latin spatial indices $\{i, j\}$.

A coordinate transformation Λ^{α}_{μ} can be used to obtain the new 4-velocity and spacetime metric from the TOV ones,

$$u^{\mu} = \Lambda^{\mu}_{\alpha} \overline{u}^{\alpha}, \tag{5.16}$$

$$g^{\mu\nu} = \Lambda^{\mu}_{\alpha} \Lambda^{\nu}_{\beta} \overline{g}^{\alpha\beta}. \tag{5.17}$$

With the freely chosen fluid velocities we can construct the coordinate transform

$$\Lambda^{\mu}_{\alpha} = \begin{pmatrix} 1 & 0^i \\ v_a/\overline{\alpha} & \delta^i_a/W \end{pmatrix}, \tag{5.18}$$

where the new Lorentz factor $W = (1 - v^i v^i / \overline{\gamma}^{ii})^{-1/2}$.

This particular coordinate transform results in a new set of spacetime quantities

$$\alpha = W \overline{\alpha}, \tag{5.19}$$

$$\beta^i = \alpha v^i, \tag{5.20}$$

$$\gamma^{ij} = \overline{\gamma}^{ij} / W^2 + v^i v^j. \tag{5.21}$$

The contravariant 4-velocity remains unchanged by this transformation, such that $u^\mu = \overline{u}^\mu$.

5.6.1 Choosing a Velocity Field

I choose a simple (piecewise linear) triangle function in cylindrical radius $\left(r_c = \sqrt{x^2 + y^2}\right)$ and height (z). A cylindrically symmetric torus shaped velocity magnitude $V(r_c, z)$, for $|z| < Z_{height}$, can be constructed as

$$V(r_c, z) = \begin{cases} \dfrac{V_{max}}{R_{peak}} r_c \left(1 - \dfrac{|z|}{Z_{height}}\right), & r_c \leq R_{peak}, \\ \dfrac{-V_{max}}{R_{out} - R_{peak}} (r_c - R_{out}) \left(1 - \dfrac{|z|}{Z_{height}}\right), & r_c \in (R_{peak}, R_{out}), \\ 0, & r_c \geq R_{out}. \end{cases} \tag{5.22}$$

In order to mimic simulated neutron star merger remnant conditions, I choose that: $V_{max} = 0.4$, $R_{peak} = 10$, $R_{out} = 20$, and $Z_{height} = 10$. With a clockwise rotation, the fluid velocity components can be obtained with

$$v^x = \frac{V}{r_c}(-y), \tag{5.23}$$

$$v^y = \frac{V}{r_c}(x), \tag{5.24}$$

$$v^z = 0. \tag{5.25}$$

I have chosen to use piecewise linear velocity profiles due to their simplicity, and because simulation experiments have shown that the stationary fluid configuration is better maintained when the divergence of the velocity field is small. These profiles minimise the spatial region with non-vanishing velocity divergence whilst still reaching the peak velocity magnitudes in the appropriate regions to mimic the remnant velocity field.

5.6.2 Evolved Simulations with Rotation

Following the simulation described in section 5.4.1, the TOV solution profiles have been created on a Cartesian computational domain covering $(x, y, z) \in [-24, 24]^3$, with a uniform discretisation of 192^3 zones. A torus-shaped velocity field (as described in section 5.6.1) has been imposed, with the corresponding coordinate and

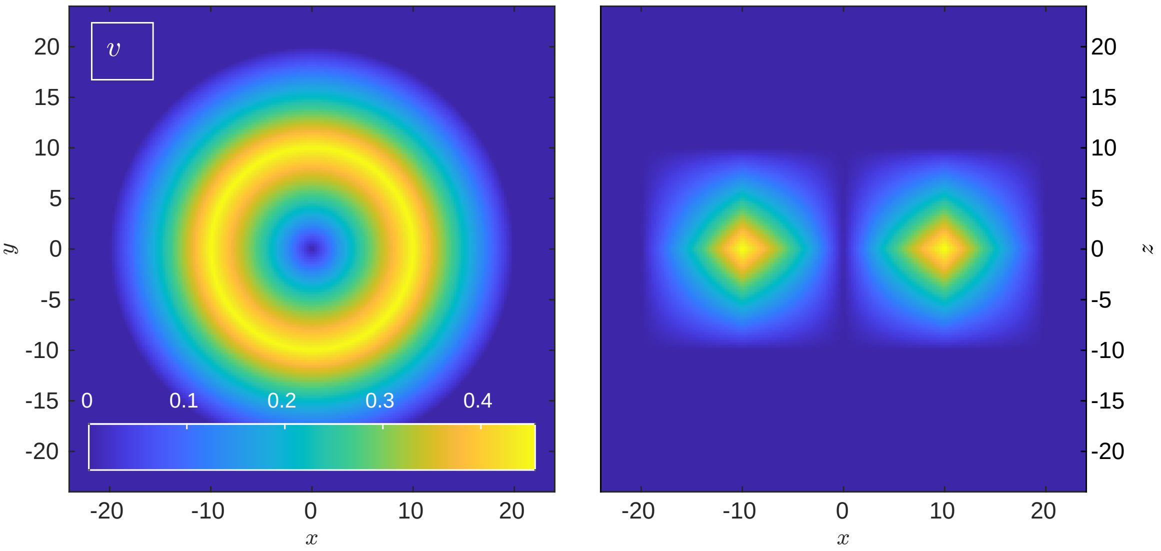

Figure 5.10: Initial neutron star model (with rotation) fluid velocity magnitude profile. Left panel: $z = 0$ slice. Right panel: $y = 0$ slice.

spacetime quantity transformations (as introduced in section 5.6). The resulting initial velocity magnitude profile is illustrated in figure 5.10. The diamond shaped structure is a result of the triangular profiles in r_c and z.

This initial state has been evolved to a time of $t = 500$ in geometric units. The evolved rest mass density is shown in figure 5.11, along with its deviation from the initial state, and the evolved velocity magnitude is shown in figure 5.12. The differentially rotating fluid results in shearing velocities, and the Cartesian discretisation induces perturbations which lead to turbulent vortices. This turbulent behaviour is visible in the rest mass density field around the TOV solution's surface only where the fluid velocities are largest, outside of the toroidal region the evolution remains smooth. Even with these detailed features the deviation from initial state is still around 10%, suggesting that the code is capable of accurately evolving a rotating fluid within a rotating spacetime, whilst maintaining hydrostationary behaviour.

5.7 Evolved Simulation with Both Radiation and Rotation

The TOV setup including both radiation and rotation has been evolved to a time of $t = 500$ in geometric units. The evolved rest mass density and radiation energy density profiles are shown in figure 5.13, along with their deviation from the initial

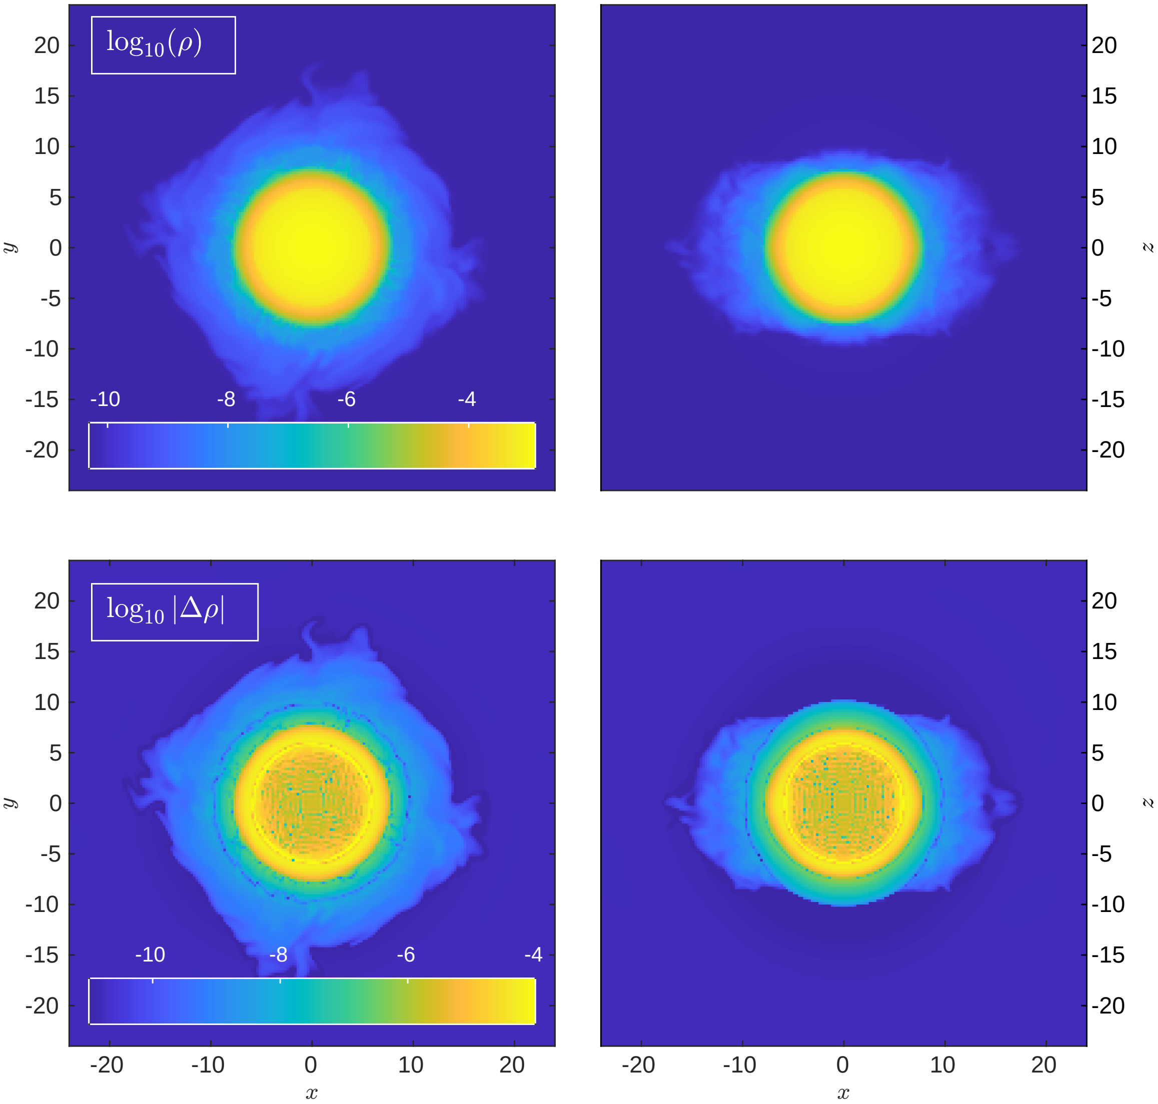

Figure 5.11: Evolved TOV model with rotation profiles. The top row shows the evolved neutron star model (with rotation) log-scale rest mass density profile. Left panels: $z = 0$ slice. Right panels: $y = 0$ slice. The bottom row shows the log-scale deviation from the initial state. The quantities are shown at time $t = 500$ in geometric units.

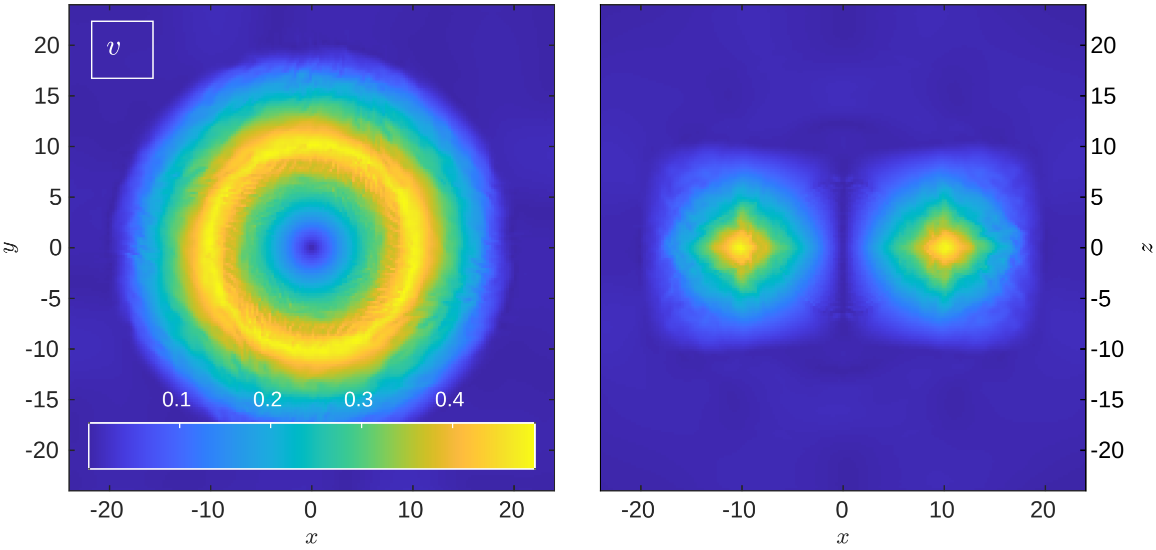

Figure 5.12: Evolved neutron star model (with rotation) fluid velocity magnitude profile. Left panel: $z = 0$ slice. Right panel: $y = 0$ slice. The quantities are shown at time $t = 500$ in geometric units.

states. We can see that the spherical symmetry, magnitudes, and radial profiles are approximately maintained. These results exhibit features from both the previous radiation and rotation simulations. The radiation field has evolved to smooth out steep gradients, balance against the spacetime curvature, and exchange energy with the fluid. The rotation and shearing motion has resulted in turbulent vortices around the surface of the TOV solution.

5.8 Central Rest Mass Density

The previous qualitative analyses in sections 5.4.1, 5.5.1, 5.6.2, and 5.7 discuss features in 2D slices of profiles at a single evolved time allowing us to conclude that deviations from a stationary configuration are small. I now conduct detailed quantitative comparisons using the time evolution of a single point in space. Our focus is the long term evolution of central rest mass density (as is the standard quantitative diagnostic used in [102, 103], and in the same theme as in [60]) which allows us to quantify the dissipation of the numerical schemes used.

The central rest-mass density is output for each time-step. This is a convenient diagnostic for observing how the fluid is evolving in time. This diagnostic can also reveal how changes such as varying the opacity values, or modifying the numerical fluxes affects the evolution of the fluid fields without relying on memory intensive

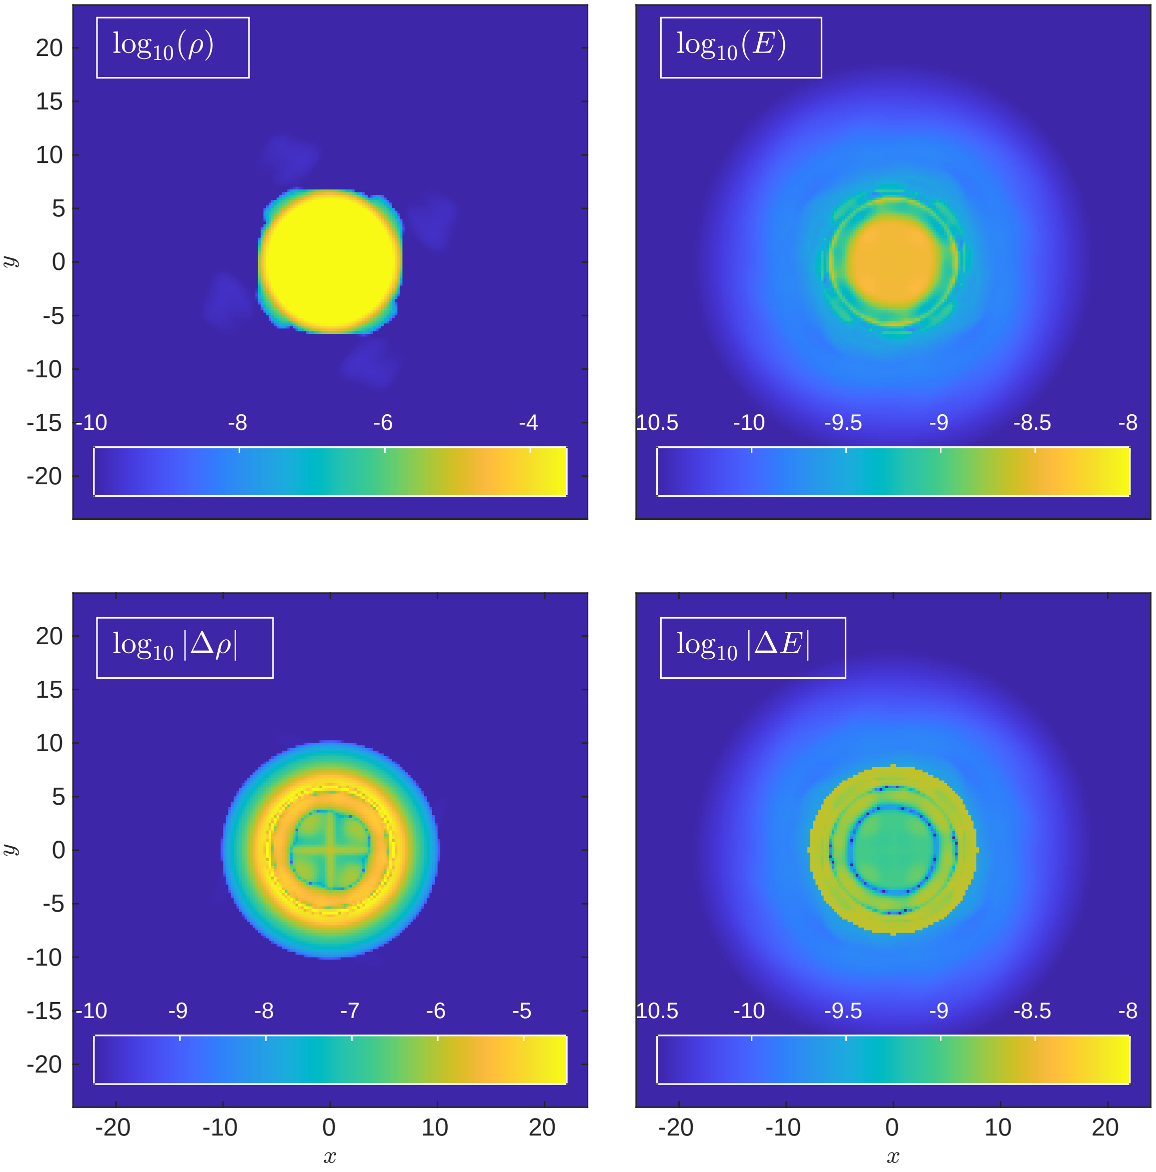

Figure 5.13: Evolved TOV model with radiation and rotation profiles. The top row shows evolved neutron star model (with both radiation and rotation) log-scale rest mass density and radiation energy density profiles. The bottom row shows the absolute deviations from initial profiles. The quantities are shown at time $t = 500$ in geometric units.

three-dimensional outputs.

For each of the preceding simulations in this chapter (5.4.1, 5.5.1, 5.6.2, and 5.7) the evolution of the relative deviation of central rest mass density from initial value (which is calculated as $|\rho_c(t) - \rho_c(0)|/\rho_c(0)$) is shown in figure 5.14. The oscillations in relative deviation are small, with magnitudes of order 10^{-3}. They oscillate around a stable value with negligible long term growth or decay. The evolved solution finds a stable discretised configuration which is not necessarily the same as the continuous equilibrium due to grid resolution and reconstruction schemes. This further suggests that the code is capable of accurately maintaining stable configurations with a physical equation of state and a curved spacetime obtained from a numerical TOV solver (not an analytic spacetime profile such as the Schwarzschild solution which was used in section 4.5).

The numerical reconstruction schemes I use have the effect of clipping local extremal values. This can introduce errors which move the system away from a strict hydrostationary configuration, and thus result in bulk or oscillatory motion. The standard TOV central rest mass density oscillations reveal the fundamental mode (and overtones) of the TOV solution's hydrodynamic system. The TOV with radiation shows the fundamental modes of the full fluid and radiation coupled system, introducing an additional channel within which energy and momentum can be exchanged. The fluid can emit and absorb a small amount of energy into and out of the radiation field. This can act to smooth the fluid oscillations, resulting in the loss of higher frequency oscillations, and increasing the oscillation magnitude at low frequencies. The rotating TOV case has the same system as the standard hydrodynamic setup, however there is an added source of rotational motion. This removes some lower frequency oscillations, but introduces new higher frequencies as the coordinates pass through the mesh in a non-grid-aligned way.

5.9 AMR Runtime Performance

Three dimensional simulations stand to benefit the most from adaptive mesh refinement, as the alternative approach of increasing spatial resolution throughout the entirety of the domain results in cubic growth in the number of grid cells. In this section I evolve the standard TOV solution described in section 5.4.1, whilst varying the number of AMR levels. The inner region ($R \leq 8$) is maintained at a spatial resolution of $\Delta x = 0.25$, as the outer regions are coarsened. An example setup is

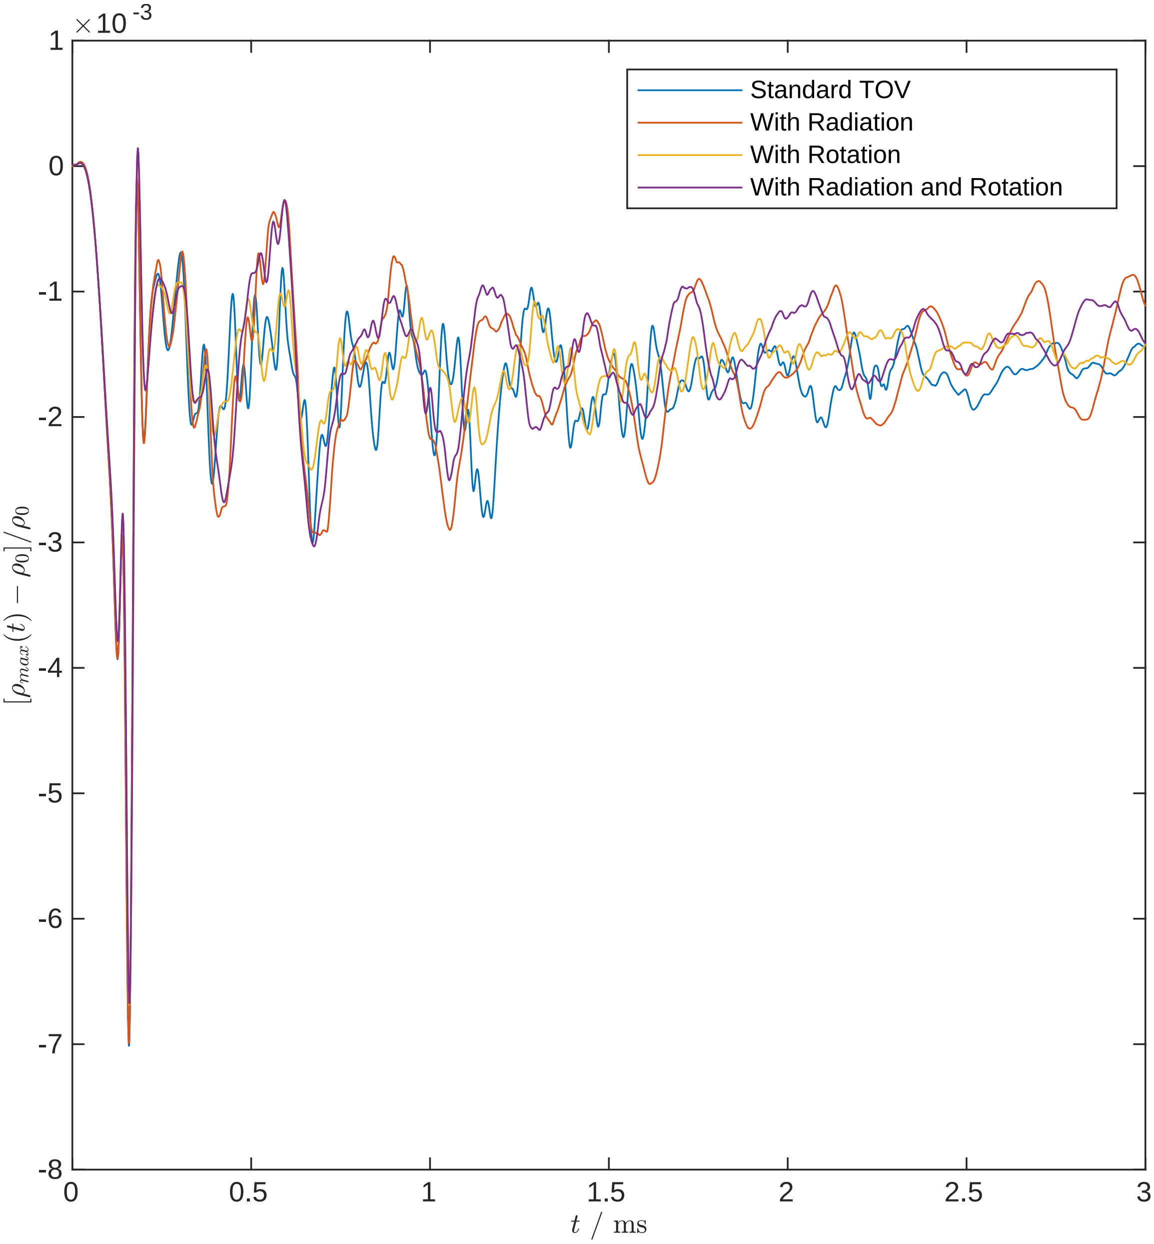

Figure 5.14: Central rest mass density evolution for the standard, with radiation, and with rotation TOV simulations. The quantities are shown up to time $t = 3$ ms.

Levels of refinement	0	1	2	3
Compute time / min	9332	1499	1362	1314
L_1 deviation	7.837×10^{-8}	2.517×10^{-7}	3.053×10^{-7}	3.068×10^{-7}
L_2 deviation	5.723×10^{-7}	1.572×10^{-6}	1.833×10^{-6}	1.843×10^{-6}

Table 5.1: Table of compute times required to reach simulation time of $t = 10^3$ in geometric units and deviation from initial state for a range of levels of refinement.

shown in figure 5.15 with three levels of refinement.

The evolution is repeated with zero, one, two, and three levels of factor-2 refinement. As a metric I use the compute time taken to reach a simulation time of $t = 10^3$ in geometric units, and measure any increases in error at this time with a volume integrated deviation from the initial condition using the L_1 and L_2 norms (as introduced in (4.1)). The results are shown in table 5.1 and figures 5.16, 5.17, and 5.18.

In this setup there is a large compute time benefit to using one level of mesh refinement (with relatively small increase in error). However there are diminishing returns for adding more levels (further coarsening the outer regions). This could be due to the outer regions being computationally cheap to evolve as it is (due to fast convergence of primitive recovery iterations compared to in the central region), and increasing efficiency here yields little net gain in calculation time. However this confirms that I could use additional levels of refinement to increase the domain size which will be important for the full merger simulation. These results also show that the flexibility of AMReX in producing the relatively complex patch structure seen in figure 5.15 does not negatively impact the runtime. This will be important in simulations with more complex small-scale features such as to be expected in a neutron star merger remnant.

5.10 Summary

In this chapter I have created setups designed to mimic a neutron star merger remnant. This has involved using a three parameter tabulated equation of state which is more physical than the simpler Gamma-law models employed in section 4. I have used the TOV equations (5.1) to construct stationary fluid and spacetime configurations with a central rest mass density in the range expected for a neutron star merger remnant. The resulting spacetime is more general than the analytic

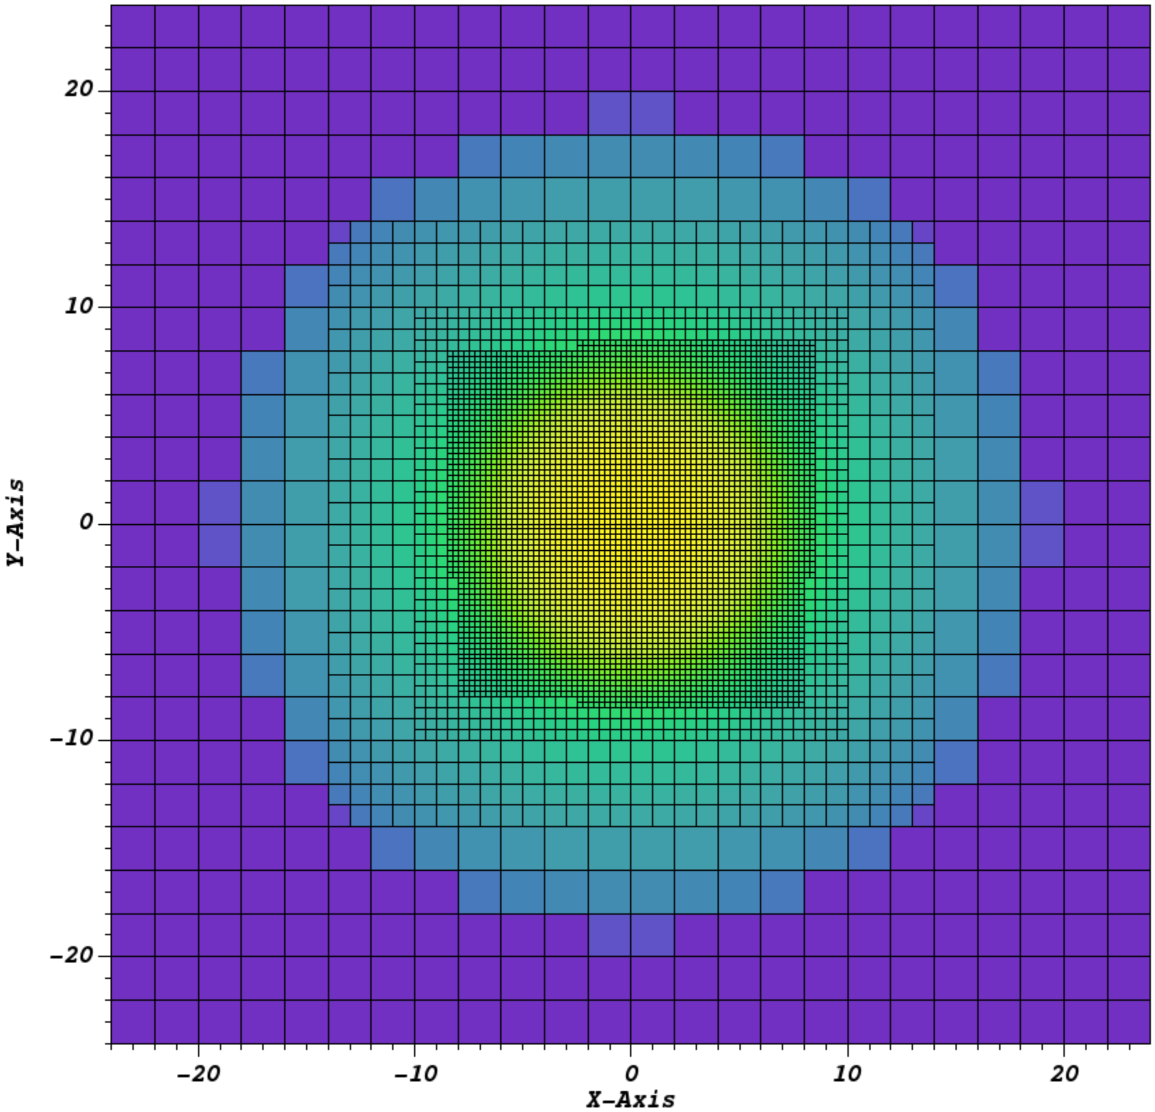

Figure 5.15: Block-structured mesh refinement covering the region with a radius $r = 8$ in geometric units. The log-scale rest mass density field is shown to highlight that the high density core is finely resolved.

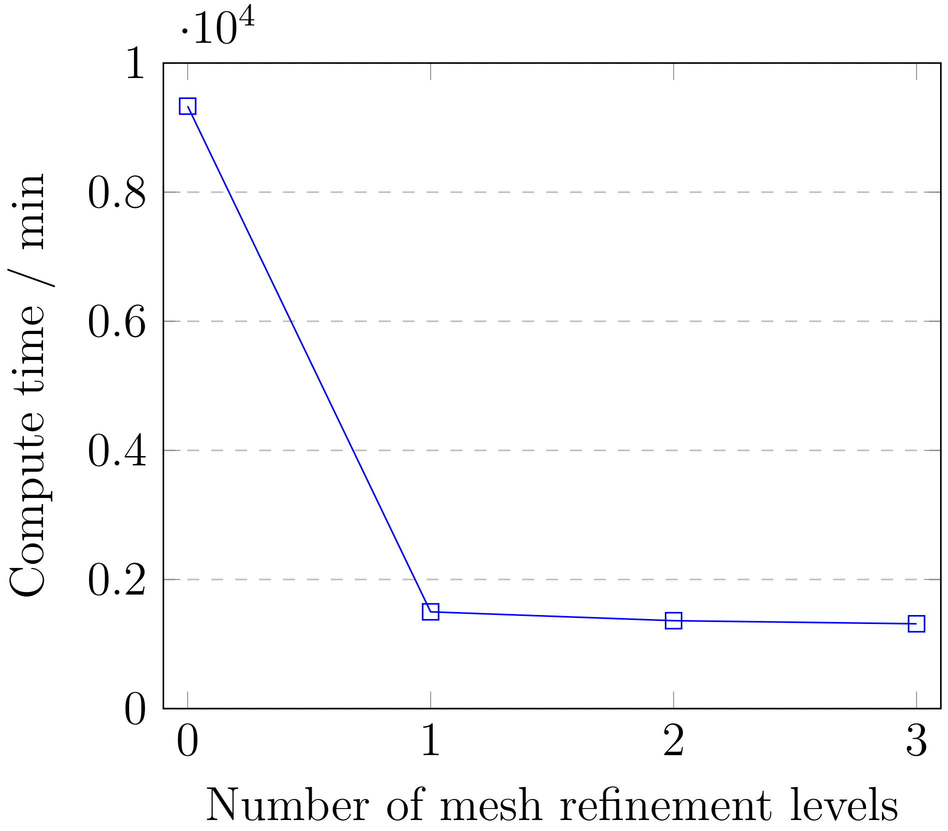

Figure 5.16: Plot of compute time used to reach simulation time $t = 10^3$ for a range of levels of refinement.

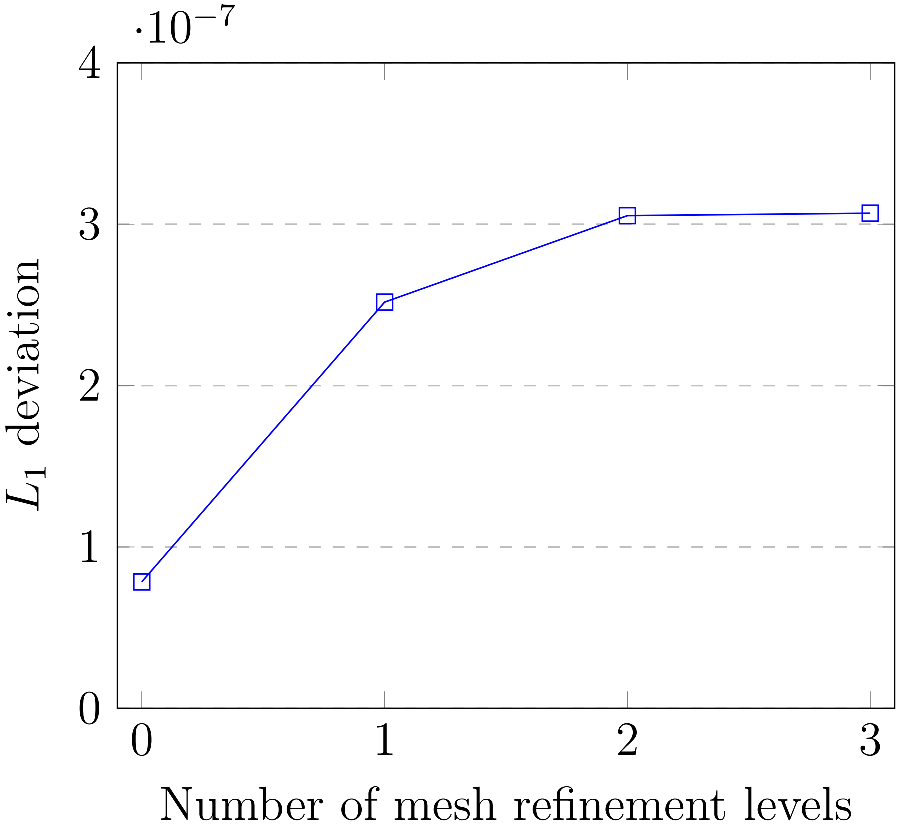

Figure 5.17: Plot of the L_1 deviation from the initial state at time $t = 10^3$ in geometric units, for a range of levels of refinement.

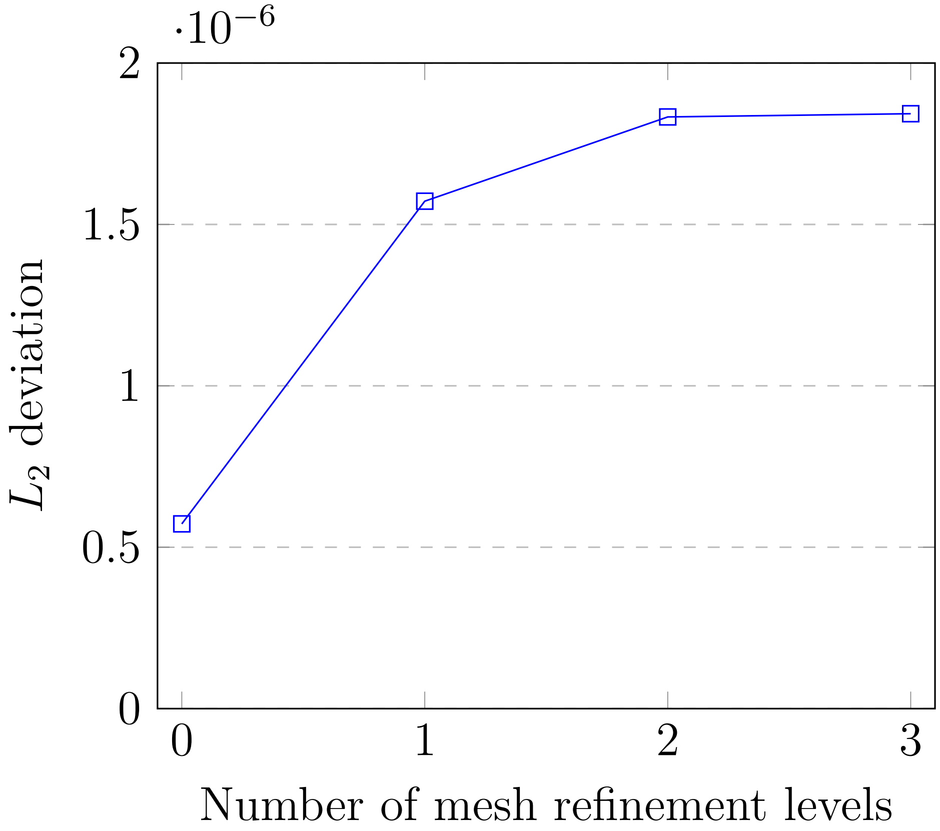

Figure 5.18: Plot of L_2 deviation from the initial state at time $t = 10^3$ in geometric units, for a range of levels of refinement.

solution used in section 4.5. I have implemented a radiation field according to physical emission and opacity parameters, and added relativistic fluid rotation according to the velocity profile of a neutron star merger simulation. This rotation implementation has involved a coordinate transform which resulted in non-trivial spacetime metric components, allowing me to validate that the code is capable of using even more general spacetime quantities. These TOV solutions have been evolved for multiple milliseconds (requiring thousands of computational time steps) and the features and deviations from initial configuration have been discussed. The central rest mass density evolution in each case has been shown, revealing the expected oscillations, and how they are affected by the presence of a radiation field and rotational velocities. I conclude that the code is indeed capable of accurately evolving a physical fluid at relativistic velocities within a general curved background spacetime, whilst interacting with a radiation field. I then experiment with making use of adaptive mesh refinement, and quantify the effect this has on calculation time. I conclude that it would be beneficial to make use of AMR to increase the spatial domain size whilst simultaneously keeping high spatial resolution around the compact remnant object.

Chapter 6

Neutron Star Merger Remnant

6.1 Introduction

The focus of this research has been on the effects that photon transport has on the evolution of a binary neutron star merger remnant. These include how the transport of radiation affects the cooling rate and energy loss of the remnant, and how energy is redistributed around the system, thus affecting its angular momentum. As the angular momentum is affected so too are the gravitational waves emitted by the rotating remnant. Photon transport plays an important role in the generation of observable signals. The luminosity, spectrum, and timings of the electromagnetic wave signals depend directly on the photon transport.

Further effects of photon transport on a remnant's evolution include the remnant's collapse to a black hole. I will not be modelling this due to my use of the stationary spacetime approximation, however it further highlights the usefulness of investigating radiation fields in relativistic scenarios. As energy is exchanged from the fluid to the radiation field, or as energy is removed from the system via radiative cooling, the pressure will be reduced, and thus may not be sufficient to support against collapse. Furthermore the changes to the angular momentum of the system due to photon transport could reduce the remnant's rotational support against collapse. As the differentially rotating remnant transitions towards uniform rotation its rotational support is reduced. This collapse could be noticeable in the observable gravitational wave signal as it would be cut off when the black hole is formed, and also in the electromagnetic signal as the surrounding material could be accreted towards the recently formed black hole remnant and lead to a short gamma-ray burst as observed in GRB 170817A [23].

I make use of a precalculated binary neutron star merger simulation provided by P. Hammond [20] using the Einstein Toolkit code. The required hydrodynamic and spacetime quantities have been reconstructed on an AMReX mesh, and the evolution is continued within my radiation hydrodynamics code (as detailed in chapter 3 and validated in chapters 4 and 5).

The Einstein Toolkit differs from my radiation hydrodynamics code in that it can evolve the spacetime quantities, however it does not have the capability to evolve a radiation field. I make use of the Cowling approximation as we assume that the spacetime is evolving sufficiently slowly. This has the benefit of avoiding issues which arise when evolving spacetime quantities across mesh refinement boundaries. The AMReX mesh is able to refine around much smaller spatial regions than the Einstein Toolkit as it can make use of many small patches. The large patches used in the Einstein Toolkit minimise the surface of refinement boundaries, allowing for accurate gravitational wave calculations. I expect to see a speedup in runtime due to the efficiencies gained from the flexibility in using many small AMR patches, as shown in chapter 5, and from avoiding spacetime evolution calculations.

In this chapter I make use of a pre-calculated binary neutron star merger remnant (soon after merging) and continue its evolutions in my code. I then look at approaches to including a radiation field onto the remnant initial condition, and how this could affect the evolution. I also explore the possibility to improve numerical calculation by applying a coordinate transform with the intent of reducing the numerical fluxes. Lastly I discuss how these modifications could affect oscillations of the central mass density.

6.2 Merger Simulation

A binary neutron star merger has been simulated by P. Hammond [26]. A pair of equal mass, corotating neutron stars (from TOV solutions) in approximate hydrostationary equilibrium (calculated with the LORENE code [104]), initialised in β-equilibrium, with a uniform fluid temperature, have been hydrodynamically evolved through inspiral to merger using the Einstein Toolkit code. Although the neutron stars are initialised in β-equilibrium there are no β-reactions occurring during the evolution. The electron number density has been advected along with the fluid allowing the material to become out of β-equilibrium. This simulation has made use of the tabulated three parameter equation of state SFHx (with electrons) [40], from the CompOSE library (as introduced in section 5.2). The spacetime has been evolved numerically using the BSSNOK formulation [105] with the aim of extracting the gravitational wave signals. The remnant object is hot and differentially rotating. This calculation did not involve any radiation field or radiative transfer. The spatial domain is discretised with a 3D Cartesian grid spanning $[-1584, 1584]^3$ (in geometric units) with eight levels of factor two mesh refinement and a coarsest cell width of 36 (also in geometric units).

6.3 Implementation into the AMReX Code

At some point after merger when the spacetime evolution is sufficiently slow, the fluid and spacetime quantities are transferred over to an AMReX mesh for further evolution within the radiation hydrodynamics code. This code assumes the Cowling approximation and does not evolve the spacetime quantities, the spacetime from the merger simulation is held stationary. This greatly reduces the computational expense of continuing the merger remnant evolution. The code's radiation wavespeed restricted timestep is comparable to the hydrodynamic timestep when modelling neutron star fluid. By utilising a large number of small patches, the adaptive mesh refinement is capable of refining much smaller spatial regions as spacetime evolution across resolution boundaries will not be an issue.

A linking code was required to construct the AMReX mesh and populate it with values extracted from the Einstein Toolkit mesh. The radiation hydrodynamics code does not take advantage of the symmetry of the problem in the same way that the Einstein Toolkit does. The spatial domain was duplicated four times over to reproduce reflective and rotational symmetries of the problem. The code required a number of necessary capabilities in order to continue the evolution of the merger remnant. In order to be compatible with the three parameter EoS and the associated algorithm for recovering the primitive quantities from the conserved quantities, the code requires access to rest mass density, fluid temperature, and electron fraction fields.

The energy evolution equations include source terms involving spacetime quantities. With a general remnant setup it is inaccurate to treat the spacetime as truly stationary. I can include this effect by making use of the extrinsic curvature in the energy equation geometric source terms. As shown in section 2.4 the relativistic energy evolution equation is

$$\partial_\mu \left(\alpha^2 \sqrt{\gamma} T^{\mu 0}\right) = -\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\mu (n_\nu). \tag{6.1}$$

The geometric source term can be expanded as shown in [35] to give

$$-\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\mu (n_\nu) = \sqrt{\gamma} \left(\alpha T^{ik} K_{ik} - \alpha T^{j0} \partial_j \alpha\right). \tag{6.2}$$

Here $K_{\mu\nu} = -(\delta^\lambda_\mu + n^\lambda n_\mu)\nabla_\lambda(n_\nu)$ is the extrinsic curvature. It is orthogonal to the time-like normal such that $(K_{\mu\nu} n^\mu = 0_\nu)$. Thus its temporal components are $K_{i0} = \beta^i K_{ij}$ and $K_{00} = \beta^i \beta^j K_{ij}$. The remnant initial data from the Einstein Toolkit simulation gives us access to the extrinsic curvature quantities K_{ij}. These can then be copied into the AMReX mesh along with the fluid and other spacetime quantities. In truly stationary spacetime cases, this source term can be written in terms of the energy momentum tensor and spacetime metric components only, as

$$-\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\mu (n_\nu) = \sqrt{\gamma} \left(\frac{1}{2} T^{ik} \beta^j \partial_j \gamma_{ik} + T^j_i \partial_j \beta^i - \alpha T^{j0} \partial_j \alpha\right). \tag{6.3}$$

6.4 Remnant Initial Data

The data from the Einstein Toolkit simulation, roughly 8ms (or 1600 in geometric units) after merger, is taken to be approximately hydrostationary and used as initial data for the radiation hydrodynamic simulation of the remnant within the AMReX code. It is centred at coordinates $(x, y, z) = (0, 0, 0)$. The initial quantity profiles are shown in figures 6.1, 6.2, 6.3, and 6.4. These show slices through $z = 0$ and $y = 0$. The cell widths on the coarsest grid level are 36 in geometric units. Eight AMR levels are used in total, each with a refinement factor of two. The computational grid created using the AMReX framework is illustrated in figure 6.5, showing a slice through $z = 0$. We can see that the remnant consists of a dense and hot central object, which is highly spatially resolved. It has a rotating torus of surrounding cooler and lower density material. The remnant and torus are neutron rich (indicated by the low electron fraction), within a cold and rarefied background atmosphere of symmetric nuclear matter.

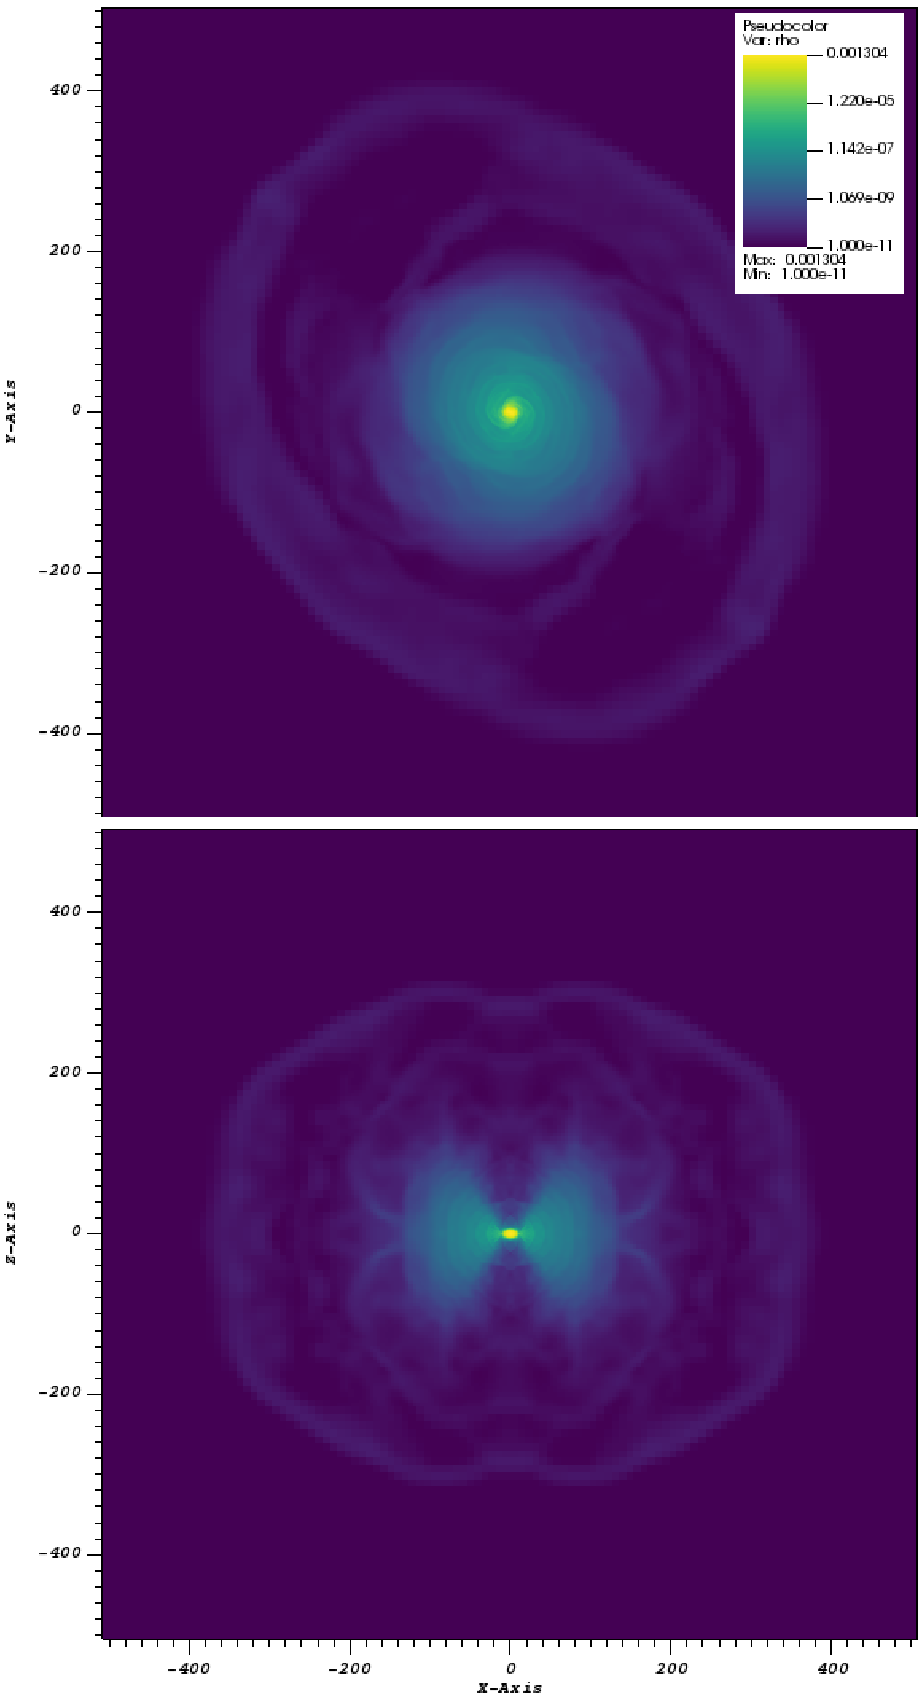

Figure 6.1: Remnant initial log-scale rest mass density profile. Top panel: $z = 0$ slice. Bottom panel: $y = 0$ slice.

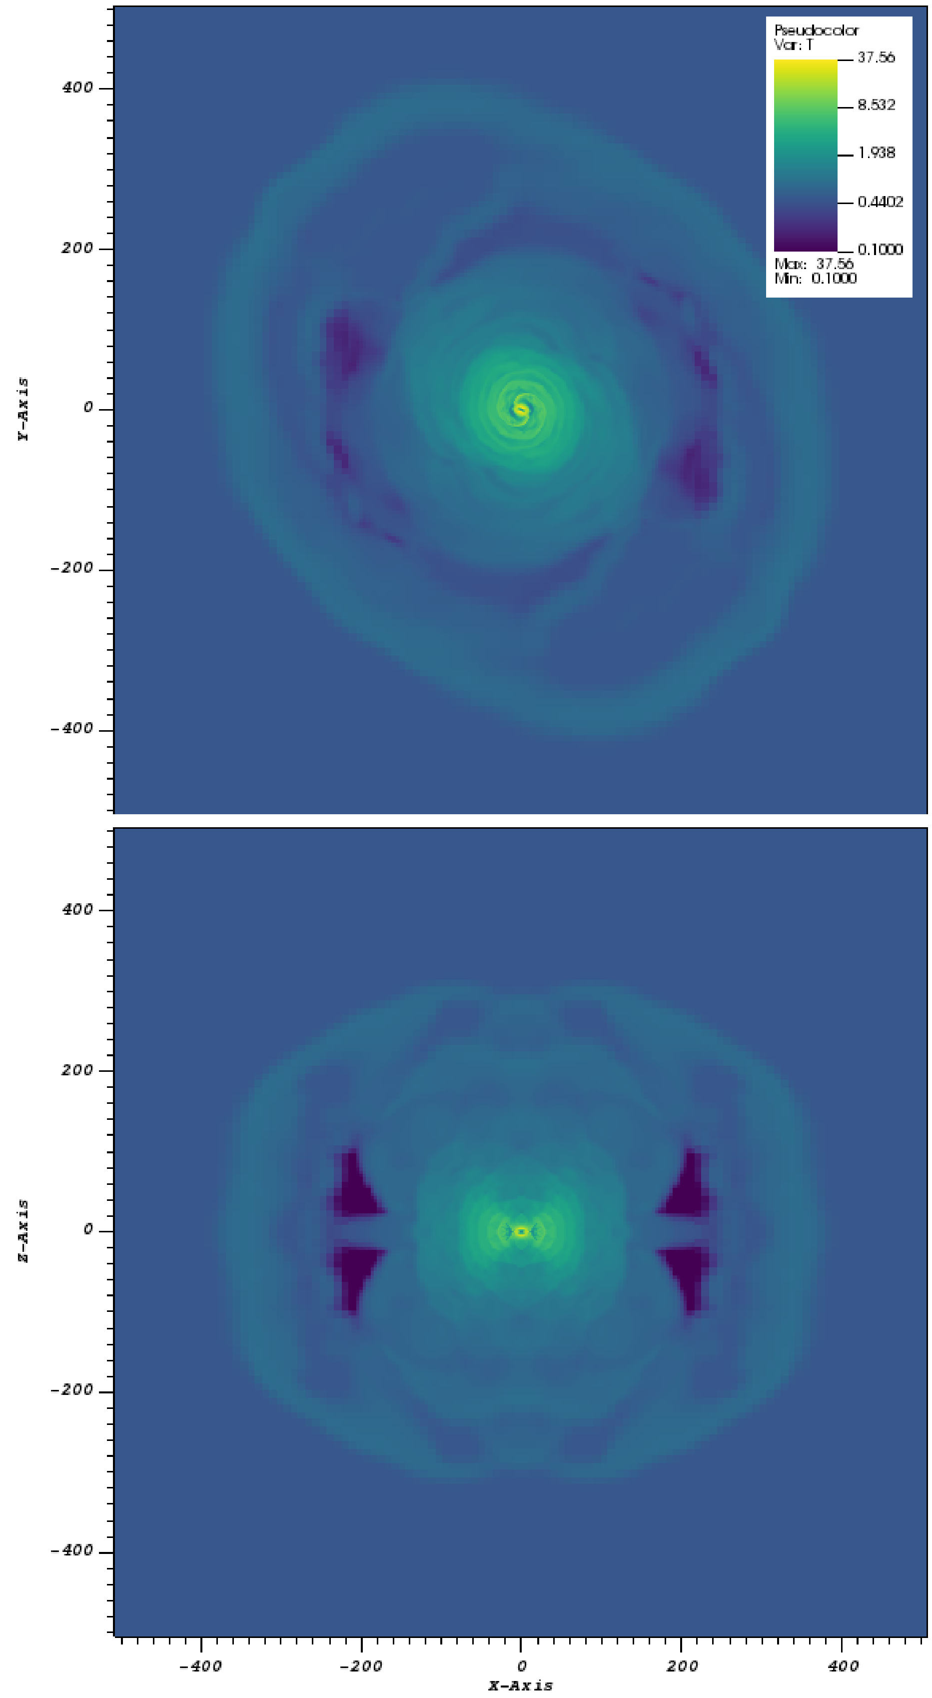

Figure 6.2: Remnant initial log-scale fluid temperature profile. Top panel: $z = 0$ slice. Bottom panel: $y = 0$ slice.

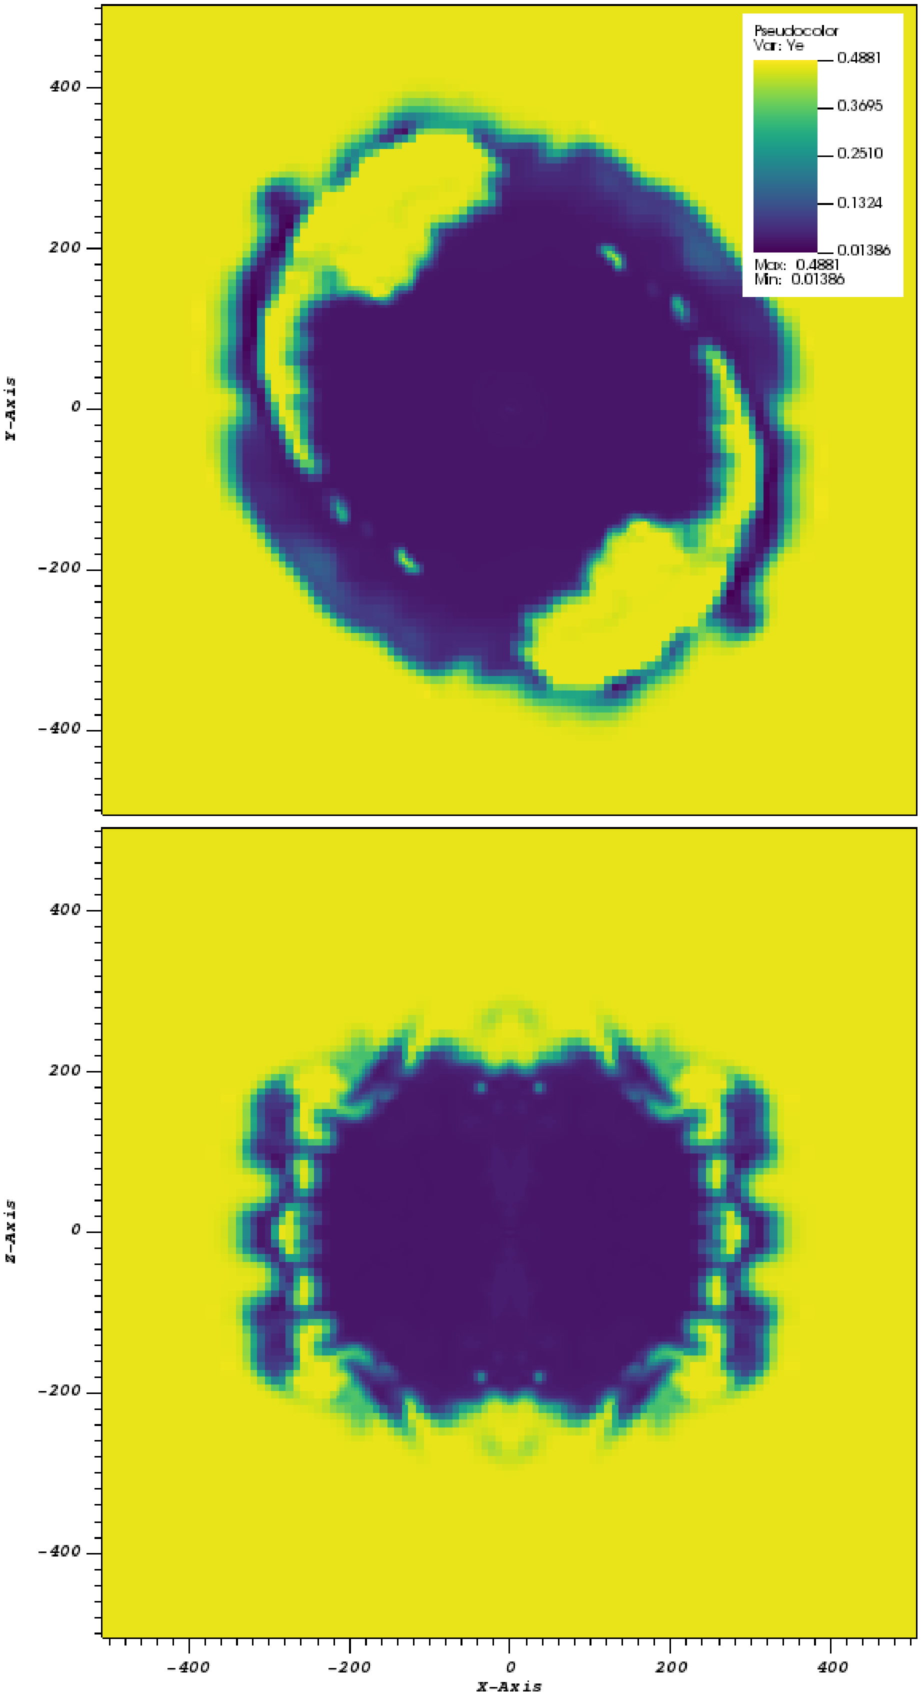

Figure 6.3: Remnant initial electron fraction profile. Top panel: $z = 0$ slice. Bottom panel: $y = 0$ slice.

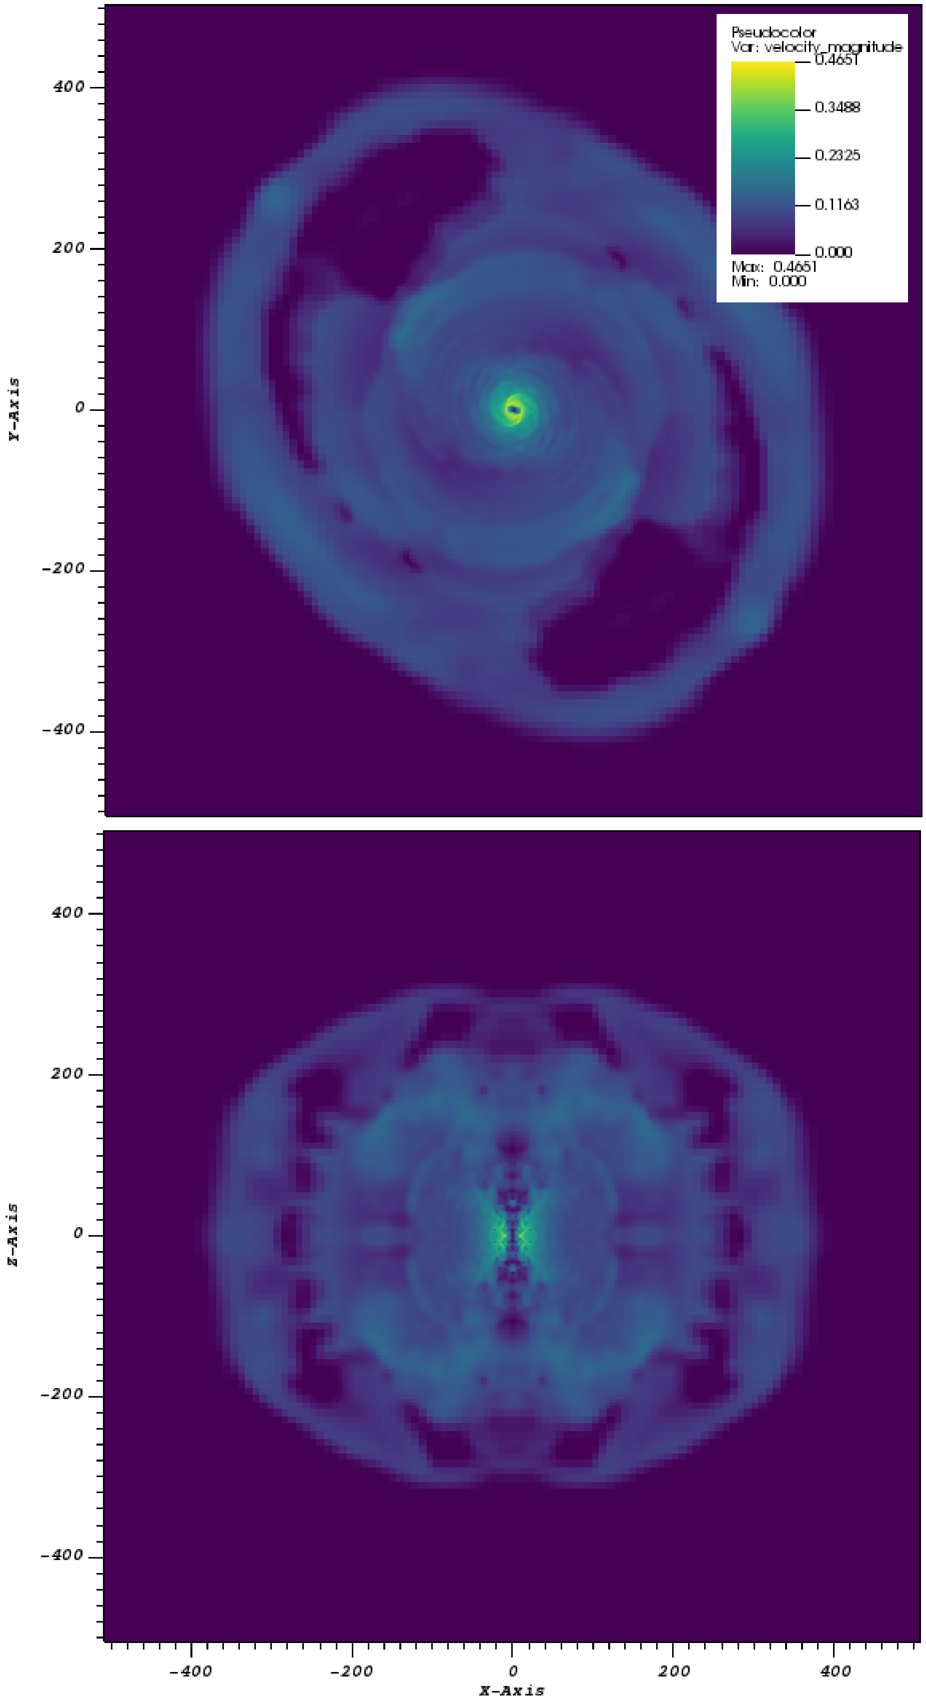

Figure 6.4: Remnant initial fluid velocity magnitude profile. Top panel: $z = 0$ slice. Bottom panel: $y = 0$ slice.

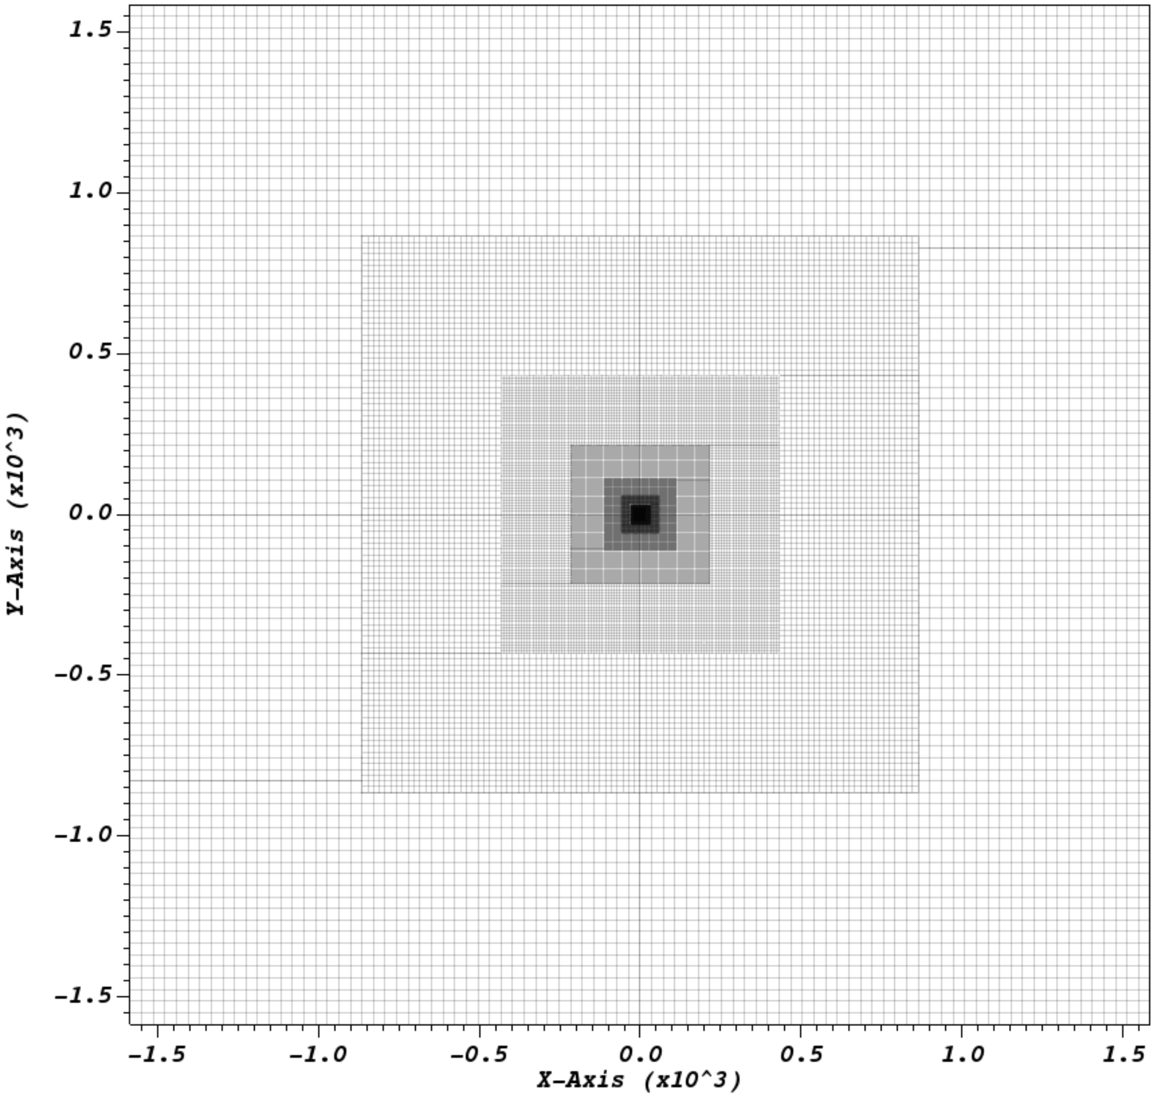

Figure 6.5: The remnant simulation employs a numerical mesh with seven levels of factor-2 refinement. It uses the box-in-box approach. The central compact object is highly spatially resolved compared to the background atmosphere.

6.5 Remnant Evolution

The remnant initial state has been evolved for a further $t = 8 \times 10^3$ in geometric units, which corresponds to around 40 ms. The evolved rest mass density, fluid temperature, electron fraction, and fluid velocity magnitude profiles are shown in figures 6.6, 6.7, 6.8, and 6.9 using $z = 0$ and $y = 0$ slices. Zoomed in profiles for fluid temperature and velocity magnitude are shown in figure 6.10.

We can see that the central compact object and its peak rest mass density have been maintained. The torus of surrounding material is comparable but has qualitatively evolved – it is now "puffed up" in the z direction. The outer shell of material has expanded, with significant outflow away from the poles of the remnant. The peak temperature has increased from around 38 MeV to about 70 MeV, whilst the peak fluid velocity has decreased from 0.46 to 0.24 in geometric units. The electron fraction has evolved to substantially follow the fluid, with the most noticeable changes being the visible ejection of symmetric matter along the poles. The velocity of the (low density) matter ejected along the poles is high compared to the rest of the matter in the system, but at less than 25% of the speed of light is still not extremely relativistic.

The qualitative behaviour of the remnant post-merger is not dissimilar to existing cases in the literature such as those shown by [26]. In particular, the ejection of matter from the system is expected to be crucial in the potential formation of a jet and a kilonova. In this scenario, the matter along the poles needs to be rapidly ejected to allow the jet to form. However, a fully self-consistent simulation of jet formation requires the inclusion of magnetic fields, so any conclusions drawn here are necessarily incomplete.

The detailed evolution of the temperature and velocity profiles will depend on the local features of the simulation, the transport of energy-momentum, the precise microphysics included, and the numerical accuracy. The temperatures seen here are within the range of those seen in the literature [26], but the peak temperatures are on the high side. The qualitative features (cold spots within the "cores" of the original merged neutron stars, the surrounding hot region that will see the largest change in the radiation opacity, and the spiral arm features involving internal shocks due to the outflowing matter) are comparable to other simulations [26].

With the qualitative features now shown for a purely hydrodynamical simulation, I can move on to introducing a radiation field.

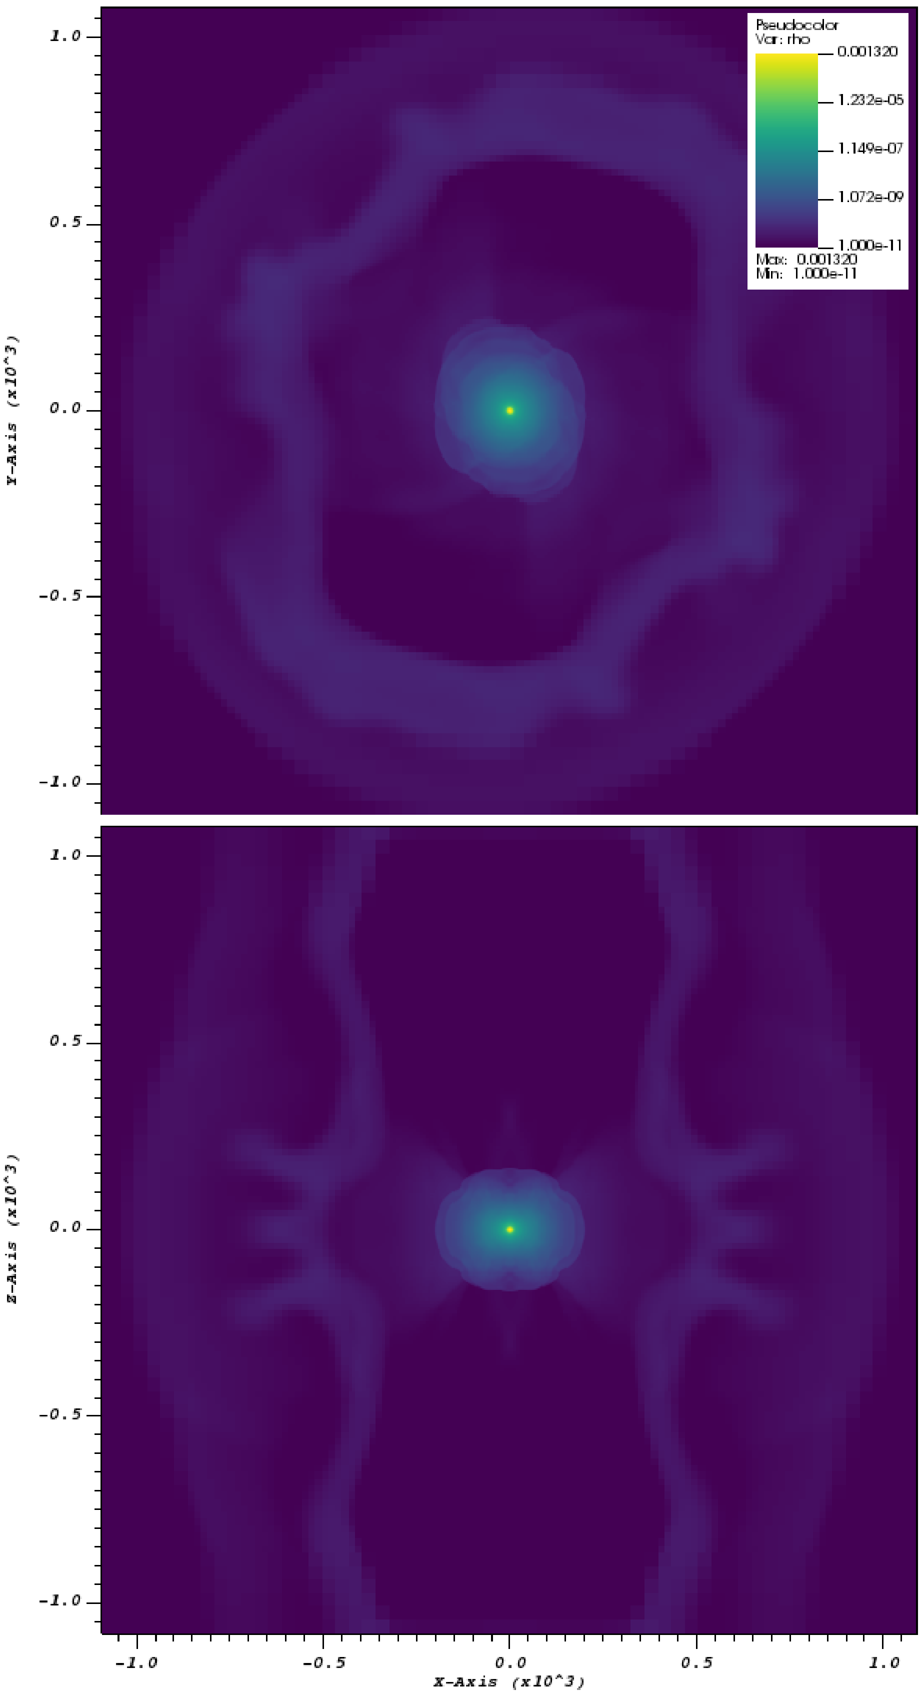

Figure 6.6: Log-scale rest mass density profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

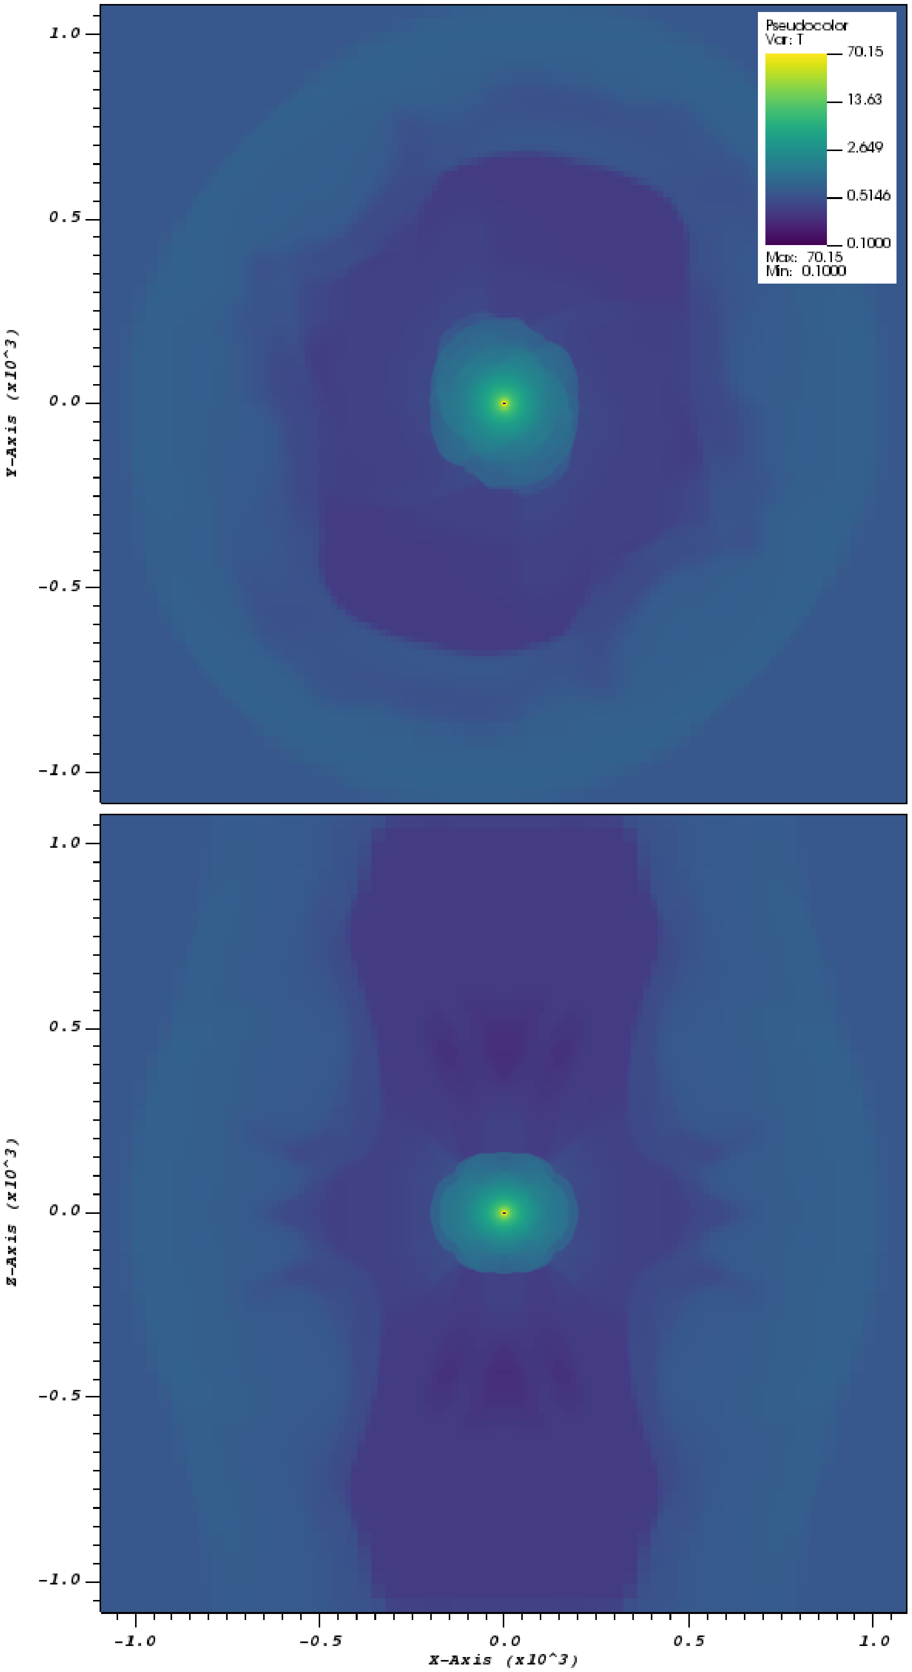

Figure 6.7: Log-scale fluid temperature profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

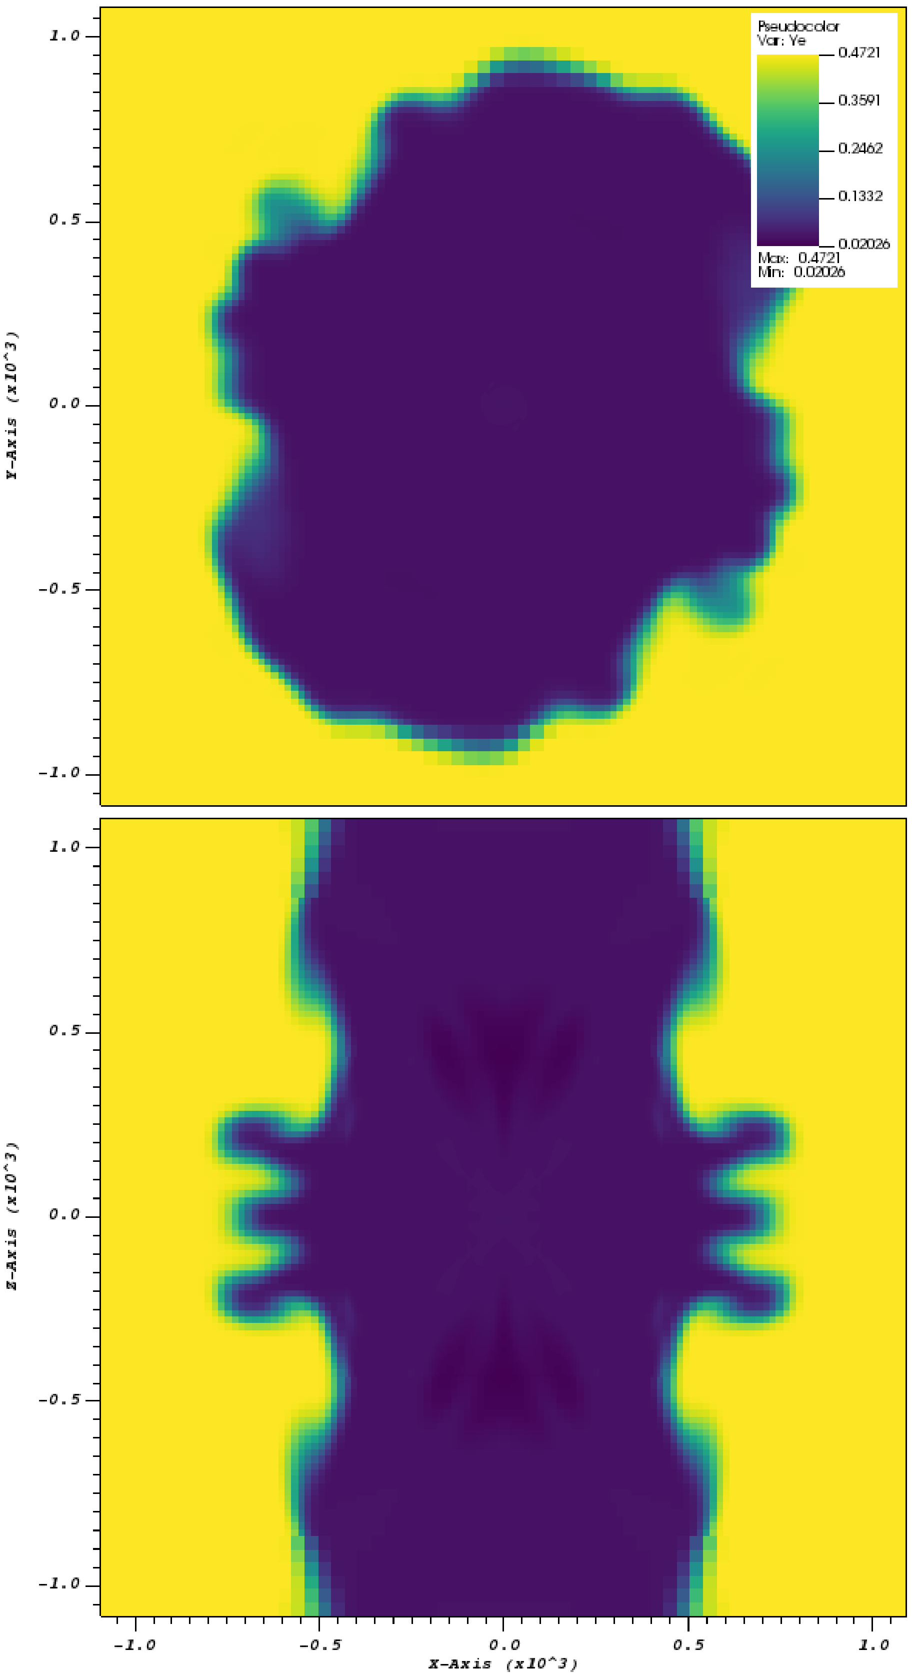

Figure 6.8: Electron fraction profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

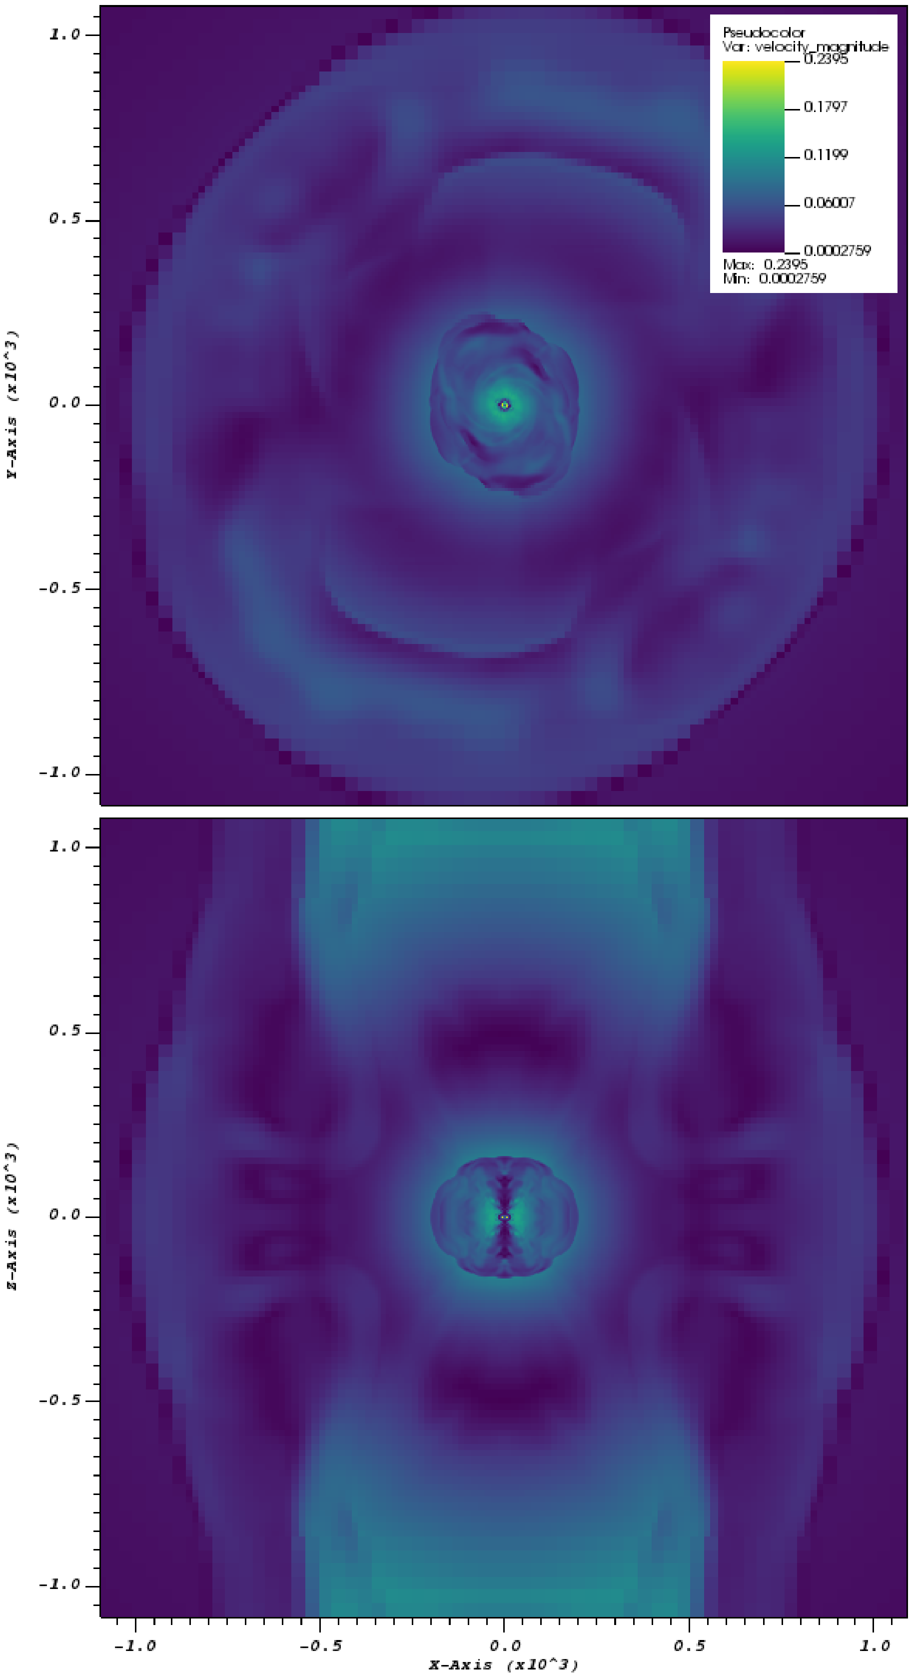

Figure 6.9: Fluid velocity magnitude profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

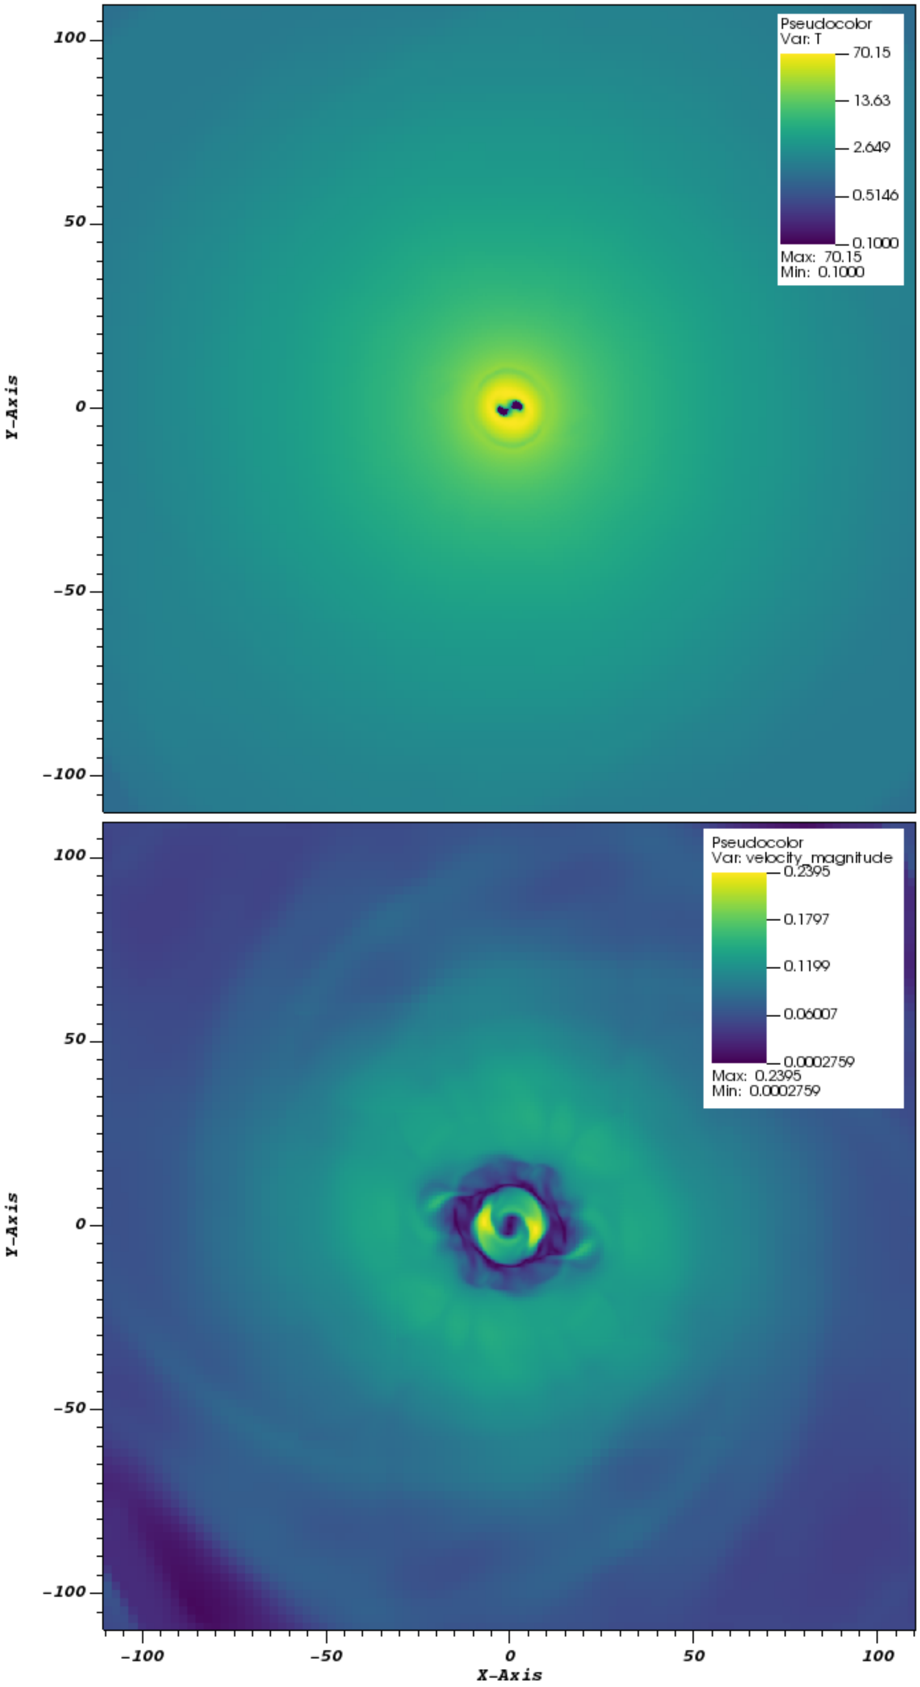

Figure 6.10: Slices through $z = 0$ showing zoomed-in evolved remnant profiles after 40 ms. The top panel shows log-scale fluid temperature. The bottom panel shows fluid velocity magnitude.

6.6 Radiation Initialisation

The remnant initial condition provides us with both the spacetime and fluid profiles. I aim to avoid modifying these when introducing a radiation field. As discussed with the TOV model in section 5.5 I have the options of adding a radiation field in local thermodynamic equilibrium with the fluid temperature, or adding the radiation pressure to balance against the spacetime forces - constructing a mechanical equilibrium. Having a single radiation field and two separate conditions to satisfy appears to be an over-constrained problem. From numerical experiments I have seen that a non-mechanically stable radiation field will re-arrange itself towards a stable distribution without issues (such as violent waves or overshooting expected limits). A non LTE radiation field in high opacity regions will quickly approach the LTE state, and in low opacity regions will slowly approach the LTE state. A practical solution is to satisfy only the most demanding and problematic condition in each region, with the expectation that the unsatisfied conditions will not cause significant qualitative changes to the bulk evolution. Therefore I initialise the radiation field in LTE in regions with high opacity (greater than some threshold), and include no radiation field whatsoever in low opacity regions. Here the fluid is allowed to slowly emit radiation until it reaches LTE over a long duration, and the radiation field is allowed to smoothly redistribute to reduce any non-mechanical equilibrium effects.

In the high opacity region a radiation field is calculated with an energy density in LTE with the fluid's radiation emission. This is determined from the fluid temperature and is assumed to have had enough time to relax to the thermodynamic equilibrium state: $E = a_{rad}T^4$. I initialise this radiation field co-moving with the fluid, with a vanishing net radiation flux in the fluid frame, and thus an isotropic pressure tensor.

At the temperatures reached in this remnant this results in a radiation field with a small fraction of the fluid's internal energy density. It is likely that any energy exchange between the fluid and radiation would not be overly disruptive to the remnant's hydrodynamic stability, whilst still having some noticeable effect on its evolution. Some disruption is to be expected as I am adding in the radiation energy on top of the fluid energy already present in the domain, breaking any strict equilibrium between the total pressure gradient and the background spacetime.

It is preferable to have small radiation energy amplitudes away from the central merger object to prevent any atmospheric fluctuations or artefacts significantly affecting the radiation evolution.

6.6.1 Opacity Implementation

I use the frequency averaged functional form opacity formula from [42] introduced in section 2.8.3,

$$\chi = 1.928 \times 10^{17} \; Y_e^2 \; \rho^2 \; T^{-7/2}. \tag{6.4}$$

I do not include scattering opacities in the simulation, as they are many orders of magnitude smaller than the absorption opacity, and would have a negligible effect on the evolution.

With this opacity formula the opacities in the merger simulation range between $6.8 \times 10^{-8} L^{-1}$ in the background atmosphere, and $3.0 \times 10^{10} L^{-1}$ in the high density core region, where L is the appropriate length scale. The threshold above which radiation is initialised in LTE with the fluid is therefore chosen as $1 \times L^{-1}$. In the most extreme case at the centre of the remnant, the large opacity scales the coupling source term, and would require an explicit numerical time step of around 10^{-11} (in geometric units) to accurately resolve. The LTE overshoot prevention method allows the use of much larger time steps dependent only on the grid resolution and physical wave-speeds involved.

The artificial atmosphere required for stable numerical hydrodynamics, which avoids having regions of near-zero mass density within the computational domain, is prevented from artificially increasing the opacity. This could potentially cause energy and momentum exchange issues. In regions where the artificial atmosphere is used ($\rho \lesssim 10^{-10}$ in geometric units) the opacity is set to zero. Otherwise the opacity would remain artificially high, where I would expect it to decrease significantly when away from the remnant material.

6.6.2 Remnant Evolution with Radiation

The remnant initial state with the introduction of a radiation field as described above has been evolved for a further $t = 8 \times 10^3$ in geometric units, which corresponds to around 40 ms. The evolved rest mass density and fluid temperature profiles are shown in figures 6.13 and 6.14, with zoomed in profiles shown in figure 6.15. The evolved opacity, radiation energy density, and radiation flux magnitudes are shown in figures 6.16, 6.17, and 6.18.

Compared to the evolution of the purely hydrodynamic case shown in figures 6.6-6.10

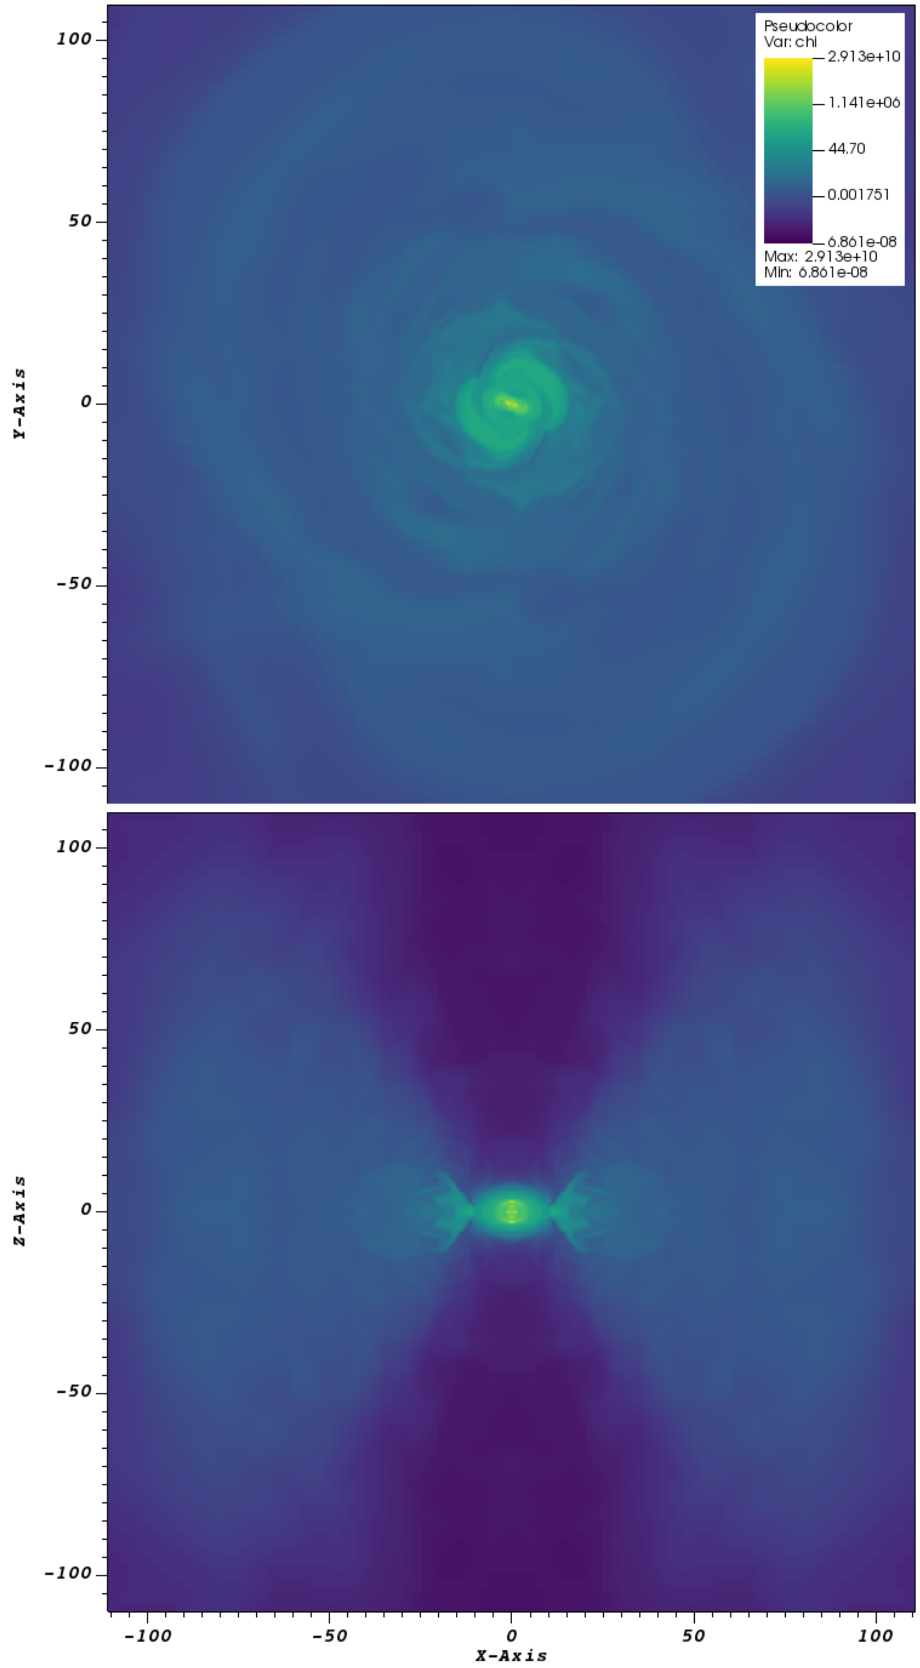

Figure 6.11: Remnant log-scale initial opacity profile. Top panel: $z = 0$ slice. Bottom panel: $y = 0$ slice.

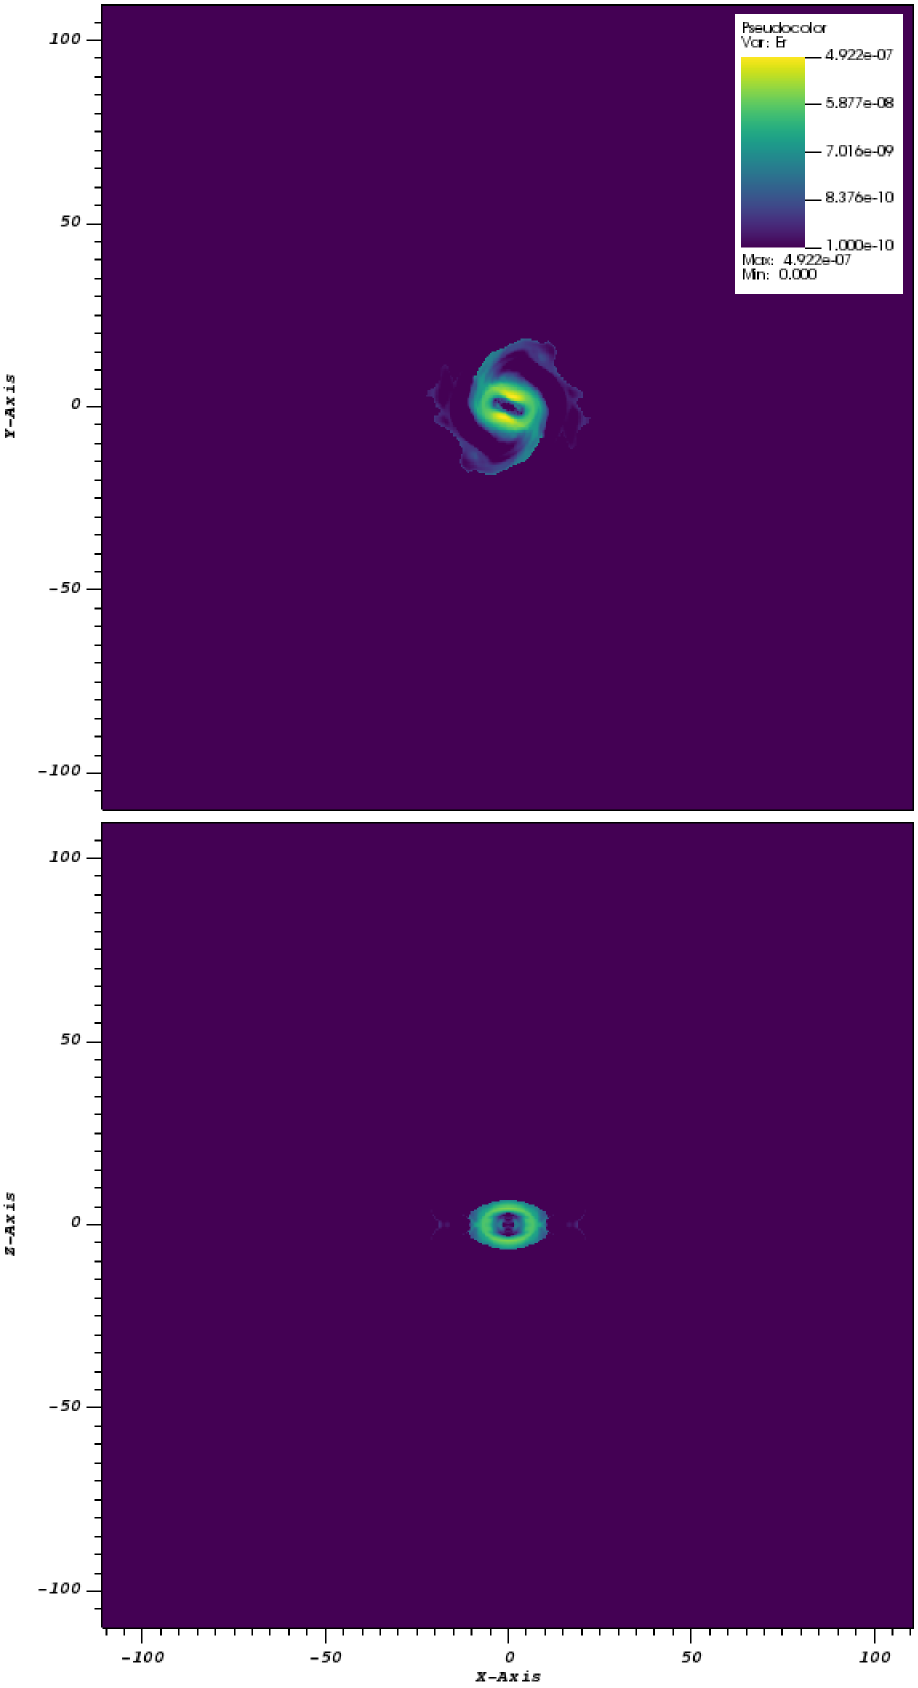

Figure 6.12: Remnant log-scale initial radiation energy density profile. Top panel: $z = 0$ slice. Bottom panel: $y = 0$ slice.

we can see that the core, high density neutron star remnant is qualitatively similar, but the surrounding material is substantially different. The outer layers of material have rapidly dispersed, and all quantities have been “smoothed out”. This seems to be driven by the evolution of the radiation field: the exchange of energy-momentum with the radiation, leading to isotropic emission and transfer of energy, has a substantial impact on the lower density (and hence lower opacity) matter.

In the profile shown in 6.15 we see that the radiation has not yet managed to smooth out the cold spots at very high density (and high opacity) that correspond to the original neutron star cores. The peak temperature has dropped (it is around 48 MeV, compared to the 70 MeV reached in the purely hydrodynamic case), but the timescale on which radiation can evolve, propagate, and diffuse in this spacetime region is considerably slower. The peak opacity has decreased by a factor of around 300. This is driven in part by the change in temperature, but is also sensitive to low-density features of the equation of state. This change in the opacity is however not sufficient to accelerate the diffusion of radiation into the core. The peak radiation energy density has increased from around 5×10^{-7} to around 1.3×10^{-6}. This corresponds very closely with the increase in fluid temperature. The shape of the radiation energy density field also corresponds closely with the fluid temperature profile. This suggests that the radiation field has remained near to local thermodynamic equilibrium with the fluid, as is expected in the high opacity central region where they are strongly coupled. The radiation flux profiles highlight the nearly isotropic behaviour of the radiation fields, but do show how the high opacity regions inhibit the propagation of radiation energy-momentum.

6.7 Low Advection Transformation

The rotational behaviour of the hydrodynamical evolution of the remnant is strong, and dominates over the isotropic behaviour in the purely hydrodynamic case. As I am simulating on a Cartesian mesh, we have seen in chapter 5 that the motion of a fluid with respect to the mesh or the coordinates can lead to numerical inaccuracy. I have used the coordinate transformation introduced in section 5.6.2 to minimize the apparent rotation of the matter with respect to the computational coordinates.

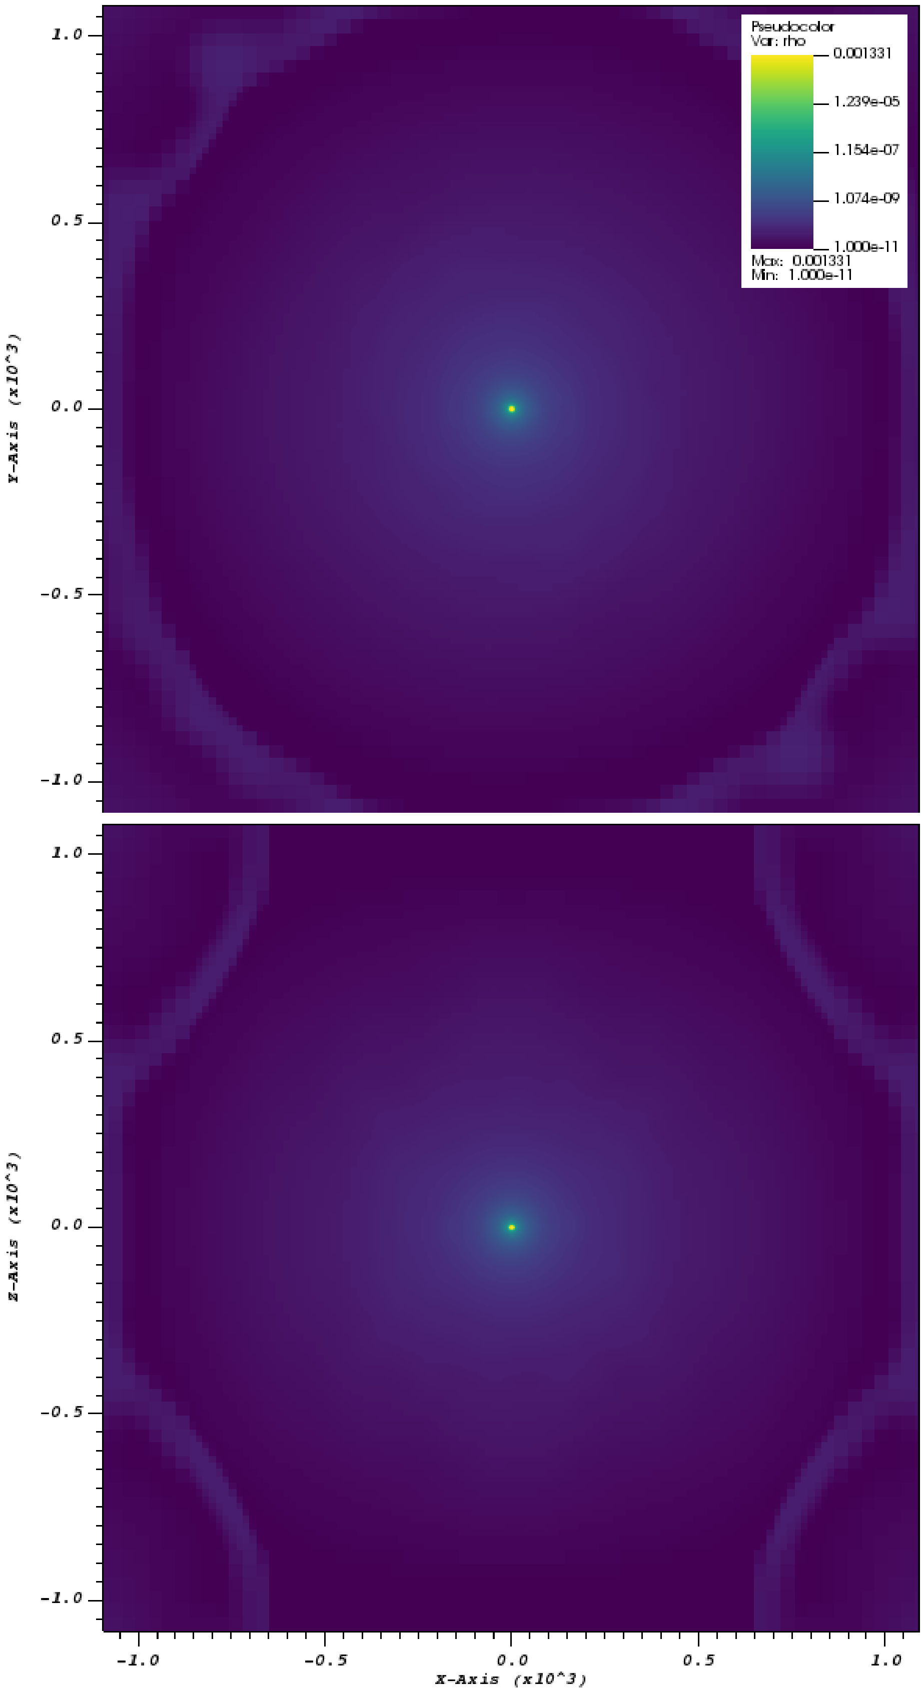

Figure 6.13: Log-scale rest mass density profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

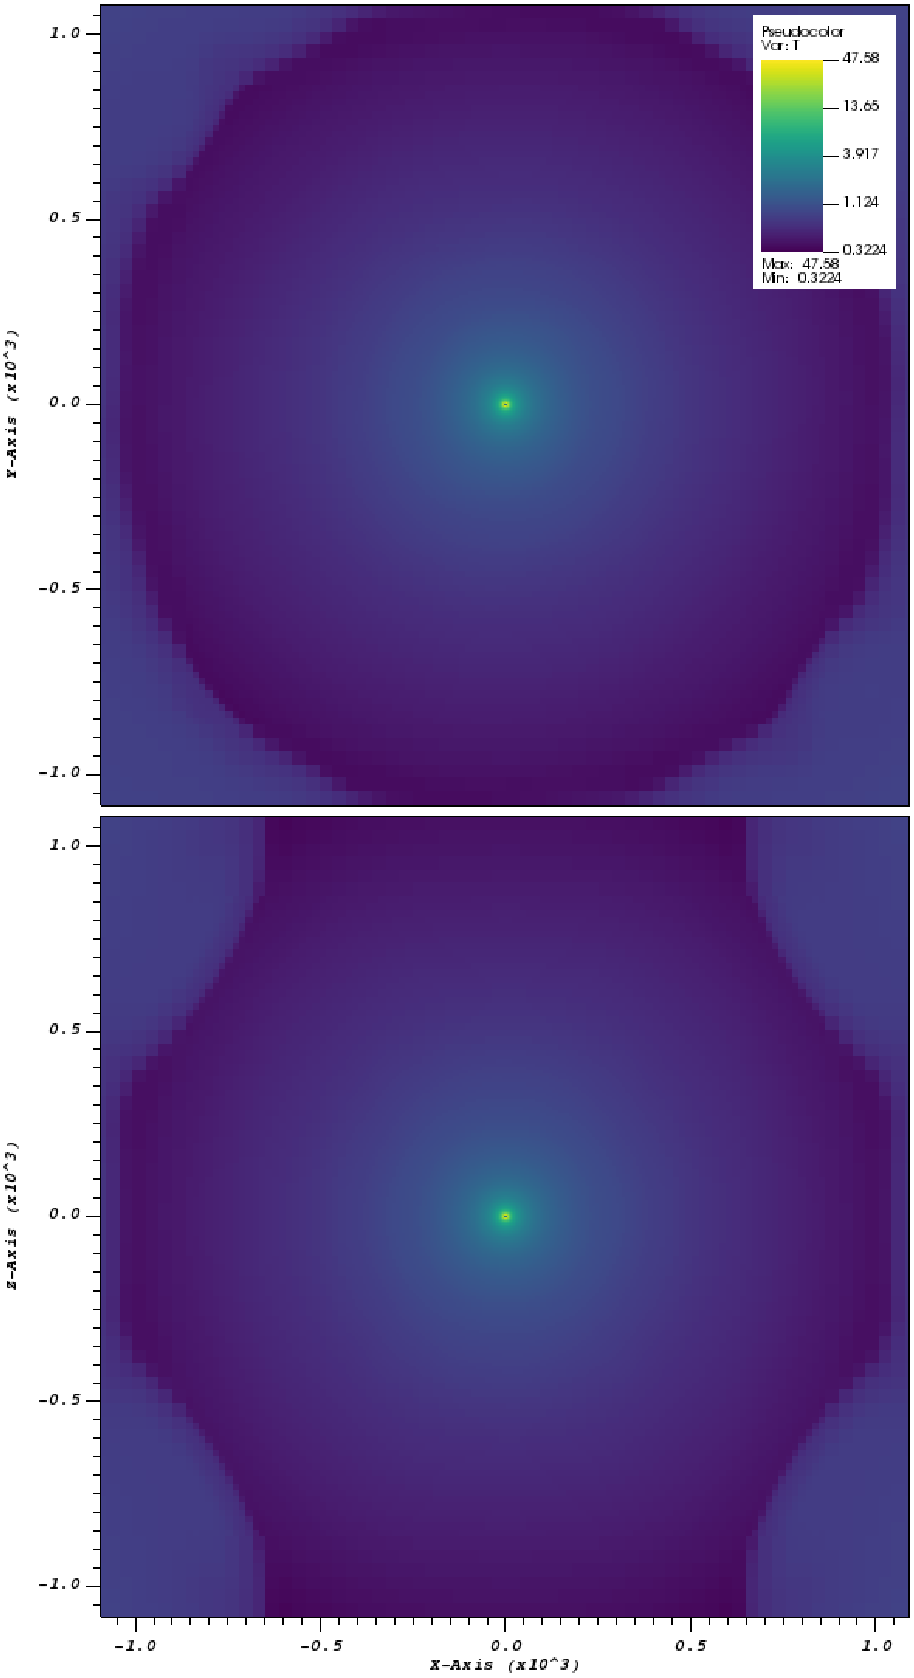

Figure 6.14: Log-scale fluid temperature profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

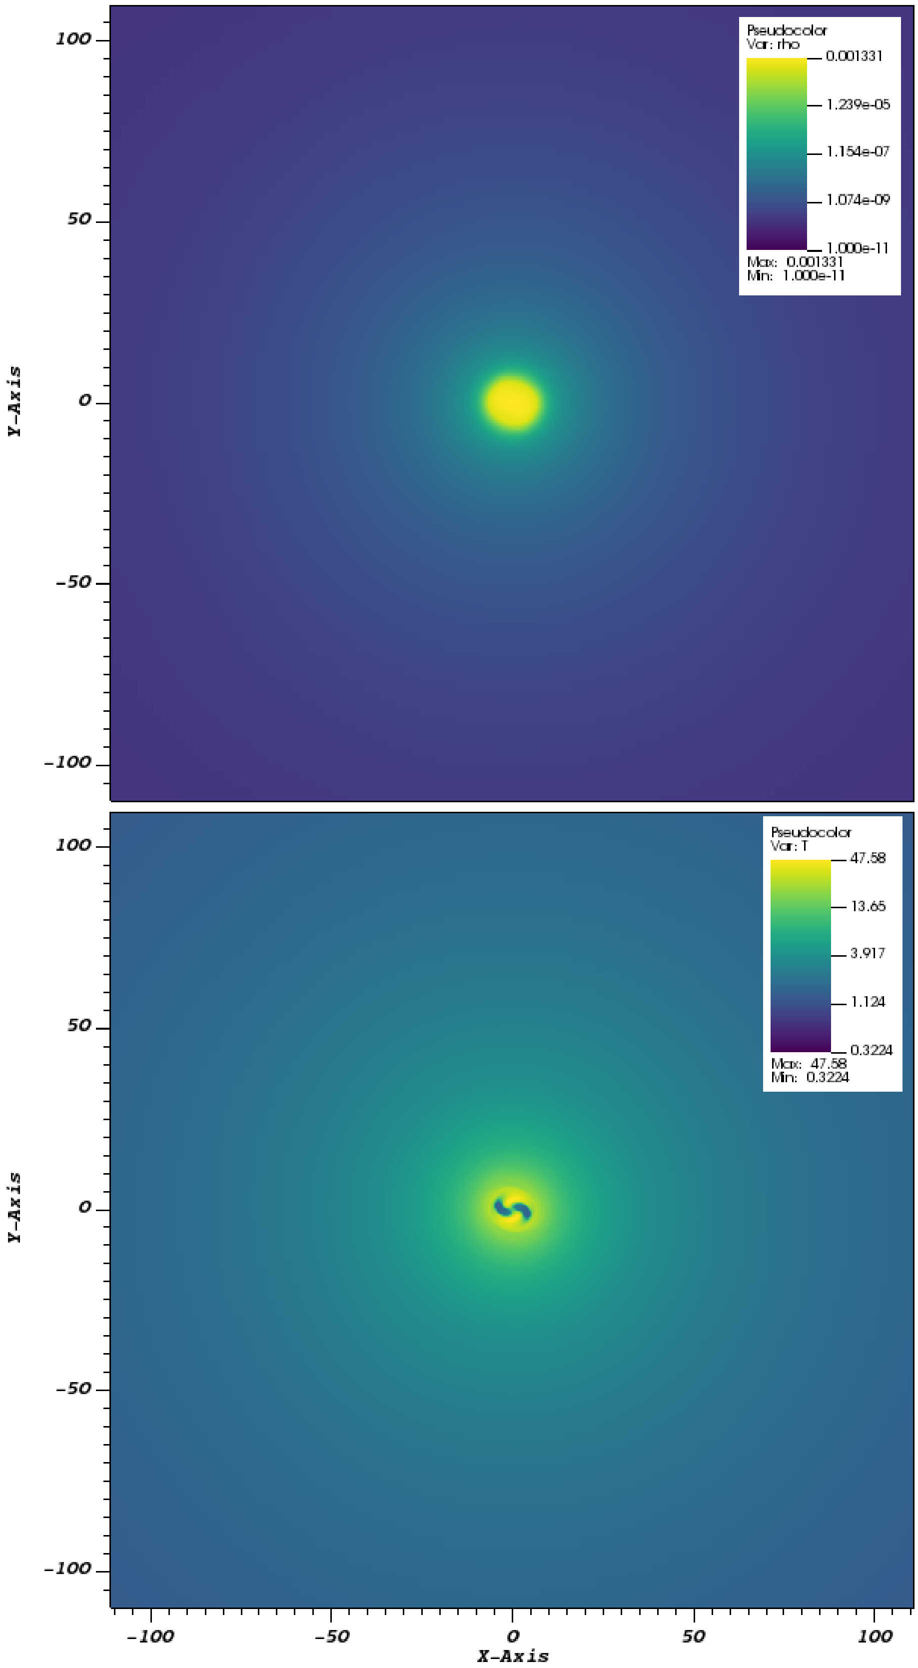

Figure 6.15: Slices through $z = 0$ showing zoomed-in evolved remnant profiles after 40 ms. The top panel shows log-scale rest mass density. The bottom panel shows log-scale fluid temperature.

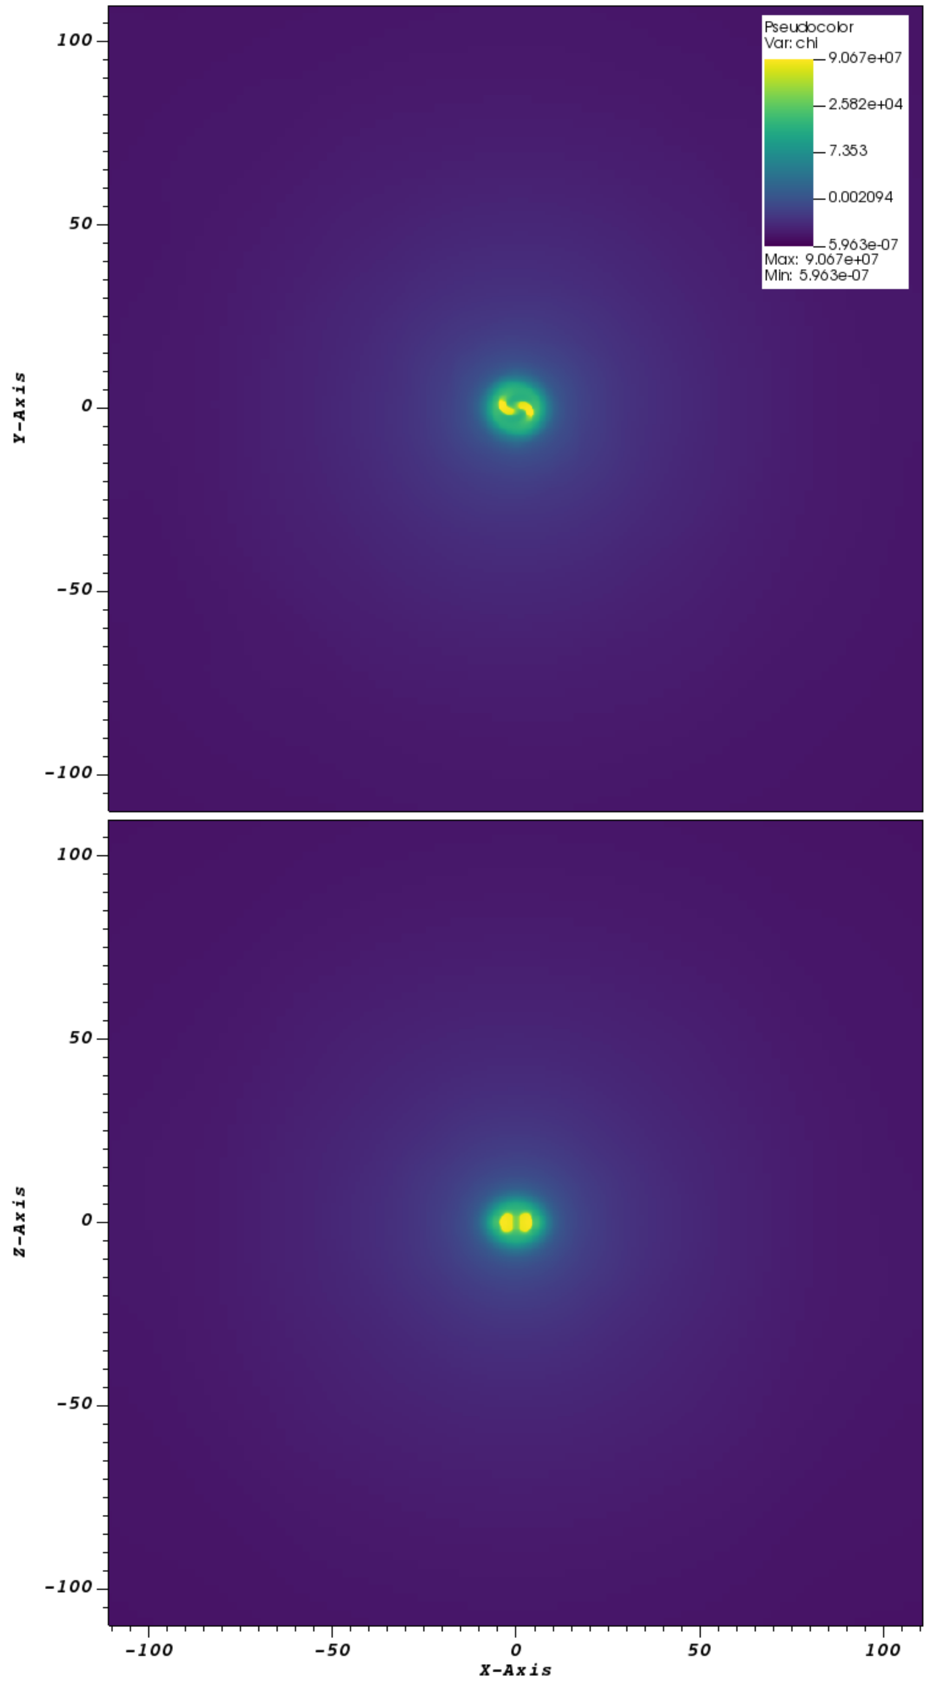

Figure 6.16: Log-scale opacity profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

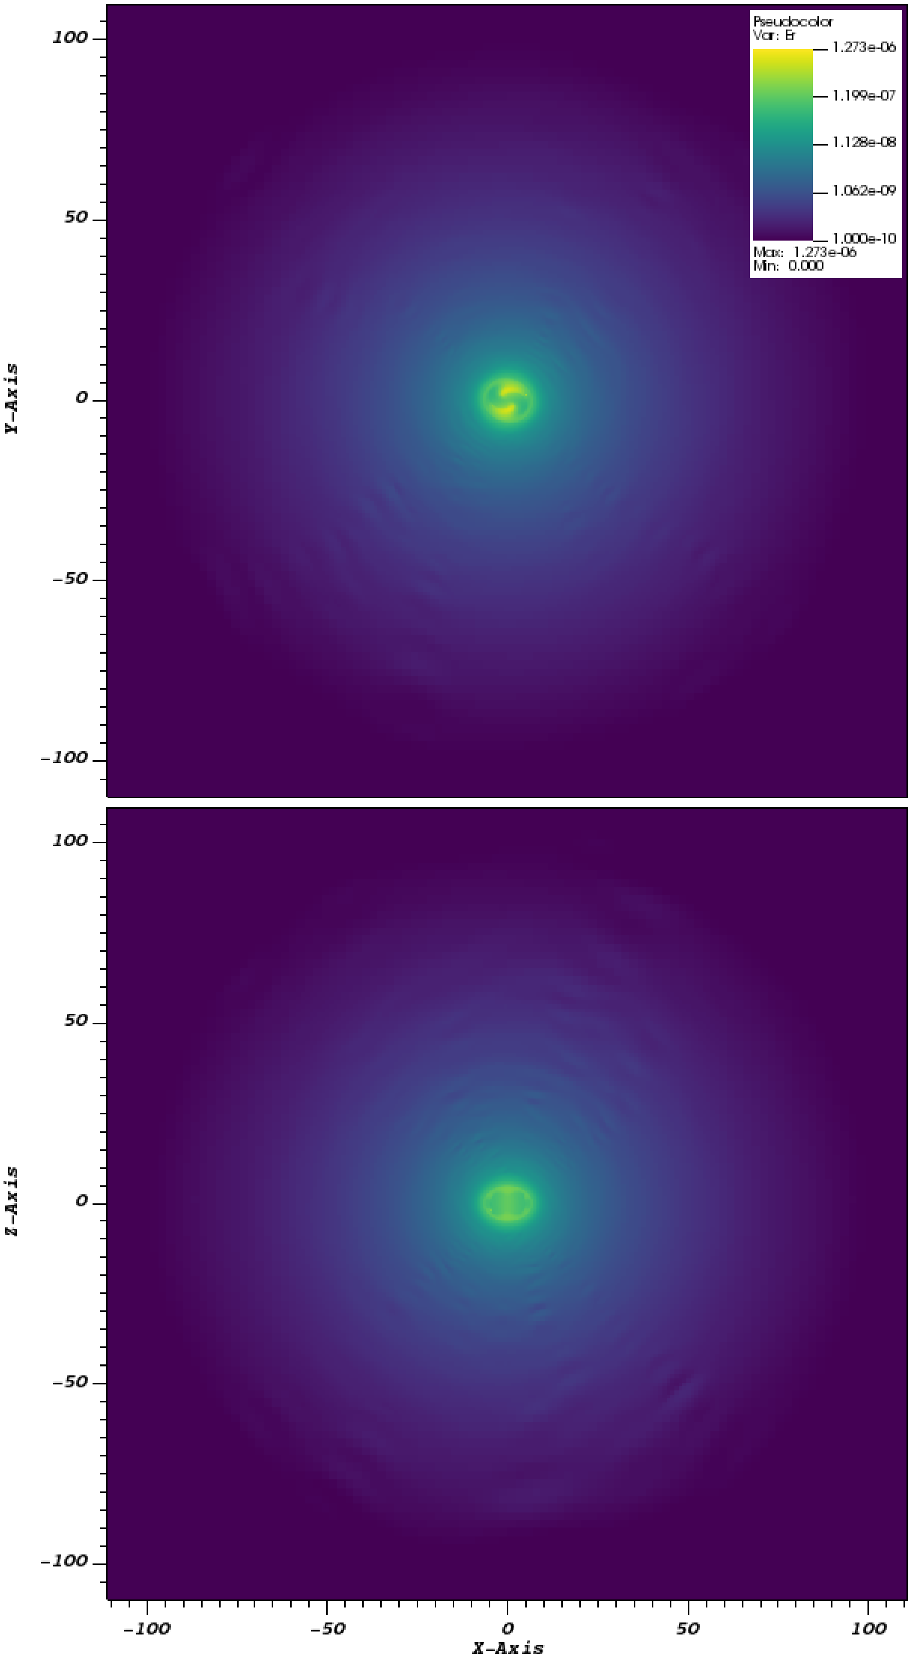

Figure 6.17: Log-scale radiation energy density profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

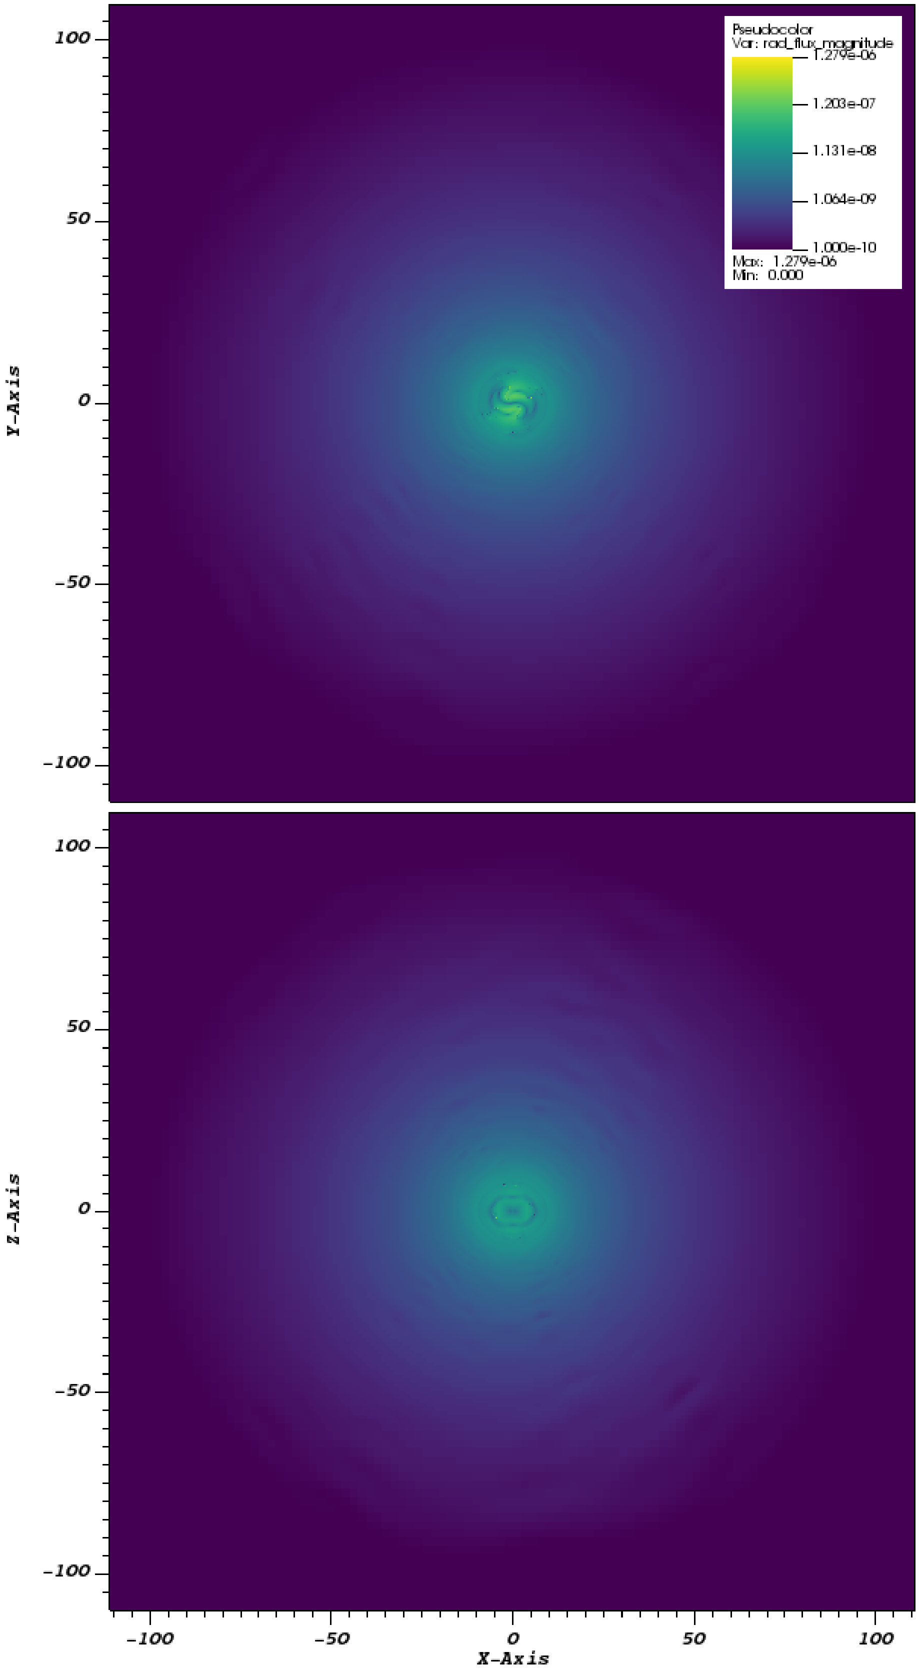

Figure 6.18: Log-scale radiation flux magnitude profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

Taking the coordinate transform

$$\Lambda^{\mu}_{\alpha} = \begin{pmatrix} 1 & \lambda^{i} \\ 0_{a} & \delta^{i}_{a} \end{pmatrix} \tag{6.5}$$

I can transform the contravariant fluid 4-velocity and spacetime metric to a new coordinate system

$$u^{\mu} = \Lambda^{\mu}_{\alpha}\overline{u}^{\alpha}, \tag{6.6}$$

$$g^{\mu\nu} = \Lambda^{\mu}_{\alpha}\Lambda^{\nu}_{\beta}\overline{g}^{\alpha\beta}. \tag{6.7}$$

Here overlined quantities denote those in the Einstein Toolkit coordinate system.

I can now obtain the new transformed spacetime quantities and fluid velocity

$$\alpha = \overline{\alpha}, \tag{6.8}$$

$$\beta^{i} = \overline{\beta}^{i} - \lambda^{i}, \tag{6.9}$$

$$\gamma^{ij} = \overline{\gamma}^{ij}, \tag{6.10}$$

$$v^{i} = \overline{v}^{i}. \tag{6.11}$$

The lapse function, spatial metric, and fluid velocities remain unchanged, however this transformation allows us to modify the shift vector. This results in a transformed contravariant 4-velocity of the form

$$u^{i} = W\left[\overline{v}^{i} - \left(\overline{\beta}^{i} - \lambda^{i}\right)/\overline{\alpha}\right]. \tag{6.12}$$

If we choose for the coordinate transform to use $\lambda^{i} = \overline{\beta}^{i} - \overline{\alpha v}^{i}$, then the new shift vector becomes $\beta^{i} = \overline{\alpha v}^{i}$, and the 4-velocity vanishes $u^{i} = 0^{i}$. This would have the effect of removing the advective flux terms from the evolution equations, in the remnant initial condition setup. As the fluid velocities evolve due to non-advective terms (such as geometric and reaction source terms and pressure gradients), this zero-advection result will no longer be exact.

As in the TOV rotation approach, I can construct a cylindrically symmetric torus shaped approximation $\lambda^{i} \approx \overline{\beta}^{i} - \overline{\alpha v}^{i}$, leading to $u^{i} \approx 0^{i}$. This reduces, on average, the size of the 4-velocity components. This thus reduces the impact of the advective flux terms. If we assume an approximately axially symmetric motion to the remnant, which is expected, then this advective flux reduction will continue to hold beyond the initial condition.

From the remnant initial data, I use the parameters $R_{peak} = 10$, $R_{out} = 60$, $\lambda_{peak} = 0.3$, and $Z_{height} = 50$. I can then construct the toroidal function λ with the rotational profile

$$\lambda(r_c, z) = \begin{cases} \dfrac{\lambda_{peak}}{R_{peak}} r_c \left(1 - \dfrac{|z|}{Z_{height}}\right), & r_c \leq R_{peak}, \\ \dfrac{-\lambda_{peak}}{R_{out} - R_{peak}} (r_c - R_{out}) \left(1 - \dfrac{|z|}{Z_{height}}\right), & r_c \in (R_{peak}, R_{out}), \\ 0, & r_c \geq R_{out}. \end{cases} \tag{5.22}$$

The components can then be obtained with

$$\lambda^x = \frac{\lambda}{r_c}(-y), \tag{6.13}$$

$$\lambda^y = \frac{\lambda}{r_c}(x), \tag{6.14}$$

$$\lambda^z = 0. \tag{6.15}$$

We must also apply this coordinate transform to the extrinsic curvature tensor $K^{\mu\nu} = \Lambda^{\mu}_{\alpha} \Lambda^{\nu}_{\beta} \overline{K}^{\alpha\beta}$.

6.7.1 Remnant Evolution with Low Flux Transformation

The remnant initial state with the flux-reducing coordinate transformation, as described above, has been evolved for a further 40 ms. The evolved rest mass density and fluid velocity magnitude are shown in figures 6.19 and 6.20. Zoomed in rest mass density, fluid temperature, and fluid velocity magnitude profiles are shown in figures 6.21 and 6.22.

As expected, the large scale qualitative behaviour is comparable to the original purely hydrodynamical simulation shown in figures 6.6-6.10. The key differences can be seen by looking at the details of the zoomed profiles, particularly the velocity magnitudes comparing figure 6.10 with figure 6.22. We can see that the coordinate transformation allows for the details of the spiral arm evolution to be much more sharply captured.

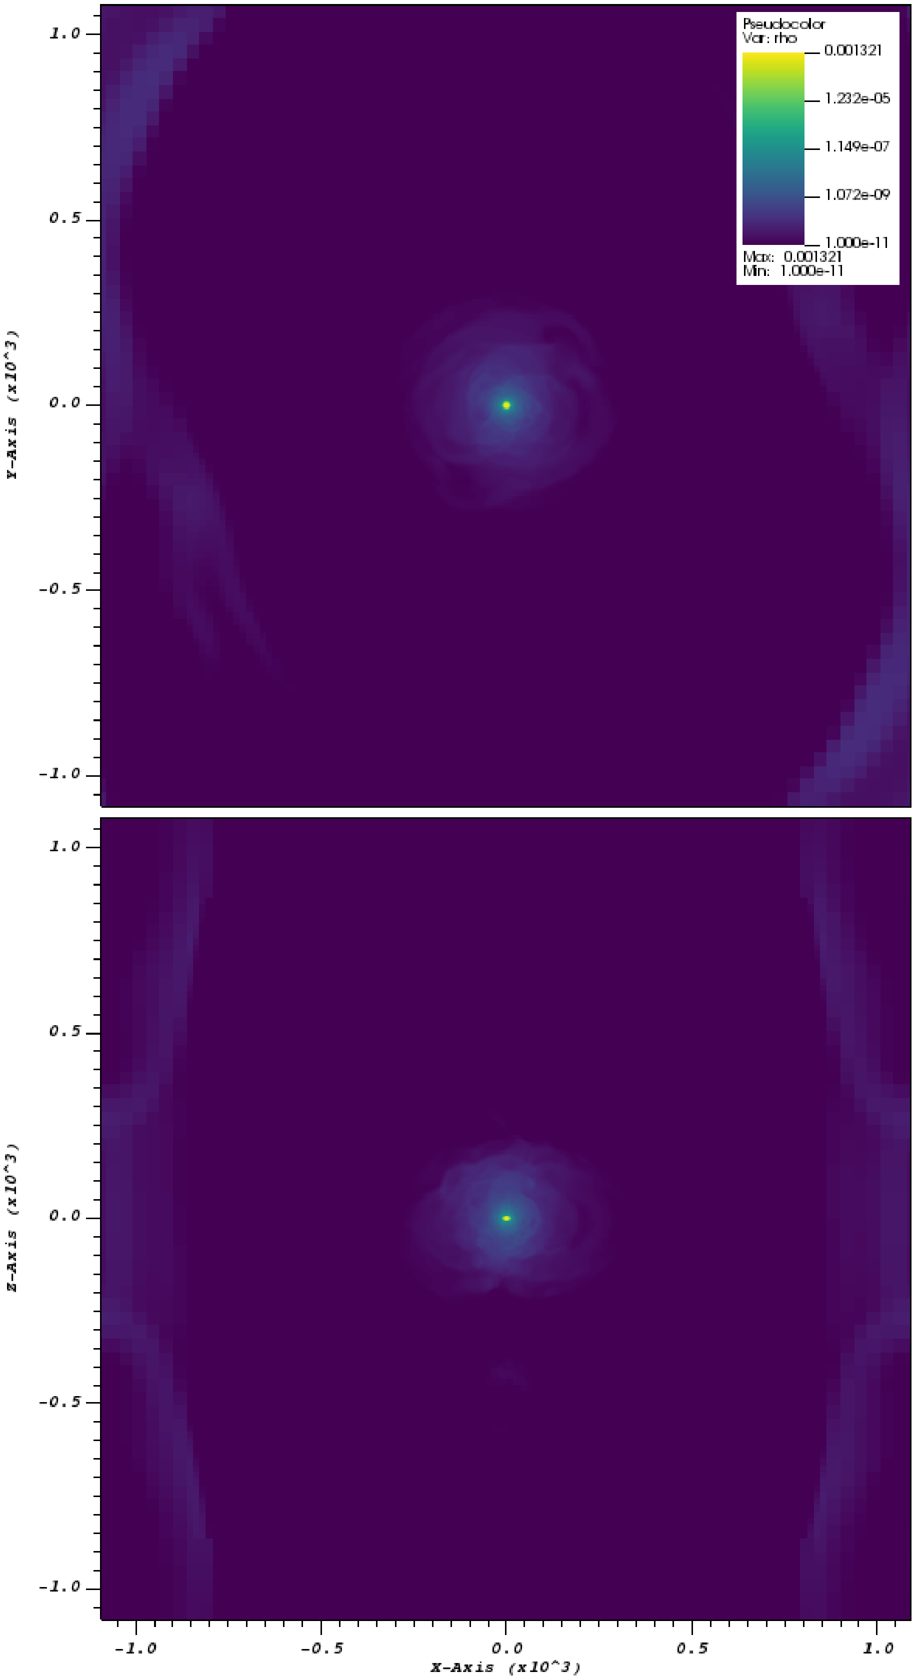

Figure 6.19: Log-scale rest mass density profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

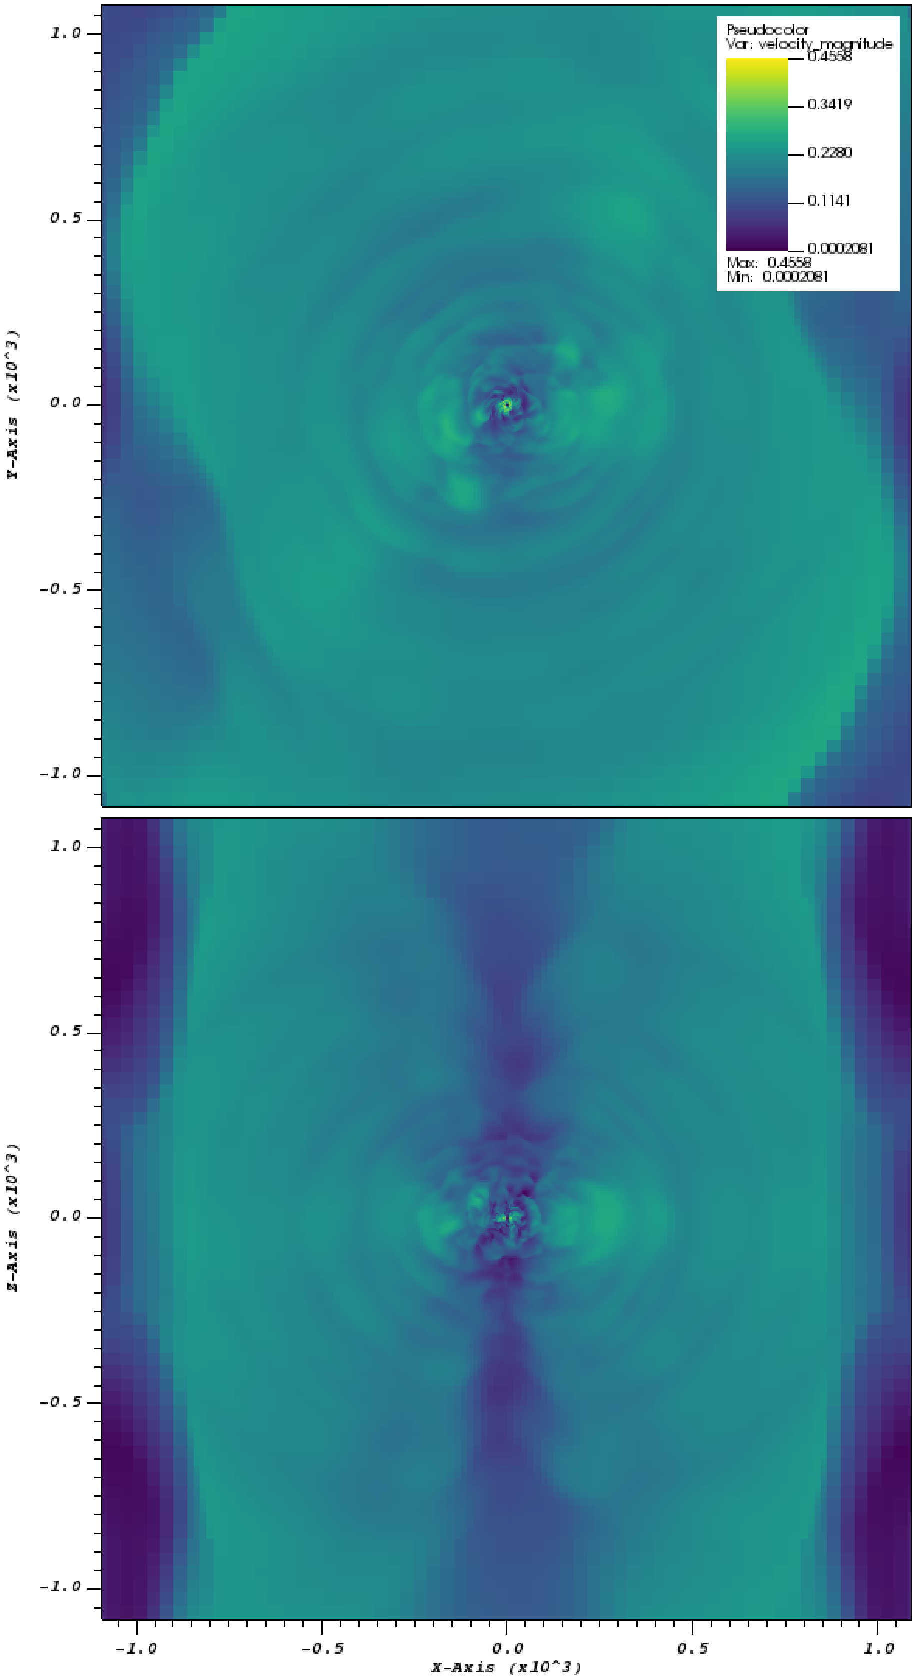

Figure 6.20: Log-scale fluid velocity magnitude profiles for the evolved remnant after a time of 40 ms. The top panel shows the $z = 0$ slice, and the bottom panel shows the $y = 0$ slice.

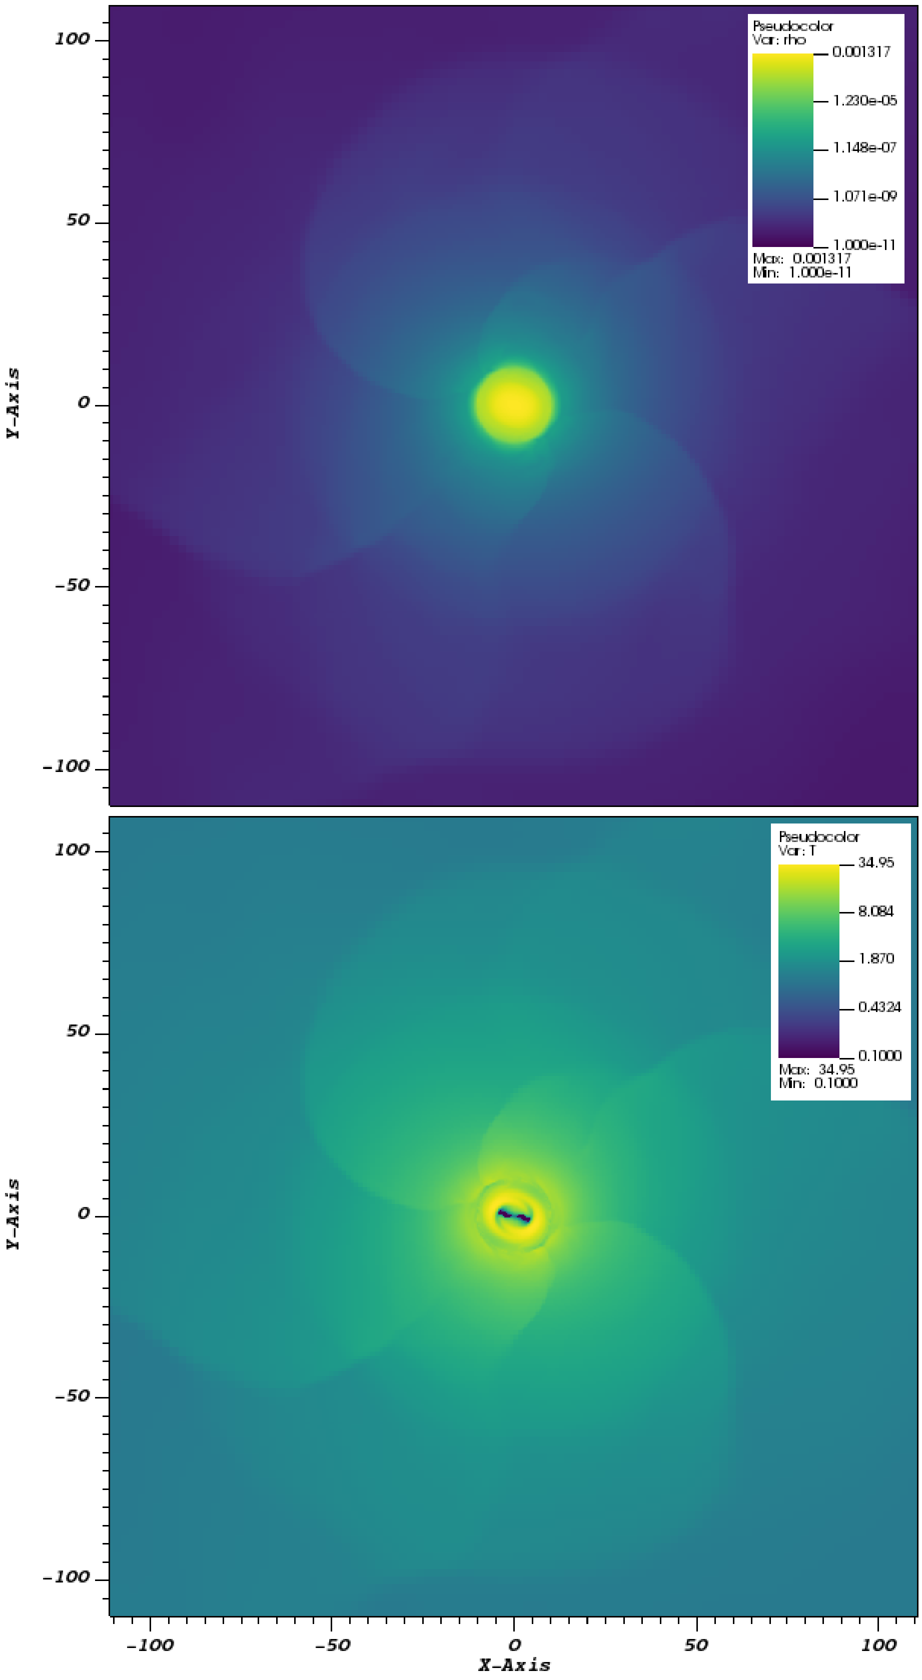

Figure 6.21: Slices through $z = 0$ showing zoomed-in evolved remnant profiles after 40 ms. The top panel shows log-scale rest mass density. The bottom panel shows log-scale fluid temperature.

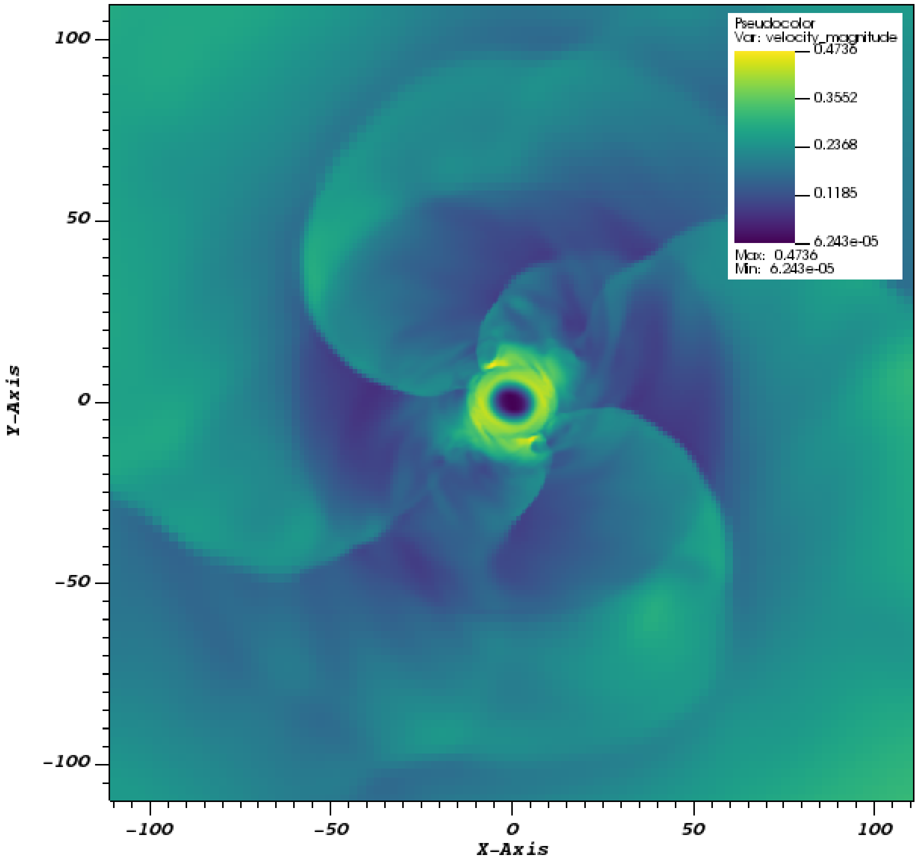

Figure 6.22: Slices through $z = 0$ showing the zoomed-in evolved remnant fluid velocity magnitude profile after 40 ms.

6.8 Remnant Central Rest Mass Density

As in the TOV models in chapter 5, and as particularly discussed in section 5.8, the evolution of the maximum density is useful for discussing the qualitative behaviour of the system. In this case, where the remnant solution is approximately but not exactly hydrostationary, there is no reason to believe that the true solution is that the maximum density should be constant. However, the *effective* behaviour seen in section 5.8 should still hold: the solution is approximately given by a hydrostationary background with a (possibly large) perturbation, leading to an approximately constant maximum density profile superimposed with oscillations driven by the fundamental mode of the remnant. There may in addition be long-term secular behaviour due to changes in pressure support in the core, which could indicate the eventual collapse to a black hole.

In figure 6.23 the maximum density evolution is tracked for the three different cases evolved in this chapter.

We see that the result is approximately stable in all cases. The discretisation effects have not lead to significant inflation or collapse of the remnant, instead they introduce small amplitude oscillations which decay over time, as was seen in section 5.8. The purely hydrodynamic simulations are evolving to very similar equilibrium solutions. The advantages of the use of the rotational coordinate transform are clear: the improved numerical accuracy allows the fundamental oscillation mode about the approximate equilibrium to be clearly seen. Advective fluxes due to non-grid-aligned 4-velocities can introduce numerical errors, transforming the coordinates to reduce the 4-velocities can work to reduce these errors. The remaining fluxes will be due to pressure gradients which largely balance against the geometric source terms in approximately stationary configurations. Here we see that this coordinate transform has the effect of reducing early time high and low frequency perturbations, whilst introducing steady frequency small oscillations which decay over time.

The addition of radiation to the system gives another channel for energy-momentum exchange. Again, the smoothing effect of this additional channel is clearly seen. However, the combined effect of the additional radiation energy and the impact on the low-density fluid matter leads to an increase in the core density which only decays over a long time period.

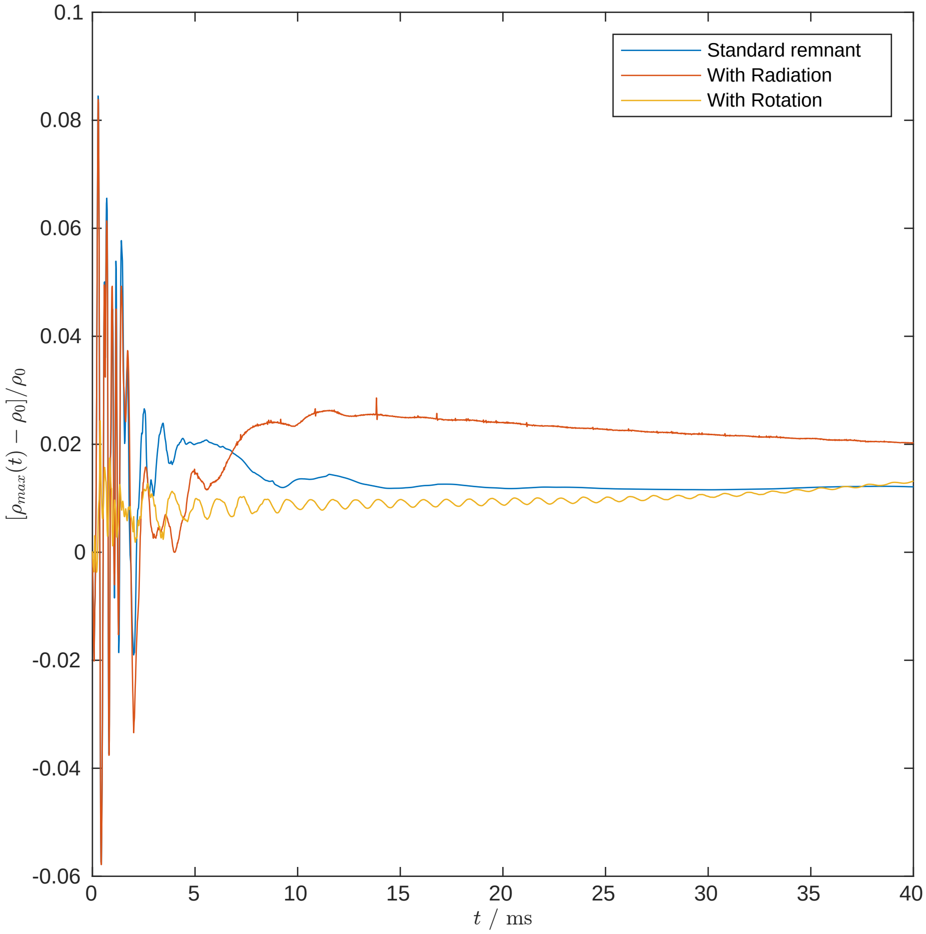

Figure 6.23: The maximum rest mass density of the remnant as evolved in the three situations discussed here. We see that the result is approximately stable in all cases. The purely hydrodynamical simulations are evolving to very similar equilibrium solutions. The advantages of the use of the rotational coordinate transform are clear: the improved numerical accuracy allows the fundamental oscillation mode about the approximate equilibrium to be clearly seen. The addition of radiation to the system gives another channel for energy-momentum exchange. Again, the smoothing effect of this additional channel is clearly seen. However, the combined effect of the additional radiation energy and the impact on the low-density fluid matter leads to an increase in the core density that only slowly reduces.

6.9 Summary

In this chapter I have used the efficient general relativistic radiation hydrodynamics code developed and described in chapters 3-5 to evolve a neutron star remnant. This has allowed us to explore the impact of including photon radiation as an additional aspect of the physical model. I have also explored the use of a novel coordinate choice to improve the numerical accuracy of the rotating fluid remnant.

We have seen that, even with limited computational resources, using a flexible and efficient mesh refinement framework allows us to accurately evolve the complex multiphysics problem of a neutron star remnant. The improved accuracy from the novel rotating coordinate choice is most clearly seen in the capturing of the fundamental mode oscillations in section 6.8. This highlights that the approach of constructing numerical methods and techniques targeted to specific astrophysical situations can lead to substantial efficiency gains.

We have also seen the qualitative impact of the addition of photon radiation through the results in section 6.6. The smoothing effect of the additional radiation is also clearly seen in the impact on the maximum density oscillations.

The qualitative impact of the physical model choices (such as the inclusion of radiation) and precise numerical choices (such as the use, or not, of the rotating coordinates) lead to the question of the robustness and precision of the results computed here. This will be tackled more precisely in the next chapter.

Chapter 7

Uncertainty Quantification

7.1 Introduction

In chapter 6 I performed general relativistic radiation hydrodynamic simulations of a neutron star merger remnant. To do this I had to make choices. There are physical choices, such as which equation of state to use, or which radiation treatment to impose. Then there are "purely numerical" choices, such as at which time to transfer the data from the Einstein Toolkit grid to the AMReX grid, or which numerical solvers to use. We need to know how robust our results are to variations in these choices.

Our setup consists of a system of partial differential equations (the balance laws introduced in section 2.7), a set of boundary conditions, multiple closure options, and precise initial conditions to describe each quantity at the start time. Each of these aspects can be parameterised but they are not fully constrained. Physical observations can be used to constrain the closure schemes used. For example the measurement of gravitational waves can be used to constrain the fluid equation of state. There is *uncertainty* in the numerical setup which we use, particularly in our choice of the initial conditions. Small changes in these parameters can propagate, qualitatively affecting the solutions. The detailed study of this is called *uncertainty quantification*. Formally, uncertainty quantification is the modelling and approximation of the propagation of uncertainty in the solution due to uncertainty in the inputs [106].

A significant source of uncertainty is in our implementation of the radiation field. The initialisation of the radiation field is a purely computational choice, not a physical process which could be constrained by observation. We choose for it take the form of being in local thermodynamic equilibrium with the fluid, and we choose the time at which these fluid temperatures are used. The temperatures are still evolving soon after merger. We make use of this ad-hoc approach as we do not have a radiation field available from the binary neutron star inspiral and merger simulation.

Most fluid simulations involve some degree of turbulent behaviour. During the neutron star merger process both shearing motion and the propagation of shock waves significantly influence the evolution of the system [107]. As we saw in section 6 turbulence is present in remnant evolution. Shocks and shearing motion in the spiral arms induce Kelvin-Helmholtz-like vortices. Kelvin-Helmholtz instabilities involve both discontinuities and small-scale turbulent behaviour, hence in this chapter I use this an appropriate model to begin studying uncertainties in shock-turbulence interaction.

It has been shown by [106] that in the Newtonian regime shocks and turbulence can lead to complications when quantifying uncertainty. For example pointwise properties do not converge with resolution, but summary statistics (such as the mean and variance) can. To formalise this, the idea of a “statistical solution” has been introduced (see, for example, [108]), which looks at the numerical convergence of the distribution of quantities.

The quantification of uncertainty in Kelvin-Helmholtz instabilities has been investigated in Newtonian hydrodynamics by [108]. In this chapter this novel research explores how these results extend to relativistic hydrodynamics and relativistic radiation hydrodynamics (this is introduced in [19]). I expect for the Newtonian conclusions to carry over: that pointwise quantities will not convergence, but that their distribution (and summary statistics of the distribution) will. I explore what observables I can make use of, and if the results are robust, meaningful, and useful.

I have also extended the Newtonian investigation to look at non-local quantities such as those involving differential and integral operators. One example is the vorticity – the curl of the velocity field. Differential operators appear in magnetic fields and in dissipative non-ideal hydrodynamic treatments such as bulk and shear viscocities and heat fluxes (see, for example, [26]). An alternative observable would be light curve calculations, which depend on the optical depth. This is a non-local quantity involving an integral operator. I investigate whether or not these non-local quantities converge with spatial resolution. It is expected that gradients may steepen as spatial resolution increases, so differential operators may not converge pointwise or with summary statistics. Conversely integral operators may be insensitive to spatial resolution, so we may expect both pointwise and statistical quantities to converge.

In this chapter I investigate the statistical behaviour of computational simulations of Kelvin-Helmholtz instabilities in the relativistic hydrodynamic and relativistic radiation hydrodynamics regimes. I quantify the convergence behaviour of individual scenarios as resolution is increased. I also quantify the convergence behaviour of the mean and variance of a set of solutions as resolution is increased. I investigate whether the same statistical results hold for relativistic hydrodynamics, and how introducing radiative transfer affects the solutions.

7.2 Relativistic Kelvin-Helmholtz Instability

A dense cool fluid moving relative to a hot rarefied fluid is a common occurrence in astrophysical phenomena. Examples include the interstellar medium, circumgalactic medium, intracluster medium, supernova remnants and superbubbles, cosmic filaments, galactic winds, protoplanetary disks, protostellar jets and jets from active galactic nuclei. Burning and energy release in turbulent media takes place in stellar interiors and in supernovae [15]. These fluids are often in hydrostatic and local thermodynamic equilibrium. A mixing layer occurs at the shearing interface in these fluids leading to a Kelvin-Helmholtz instability (KHI).

In phenomena where KHI occurs it is very unlikely that the initial conditions are known to an accuracy where the resulting turbulence can be precisely reproduced with physics models. However if numerous simulations are calculated, with initial conditions sampling the space of likely initial conditions, then the distribution of results allows us to predict the expected behaviour with a quantified uncertainty.

In the Newtonian limit of hydrodynamics it has been shown that for numerical simulations of KHI, a conserved quantity does not converge locally with increasing spatial resolution [106]. However when numerous simulations are run with varying perturbations in initial conditions leading to varying turbulent motion, then the mean and variance of the quantity taken over the numerous samples does indeed converge with increasing spatial resolution. This same result is seen with algebraic functions of the conserved quantities.

The numerical experiments in this chapter analyse the rest-mass density field as an example of the statistical properties of algebraic functions of conserved quantities.

The properties of integral and differential operations on conserved quantities are also investigated.

For a detailed investigation I conduct uncertainty quantification with flat spacetime in two spatial dimensions. I set up Kelvin-Helmholtz initial conditions and extend the work in [106] to special relativistic regimes. My approach differs to the standard in relativistic hydrodynamics which uses a smeared interface with perturbations in the velocity field. I make use of a sharp interface with perturbations in the interface location. I have used relatively large perturbation magnitudes, however [106] shows that the resulting instabilities are insensitive to this amplitude.

7.2.1 Problem Definition

I set up a square spatial domain with $x \in [0, 1]$, $y \in [-1/2, 1/2]$, and with periodic boundary conditions. The domain is split into three horizontal sections. The fluid in the top and bottom sections initially flows towards the right, and the fluid in the middle section flows towards the left. Small perturbations are made to the interface positions. The shearing velocities at the perturbed interfaces induce vortices which grow into turbulent flow. The initial rest mass density and horizontal relativistic fluid velocity profiles are given by

$$\rho(x, y) = \begin{cases} 1, & y^-(x) < y < y^+(x), \\ \frac{1}{10}, & \text{else}, \end{cases} \tag{7.1}$$

$$v^x(x, y) = \begin{cases} -\frac{1}{2}, & y^-(x) < y < y^+(x), \\ \frac{1}{2}, & \text{else}. \end{cases} \tag{7.2}$$

Here $y^{\pm}(x)$ are the perturbed interface locations. The initial fluid pressure is given by $P = 1$, the vertical fluid velocity is $v^y = 0$, and an adiabatic index $\Gamma = 4/3$ is used.

This set-up is similar to [109] in quantity values, however the domain size has been reduced to increase interaction between the upper and lower resulting vortices. The interface perturbations are as in [108] using a perturbation in position rather than in vertical fluid velocity. The upper and lower interface perturbations are each a linear combination of sinusoidal waves, with amplitudes and phase shifts generated using a pseudo random number generator [108]

$$y^{\pm}(x) = \pm\frac{1}{4} + \frac{1}{10}\sum_{n=1}^{10} a_n^{\pm}\cos(b_n^{\pm} + 2\pi n x). \tag{7.3}$$

The amplitudes $(a_n^{\pm})$ are scaled to sum to unity, and the phase shifts are linearly transformed to the interval $b_n^{\pm} \in (-\pi, \pi]$. The coefficients a_n, b_n are drawn from a uniform random distribution.

Figure 7.1 shows a visualisation of the initial set-up of the mass-density and velocity fields.

7.3 Evolved Solutions

The Kelvin-Helmholtz instability setup is evolved to a time $t = 2$ in geometric units. The evolved rest mass density field is shown in figure 7.2. In this case the spatial domain used a uniform discretisation with 512^2 zones. Although still at early time we can see qualitative features of interest such as shock waves, the rollup of vortices, and turbulent behaviour. It shows that the numerical code can capture and simulate the qualitative features expected of this type of hydrodynamic instability. The discontinuities in the rest mass density field within the higher density layer are stronger than in Newtonian cases due to special relativistic effects.

7.3.1 Pointwise Convergence

I have run a series of Kelvin-Helmholtz instability simulations using identical initial condition perturbations with a range of spatial resolutions. Figure 7.3 shows the rest mass density profile for spatial resolutions using 16^2, 32^2, 64^2, 128^2, 256^2, and 512^2 zones. We can see that the large scale qualitative features are comparable. There are internal shocks, and vortices appear to roll up at similar locations. However as the resolution increases, tighter and smaller vortices are able to form leading to more detailed turbulent behaviour.

A quantitative check for convergence is required. Here I will be using the L_2-norm (as introduced in section 4.1) of the difference in rest-mass density profiles between two simulations at different resolutions. To allow a numerical comparison each lower resolution solution is up-scaled to match the higher resolution solution assuming piecewise constant data. For our two dimensional discretised solutions this error metric takes the form

$$L_2\left(q^{lo}\right) = \left(\frac{1}{N^2}\sum_{i=1}^{N}\sum_{j=1}^{N}\left(q_{ij}^{hi} - q_{ij}^{lo}\right)^2\right)^{1/2} . \qquad (7.4)$$

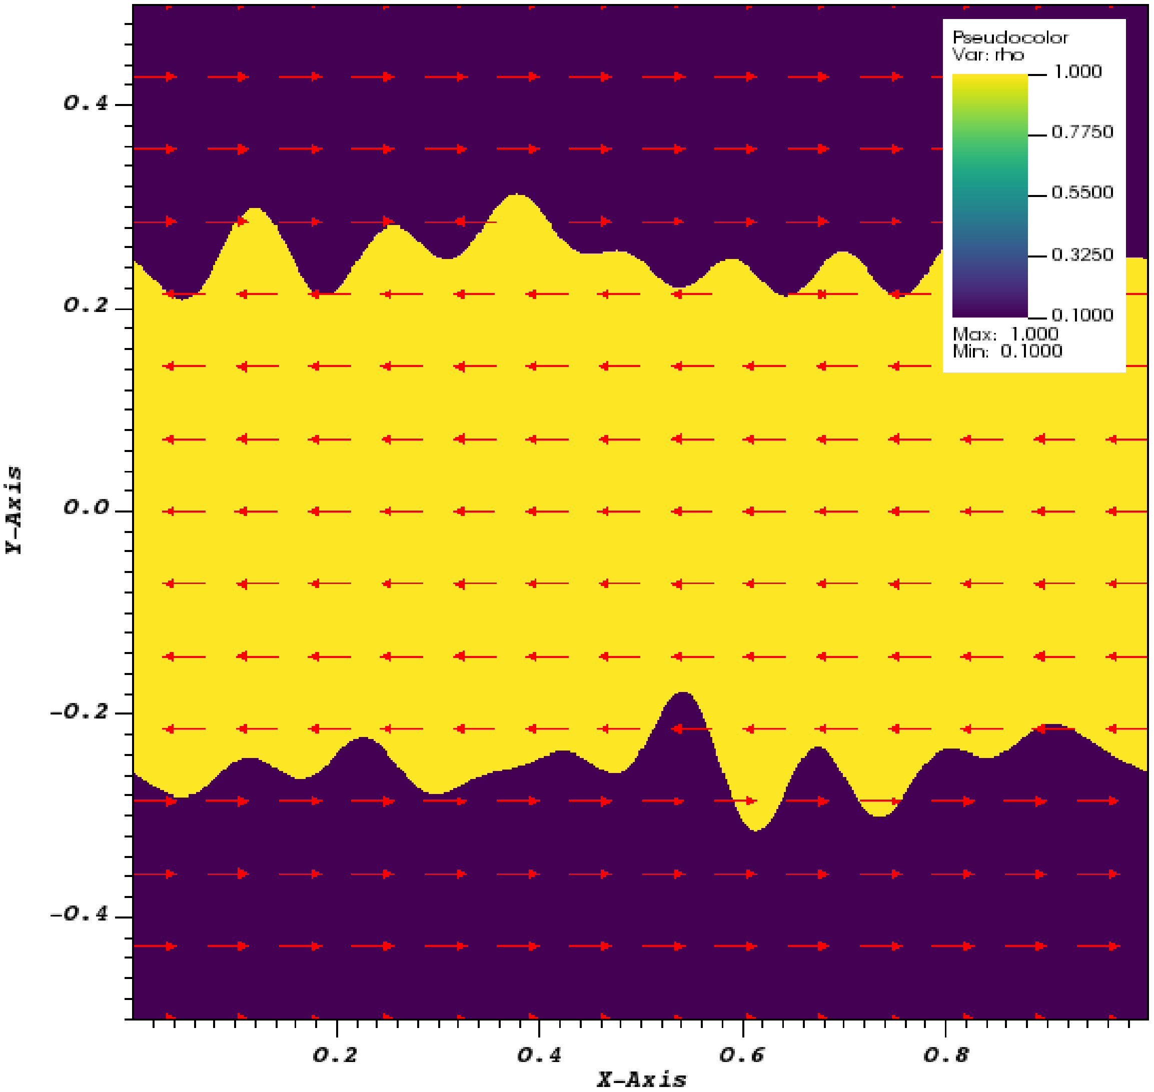

Figure 7.1: Initial set-up of a relativistic Kelvin-Helmholtz instability, showing rest-mass density and fluid 3-velocity vectors. We can see the two uniform states separated by a random number based wave interface as in [108].

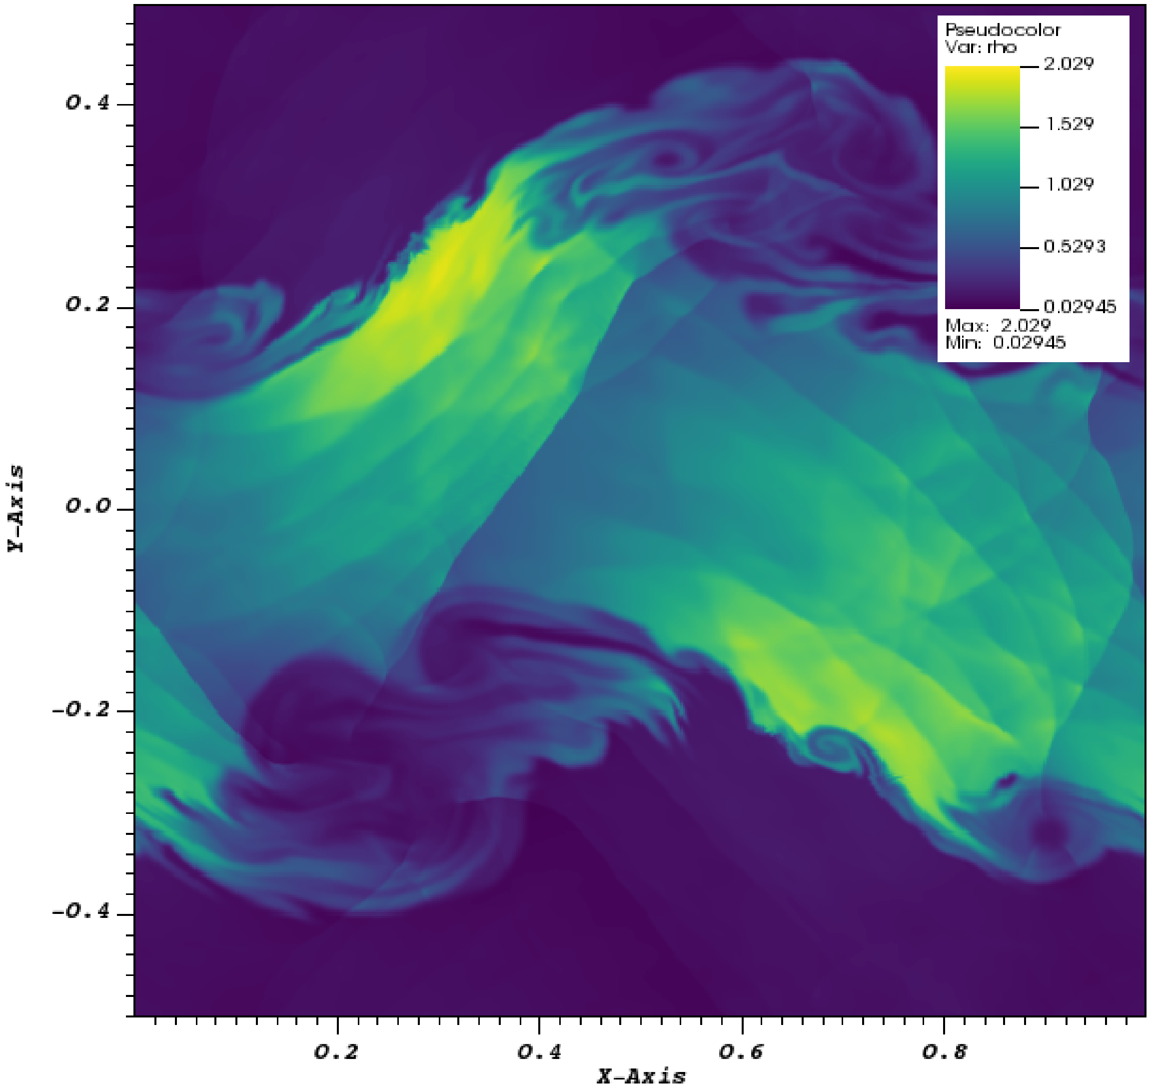

Figure 7.2: The rest mass density profile of the relativistic Kelvin-Helmholtz instability evolved to time $t = 2$ in geometric units. We can see the formation of vortices and turbulent behaviour. This is a relativistic version of the results in [108].

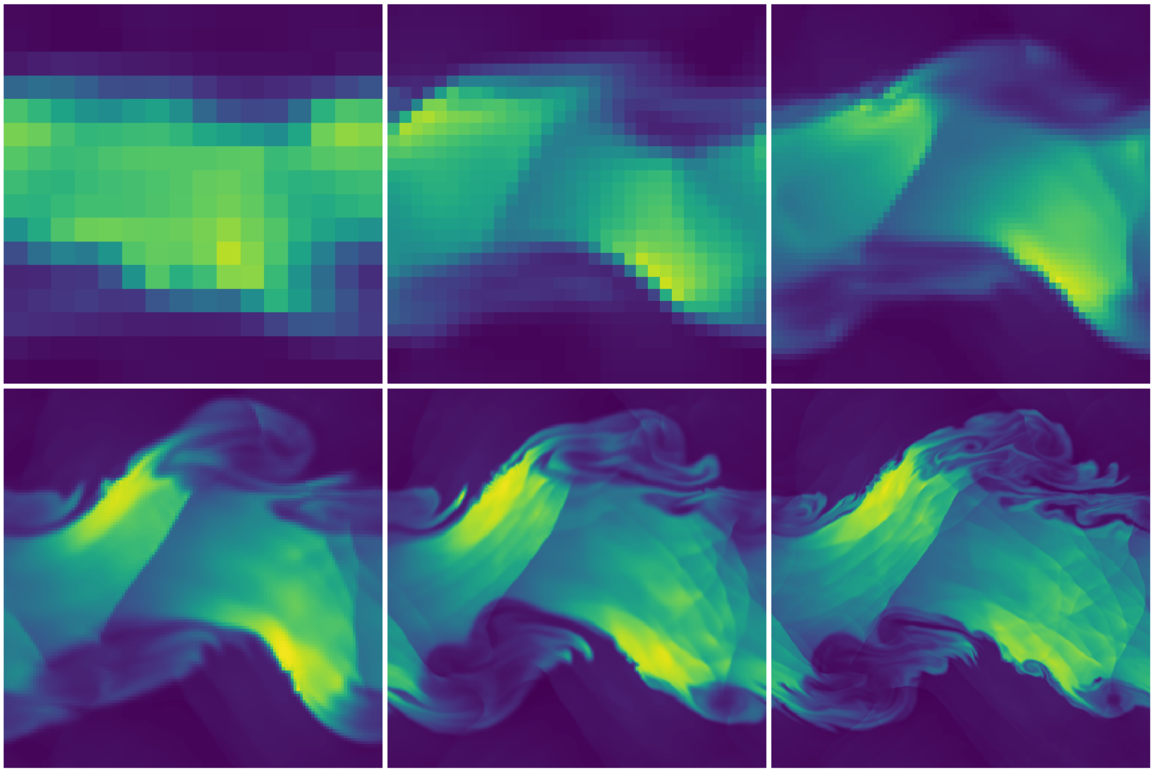

Figure 7.3: The rest mass density profiles for a series of relativistic Kelvin-Helmholtz instabilities evolved to time $t = 2$ in geometric units. Each simulation had identical initial condition definitions, however spatial resolutions varied from 16^2 to 512^2.

Here q^{hi} is the higher resolution solution, and q^{lo} is the lower. For this study I compare to the highest resolution available (512^2), treating it as the best case solution. If this error metric decreases as the spatial resolution increases then we have Cauchy convergence.

Figure 7.4 shows the L_2-norm of the difference in rest-mass density profiles at spatial resolutions 16^2 to 256^2. This has been repeated for a set of 32 different initial condition perturbations. As can be seen in figure 7.4 none of the individual realisations of initial condition perturbations are converging with increasing spatial resolution. Accordingly the mean of the set of L_2-norms is also not converging. These results have previously been shown for the Newtonian hydrodynamic regime in [108].

At first glance, this is catastrophic for our numerical simulations. It indicates that there is no continuum limit for the quantities, and suggests the problem is not well posed. As discussed in, for example, [108], well-posedness and the utility of numerical solutions for these problems, can be recovered by considering their statistical properties.

7.4 Analysing a Distribution of Solutions

The Newtonian results, as extended to the relativistic regime above, show that pointwise convergence can be lost in shock-turbulence interaction simulations. However, the idea of statistical solutions is to look at the convergence of the distribution of quantities, and the summary statistics of those distributions. This requires computing many samples, and hence many simulations.

With a 3+1 dimensional simulation code numerous 2+1 simulations can be calculated simultaneously. I run a separate instance of the problem on each slice in (for example) the z-direction. The initial conditions for each slice are generated using the z-index as the random number generator seed. Effectively the z-index has become a proxy-parameter for the initial condition perturbation. To prevent the separate simulations from interacting with each other the fluxes in the z-direction are set to zero ($\underline{F}^z = 0$). Figure 7.5 shows the rest-mass density profiles for a simulation block with 32 slices for a series of relativistic Kelvin-Helmholtz instabilities evolved to a time of $t = 2$ in geometric units.

At each (x_i, y_j) location there exists a distribution of values over the number of samples calculated. I can calculate the mean and variance of a quantity (q) over these

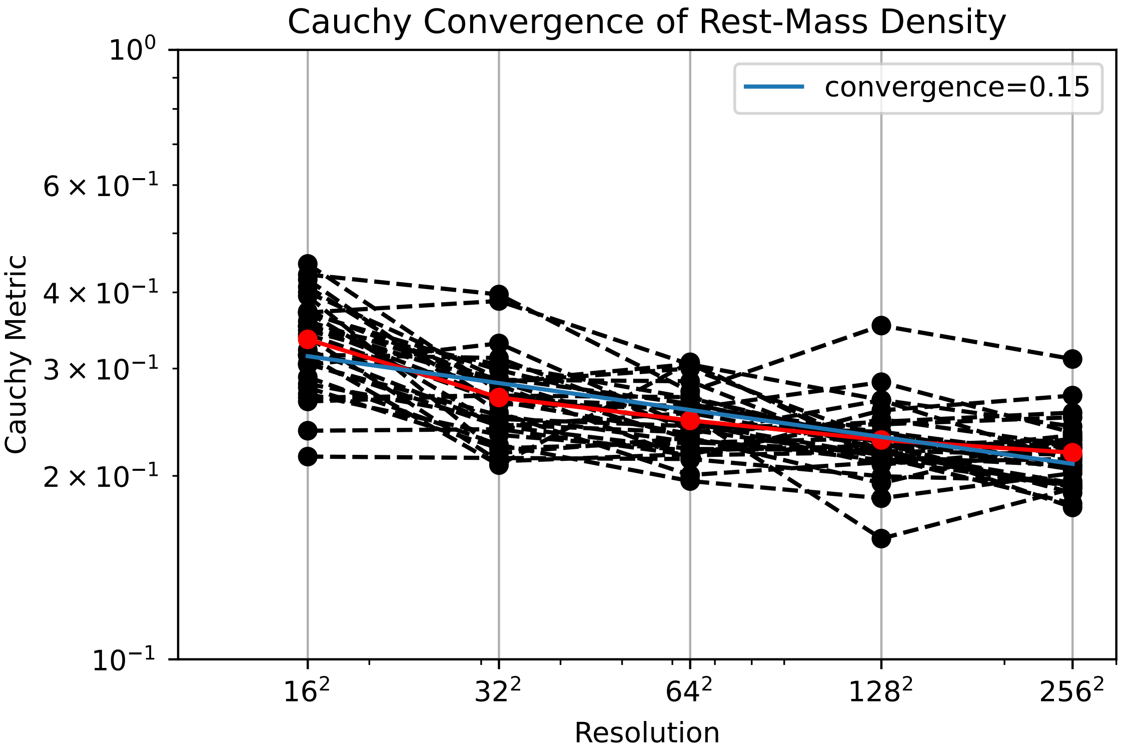

Figure 7.4: The L_2-norm of the difference in rest-mass density profiles at spatial resolutions 16^2 to 256^2. This has been repeated for a set of 32 different initial condition perturbations (dashed black lines), with the mean of the L_2-norms (solid red line), and the linear regression (solid blue line). This figure suggests that individual realisations do not converge to a continuum limit. This is consistent with the results in [108].

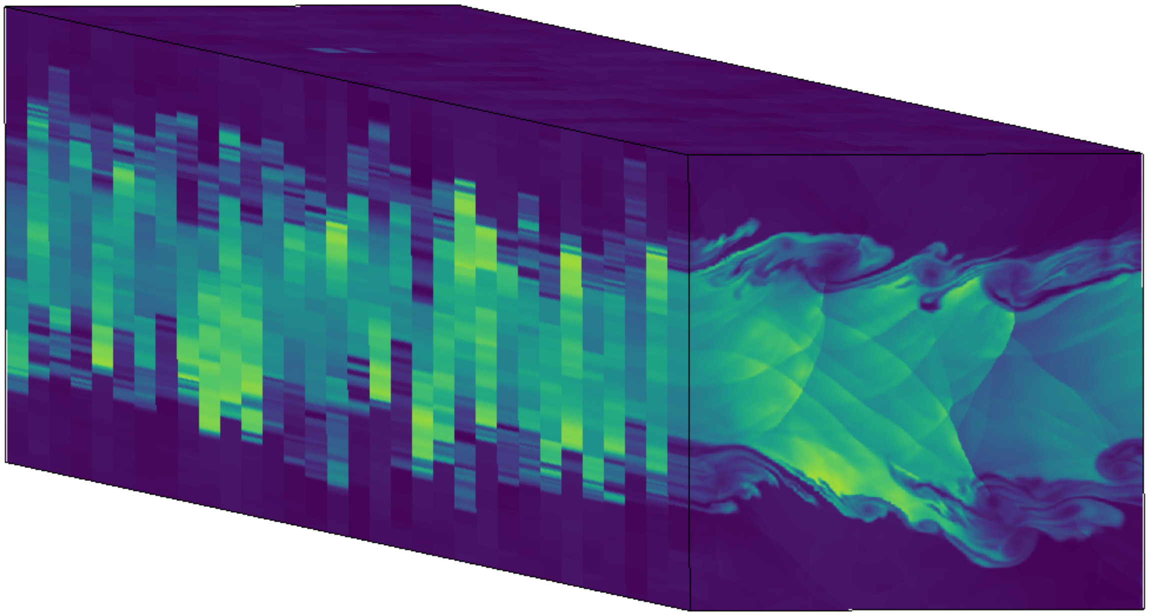

Figure 7.5: A block of relativistic Kelvin-Helmholtz instability simulations, with each slice a separate realisation of initial condition perturbations, evolved to time $t = 2$ in geometric units.

numerous samples as

$$\overline{q} = \frac{1}{K}\sum_{k=1}^{K} q_k, \tag{7.5}$$

$$q_{var} = \frac{1}{K}\sum_{k=1}^{K}(q_k - \overline{q})^2. \tag{7.6}$$

Here the mean and variance are calculated over K samples.

Figure 7.6 shows the mean and variance of the rest-mass density profile over a set of relativistic Kelvin-Helmholtz instability simulations evolved to time $t = 2$ in geometric units. As can be seen in figure 7.6 the asymmetries in the x-direction are being washed out, converging towards symmetric solutions as the size of the simulation sample set increases. This has previously been shown in the Newtonian hydrodynamic regime [108].

Figure 7.7 suggests that the mean of the rest-mass density profiles converges towards some continuum limit as the number of samples increases. Here I have used the mean rest mass density profile over 256 samples to approximate the best case solution. This does *not* suggest (yet) that we have recovered well-posedness of the solution, as that would require looking at convergence towards the continuum limit (in both space and

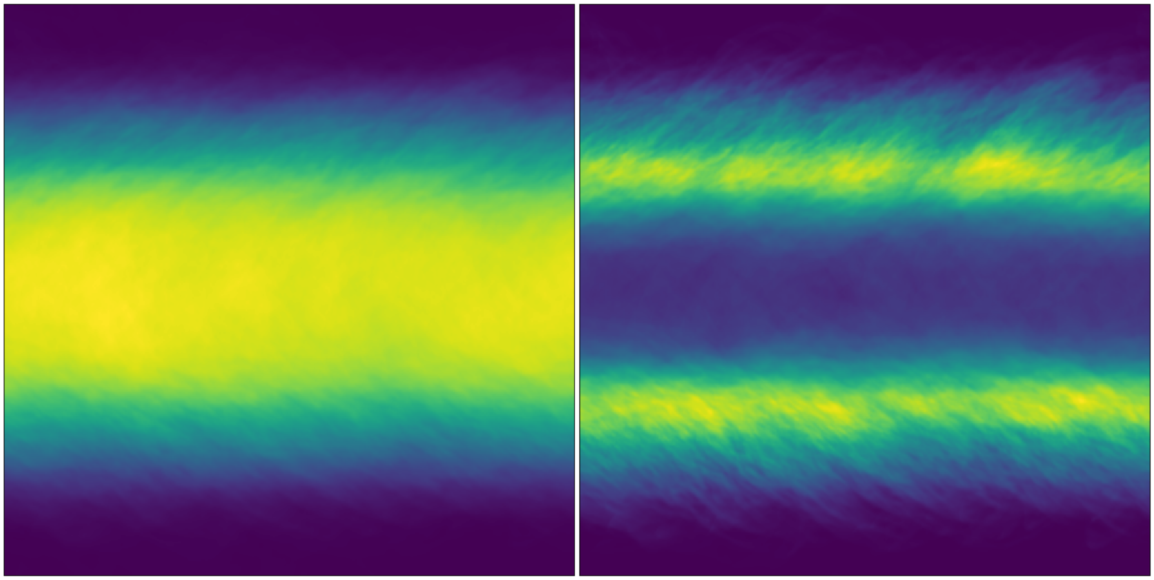

Figure 7.6: The mean (left) and variance (right) of rest-mass density profiles from a set of relativistic Kelvin-Helmholtz simulations, each with distinct initial condition perturbations, evolved to time $t = 2$ in geometric units. This figure shows that statistical properties such as mean and variance appear to be robust and tend towards independence on x-position. This is consistent with the Newtonian case shown in [108].

time). This instead suggests that the statistical limit of an infinite number of random samples (the continuum limit in parameter space) is well behaved.

7.5 Convergence of a Distribution of Solutions with Spatial Resolution

By perturbing the initial condition by a small amount we can obtain a distribution of solutions. I am interested in how this distribution of solutions behaves as the spatial resolution is increased.

Again I take the highest resolution results available as the best case approximation to the continuum limit. I calculate the L_2-norm of the difference in mean and variance of the rest-mass density between lower resolutions and the highest, with the lower resolutions up-scaled assuming piecewise-constant data.

Figure 7.8 shows that as spatial resolution increases, both the mean and the variance of the rest-mass density over numerous instances do indeed converge towards a continuum limit. This has previously been shown in the Newtonian hydrodynamic

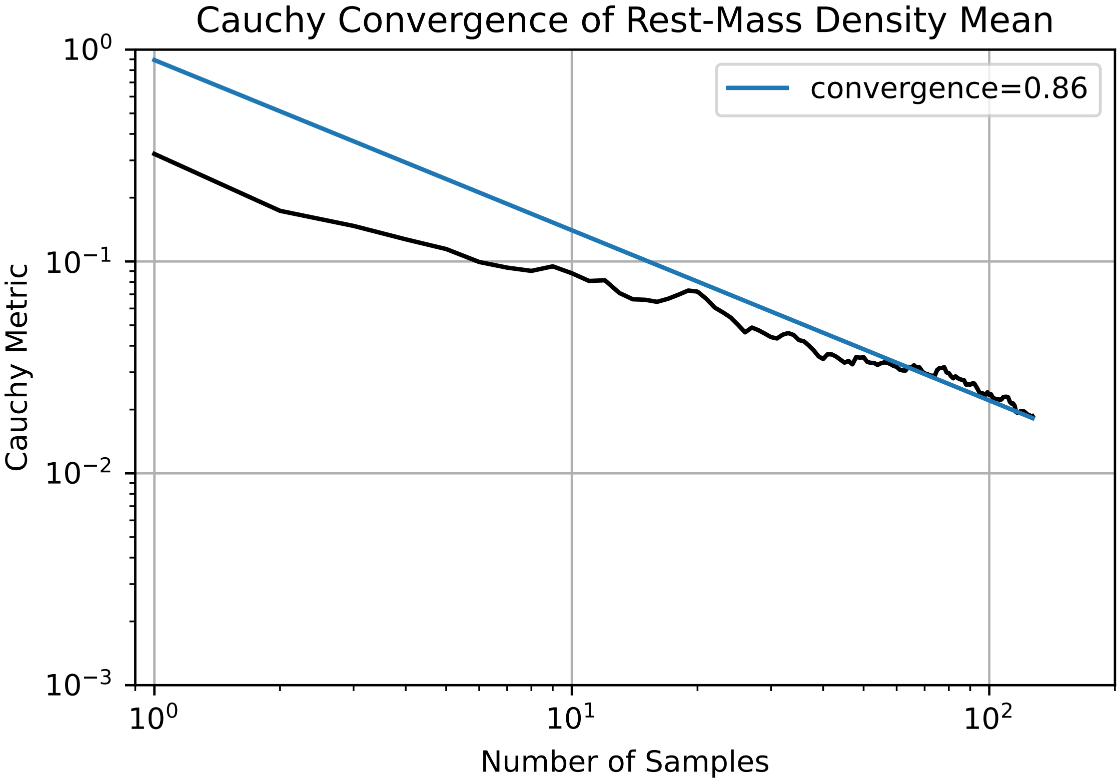

Figure 7.7: The mean profile converges to some continuum limit (in parameter space) as the number of samples increases.

regime by [108].

Recall the conclusions drawn from figure 7.4: this shows that the mean of the L_2-norm of the difference in the rest-mass density does not decrease with spatial resolution. Therefore the sequence of means of errors at different resolutions is not *strongly* convergent. However figure 7.8 shows that the L_2-norm of the difference in the mean rest-mass density does decreases with spatial resolution. This sequence of the difference of the means of the errors is *weakly* convergent. The order of operations is important implying weak convergence and not strong convergence.

7.6 Wasserstein Metric

The one-point Wasserstein metric (also known as the earth mover's distance) is a measure of the difference between two probability distributions. The metric corresponds to the cheapest way to rearrange one distribution into the shape of the other. This is a more complete comparison of probability distributions than just comparing means and variances since the entire distribution is considered, not just the first moments. The Wasserstein metric has previously been used to check the (weak) convergence with spatial resolution of the distribution of rest-mass density fields for Newtonian Kelvin-Helmholtz simulations [106, 108].

Figure 7.9 shows that the Wasserstein metric is converging towards a continuum limit as the spatial resolution increases. This suggests that for the range of initial conditions used, there is a continuum limit for the distribution of solutions, and the numerical method is finding it.

Special relativistic hydrodynamics leads to consistent conclusions as with the Newtonian results by [106, 108]. The pointwise quantities do not converge with spatial resolution but the distribution of solutions, including the summary statistics such as the mean and variance, does weakly converge towards a continuum limit as the spatial resolution increases.

7.7 Relativistic Radiation Hydrodynamic

I now aim to check if these conclusions extend to radiation hydrodynamics and whether the radiation parameters (for example, the strength of the opacity) can be distinguished in a statistical sense.

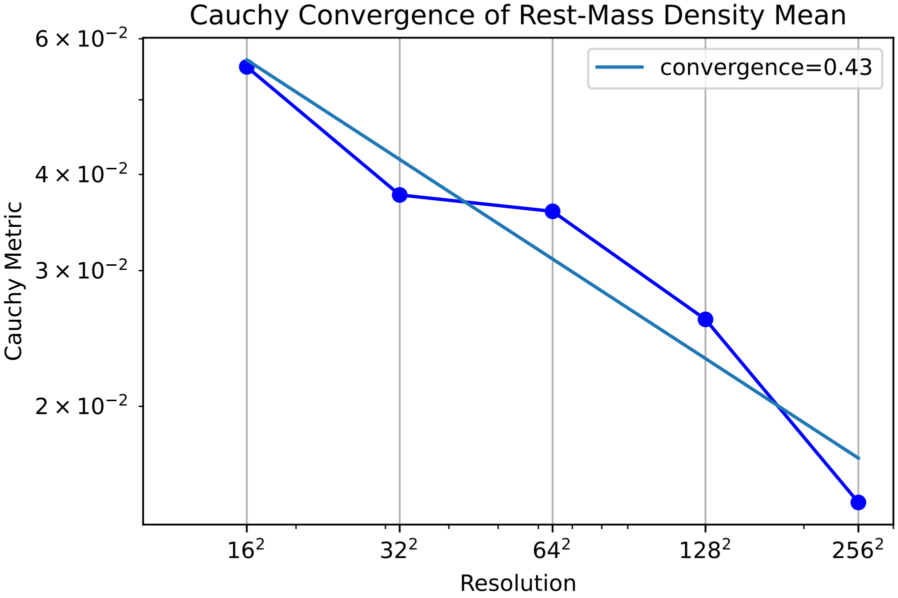

(a) Mean

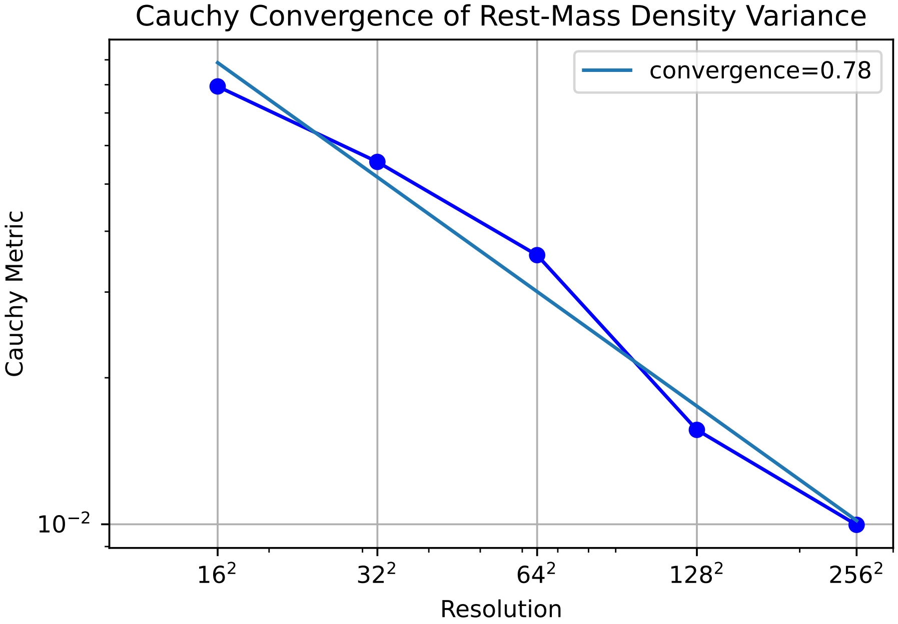

(b) Variance

Figure 7.8: The L_2-norms of the difference in mean (top) and variance (bottom) of rest-mass density from sets of distinct relativistic Kelvin-Helmholtz simulations at spatial resolutions 16^2 to 512^2, with linear regressions. This figure shows that statistical properties such as mean and variance do converge to a continuum limit. This is consistent with the Newtonian regime shown in [108].

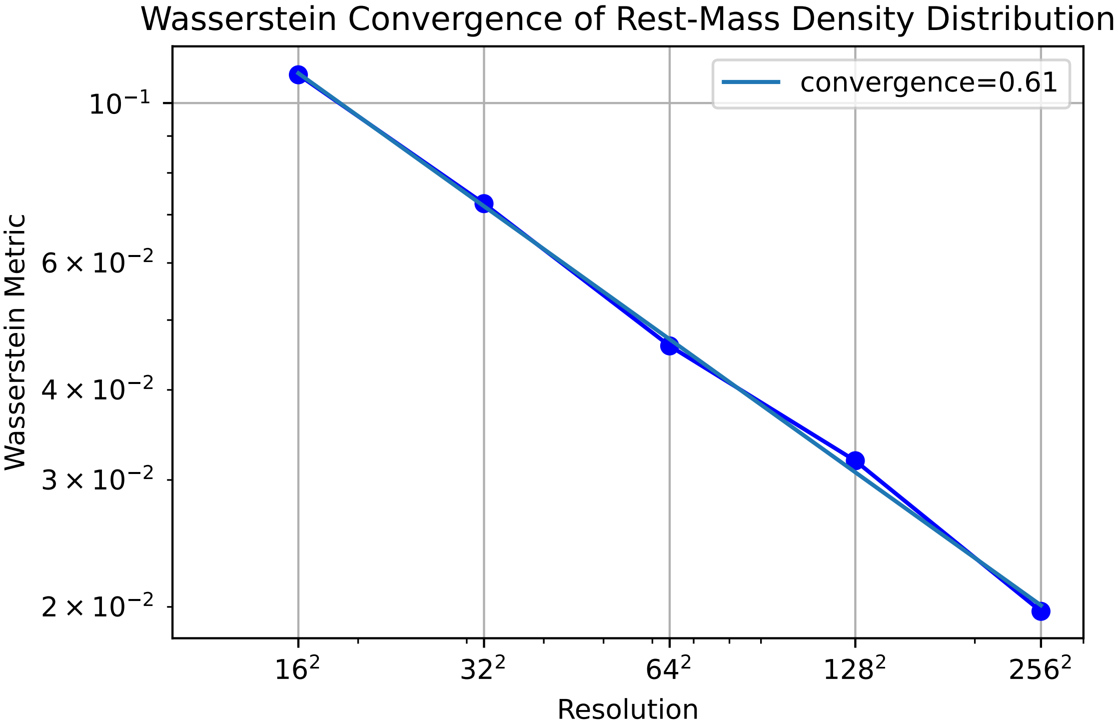

Figure 7.9: The Wasserstein metric of the difference between rest-mass density distributions for lower and higher resolution simulations. We can see that as the spatial resolution increases the distribution of rest-mass density values converges to a continuum limit. This is consistent with the Newtonian regime shown in [108].

A radiation field is introduced to the previously defined hydrodynamic Kelvin-Helmholtz initialisation. This radiation field is in local thermal equilibrium with the fluid ($E = a_{rad}T^4$) leading to

$$E = \begin{cases} 1, & y^-(x) < y < y^+(x), \\ 10^4, & \text{else.} \end{cases} \tag{7.7}$$

It has zero net flux with respect to the fluid ($F^x = F^y = 0$). I vary strength of the radiation hydrodynamic coupling by considering an optically thin case with $\chi_a = 10^{-4}\rho$, and an optically thick case with $\chi_a = 10^{-3}\rho$.

As with the hydrodynamic simulations the radiation hydrodynamic initial conditions are evolved to time of $t = 2$ in geometric units. I obtain a set of solutions in each case by initialising with a range of interface perturbations.

Figure 7.10 shows the rest mass density profile for a single realisation (on the left column), the mean solution (the centre column), and variance of the solutions (the right column). The different cases shown are the purely hydrodynamic (the top row), optically thin (the centre row), and and optically thick (the bottom row) relativistic radiation hydrodynamic Kelvin-Helmholtz simulations. The purely hydrodynamic case can be thought of as having zero radiation hydrodynamic coupling.

The differences in the solution due to the introduction of the radiation field, both weakly coupled and strongly coupled, can be seen in the individual realisations, the means, and the variances. In the radiation hydrodynamic simulations there has been more motion in the y-direction than in the purely hydrodynamic case. Additionally the central horizontal strip of larger rest-mass density is beginning to split into upper and lower bands with a region of lower rest-mass density separating them. This effect is more severe in the larger opacity case since the strength of the radiation hydrodynamic coupling is increased.

I am interested in whether the mean and variance of the rest-mass density converges with increasing spatial resolution even when there is radiation hydrodynamic coupling. Recalling the process used for figure 7.8 the highest resolution case is taken as a best approximation to the continuum limit for the mean and variance. I calculate the L_2-norm of the difference in mean and variance of rest-mass density profiles between lower resolutions and the highest, with the lower resolutions being up-scaled assuming piecewise-constant data.

Figures 7.11 and 7.12 show that as the spatial resolution increases both the mean and

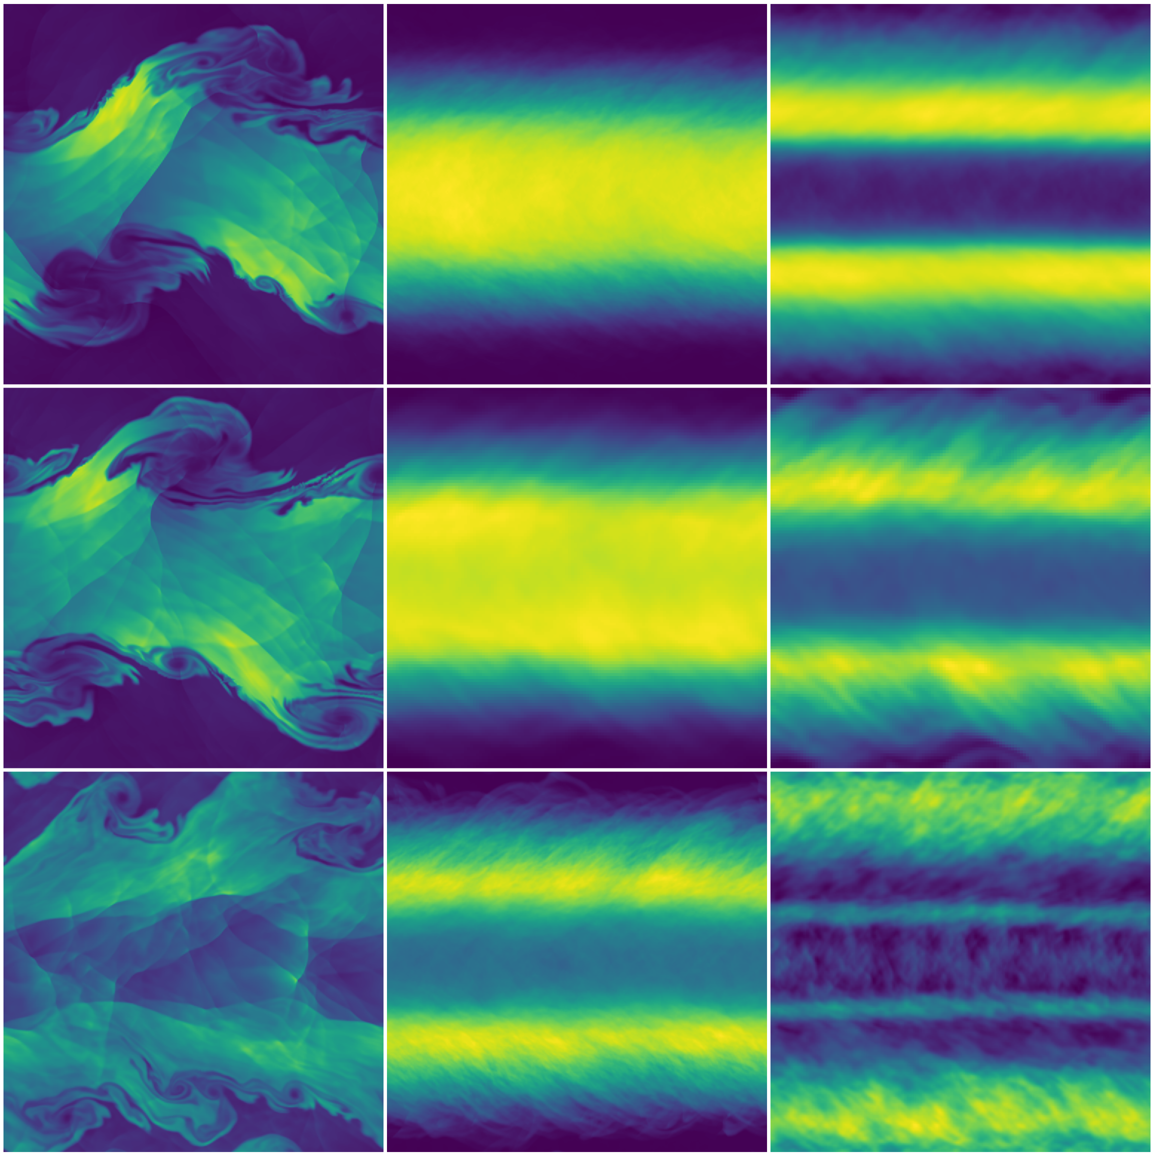

Figure 7.10: Single realisations (left column), means (central column), and variances (right column) of rest-mass density from purely hydrodynamic (top row), weakly coupled radiation hydrodynamic (middle row), and strongly coupled radiation hydrodynamic (bottom row) relativistic Kelvin-Helmholtz simulations, evolved to time $t = 2$ in geometric units. This figure shows that different physical models produce different results, and that features which are hard to distinguish between models when comparing individual realisations are clearly distinguishable in the statistical properties.

variance of rest-mass density do indeed weakly converge for both the optically thin and the optically thick cases.

As in section 7.6 I check to see if the distribution of solutions is converging and well behaved for the optically thin and thick radiation hydrodynamic KHI simulations. The results shown in figure 7.13 suggest that they indeed do.

In order to quantitatively confirm that introducing a radiation field with non-zero fluid coupling modifies the solution, even as spatial resolution is increased, the Wasserstein metric of the difference between the mean rest-mass density distributions from radiation hydrodynamic cases and the purely hydrodynamic case is calculated for a range of spatial resolutions. This is shown in figure 7.14 for both the optically thin and thick cases.

The lack of convergence here confirms that the radiation hydrodynamic solutions are indeed converging towards a different distribution to the purely hydrodynamic case. The optically thick case has a consistently larger Wasserstein metric than the optically thin case. This suggests that the thick case is *more* different to the purely hydrodynamic case than the thin case is. This indicates a monotonic dependence on the opacity: $\chi_a = 0$, $10^{-4}\rho$, and $10^{-3}\rho$ for the non-coupled, optically thin, and the optically thick cases respectively. We can clearly distinguish between purely hydrodynamic and coupled radiation hydrodynamic models for a range of opacities. This approach could be extended to identify the smallest change in opacity which could be detected with this quantitative method.

7.8 Non-Local Operators

In this section I investigate the statistical properties of non-local quantities, such as those involving differential and integral operators. As discussed in the introduction to this chapter, differential and integral operators are relevant to neutron star merger remnant evolution as they can appear in non-ideal hydrodynamics and in optical depth calculations. Here I use the vorticity (the curl of the fluid velocity field) and the density gradient magnitude as representative differential operators.

7.8.1 Differential Operators

A fluid's vorticity can be used to help characterise its turbulent flow. It is calculated using a differential operator on the primitive fluid quantities, namely the curl of the

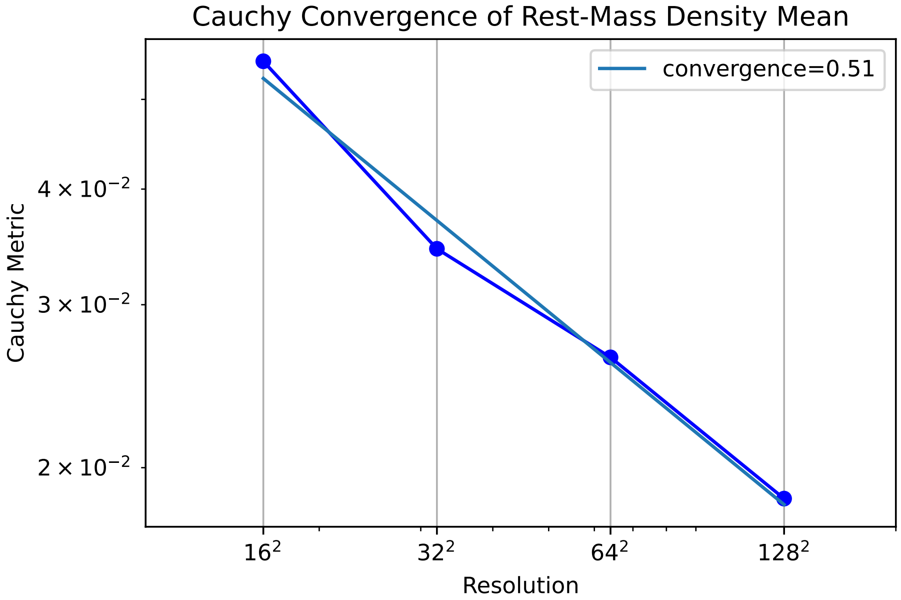

(a) Convergence of mean in the optically thin case.

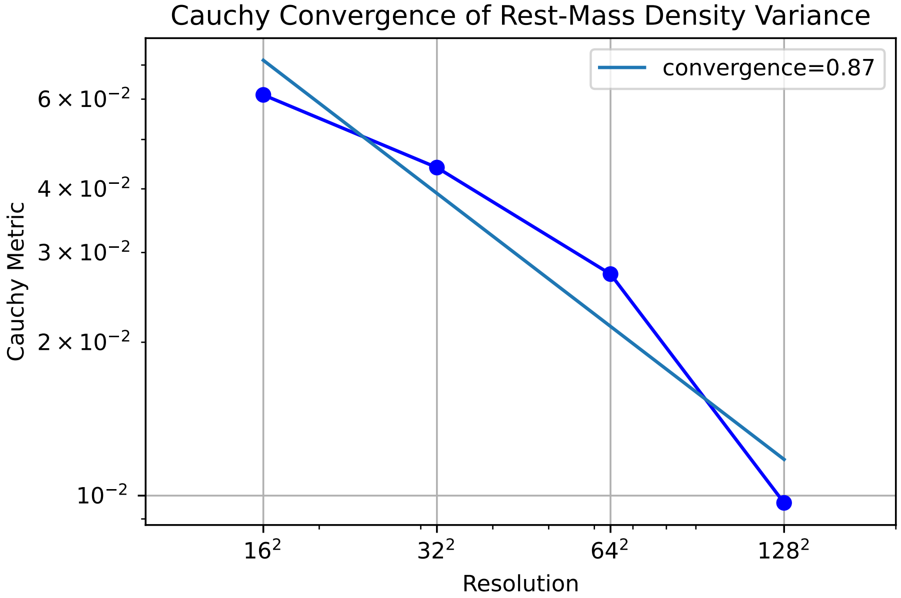

(b) Convergence of variance in the optically thin case.

Figure 7.11: Two-norms of the difference in mean (top panel) and variance (bottom) of rest-mass density from sets of distinct relativistic optically thin coupled radiation hydrodynamic Kelvin-Helmholtz simulations at spatial resolutions 16^2 to 128^2, with linear regressions. This figure shows that statistical properties still converge to a continuum limit, even for relativistic radiation hydrodynamics.

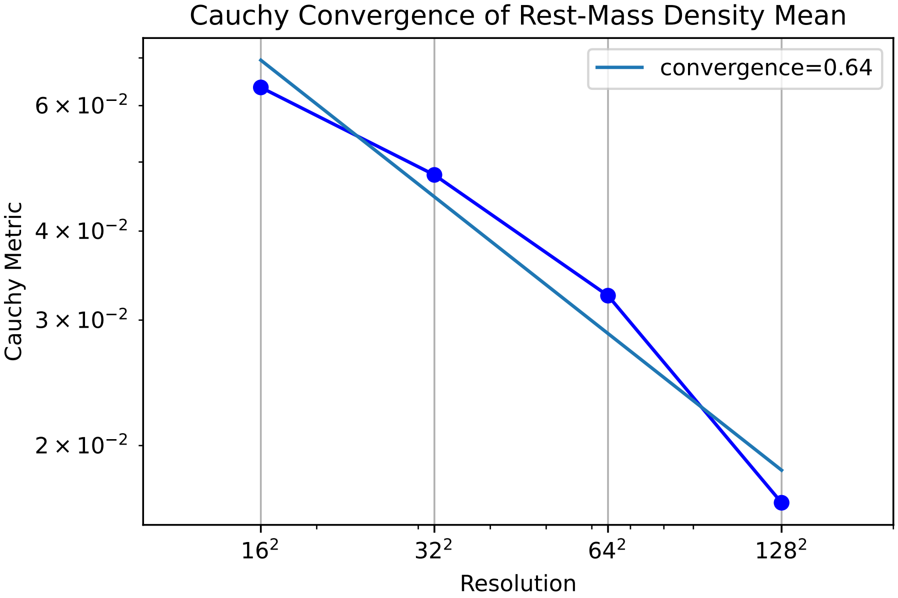

(a) Convergence of mean in the optically thick case.

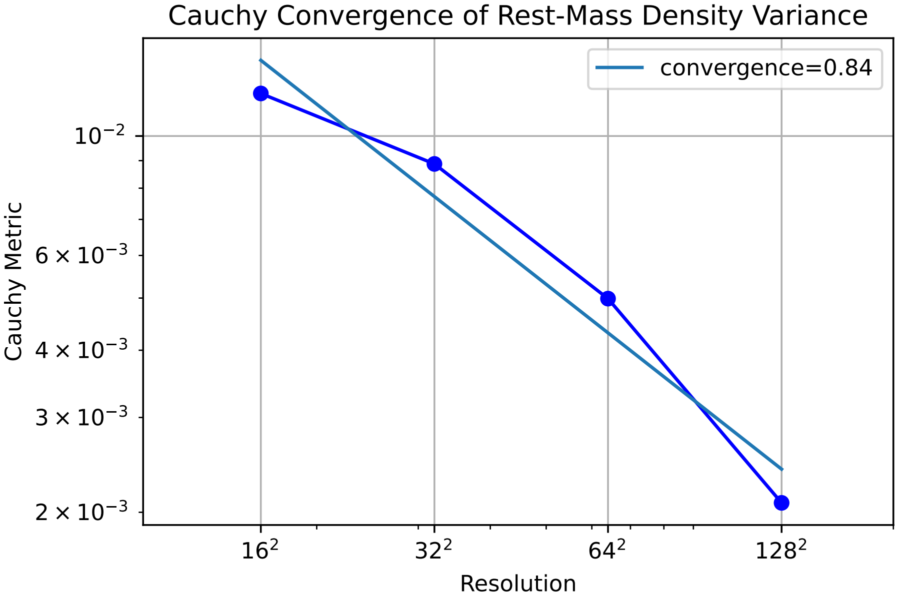

(b) Convergence of variance in the optically thick case.

Figure 7.12: Two-norms of the difference in mean (top panel) and variance (bottom) of rest-mass density from sets of distinct relativistic optically thick coupled radiation hydrodynamic Kelvin-Helmholtz simulations at spatial resolutions 16^2 to 128^2, with linear regressions. This figure shows that statistical properties still converge to a continuum limit, even for strongly coupled relativistic radiation hydrodynamics.

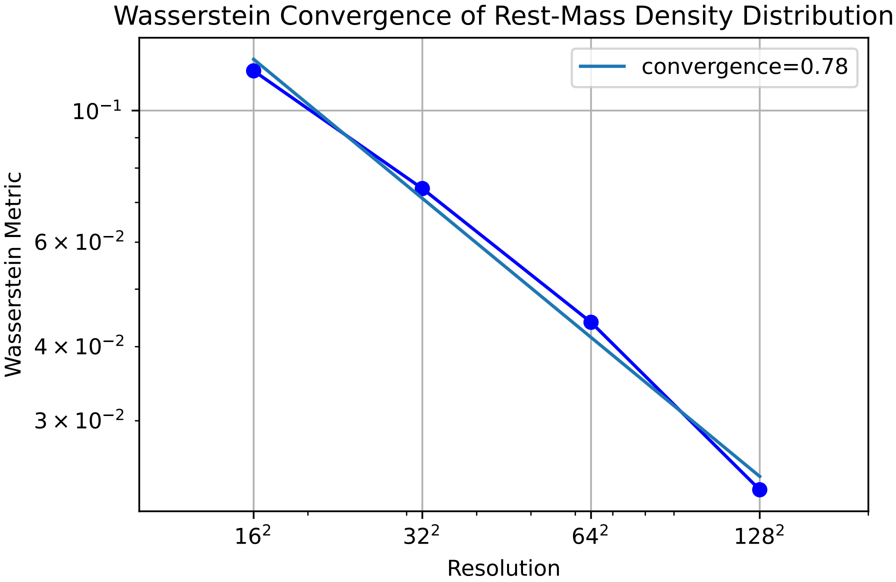

(a) Weak Coupling

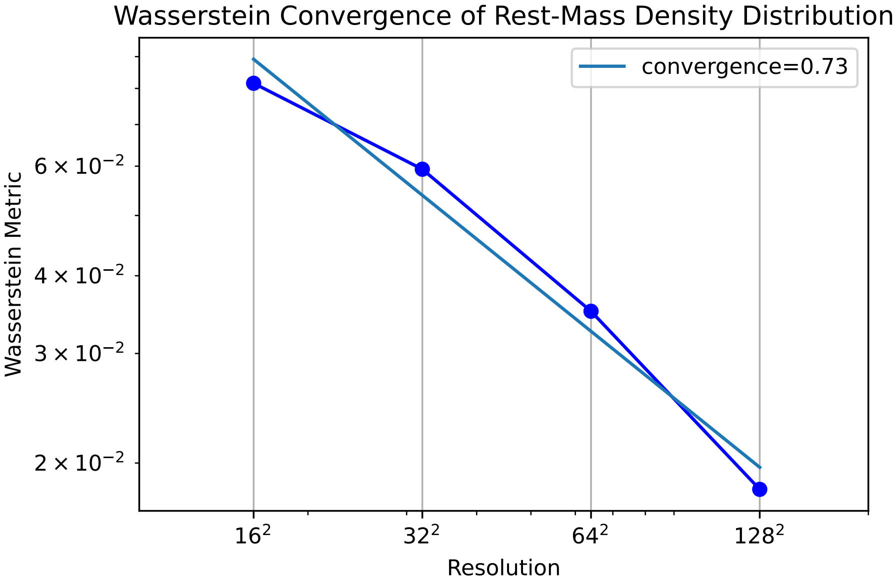

(b) Strong Coupling

Figure 7.13: Wasserstein convergence of the difference in rest-mass density distributions for different spatial resolutions for radiation coupled hydrodynamic simulations. The optically thin case is shown in the top panel, and the optically thick case on the bottom. This figure shows that the distributions still converge to a continuum limit, even for strongly coupled relativistic radiation hydrodynamics.

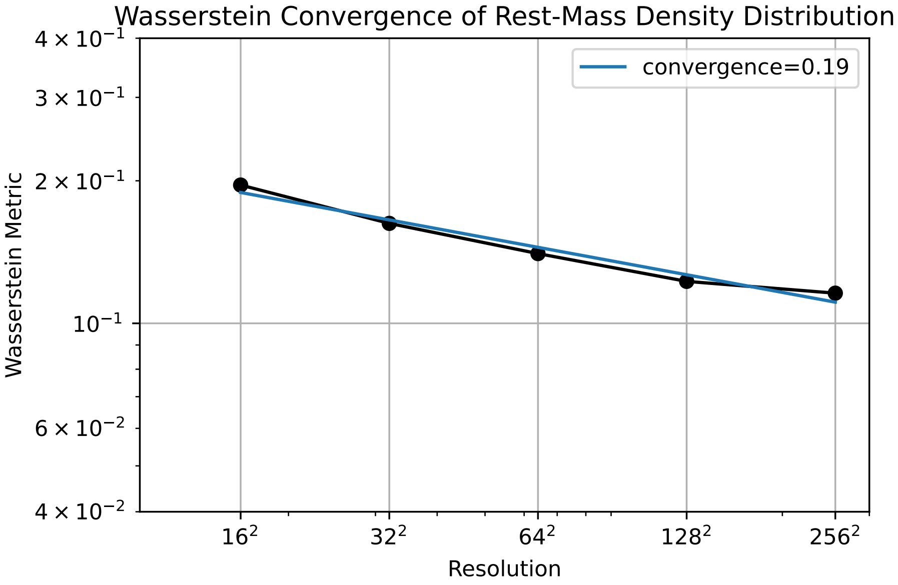

(a) The optically thin case.

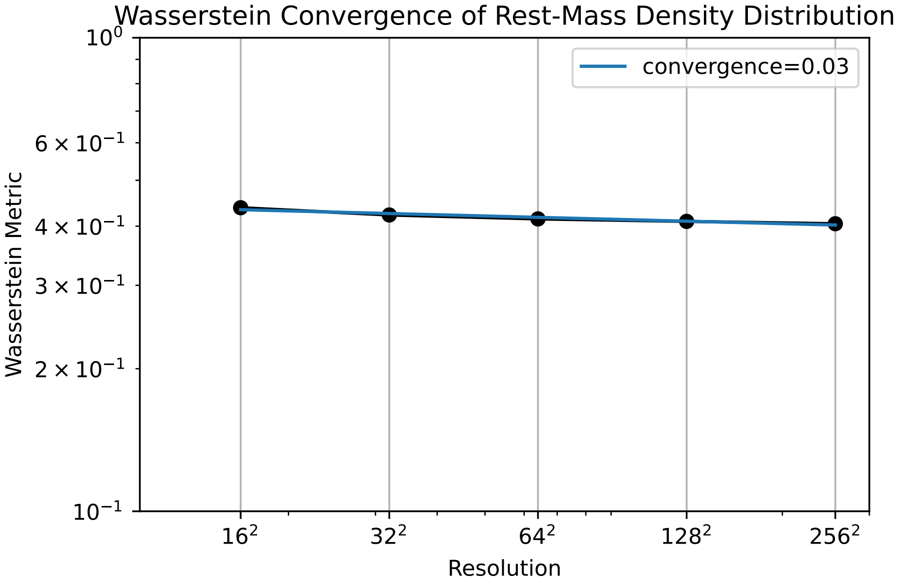

(b) The optically thick case.

Figure 7.14: The Wasserstein metric of the difference in mean rest-mass density distributions between optically thin (top panel) and optically thick (bottom panel) radiation hydrodynamic with the purely hydrodynamic relativistic Kelvin-Helmholtz simulations at spatial resolutions 16^2 to 256^2. A linear regression is shown in blue. The lack of even weak convergence suggest that the different physics models used, and the differences in their parameters can be clearly distinguished in the statistical properties.

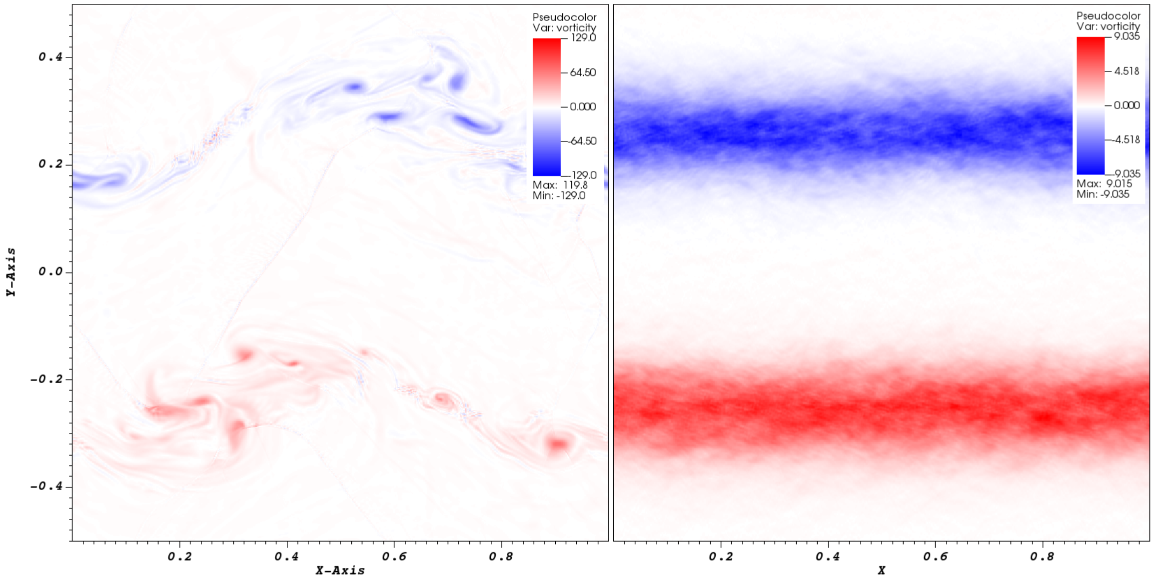

Figure 7.15: A single realisation (in the left panel) and the mean of a set of distinct simulations (on the right) of the vorticity field resulting from relativistic Kelvin-Helmholtz simulations evolved to time $t = 2$ in geometric units. We can see that the upper and lower shearing mixing layers between the two initial states contain vortices with opposing rotational directions. Taking the mean over numerous simulations approaches independence in the x-dimension, just as in the rest-mass density field shown in 7.6.

fluid velocity field ($w = \nabla \times v$). For a 2D velocity field (here I use the x-y plane) only the z-component of the vorticity can be non-zero, given by $w^z = \partial_x v^y - \partial_y v^x$.

In the left panel of figure 7.15 we see a single realisation of the vorticity field for the relativistic hydrodynamic Kelvin-Helmholtz instability simulation introduced in section 7.2.1. We can identify qualitative features such as the vorticity being strongest in the mixing layer, and the vortices corresponding to the centres of the roll up features visible in figure 7.2. The vortices in the upper and lower mixing layers have opposing rotational directions, as a result of following the rotational fluid flow. The shocks present in the central higher density layer are difficult to distinguish in the vorticity profile. The right panel of figure 7.15 shows the mean vorticity field taken over a number of individual solutions. We can see that this quantity approaches an independence in the x-dimension, similar to the mean rest-mass density field shown in 7.6.

Figure 7.16 shows that the mean over a set of realisations of the vorticity field struggles to converge towards a continuum limit as the spatial resolution is increased. This is to be expected as finer spatial resolutions allow for tighter vortices before the

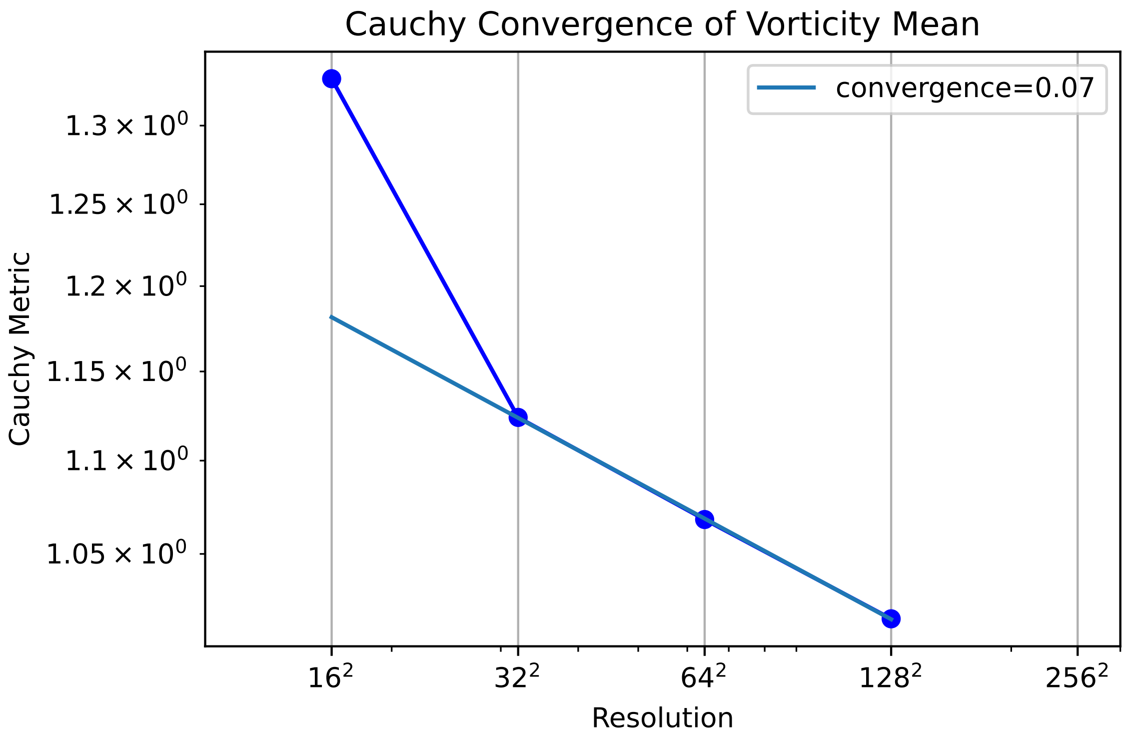

Figure 7.16: The L_2-norms of the difference in mean vorticity profiles from a set of relativistic hydrodynamic KHI simulations for a range of spatial resolutions.

fluid becomes averaged over a single computational cell, and hence allows for more extreme vorticity. The finer resolution allows for the existence of more extreme local maxima and minima, whilst preserving volume averaged quantities, thus allowing for ever steepening spatial gradients.

This can be qualitatively discussed by considering what the continuum limit of the vorticity at a shock means, or would converge to. The shock itself is defined distributionally as a discontinuity. The vorticity, as a derivative of that, would "look like a δ function" at the shock. This should not be expected to converge to a pointwise solution. The results shown here indicate that a statistical distribution of δ function-like behaviour also should not be expected to converge pointwise, even in a weak sense.

To check that this is not an artefact of the curl, I look at a different type of differential operator. The magnitude of the gradient of the rest-mass density is also a differential operator, and is given by $|\nabla\rho|$. For a 2D problem this reduces to $[(\partial_x\rho)^2 + (\partial_y\rho)^2]^{1/2}$.

In the left panel of figure 7.17 we see a single realisation of the density gradient field for the relativistic hydrodynamic Kelvin-Helmholtz instability simulation introduced

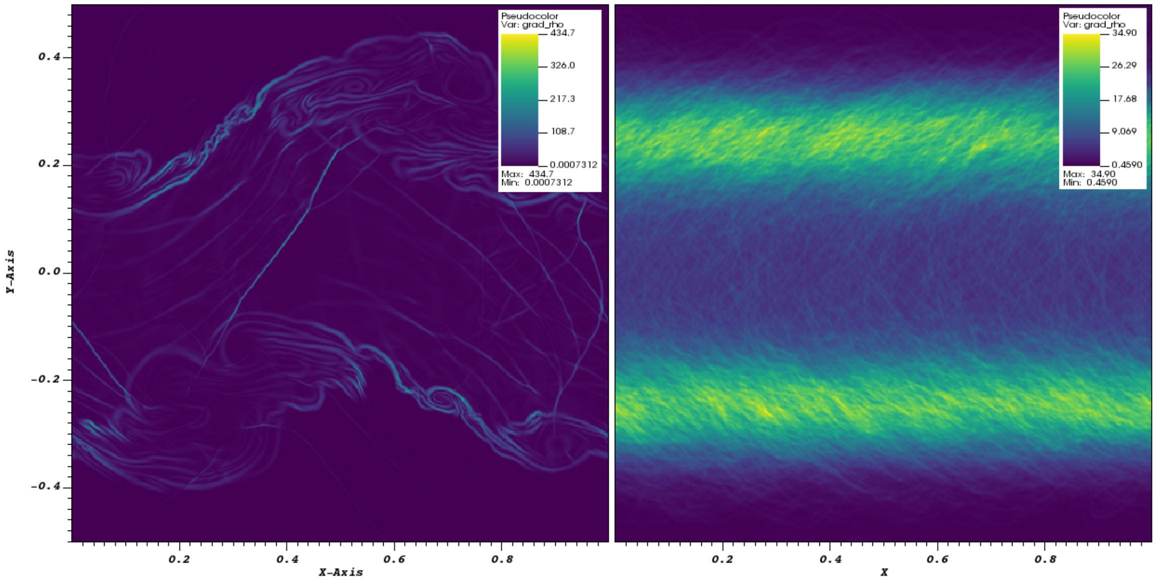

Figure 7.17: Single realisation (left panel) and mean from a set of distinct simulations (right panel) of the magnitude of the spatial gradient of rest-mass density, evolved to time $t = 2$ in geometric units. We can see how the initial discontinuity in rest-mass density has rolled up to follow the vortex flow. New discontinuities have been introduced in between the mixing layers due to internal interactions. Taking the mean over numerous simulations the average profile approaches independence in the x-location, as in the rest-mass density field shown in 7.6. This approach is very slow due to the sharpness of the the gradient magnitudes.

in section 7.2.1. This highlights steep changes in the rest mass density field such as the discontinuities found around shocks. We can identify the steep changes in the mixing layers, and internal shocks throughout the solution. The right panel of figure 7.17 shows the mean density gradient field taken over a number of realisations. The small scale details in this profile suggest that a very large number of samples would be required to obtain a smooth mean solution. Furthermore such a smooth mean solution would exhibit independence of the x-position, just as with the mean curl profile and the rest mass density profiles in figure 7.6.

In figure 7.18 we can see that, as with the vorticity, the mean over a set of realisations of density gradient struggles to converge towards a continuum limit as the spatial resolution is increased. This is also to be expected as finer spatial resolution allows for the existence of increasingly sharp gradients. Just like in the vorticity analysis, the δ function-like limit would not be expected to converge to a pointwise solution. These results indicate that a distribution of resolved density gradient profiles do not (even weakly) converge to a statistical solution.

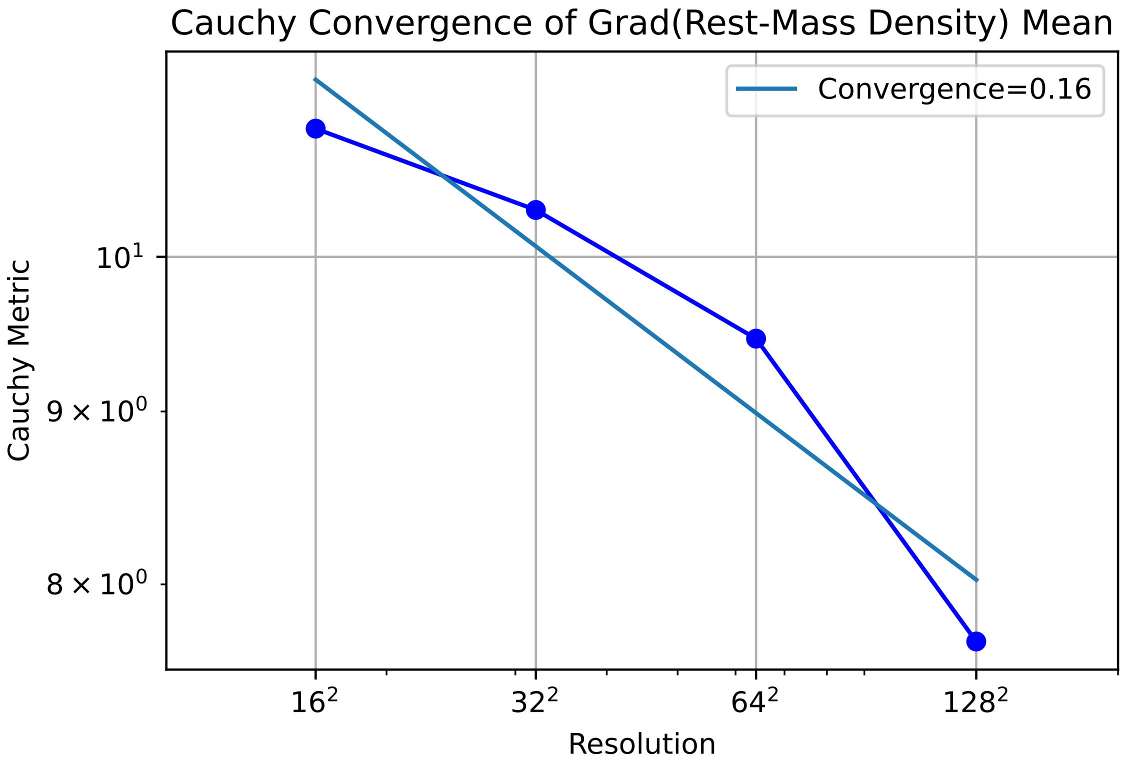

Figure 7.18: The L_2-norms of the difference in mean density gradient magnitude profiles from a set of relativistic hydrodynamic KHI simulations for a range of spatial resolutions.

7.8.2 Integral Operators

I am interested in whether non-local quantities involving integral operators such as the optical depth converge towards a continuum limit as the spatial resolution is increased. The optical depth of a volume of fluid is useful in determining its rate of radiative cooling, and how a radiation field evolves within the fluid due to radiative transfer. The optical depth can be obtained by integrating the opacity along a line of sight. It is non-locally defined and involves an integral operator as opposed to a differential operator, as discussed in section 7.8.1. The optical depth is given by

$$\tau = \int_{y_0}^{y_1} \chi_a dy. \tag{7.8}$$

For this analysis I choose the opacity to equal the rest-mass density: $\chi_a = \rho$.

The optical depths are calculated by integrating the rest-mass density fields over the y-direction: $\tau_k(x_i) = \int_0^1 \rho_k(x_i, y)dy$. At each x-position there is a distribution of optical depths, each corresponding to a uniquely generated initial condition.

The L_2-norm of the difference in mean optical depths has been used to quantify the effect of increasing the spatial resolution on the distribution $\tau_k(x_i)$ (where the lower resolution simulations have been upscaled to allow numerical comparison). This has been reproduced for the optically thin and optically thick coupled radiation hydrodynamic models. Figure 7.19 shows that in both cases the mean optical depth does indeed weakly converge towards a continuum limit as the spatial resolution increases.

7.9 Summary

All numerical simulations of neutron stars, including our evolutions of neutron star merger remnants in chapter 6, have the same issue: key quantities are uncertain. In some cases the aim is to constrain these uncertainties by experiment (the microphysical equation of state being the key quantity here). In other circumstances there is very little chance of constraining the property experimentally, or the property is purely an artefact of the numerical simulation process (such as aspects of the initial-boundary conditions). In these circumstances we need to quantify the systematic uncertainties introduced in the numerical modelling process.

Statistical results from Newtonian hydrodynamic turbulence simulations have

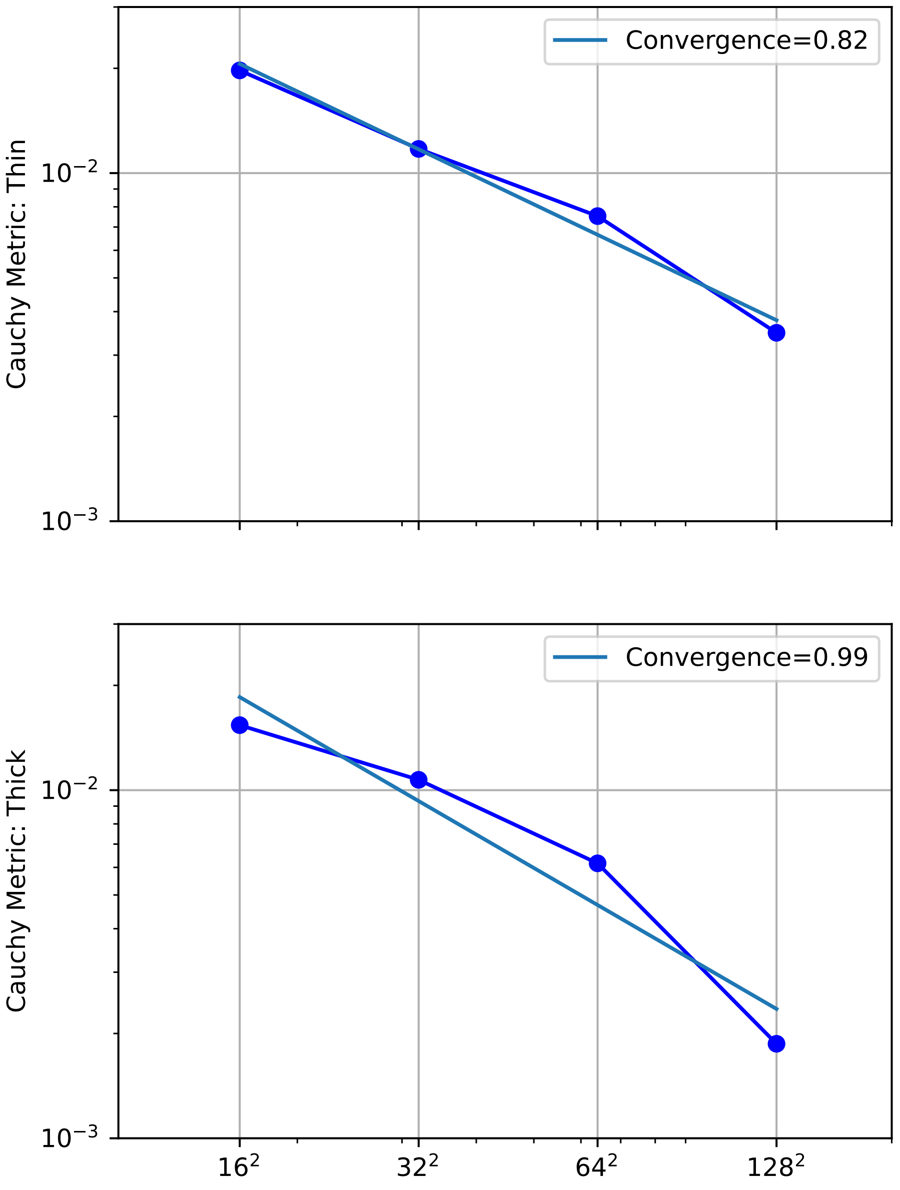

Figure 7.19: The L_2-norm of the difference in mean optical depth optically thin (top) and optically thick (bottom) relativistic Kelvin-Helmholtz simulations for a range of spatial resolutions. Linear regressions are shown with blue lines.

suggested that there are fundamental problems with uncertainty quantification in the case of shock turbulence interaction. In particular, it is possible that quantities may not converge pointwise. The suggested solution is to consider the (weak) convergence of statistical properties. These do converge in Newtonian cases, rescuing well-posedness. This chapter has shown that precisely the same issues and solution methods appear for relativistic hydrodynamic and relativistic radiation hydrodynamics simulations.

I have shown that for an individual realisation of initial condition perturbations, which grow into turbulent flow, the numerical solution evolved with relativistic hydrodynamics does not converge with increasing spatial resolution. I have shown that when a sufficient number of solutions are calculated, each generated from randomly sampled initial condition perturbations, then the mean and variance do indeed weakly converge with increasing spatial resolution. Thus the expected distribution of solutions arising from a distribution of initial conditions can be calculated.

These experiments have been reproduced involving radiation fields, including both optically thin and optically thick cases. I have shown that the mean of radiation hydrodynamic solutions converge to distinct solutions to the purely hydrodynamic case.

Introducing radiation modifies the fluid evolution by inducing additional motion in the y-direction. The horizontal region of larger rest mass density splits into upper and lower bands, with a region of lower density separating them. This effect becomes more exaggerated for the larger opacity as this increases the strength of the coupling between the radiation and fluid. This allows us to state a minimum opacity for which we can genuinely see differences between purely hydrodynamic and radiation coupled hydrodynamic systems. This observable difference manifests in both the single simulation results and the averaged results, indicating that the differences in solution caused by introducing radiation transport are independent of the perturbations in the initial condition. There are clear distinctions between the solutions of purely hydrodynamic, weakly coupled radiation hydrodynamic, and strongly coupled radiation hydrodynamic simulations.

The statistical properties of non-local operators, such as differential and integral operators of conserved quantities have been investigated. The fluid's vorticity is composed of spatial gradients of the fluid velocity field. Neither the individual realisations of the vorticity, nor the mean over a set of realisations of the vorticity converges as the spatial resolution is increased. This is to be expected as finer resolutions allow for tighter vortices before the fluid becomes averaged over the smallest computational cell, and hence allows for more extreme vorticity. The finer resolution allows for the existence of more extreme local maxima and minima, whilst preserving volume averaged quantities, thus allowing ever steepening spatial gradients. The magnitude of the gradient of the rest-mass density is also a differential operator, and as expected also exhibits this non-convergent statistical property. We cannot make any statistical conclusions about non-locally defined differential quantities such as vorticity and rest-mass density gradient: they are not considered observable quantities.

The optical depth of a volume of fluid can be obtained by integrating its opacity along a line of sight. It is non-locally defined and involves an integral operator as opposed to a differential operator. The optical depth appears to converge towards a continuum limit with increasing spatial resolution when analysing the mean over a number of samples. This successful convergence test works for the purely hydrodynamic and the coupled radiation hydrodynamic models for a range of opacities. This suggests that integral non-local operators are to be preferred when using the results of nonlinear numerical simulations as experimental observables.

Chapter 8

Conclusions

In this thesis I have investigated the impact of photon radiation on neutron star merger remnants. The numerical simulations I have conducted suggest that the presence of a radiation field can cause the rapid dispersion of material surrounding a remnant object. This could have repercussions for the material available to accrete onto the remnant at a later time. Additionally, discretised advective fluxes can be significant source of error in numerical fluid simulations, in certain situations, for example with approximately hydrostationary configurations, a coordinate transform can be applied to reduce this effect. Through the implementation of a new general relativistic radiation hydrodynamics code I have looked at the importance of flexible and efficient mesh refinement and numerical methods on the results. I have also looked at uncertainty quantification, seeing how the uncertainties in (for example) the initial data can have a qualitative impact on the observables.

Chapters 3-5 have concentrated on the introduction and validation of the new code. The important contributions are the use of the AMReX framework, and the "ground-up" reconsideration of the methods to be used in the context of neutron star merger remnants. We have seen that within a modern framework it is possible to build a complex multiphysics code for efficiently simulating a specific astrophysical phenomenon. Whilst frameworks such as the Einstein Toolkit allow a much broader range of problems to be tackled, the use of fixed spacetimes, and the careful choice of numerical methods, allow for substantial numerical efficiencies. This has been quantified in chapters 4 and 5.

I emphasise again that the code designed and implemented here differs from other state of the art codes in the field. We can compare with GRaM-X [3] which is a recent AMReX based GRMHD code which uses the standard approaches from the Einstein Toolkit but makes use of the AMReX functionality to do the grid management. GRaM-X is designed, as the Einstein Toolkit was, to be a general purpose code. It primarily uses AMReX in order to get the best performance on modern exascale machines. In particular it relies on AMReX's GPU functionality but still uses numerical methods with broad stencils and evolves the spacetime as well as the matter. In our case I have concentrated on using AMReX for its flexibility and efficient patch structures. I make additional efficiency gains by also targeting a specific astrophysical case, which also allows us to make efficient choices of numerical methods and to use a fixed spacetime.

An illustration of the utility of this approach is given by the low computational cost of the merger remnant simulations shown in chapter 6. The astrophysical conclusions of the impact of radiation on the post-merger remnant would be considerably more costly and complex to compute using, for example, the standard Einstein Toolkit approach.

Efficiency is also important in the discussion of uncertainty quantification in chapter 7. The need to run many samples to quantify the uncertainty in parameter space means that computational efficiency is essential. Building on the efficiency gains, I have extended the results of Newtonian work on statistical solutions and convergence in two ways. Firstly I have shown that the strong (pointwise) convergence of solutions to relativistic (radiation) hydrodynamic equations with shock-turbulence interaction may fail, but that this can be rescued using weak (statistical) convergence. Secondly, I have extended this work to look at the convergence of non-local operators of the evolved variables, showing that non-local integral operators do converge cleanly in all cases, but that non-local differential operators fail to converge even weakly.

Further questions remain to be investigated from this work. Here I highlight just three.

Firstly from an astrophysics point of view the main source of radiation emission in neutron star mergers is currently believed to be neutrinos rather than photons. The implementation of neutrinos is conceptually similar but practically more complex, hence our focus on photon radiation here. Extending to neutrinos would be important to compare with astrophysical observations. It would require more complex absorption and emission terms and likely more complex numerical methods.

Secondly from a numerical point of view the treatment of stiff source terms should be investigated in more detail. The approach used in chapter 3 prevents overshoots in the evolution of energy and temperature. However for more complex situations such as the inclusion of neutrinos it is likely that more accurate methods will be required. Possible approaches include semi-implicit methods (as in [110]) and spectral deferred corrections (as in [111]).

Finally from a mathematical point of view we need to understand how to efficiently extend the uncertainty quantification work in chapter 7 to functional uncertainties. Here the important case is the uncertainties in the equation of state (or similar closure relations). Constraining the properties of the equation of state is a fundamental goal of multi-messenger astronomy. The current work shows that quantifying the uncertainty with which the equation of state can be constrained is, even in principle, a very complex process.

References

[1] N. Andersson, *Gravitational-wave astronomy: exploring the dark side of the universe*. Oxford Graduate Texts, 2019.

[2] F. Löffler, J. Faber, E. Bentivegna, T. Bode, P. Diener, R. Haas et al., *The Einstein Toolkit: a community computational infrastructure for relativistic astrophysics*, *Classical and Quantum Gravity* **29** (2012) 115001.

[3] S. Shankar, P. Mösta, S. R. Brandt, R. Haas, E. Schnetter and Y. de Graaf, *GRaM-X: A new GPU-accelerated dynamical spacetime GRMHD code for exascale computing with the Einstein Toolkit*, 2022.

[4] B. W. Carroll and D. A. Ostlie, *An introduction to modern astrophysics*, 1996.

[5] C. D. Ott, E. Abdikamalov, P. Mösta, R. Haas, S. Drasco, E. P. O'Connor et al., *General-relativistic simulations of three-dimensional core-collapse supernovae*, *The Astrophysical Journal* **768** (2013) 115.

[6] P. Haensel, A. Y. Potekhin and D. G. Yakovlev, *Neutron stars 1: Equation of state and structure*, vol. 326. Springer Science & Business Media, 2007.

[7] S. Chandrasekhar, *The highly collapsed configurations of a stellar mass (second paper)*, *Monthly Notices of the Royal Astronomical Society* **95** (1935) 207–225.

[8] T. Fischer, M. Hempel, I. Sagert, Y. Suwa and J. Schaffner-Bielich, *Symmetry energy impact in simulations of core-collapse supernovae*, *The European Physical Journal A* **50** (2014) 1–18.

[9] A. Burrows, *Supernova explosions in the universe*, *Nature* **403** (2000) 727–733.

[10] J. A. Faber and F. A. Rasio, *Binary neutron star mergers*, *Living Reviews in Relativity* **15** (2012) 1–83.

[11] D. R. Lorimer, *Binary and millisecond pulsars*, *Living Reviews in Relativity* **11** (Nov., 2008) .

[12] B. P. Abbott, R. Abbott and T. D. Abbott, *Multi-messenger observations of a binary neutron star merger*, *The Astrophysical Journal Letters* **848** (Oct., 2017) L12.

[13] M. Nicholl, E. Berger, D. Kasen, B. Metzger, J. Elias, C. Briceño et al., *The electromagnetic counterpart of the binary neutron star merger LIGO/Virgo GW170817. III. Optical and UV spectra of a blue kilonova from fast polar ejecta*, *The Astrophysical Journal Letters* **848** (2017) L18.

[14] B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams et al., *GW170817: observation of gravitational waves from a binary neutron star inspiral*, *Physical Review Letters* **119** (2017) 161101.

[15] D. B. Fielding, E. C. Ostriker, G. L. Bryan and A. S. Jermyn, *Multiphase gas and the fractal nature of radiative turbulent mixing layers*, *The Astrophysical Journal Letters* **894** (2020) L24.

[16] B. D. Metzger, *Kilonovae*, *Living reviews in relativity* **20** (2017) 3.

[17] N. Andersson and G. L. Comer, *Relativistic fluid dynamics: physics for many different scales*, *Living reviews in relativity* **24** (2021) 3.

[18] O. Zanotti, C. Roedig, L. Rezzolla and L. Del Zanna, *General relativistic radiation hydrodynamics of accretion flows–I. Bondi–Hoyle accretion*, *Monthly Notices of the Royal Astronomical Society* **417** (2011) 2899–2915.

[19] D. Radice and I. Hawke, *Turbulence modelling in neutron star merger simulations*, 2024.

[20] P. Hammond, I. Hawke and N. Andersson, *Thermal aspects of neutron star mergers*, *Phys. Rev. D* **104** (Nov, 2021) 103006.

[21] M. Planck, *The theory of heat radiation.* Blakiston, 1914.

[22] S. Matz, G. Share, M. Leising, E. Chupp, W. Vestrandt, W. Purcell et al., *Gamma-ray line emission from SN1987A*, *Nature* **331** (1988) 416.

[23] B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams et al., *Gravitational waves and gamma-rays from a binary neutron star merger:*

GW170817 and GRB 170817A, *The Astrophysical Journal Letters* **848** (2017) L13.

[24] S. Valenti, J. David, S. Yang, E. Cappellaro, L. Tartaglia, A. Corsi et al., *The discovery of the electromagnetic counterpart of GW170817: kilonova AT 2017gfo/DLT17ck*, *The Astrophysical Journal Letters* **848** (2017) L24.

[25] S. Chandrasekhar, *Radiative transfer.* Courier Corporation, 2013.

[26] P. C. Hammond, *Numerical evolution of binary neutron star mergers.* PhD thesis, University of Southampton, 2022.

[27] T. G. Cowling, *The non-radial oscillations of polytropic stars*, *Monthly Notices of the Royal Astronomical Society* **101** (1941) 367.

[28] R. L. Arnowitt, S. Deser and C. W. Misner, *The Dynamics of general relativity*, *Gen. Rel. Grav.* **40** (2008) 1997–2027, [`gr-qc/0405109`].

[29] C. W. Misner, K. S. Thorne, J. A. Wheeler and D. I. Kaiser, *Gravitation.* Princeton University Press, 2017.

[30] M. Alcubierre, *Introduction to 3+ 1 numerical relativity*, vol. 140. Oxford University Press, 2008.

[31] C. Palenzuela, *Introduction to numerical relativity*, *Frontiers in Astronomy and Space Sciences* **7** (Sept., 2020) .

[32] R. A. d'Inverno, *Introducing Einstein's relativity.* Clarendon Press, 1992.

[33] J. A. Font, *Numerical hydrodynamics and magnetohydrodynamics in general relativity*, *Living reviews in relativity* **11** (2008) 7.

[34] J. M. Martí and E. Müller, *Grid-based methods in relativistic hydrodynamics and magnetohydrodynamics*, *Living Reviews in Computational Astrophysics* **1** (2015) 3.

[35] L. Rezzolla and O. Zanotti, *Relativistic hydrodynamics.* Oxford University Press, 2013.

[36] M. Shibata, K. Kiuchi, Y.-i. Sekiguchi and Y. Suwa, *Truncated moment formalism for radiation hydrodynamics in numerical relativity*, *Progress of Theoretical Physics* **125** (2011) 1255–1287.

[37] T. Kuroda, K. Kotake and T. Takiwaki, *Fully general relativistic simulations of*

core-collapse supernovae with an approximate neutrino transport, *The Astrophysical Journal* **755** (2012) 11.

[38] J. M. Martí, J. M. Ibánez and J. A. Miralles, *Numerical relativistic hydrodynamics: Local characteristic approach*, *Physical Review D* **43** (1991) 3794.

[39] B. D. Farris, T. K. Li, Y. T. Liu and S. L. Shapiro, *Relativistic radiation magnetohydrodynamics in dynamical spacetimes: Numerical methods and tests*, *Physical Review D* **78** (2008) 024023.

[40] M. Hempel and J. Schaffner-Bielich, *A statistical model for a complete supernova equation of state*, *Nuclear Physics A* **837** (2010) 210–254.

[41] E. Murchikova, E. Abdikamalov and T. Urbatsch, *Analytic closures for M1 neutrino transport*, *Monthly Notices of the Royal Astronomical Society* **469** (2017) 1725–1737.

[42] S. L. Shapiro and S. A. Teukolsky, *Black holes, white dwarfs, and neutron stars: The physics of compact objects.* John Wiley & Sons, 2008.

[43] E. Toro, *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction.* Applied mechanics: Researchers and students. Springer, 1999.

[44] 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** (2004) 677–694.

[45] D. Mihalas and B. W. Mihalas, *Foundations of radiation hydrodynamics.* Courier Corporation, 2013.

[46] A. J. Wright and I. Hawke, *A resistive extension for ideal magnetohydrodynamics*, *Monthly Notices of the Royal Astronomical Society* **491** (2020) 5510–5523.

[47] L. Baiotti and L. Rezzolla, *Binary neutron star mergers: a review of Einstein's richest laboratory*, *Reports on Progress in Physics* **80** (2017) 096901.

[48] L. Kidder, M. Scheel, S. Teukolsky, E. Carlson and G. Cook, *Black hole evolution by spectral methods*, *Physical Review D* **62** (Sept., 2000) .

[49] R. LeVeque, *Numerical methods for conservation laws. Lectures in mathematics*

ETH Zürich, department of mathematics research institute of mathematics, 1992.

[50] M. J. Berger and J. Oliger, *Adaptive mesh refinement for hyperbolic partial differential equations*, *Journal of Computational Physics* **53** (1984) 484–512.

[51] K. Sverdrup, N. Nikiforakis and A. Almgren, *Highly parallelisable simulations of time-dependent viscoplastic fluid flow with structured adaptive mesh refinement*, *Physics of Fluids* **30** (2018) .

[52] A. Almgren, J. Bell and M. Zingale, *Maestro: A low mach number stellar hydrodynamics code*, in *Journal of Physics: Conference Series*, vol. 78, p. 012085, IOP Publishing, 2007.

[53] A. Almgren, M. B. Sazo, J. Bell, A. Harpole, M. Katz, J. Sexton et al., *Castro: A massively parallel compressible astrophysics simulation code*, *Journal of Open Source Software* **5** (2020) 2513.

[54] W. E. Schiesser, *The numerical method of lines: integration of partial differential equations.* Elsevier, 2012.

[55] C. B. Laney, *Computational gasdynamics.* Cambridge university press, 1998.

[56] C.-W. Shu and S. Osher, *Efficient implementation of essentially non-oscillatory shock-capturing schemes*, *Journal of computational physics* **77** (1988) 439–471.

[57] D. I. Ketcheson, S. Gottlieb and C. B. Macdonald, *Strong stability preserving two-step Runge–Kutta methods*, *SIAM Journal on Numerical Analysis* **49** (2011) 2618–2639.

[58] S. Gottlieb and D. Ketcheson, *Chapter 21 - Time discretization techniques*, in *Handbook of Numerical Methods for Hyperbolic Problems* (R. Abgrall and C.-W. Shu, eds.), vol. 17 of *Handbook of Numerical Analysis*, pp. 549–583. Elsevier, 2016. DOI.

[59] A. Harten, *High resolution schemes for hyperbolic conservation laws*, *Journal of computational physics* **135** (1997) 260–278.

[60] W. Kastaun, *High-resolution shock capturing scheme for ideal hydrodynamics in general relativity optimized for quasistationary solutions*, *Physical Review D* **74** (2006) 124024.

[61] P. Mösta, B. C. Mundim, J. A. Faber, R. Haas, S. C. Noble, T. Bode et al.,

GRHydro: a new open-source general-relativistic magnetohydrodynamics code for the Einstein Toolkit, *Classical and Quantum Gravity* **31** (Nov., 2013) 015005.

[62] L. Baiotti, I. Hawke, P. J. Montero, F. Löffler, L. Rezzolla, N. Stergioulas et al., *Three-dimensional relativistic simulations of rotating neutron-star collapse to a Kerr black hole*, *Physical Review D* **71** (Jan., 2005) .

[63] B. Fornberg and J. Zuev, *The Runge phenomenon and spatially variable shape parameters in RBF interpolation*, *Computers & Mathematics with Applications* **54** (2007) 379–398.

[64] W. J. Thompson, *Fourier series and the Gibbs phenomenon*, *American journal of physics* **60** (1992) 425–429.

[65] B. Van Leer, *Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method*, *Journal of computational Physics* **32** (1979) 101–136.

[66] S. K. Godunov, *A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics*, *Matematicheskii Sbornik* **89** (1959) 271–306.

[67] M. Berger, M. Aftosmis and S. Muman, *Analysis of slope limiters on irregular grids*, in *43rd AIAA Aerospace Sciences Meeting and Exhibit*, p. 490, 2005.

[68] R. J. LeVeque, *Finite volume methods for hyperbolic problems*, vol. 31. Cambridge university press, 2002.

[69] D. Zhang, C. Jiang, L. Cheng and D. Liang, *A refined r-factor algorithm for TVD schemes on arbitrary unstructured meshes*, *International Journal for Numerical Methods in Fluids* **80** (2016) 105–139.

[70] Y.-T. Zhang and C.-W. Shu, *ENO and WENO schemes*, in *Handbook of numerical analysis*, vol. 17, pp. 103–122. Elsevier, 2016.

[71] C.-W. Shu, *Essentially non-oscillatory and weighted essentially non-oscillatory schemes*, *Acta Numerica* **29** (2020) 701–762.

[72] G. Gerolymos, D. Sénéchal and I. Vallet, *Very-high-order WENO schemes*, *Journal of Computational Physics* **228** (2009) 8481–8524.

[73] E. F. Toro, M. Spruce and W. Speares, *Restoration of the contact surface in the HLL-Riemann solver*, *Shock waves* **4** (1994) 25–34.

[74] J. M. Martí and E. Müller, *Numerical hydrodynamics in special relativity*, *Living Reviews in Relativity* **6** (2003) 7.

[75] V. Schneider, U. Katscher, D. Rischke, B. Waldhauser, J. Maruhn and C.-D. Munz, *New algorithms for ultra-relativistic numerical hydrodynamics*, *Journal of Computational Physics* **105** (1993) 92–107.

[76] A. Harten, P. D. Lax and B. v. Leer, *On upstream differencing and Godunov-type schemes for hyperbolic conservation laws*, *SIAM review* **25** (1983) 35–61.

[77] 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** (1997) 221.

[78] A. Mignone and G. Bodo, *An HLLC Riemann solver for relativistic flows—I. Hydrodynamics*, *Monthly Notices of the Royal Astronomical Society* **364** (2005) 126–136.

[79] R. Courant, K. Friedrichs and H. Lewy, *On the partial difference equations of mathematical physics*, *IBM journal of Research and Development* **11** (1967) 215–234.

[80] S. C. Noble, C. F. Gammie, J. C. McKinney and L. Del Zanna, *Primitive variable solvers for conservative general relativistic magnetohydrodynamics*, *The Astrophysical Journal* **641** (2006) 626.

[81] W. Kastaun, J. V. Kalinani and R. Ciolfi, *Robust recovery of primitive variables in relativistic ideal magnetohydrodynamics*, *Physical Review D* **103** (Jan., 2021) .

[82] L. R. Weih, H. Olivares and L. Rezzolla, *Two-moment scheme for general-relativistic radiation hydrodynamics: A systematic description and new applications*, *Monthly Notices of the Royal Astronomical Society* **495** (2020) 2285–2304.

[83] W. Zhang, A. Myers, K. Gott, A. Almgren and J. Bell, *AMReX: Block-structured adaptive mesh refinement for multiphysics applications*, *The International Journal of High Performance Computing Applications* **35** (2021) 508–526.

[84] M. J. Berger and P. Colella, *Local adaptive mesh refinement for shock hydrodynamics*, *Journal of computational Physics* **82** (1989) 64–84.

[85] R. Farber, "AMReX: A performance-portable framework for block-structured adaptive mesh refinement applications." `https://www.exascaleproject.org/highlight/amrex-a-performance-portable-framework-for-block-structured-adaptive-mesh-refinement-applications/`, 2023.

[86] J. Bell, M. Berger, J. Saltzman and M. Welcome, *Three-dimensional adaptive mesh refinement for hyperbolic conservation laws*, *SIAM Journal on Scientific Computing* **15** (1994) 127–138.

[87] I. Hawke, "GitHub repository: slides/et-2022." `https://github.com/IanHawke/slides/tree/f05e97f89193baac021d3aa44eecde691754469b/et-2022/figures`, 2022.

[88] A. S. Almgren, V. E. Beckner, J. B. Bell, M. S. Day, L. H. Howell, C. C. Joggerst et al., *CASTRO: A new compressible astrophysical solver. I. Hydrodynamics and self-gravity*, *The Astrophysical Journal* **715** (may, 2010) 1221.

[89] P. Pacheco, *Parallel Programming with MPI.* Elsevier Science, 1997.

[90] L. Dagum and R. Menon, *OpenMP: an industry standard API for shared-memory programming*, *Computational Science & Engineering, IEEE* **5** (1998) 46–55.

[91] Z. Lan, V. E. Taylor and G. Bryan, *A novel dynamic load balancing scheme for parallel systems*, *Journal of Parallel and Distributed Computing* **62** (2002) 1763–1781.

[92] J. Tao, "Cactus runs on 131,072 cores on Blue Gene/P at ANL." `https://www.cactuscode.org/media/news/BGP-131072/index.html`.

[93] E. Schnetter, "Using CarpetX: A guide for early adopters." `https://einsteintoolkit.org/seminars/2021_03_18/CarpetX.pdf`.

[94] W. Zhang, A. Almgren, M. Day, T. Nguyen, J. Shalf and D. Unat, *Boxlib with tiling: An amr software framework*, .

[95] M. Zingale, A. S. Almgren, M. G. Barrios Sazo, V. E. Beckner, J. B. Bell, B. Friesen et al., *Meeting the challenges of modeling astrophysical*

thermonuclear explosions: Castro, Maestro, and the AMReX astrophysics suite, *Journal of Physics: Conference Series* **1031** (may, 2018) 012024.

[96] W. Zhang and A. I. MacFadyen, *RAM: A relativistic adaptive mesh refinement hydrodynamics code*, *The Astrophysical Journal Supplement Series* **164** (2006) 255.

[97] A. Sadowski, R. Narayan, A. Tchekhovskoy and Y. Zhu, *Semi-implicit scheme for treating radiation under M1 closure in general relativistic conservative fluid dynamics codes*, *Monthly Notices of the Royal Astronomical Society* **429** (2013) 3533–3550.

[98] C. Roedig, O. Zanotti and D. Alic, *General relativistic radiation hydrodynamics of accretion flows–II. Treating stiff source terms and exploring physical limitations*, *Monthly Notices of the Royal Astronomical Society* **426** (2012) 1613–1631.

[99] J. M. Lattimer, C. Pethick, M. Prakash and P. Haensel, *Direct URCA process in neutron stars*, *Physical review letters* **66** (1991) 2701.

[100] A. W. Steiner, M. Hempel and T. Fischer, *Core-collapse supernova equations of state based on neutron star observations*, *The Astrophysical Journal* **774** (Aug., 2013) 17.

[101] I. Hawke and F. Loeffler, "EinsteinToolkit documentation: TOVSolver." `https://einsteintoolkit.org/thornguide/EinsteinInitialData/TOVSolver/documentation.html`, 2024.

[102] F. Cipolletta, J. V. Kalinani, B. Giacomazzo and R. Ciolfi, *Spritz: A new fully general-relativistic magnetohydrodynamic code*, *Classical and Quantum Gravity* **37** (2020) 135010.

[103] S. Rosswog, *Modelling astrophysical fluids with particles*, *Proceedings of the International Astronomical Union* **16** (2020) 382–397.

[104] E. Gourgoulhon, P. Grandclément, K. Taniguchi, J.-A. Marck and S. Bonazzola, *Quasiequilibrium sequences of synchronized and irrotational binary neutron stars in general relativity: Method and tests*, *Physical Review D* **63** (Feb., 2001) .

[105] T. W. Baumgarte and S. L. Shapiro, *Numerical integration of Einstein's field equations*, *Physical Review D* **59** (1998) 024007.

[106] R. Abgrall and S. Mishra, *Uncertainty quantification for hyperbolic systems of*

conservation laws, in *Handbook of Numerical Analysis*, vol. 18, pp. 507–544. Elsevier, 2017.

[107] D. Radice, *Advanced numerical approaches in the dynamics of relativistic flows.* PhD thesis, Technische Informationsbibliothek und Universitätsbibliothek Hannover (TIB), 2013.

[108] U. S. Fjordholm, K. Lye, S. Mishra and F. Weber, *Statistical solutions of hyperbolic systems of conservation laws: Numerical approximation*, *Mathematical Models and Methods in Applied Sciences* **30** (2020) 539–609, [`https://doi.org/10.1142/S0218202520500141`].

[109] D. Radice and L. Rezzolla, *THC: A new high-order finite-difference high-resolution shock-capturing code for special-relativistic hydrodynamics*, *Astronomy & Astrophysics* **547** (2012) A26.

[110] D. Radice, S. Bernuzzi, A. Perego and R. Haas, *A new moment-based general-relativistic neutrino-radiation transport code: Methods and first applications to neutron star mergers*, *Monthly Notices of the Royal Astronomical Society* **512** (Mar., 2022) 1499–1521.

[111] M. Zingale, Z. Chen, E. T. Johnson, M. P. Katz and A. S. Clark, *Strong coupling of hydrodynamics and reactions in nuclear statistical equilibrium for modeling convection in massive stars*, `2403.14786`.