THE UNIVERSITY *of* EDINBURGH

Examining the effects of magnetic fields in neutron star mergers through numerical simulations

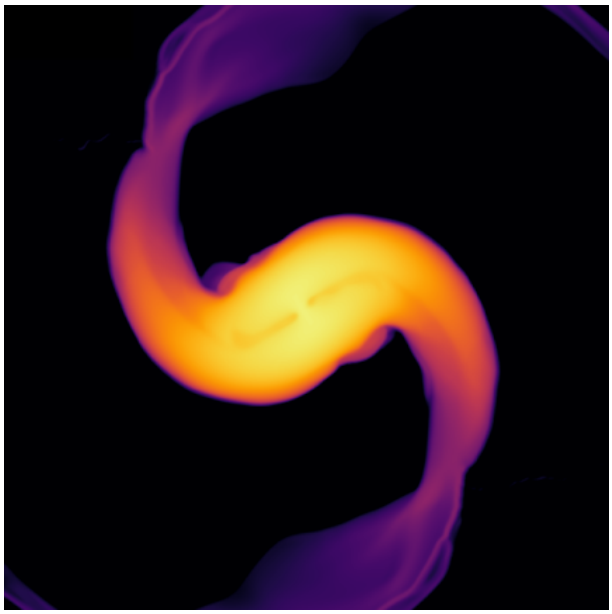

Charlotte Summers

Doctor of Philosophy
University of Edinburgh
November 2022

Declaration

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text.

(*Charlotte Summers*)

Acknowledgements

Firstly, I would like to thank my supervisor, Max Ruffert, for all his advice, enlightening discussion and the many cups of green tea we have enjoyed during our meetings. For her always excellent advice and support and encouragement just when I needed it most, I would also like to thank Anna Lisa Varri.

For the many, many tea time crosswords that brightened my afternoons during lockdown I would thank Emily Roff, Bruno Barton-Singer, Kai Hutenberg and Bethany Warburton. For new skills, many laughs and more than a few bruises, I would also like to thank the Edinburgh University Jiu Jitsu club and the Maths PhD women's football team.

For your constant support through all the ups and downs, I am forever grateful to my mum, Debbie Harris.

And finally, for nearly 10 years of everyday joy no matter how far apart we are, I would like to thank Guy Boyde, I certainly wouldn't have made it this far without you.

Technical Acknowledgements

The FLASH code framework that forms the basis of this work was developed in part by the DOE NNSA- and DOE Office of Science-supported Flash Center for Computational Science at the University of Chicago and the University of Rochester. In addition, the visualisation and analysis work utilised the yt-project Python package [Turk et al., 2011].

The majority of the simulations presented here were performed using the ARCHER2 UK National Supercomputing Service. Additional simulations were performed using the Cirrus UK National Tier-2 HPC Service at EPCC (http://www.cirrus.ac.uk) funded by the University of Edinburgh and EPSRC (EP/P020267/1). Analysis of the simulation outputs and other development work made use of computing resources provided by the Edinburgh Compute and Data Facility (ECDF).

Lay Summary

Gamma ray bursts (GRBs) are flashes of gamma ray radiation, first observed in the late 1960s by a US military satellite program designed to monitor for gamma ray signatures of nuclear weapons testing. With only a handful of observations initially, the source of the bursts was poorly understood and dozens of different theories for their origin were proposed. Over the following decades, further observations by purpose-built gamma ray telescopes revealed that such events occur on a daily basis and that they originate from highly energetic events in other galaxies. This new information narrowed the range of likely theories and eventually a prominent theory emerged, which suggested that some GRBs could be produced by the collision of two neutron stars. In 2017, the observation of gravitational waves from a neutron star merger that coincided with a GRB [Abbott et al., 2017a] provided the first direct evidence in support of this, but the details of what conditions are required in a merger in order for it to produce a GRB are still not fully understood.

Neutron star mergers involve extreme physical conditions that are far beyond what can be recreated in a laboratory, therefore in order to study them we use numerical simulations. In this work we created simulations of neutron star mergers to examine what effect the initial magnetic field of the neutron stars has on the ability of the merger to produce a GRB. In a second project we also created simulations of a black hole and magnetized accretion disc, which can be used to study the remnant of a neutron star merger. Our simulations use Newtonian physics with corrections in order to replicate relativistic effects, as this is much less computationally expensive then using a fully relativistic treatment. In both projects we compare our non-relativistic simulations with fully relativistic results in order to see to what extent these lower computational cost simulations can be used to explore effects and parameters that would be otherwise too expensive to investigate.

Abstract

In this thesis, we present simulations of merging binary neutron stars, carried out using the publicly available FLASH code framework. These are 3D Newtonian magnetohydrodynamic simulations, in which we have included gravitational wave effects through the use of a source term. We trial different implementations of this source term and discuss the results. We then use this model to investigate the role of magnetic fields in binary neutron star mergers. We endow each neutron star with a dipolar magnetic field and examine how the orientation of the dipole affects the strength and structure of the magnetic field during the merger. This has important implications for the ability of the merger remnant to produce a short gamma ray burst jet.

In a second project, we simulate a magnetized accretion torus surrounding a black hole. We implement a model black hole in the FLASH code framework using a Pseudo-Newtonian potential to reproduce features, such as the innermost stable circular orbit, which are important to accretion disc studies. We compare the results of our magnetized accretion disc simulations with similar studies, finding broad agreement with accretion rates and the general structure of the magnetic field.

Contents

Chapter 1

Introduction

1.1 Gamma Ray Bursts (GRBs) and neutron star mergers

Gamma ray bursts (GRBs) are flashes of gamma ray radiation, first observed in the late 1960s. These initial observations were made by the Vela satellites, a US military program designed to monitor for gamma ray signatures of nuclear weapons testing. With only a handful of events, relatively little could be learned from these early observations and a wide range of theories were posited for their origin in the following two decades. The launch of the Compton gamma-ray observatory in 1991, which carried the Burst and Transient Source Explorer (BATSE) instrument vastly increased the number of observed bursts and as well as giving on-sky positions for the observed bursts. This crucially showed that the burst distribution on the sky was isotropic, strongly suggesting a cosmological rather than galactic origin for the sources and narrowing the range of possible progenitor systems. Further to this, the BATSE data showed that the distribution of the duration of GRBs is bimodal, with peaks at 0.2s and 20s and a separation line between the two at around 2s [Kouveliotou et al., 1993]. The duration of the burst is described by the quantity T_{90}, which is defined as the interval between the times at which 5% and 95% of the total gamma ray fluence is detected. This lead to the identification of two subclasses: short GRBs (sGRBs) with T_{90} <2s and long GRBs with T_{90} >2s. A further distinction between the classes is that short bursts have typically have harder spectra than long bursts [Kouveliotou et al., 1993], meaning that more of the power emitted in the burst is at higher energies.

Long GRBs have been found to be associated with Type Ic supernovae, suggesting that they are generated in the collapse of a massive star alongside the supernova [Hjorth et al., 2003], whereas sGRBs showed no such associations. Further to this, the identification of bursts with their host galaxies showed that lGRBs were found to occur in late-type, star-forming galaxies, supporting the connection to supernovae of young massive stars [Bloom et al., 2002] [Fruchter et al., 2006]. On the other hand, sGRBs are found in both early and late-type galaxies and can occur at a much greater distance from their host galaxy than lGRBs [Fong and Berger, 2013]. Merging binary neutron stars and black holes

emerged as a possible progenitor for sGRBs, as they are expected to be less strongly associated with star forming regions. Furthermore, during the formation of the binary, velocity kicks from the supernovae that form the neutron stars could eject the whole binary from its host galaxy, providing an explanation for the large distance offsets observed. Finally, the very rapid variability seen in sGRBs limits the size of the emitting object through causality, therefore the small physical scale of a merger provided further support to the theory. This link between binary neutron star mergers and sGRBs was finally solidified by the coincident detection of a gravitational wave signal from merging binary NSs (GW170817) [Abbott et al., 2017a] and the sGRB 170817a [Abbott et al., 2017b].

In providing extremely valuable data, this gravitational wave event also opened up a host of new avenues to investigate the details of how sGRBs and other electromagnetic emission is produced during the merger. Neutron star mergers are extremely complex systems, combining strong gravity, magnetic fields and the nuclear physics of extremely dense matter. As such they are all but impossible to study in a purely theoretical way, therefore we must turn to numerical simulations in order to gain better understanding of questions such as how they can produce GRBs. For this in particular, magnetic fields are a crucial feature, since they provide the most promising avenue for the production of a sGRB jet. It is theorised that a strong magnetic field can extract the rotational energy of the merger remnant, launching a relativistic outflow along the system's rotation axis. This requires extremely strong fields with a large scale poloidal structure, and producing a model to investigate how exactly these conditions can be met is the subject of this work.

1.2 Summary of thesis

This thesis is organised as follows. Chapter 2 discusses the astrophysical context of neutron star mergers, covering both the observational and theoretical background. Chapter 3 covers the fundamentals of numerical hydrodynamics and a survey of the available public hydrodynamics codes used for modelling astrophysical fluids. Chapter 4 presents our binary neutron star merger model using the FLASH code framework and the results of simulations exploring the effect of different magnetic field configurations. Finally Chapter 5 describes inclusion of a black hole into the FLASH code and presents some accretion disc simulations used to assess its behaviour.

Chapter 2

Astrophysical Background

The study of neutron star mergers and their related observable events (sGRBs, gravitational waves, neutrinos, kilonovae, etc.) has spanned over 30 years and produce a vast array of results. In this chapter we summarize some of the main features of observable events such as GRBs, gravitational waves and kilonovae that are related to neutron star mergers. We also present some of the key theoretical insights into mergers, and describe some of the current questions and challenges in this area.

2.1 Observations

The study of NS mergers was initially prompted by the observations of gamma ray bursts (GRBs), and the proposal that they could be responsible for certain types of GRB. Further to this, their strength as a source of gravitational waves meant that they were also of interest for the first generations of gravitational wave detectors. Here we review the current understanding of observations relevant to the study of NS mergers, including the multimessenger event GW170817/GRB170817a/AT2017gfo, which provided the first 'smoking gun' evidence for the link between BNS mergers and short-GRBs.

2.1.1 Gamma Ray Bursts (GRBs)

The main observed burst of a GRB is referred to as the prompt emission. This length of this emission is used to categorise a burst as short or long, depending on whether the value of T_{90} is shorter or longer than 2s respectively. More recently it was discovered that the prompt emission can also be preceded by a period of precursor emission and/or followed by an afterglow. Each stage of emission has different characteristics and through studying them we can build a broader picture of the environment both before and after the main burst itself.

Prompt emission

The lightcurves of the GRB prompt emission are highly varied, with few overall trends in their structure. Some appear to be comprised of many shorter bursts of

emission, whereas others have a single fast-rise-exponential-decay profile. Furthermore, some sGRBs show extended gamma ray emission at lower energies following the main peak that can last for tens of seconds [Lazzati et al., 2001]. It was suggested that bursts with extended emission represented a separate population of sGRBs produced by different progenitors [Troja et al., 2008], however the lack of other differences between the two populations suggests that the progenitors are still quite similar. [Metzger et al., 2008] suggest that the progenitor for bursts with extended emission could be a binary NS merger, which then forms a proto-magnetar. The prompt emission would then be produced by accretion onto the magnetar and the extended emission by a relativistic wind that extracts the magnetar's rotational energy.

The spectrum of the prompt emission is primarily non-thermal and can be fitted with a Band function [Band et al., 1993], which consists of a power law with an exponential cut-off at low energies and a steeper power law component at high energies. However, in many GRBs a thermal component can be identified [Ryde, 2005], suggesting multiple sources of emission at work in the burst. The prompt emission photons are thought to be produced in internal shocks within the relativistic jet, where shells of matter travelling at different speeds collide [Rees and Meszaros, 1994]. The majority of this emission is observed at sub-MeV energies, however a handful bursts have been observed to produce photons at much higher energies [MAGIC Collaboration et al., 2019]. The dominant emission mechanism is thought to be synchrotron radiation produced by electrons that have been accelerated in the shock. If we consider the magnetic and kinetic energies to be close to equipartition within the shock, we expect the magnetic field strength at the radius of emission from the central engine to be of the order $B \approx 10^3$G [Daigne and Mochkovitch, 1998]. The peak synchrotron photon energy can then be approximated to be

$$E_{peak} = \frac{\hbar e B}{m_e c}\Gamma_e^2\Gamma_{bulk} \approx 500\text{keV}\left(\frac{\Gamma_e}{100}\right)^2\left(\frac{B}{1000\text{G}}\right)\left(\frac{\Gamma_{bulk}}{100}\right) \tag{2.1}$$

where Γ_e is the Lorentz factor of the electrons themselves and Γ_{bulk} is the bulk Lorentz factor of the material with respect to the observer. We see that for Lorentz factors of order 100, the synchrotron emission peaks in the expected range.

A relativistic jet was proposed as the main source of prompt GRB emission as it provided a solution to what is known as the compactness problem. The timescales on which GRBs vary are very short, implying the emitting area must be very small because the whole region must be in causal contact. Taken naively, the density implied by the GRB energies and these small scales would make the emitting region optically thick, which would result in the photons being trapped and thermalized, which does not correspond with the non-thermal spectra we typically see [Cavallo and Rees, 1978]. The contradiction is resolved by invoking the jet theory, as if the emitting region is travelling towards the observer at relativistic velocities, the relativistic beaming effect acts to focus the emitted radiation along the direction of the jet's motion. This means that the emitting

area then appears much smaller than it actually is.

Pre-cursor emission

Prior to the main emission peak, in some GRBs a burst of precursor emission is observed. This appears to be more common for long GRBs, with around 20% of bursts showing precursors [Lazzati, 2005]. For short GRBs, studies with different instruments have detected precursors for 1-10% of bursts [Troja et al., 2010] [Wang et al., 2020]. The precursor emission is fainter than the main burst and often has a different spectrum [Wang et al., 2020], suggesting there is some significant difference in the emission mechanism for the two components. Further to this, the spectra of different precursors also vary, which could imply multiple progenitors for precursor signals. The time delay between the precursor and main burst is mostly of the same order as the duration of the burst itself [Wang et al., 2020], but can be up to around 100s [Troja et al., 2010].

In long GRBs, there are two main models for precursor production: the shock breakout and fireball photosphere models. These are also both possible mechanisms for post-merger precursor emissions in sGRBs. In the shock breakout scenario, the precursor emission is generated by the breakout of the main GRB jet from surrounding matter, either the remnants of the massive star in the case of an lGRB or ejecta from the merging NSs in the sGRB case. As the jet propagates through the surrounding matter it deposits energy, creating a hot cocoon of matter, which may then be capable of producing soft gamma ray emission when it breaks out of the cooler matter [Bromberg et al., 2017]. In the context of GRBs a fireball is a radiation-dominated outflow from the central engine. The fireball photosphere model proposes that the precursor emission is generated by the fireball transitioning from optically thick to thin, allowing photons initially trapped within to escape [Li, 2007].

Further precursor emission mechanisms have been proposed for sGRBs produced in NS mergers, where the emission is generated in an event prior to the merging of the two stars. For example, in the final stages of the star's inspiral, tidal distortions will cause the crust of the stars to rupture. This is expected to trigger restructuring of the magnetic field, which could generate a reconnection flare that is then detected as the precursor emission [Thompson and Duncan, 1995]. Alternatively, interactions between the magnetospheres of the two neutron stars are expected to be capable of producing electromagnetic transients at a variety of energy scales prior to merger, including those that could be consistent with precursor gamma-ray emission [Troja et al., 2010] [Hansen and Lyutikov, 2001].

Afterglow

Many GRBs show significant afterglow emission that lasts for $\sim$1000 s, primarily in x-rays but in a smaller number of cases optical and radio emission also. This emission is thought to be generated by the external shock as the GRB jet collides with the surrounding interstellar medium. The afterglow emission for lGRBs

is substantially brighter than for sGRBs, suggesting that they are either more energetic or that the surrounding matter density is different for the two systems [Margutti et al., 2012]. The emission is characterised by an initial steeply decaying component, which is thought to be associated with the prompt emission, followed by a more slowly declining phase [Nousek et al., 2006].

To summarize, the main component of an sGRB is now broadly accepted to be produced by a relativistic jet launched from the remnant of a NS-NS or potentially NS-BH merger. Questions however remain about what conditions are required to produce such a jet, which we will discuss in more detail in Section 2.2, however further observations of spectra and lightcurves are refining our theories in this area. In contrast the emission mechanism for pre-cursors, particularly in sGRBs, is less well defined, but has the potential to tell us a great deal about the earliest stages of a merger, including insights into the tidal deformability and magnetospheres of the neutron stars. Finally, the afterglow emission can be used to infer features of the merger environment, helping to answer larger questions about binary populations, location and merger rates.

2.1.2 Kilonovae

Kilonovae are supernova-like transient events observed primarily in the optical and near-infrared. Their spectra are thermal and they reach peak luminosities of around $10^{40} - 10^{42}$ erg s^{-1}. These events were first proposed theoretically in an attempt to understand where the elements heavier than iron were produced in the universe, with the emission observed being produced by the decay of radioactive heavy isotopes [Li and Paczyński, 1998]. These elements are created by the rapid capture of free neutrons by a heavy seed nucleus, which is known as the r-process. This only occurs when the neutron capture timescale is shorter than the beta decay timescale, therefore it requires a very neutron-rich and high entropy environment.

Initially, the two most prominent theories for sites where these conditions occur are in the neutrino driven winds from core-collapse supernovae [Mathews et al., 1992] and the ejecta from NS-NS or NS-BH mergers [Lattimer and Schramm, 1974] [Symbalisty and Schramm, 1982]. However, at present simulations of core-collapse supernovae have shown that the very high neutrino fluxes mean that β interactions reduce the neutron fraction of the matter so rapidly that the conditions for the r-process are not reached for any substantial amount of time [Thompson et al., 2001] [Arcones et al., 2006] [Martínez-Pinedo et al., 2012]. At present it is still unclear from observations which of the two sites is the dominant source of r-process elements in the universe. Studies of the r-process element abundances in our galaxy [Wallner et al., 2015] and the dwarf Reticulum galaxy favour an r-process source with consists of rare, high yield events rather the frequent low yield ones [Ji et al., 2016]. This favours the NS merger model as they are both less frequent than core-collapse supernovae and expected to produce a higher r-process yield. However, the potential for natal kicks to eject NS binaries from

their formation sites poses a problem, as the r-process enriched material from the merger is unable to pollute the next generation of stars forming if the merger occurs too far away [Bonetti et al., 2019]. Furthermore, galactic chemical evolution studies have suggested that binary populations appear too low for mergers to be capable of reproducing galactic abundances of Europium [Côté et al., 2017], and so this remains an open question.

The basic model for kilonova emission consists of ejected, optically thick matter expanding outward from the central engine. It cools as it expands, and the optical depth lowers until it becomes transparent to radiation. The timescale for this transition from optically thick to thin determines the timescale on with the peak luminosity is reached, typically of the order days to weeks depending on the specific model. Since matter is likely ejected at a range of velocities, this leads to the luminosity from the faster matter peaking earlier than that from the slower ejecta. Further to this, kilonova emission can be divided into two components, blue and red, based on the composition of the matter they originate from. The blue component comes from ejecta which has higher electron fraction (corresponding to a lower neutron count) and so is unable to form many of the heaviest elements like Lanthanides, which gives it a light curve peaking at shorter wavelengths. In the case of NS mergers, this ejecta is likely to be produced along the rotation axis by ejection of matter from the interface between the two stars or from the post-merger accretion disc winds. The red component is generated by the ejecta components with the lowest electron fractions, which are rich in Lanthanides [Metzger and Fernández, 2014]. These are expected to be located in the tidal ejecta and low latitude disc winds. As a result of the distribution of the different components, kilonova may be observed with both components present, or only one depending on the viewing angle [Kasen et al., 2015].

To date there have been relatively few kilonova observations, particularly when compared with sGRBs. This is largely a result of them being much fainter, which therefore limits the distances at which they are visible. Further to this, efforts to search for kilonovae by follow-up optical and IR observations after sGRBs are hampered by the fact that if a GRB afterglow is present, it is typically much brighter and obscures the kilonova signal. [Tanvir et al., 2013] and [Berger et al., 2013] both observed an excess of near IR emission following GRB130603b, the light curve of which matched predictions for the red component of a kilonovae. More concretely, broadband electromagnetic follow-up of the GW170817 and GRB170817a event discovered a counterpart (referred to as AT2017gfo) at a full range of wavelengths from UV to radio [Abbott et al., 2017c]. The observations showed a UV-optical component that was brightest hours after the initial GW detection and decayed rapidly over the following two days, and a near-IR component that peaked in brightness later and faded over the course of a week. This optical follow-up also allowed for the identification of the host galaxy as NGC 4993 [Coulter et al., 2017]. Modelling of the lightcurve also predicted that the total mass of the ejected material was around 0.04 $M_\odot$ [Perego et al., 2017].

How much mass is ejected and where it comes from within the merger are key

to understanding the source of heavy elements in the universe. This depends on many parameters in a merger (total masses, spins, magnetic fields, EOS to name a few) and so further direct observational constraints, like those provided by AT2017gfo, will be important for understanding the relative impact of all the effects at work.

2.1.3 Gravitational Waves

In the first three observing campaigns by the Advanced LIGO-VIRGO gravitational wave detector collaboration, there have been two confirmed detections of compact object mergers with masses consistent with binary NSs and another two that are believed to be NS-BH mergers [Abbott et al., 2021]. The first of these events, the binary NS observation known as GW170817 is as yet the only event to have been observed in both gravitational waves and electromagnetic signals. In the case of the two NS-BH mergers, the mass ratios for both systems were relatively small (around 0.2) and so the NSs would have been swallowed by the BH before tidal disruption could occur, meaning it not expected that their would be an observable electromagnetic counterpart [Abbott et al., 2021]. The GW190425 event is believed to have been a binary NS merger as the component masses were within the range 1.5 - 1.9 $M_\odot$. However, with a total mass of 3.4 $M_\odot$, this system appears to be substantially different from the population of binary NSs within our galaxy that have been observed [Abbott et al., 2020]. This suggests that this system formed via a different channel to the galactic binaries. As there was no electromagnetic counterpart observation and no tidal signature in the gravitational wave signal, it cannot be ruled out that both objects were BHs, however a non-standard evolutionary channel would again be required to explain the formation of a binary BH system with these observed properties [Abbott et al., 2020].

2.1.4 GW170817

As the only binary NS merger observed in both gravitational waves and the electromagnetic spectrum, GW170817 has provided the most observational information to date on binary NS mergers. The gravitational wave signal gave the component masses to be in the range 1.1-1.9$M_\odot$, with a total mass of 2.74$M_\odot$ [Abbott et al., 2019]. Further to this, the gravitational wave signal was also used to place constraints on the radii of the neutron stars and the equation of state more broadly [Abbott et al., 2018]. The extent of the tidal deformation of the two neutron stars can be extracted from the gravitational waveform, since the tidal deformation increases the inspiral rate of the stars, which is measurable as a change in the rate of change of the gravitational wave frequency. This measured tidal deformation can then be used to estimate the radius of the neutron star [Maselli et al., 2013], which in the case of GW170817 gave radii 11.9$\pm$1.4km for both stars [Abbott et al., 2018].

Prior to this, the only major constraints possible came from calculations of the mass of neutron stars in binary orbits. This allowed equation of state models which predict maximum masses below the most massive stars observed (for

example the millisecond pulsar J0740+6620 [Cromartie et al., 2020] has one of the largest well measured masses to date at around 2.14$M_\odot$) to be ruled out. GW170817 also provided a tentative upper limit on the maximum NS mass, through the follow-up electromagnetic observations. Using the total mass of the binary, which can be determined from the gravitational wave signal along with a variety of theoretical models for the remnant, estimates for an upper bound on the maximum non-rotating neutron mass of between 2.15 $M_\odot$ and 2.3 $M_\odot$ have been obtained [Margalit and Metzger, 2017] [Shibata et al., 2017].

Most significantly, the localisation of the gravitational wave (GW) signal and coincident detection of the gamma ray signal GRB170817a provided the first direct evidence for sGRBs being produced in binary NS mergers. The GRB was detected around 2 seconds after the peak GW signal by both Fermi GBM and INTEGRAL [Abbott et al., 2017b]. Although the signal is consistent with a sGRB, it is several orders of magnitude fainter than any previously observed sGRB. This combined with long-term observations of the afterglow suggest that the GRB jet was observed 15-30° off-axis, meaning that only the surrounding mildly relativistic outflow was observed [Abbott et al., 2017b] [Margutti et al., 2017] [Alexander et al., 2018].

Further to this, in the days following the GW detection of the merger, optical and near infrared emission from the resulting kilonova was detected [Arcavi et al., 2017], with an initial peak in the UV which evolved over the following days to peak in the IR lasting around a week. This two component structure agreed well with predicted behaviour for kilonova emission [Metzger et al., 2010]. Comparison of the spectra and light curves to models predicted ejecta masses of 0.02 $M_\odot$ [Nicholl et al., 2017] and 0.04 $M_\odot$ [Kasen et al., 2017] for the blue and red components respectively. In addition, over the following weeks a full range of counterparts, from gamma rays to radio were identified [Abbott et al., 2017c]. In x-ray and radio, the afterglow emission from the sGRB was detected around 9 days after the initial gravitational wave detection [Troja et al., 2017] [Alexander et al., 2017]. This again was consistent with a sGRB generated by a highly relativistic jet viewed off-axis.

Gravitational waves have been key in establishing the link between electromagnetic transients, such as GRBs, and mergers. As more NS-NS mergers are detected, this will provide constraints on the merger rate and the population of binaries in the universe, furthering our understanding of the role of mergers in the production of r-process elements. In the future it is hoped that more sensitive detectors will be able to observe deviations in the inspiral waveform from the point mass solution, providing insights into the neutron star equation of state. Additionally, a clearer picture of the signal during and after the merger could tell us a great deal about the mass distribution during the merger and in the remnant.

2.2 Theoretical background of mergers

2.2.1 Binary formation and evolution

The first binary neutron star was discovered in 1974 [Hulse and Taylor, 1975], known as the Hulse-Taylor pulsar it is comprised of a radio pulsar and a neutron star. Precise measurements of the pulse arrival times from the pulsar allowed the orbital motion of the two stars to be tightly constrained, providing stringent tests of general relativity and giving the first indirect evidence for the existence of gravitational waves [Taylor et al., 1979]. General relativity predicts that accelerating masses will produce gravitational waves, in a similar way to accelerating charges producing electromagnetic waves in electromagnetism. Therefore, a pair of masses orbiting one another will radiate away some amount of their orbital energy, causing them to fall inward and eventually collide. The main properties of the gravitational wave emission from a system depend on the system's mass quadrupole moment. The quadrupole moment of a continuous matter distribution is defined as

$$Q_{ij} = \int \rho \left(r_i r_j - \frac{1}{3}\delta_{ij} r_k r_k \right) dV \tag{2.2}$$

where $\rho = \rho(\mathbf{r}, t)$ is the matter density. For a system of point masses this becomes

$$Q_{ij} = \sum_l m_l \left(r_{il} r_{jl} - \frac{r^2}{3} \right) \tag{2.3}$$

where r_i is the ith component of a particle's position and m is it's mass. For a circular two-body orbit, the position of each body is described by $\mathbf{r}_l = r_l \left(\cos(\omega t)\hat{\mathbf{x}} + \sin(\omega t)\hat{\mathbf{y}}\right)$, where ω is the orbital frequency of the binary and r_l is the distance from the particle to the center of mass of the system. Using this, the quadrupole moment of the binary can be written as,

$$Q_{ij} = \frac{\mu r^2}{2} \begin{pmatrix} \cos 2\omega t + \frac{1}{3} & \sin 2\omega t & 0 \\ \sin 2\omega t & \frac{1}{3} - \cos 2\omega t & 0 \\ 0 & 0 & -\frac{2}{3} \end{pmatrix} \tag{2.4}$$

where $\mu = m_1 m_2/(m_1 + m_2)$ is the reduced mass of the system and r is the separation of the two masses. Assuming spacetime to be approximately flat and that the gravitational waves are a small perturbation to this, the gravitational wave strain h_{ij} can be estimated using the quadrupole moment as

$$h_{ij} = \frac{2G}{c^4 d_L} \ddot{Q}_{ij} \tag{2.5}$$

where d_L is the luminosity distance to the observer [Einstein, 1918]. Combining this with the quadrupole moment for a two-body system we can see that the frequency of the emitted gravitational waves is twice the orbital frequency.

Continuing to use this approximation, the energy radiated from the binary by

gravitational waves is given by

$$\frac{dE_{GW}}{dt} = \frac{G}{5c^5} \dddot{Q}_{ij} \dddot{Q}_{ij}. \tag{2.6}$$

Once again, we substitute the quadrupole moment for the 2-body circular orbit in to Equation 2.6 to obtain

$$\frac{dE_{GW}}{dt} = \frac{32G}{5c^5} \omega^6 \mu^2 r^4 \tag{2.7}$$

Using Kepler's third law, we may eliminate ω to give

$$\frac{dE_{GW}}{dt} = \frac{32G^4}{5c^5} \frac{M^3 \mu^2}{r^5} \tag{2.8}$$

where M is the total mass of the binary. We can then equate this with the change in the orbital energy $\dot{E}_{Orb} = GM\mu\dot{r}/2r^2$, to give a differential equation for the orbital separation of the form

$$\dot{r} r^3 = \frac{64G^3 M^2 \mu}{5c^5} \tag{2.9}$$

Integrating this, we rewrite the resulting equation to find the merger time ($r(T_M) = 0$) in terms of the initial separation r_0 as

$$T_M = \frac{5c^5 r_0^4}{256G^3 M^2 \mu} \tag{2.10}$$

The strong dependence on separation means that the initial separation of binaries is an important factor in determining expected merger rates.

There are two main channels through which neutron star binaries are expected to form. The first (and most dominant) is from a high mass binary in which both companions undergo a supernova explosions that do not break the binary orbit. The heavier companion undergoes the standard evolution for a high mass star, eventually exploding in a supernova and forming the first compact object. Upon the secondary reaching its red giant phase, the binary undergoes a common envelope phase, where the red giant expands to cover the orbit of the compact object companion. This stage is vital to the formation of a merging binary, as the dynamical friction generated by the common envelope shrinks the orbit, giving a much tighter binary orbit when the secondary companion finally becomes the second compact object in a supernova [Bhattacharya and van den Heuvel, 1991].

The second formation channel is the dynamical formation of a binary compact object in a dense stellar environment. This involves preexisting compact objects forming a binary as a result of a dynamical encounter, which requires a very dense stellar environment such as a globular cluster or the galactic center to ensure the probability of interactions is high enough for there to be a reasonable chance of a binary forming. As a result, this channel is expected to produce substantially fewer

NS binaries, for example [Grindlay et al., 2006] found that dynamical binaries may account for only 10-30% of merging NS binaries. These estimates however rely on parameters such as the initial mass fraction of dense clusters which are poorly constrained. In contrast, black hole binaries, which may in principle be formed through the same two channels, are expected to be formed far more efficiently through dynamical encounters [Portegies Zwart et al., 1999]. This leads to the expected binary BH population to be dominated instead by dynamically formed binaries [Sadowski et al., 2008].

During the inspiral phase, the system emits gravitational waves at twice the orbital frequency, which increases as the orbit shrinks. The luminosity also grows with the decay of the orbit, and these combine to give the characteristic 'chirp' signal for mergers. In addition, tidal deformations of the stars as they get closer will alter the waveforms, and the form these changes take will allow us to probe the neutron star's tidal deformability and equation of state. Finally, the post-merger signal depends on what object is formed initially. For a NS-BH this is just a black hole, but in the NS-NS case this could be a supra- or hypermassive NS which may then collapse to a BH some time later [Blanchet, 2006]. This remnant system is the expected central engine, which produces a relativistic jet, which is observed as an sGRB. The dominant theory for this currently requires the presence of a strong magnetic field, which interacts with the central object to produce the jet in a process known as the Blandford-Znajek mechanism.

2.2.2 The Blandford-Znajek Mechanism

The Blandford-Znajek mechanism [Blandford and Znajek, 1977] is a process by which a black hole's rotational energy is extracted through interactions with a surrounding magnetic field to power a jet. It was originally developed as a theory for how AGN jets are produced by supermassive black holes, but the same principle can apply with a smaller system such as the remnant stellar mass black hole and disc following a NS-NS or NS-BH merger.

The mechanism is most intuitively understood using the membrane paradigm set out in [Thorne et al., 1986]. This is a theoretical framework developed for studying black holes from an astrophysical viewpoint. It describes a way of modelling black holes which captures the important relativistic effects near the black hole, whilst removing some of the near horizon physics that does not play an important role in astrophysical phenomena.

In the membrane paradigm, we split 4D spacetime into 3D spacelike hypersurfaces before collapsing these into a single 3D space with a global time coordinate. Physical measurements are made by a set of fiducial observers (FIDOs), which are defined as observers who are 'at rest' at a given position in space, i.e. their worldline in 4D spacetime is orthogonal to hypersurfaces of constant time. In the spacetime around a rotating black hole this means the observers must lie at fixed radius and angle from the spin axis, and will be orbiting the black hole with speed determined by their distance from it. FIDOs make physical measurements using their own

local time, which is related to a global time coordinate by the 'lapse function', α. Near the surface of a black hole, a falling observer's clock runs arbitrarily slower the closer to the horizon we get. As a result, the motion of electromagnetic fields near the horizon becomes 'sluggish', when lines further from the horizon move, those closer in drag behind them. This leads to the build up of 'fossilised' fields around the BH's horizon, containing a complete record of the history of the fields motion around the BH. These fossil fields however have very little impact on field behaviour further out from the hole. The membrane paradigm shows that the interactions of the field can be well described by introducing a conducting membrane a small distance outside the BH horizon, which will hide these complex fossil fields.

In the case of a merger remnant, a large scale external magnetic field is provided by the accretion disc and this interacts with the conducting membrane around the black hole. The overall shape of the field is expected to be roughly poloidal, however at smaller scales the field will be more tangled and chaotic. As the field is pulled with the accreting matter onto the black hole, eddy currents generated in the conducting membrane will dissipate the energy of the smaller scale fields structures, as the membrane can be shown to have non zero resistivity [Thorne et al., 1986]. Larger structures will remain as the currents are unable to dissipate them, which leads to the field being 'cleaned' by the membrane. Further to this, the rotation of the membrane and surrounding ergosphere will drag field lines around with the black hole's rotation, twisting the field lines to produce a magnetic funnel about the hole's rotation axis. Plasma will then be accelerated out along this funnel of field lines at highly relativistic speeds, producing a jet which then can power a GRB.

2.3 Numerical hydrodynamics studies

Whilst analytic models can provide insight into the inspiral and post-merger phases of a merger, the complicated dynamics present in the merger itself restrict the usefulness of such simplified models. In this phase of the merger, numerical hydrodynamics simulations have been employed to study the evolution of the system. The early simulations (dating mostly throughout the 1990s) of NS-NS mergers were performed in Newtonian and Pseudo-Newtonian gravity, using a variety of both Eulerian and SPH codes [Oohara and Nakamura, 1989] [Ruffert et al., 1996] [Zhuge et al., 1994] [Ruffert et al., 1997b] [Shibata et al., 1992]. Gravitational wave emission and its backreaction on the fluid was approximated using a variety of methods, as fully relativistic simulations were not possible. There is general agreement across the different codes on the broad features of the merger: the neutron stars are generally tidally disrupted, producing tails which become spiral arms that wind around a dense core. This core is initially highly non-axisymmetric, transitioning from a pair of cores, through a dumbbell or bar-like phase before eventually settling into an axisymmetric remnant surrounded by a disc. The lack of general relativity meant that the resulting remnant was unable to collapse to a black hole, leaving a long-lived neutron star in all cases and predicting large amounts of mass loss from this remnant to the disk. The

predicted gravitational wave emission from these simulations generally gave an accurate description of the overall GW signal: a rising frequency and amplitude 'chirp' followed by the ringdown of the central remnant, however the finer details and exact wavetrain would require a fully relativistic treatment.

Following breakthroughs in numerical general relativity (GR), which led to the first stable simulations of black holes [Pretorius, 2005] [Baker et al., 2006], the vast majority of recent merger simulations have been performed with varying forms of relativistic gravity. Whilst the major phenomenological features of the mergers are broadly similar, GR simulations predict much lower disk masses due to lower mass loss rates from the central remnant [Shibata and Uryū, 2000]. Using a relativistic gravity formalism also allows for a more accurate description of the GW emission during the merger and the remnant ringdown.

The inclusion of GR allows for the remnant to collapse to form a black hole, and so relativistic simulations have been used extensively to probe how the final product of the merger depends on the initial system parameters [Kiuchi et al., 2009]. Broadly there are three possible scenarios for the final product of a merger: a long-lived NS, a short lived hypermassive NS that is supported against collapse only temporarily by rotation and a prompt collapse to a black hole. [Kiuchi et al., 2009], for example, investigated several different initial mass configurations and showed that the merger resulted in a prompt collapse for systems with total masses greater than 2.9 $M_\odot$ irrespective of mass ratio. Further work has examined the impact of the equation of state (EOS) on the collapse timescales [Hotokezaka et al., 2011] for the central remnant, broadly finding that for a stiffer equation of state, a higher total system mass is required for prompt collapse.

Further improvements to simulations have come from the inclusion of additional microphysics. This includes the use of nuclear EOSs, neutrino emission schemes and magnetic fields. A very large range of nuclear physics derived equations of state have been used in simulations, including in early Newtonian simulations [Ruffert et al., 1996]. The lack of detectable tidal deformation signatures in the GW170817 signal placed an upper bound on the tidal deformability and by extension the EOS, however this EOS is still poorly constrained. The lack of data combined with large theoretical uncertainty about the underlying physics has led to the wide variety of potential EOSs and an extensive study of the effect they have on mergers, in an attempt to find predictions that would allow future observations to further narrow the range of possibilities [Hotokezaka et al., 2011] [East et al., 2016] [Bauswein et al., 2020] [Bauswein et al., 2012].

Additionally, neutrino leakage schemes have been used to simulate the emission of neutrinos from merging NSs, in both early Newtonian simulations [Ruffert et al., 1997a] and more recently in fully relativistic simulations [Bauswein et al., 2012] [Deaton et al., 2013] [Perego et al., 2014] [Cipolletta et al., 2021] [Sun et al., 2022] [Radice et al., 2022]. [Ruffert et al., 1997a] explored the possibility that neutrino emission in NS mergers could power a GRB, finding that a $\nu - \overline{\nu}$ annihilation driven fireball would be a factor 10-1000 too low in energy to match the energies

of observed GRBs. Their simulations showed that the neutrino luminosity grew from around 10^{50}erg/s during the merger itself, to a peak of around 10^{53} erg/s following the formation of the hot accretion torus around the central remnant. Whilst the remnant in this Newtonian simulation cannot collapse to a BH, it was speculated here that the collapse of the remnant will reduce the mass of the disc substantially, further lowering the expected neutrino emission. Neutrino emission in the disc also produced a wind, resulting in the loss of 10^{-2} - 10^{-4} $M_{\odot}$ of neutron-rich matter from the disc. This neutron-rich disc wind ejecta will likely be a source of r-process elements, contributing to the kilonova emission. Whilst these simulations show it is unlikely that neutrinos alone can produce a GRB, recent work [Mösta et al., 2020] [Sun et al., 2022] has demonstrated that neutrinos can be instrumental in clearing baryonic matter away from the poles of the merger remnant. This provides a solution to the baryon-loading problem, in which a magnetically powered GRB jet must overcome the ram pressure of the baryonic matter above the remnant's poles. In this way, the combination of initially neutrino powered outflows, followed by a magnetically powered relativistic jet is providing a promising model for sGRB generation in mergers.

2.3.1 MHD studies

A final direction in which improvements have been made to BNS merger simulations is through the inclusion of magnetic fields in simulations. The first simulations that included both GR and MHD [Anderson et al., 2008a] used very large magnetic fields (around 10^{15}G) and found that the magnetic field had very little impact on the inspiral of the stars, but it did delay the merger slightly in comparison with an unmagnetized case. They also found that for irrotational binaries, the shear layer that forms when the two stars meet is Kelvin-Helmholtz unstable, and so turbulence along this interface amplifies the magnetic field. Further field amplification occurred in the differentially rotating HMNS remnant, and that the activation of the magnetorotational instability (MRI) in the remnant transported angular momentum outward. This reduces the differential rotation of the HMNS, which causes it to collapse quicker than in a non-magnetized case [Giacomazzo et al., 2011].

Other studies have probed the production of a magnetically driven jet, which could power a GRB, like that proposed by the Blandford-Znajek mechanism. Long duration simulations of a merger remnant [Rezzolla et al., 2011] showed that a large-scale poloidal field could form around the remnant black hole, as is required by the Blandford-Znajek process. Further to this [Etienne et al., 2012] have performed a suite of GRMHD simulations, covering both BH-NS and NS-NS mergers. They have demonstrated that in both BH-NS [Paschalidis et al., 2015b] and NS-NS mergers [Ruiz et al., 2016], a large-scale poloidal field is produced and this drives a mildly-relativistic outflow (referred to as an 'incipient jet'), which is expected to be consistent with the early stages of the more strongly relativistic jet required to produce a GRB. They investigate the effect of endowing the NS with both an internal field, and one that extends beyond the surface of the star as expected for a pulsar, finding that in the BH-NS case a jet is only launched

when a external field is present [Paschalidis et al., 2015b]. In contrast, for binary NS mergers they find that both internal and external field cases are capable of launching jets [Ruiz et al., 2016]. Further to this, they examine the effect of the orientation of the initial dipole field applied to the neutron star [Ruiz et al., 2020]. They find that it has a strong impact on the lifetime of the HMNS, the time taken for a jet to be launched and the amount of matter ejected from the system. Specifically, the HMNS survives longer and the amount of matter ejected is suppressed if one or both of the initial dipoles is perpendicular to the orbital spin axis of the system, rather than being aligned with it.

Whilst current simulations have shown that the formation of a magnetically powered outflow is possible, there are still some issues with our understanding of this scenario. For example, current simulations need to use unphysically large initial magnetic fields in order to produce jets. It is known that turbulence, such as that generated along the contact interface between the two stars via the Kelvin-Helmholtz instability, amplifies the initial field [Kiuchi et al., 2015], however current simulations are unable to resolve the turbulence to the required degree to see this. It is therefore not clear whether using an initially strong field will lead to the same outcome as the expected amplification of the field during the merger. Furthermore, the strength and structure of the magnetic field within the neutron star is very poorly constrained due to the potential for the crust of the star to shield an inner field that is much larger or smaller than the external field, which could have further consequences for the production of jets.

Chapter 3

Numerical Hydrodynamics

Numerical hydrodynamics is the basis of all numerical simulations of neutron star mergers. In this chapter, we discuss the mathematical and computational procedures which allow us to numerically solve the equations of hydrodynamics and magnetohydrodynamics (MHD). We also survey some of the existing numerical hydrodynamics codes, commenting on their features and how applicable they are to neutron star mergers.

The Euler equations describe the dynamics of an adiabatic fluid with zero viscosity, and are a simplified form of the Navier-Stokes equations. In three dimensions, the equations are

$$\frac{\partial \rho}{\partial t} + \frac{\partial \rho v_j}{\partial x_j} = 0 \tag{3.1}$$

$$\frac{\partial \rho v_i}{\partial t} + \frac{\partial \rho v_i v_j}{\partial x_j} = -\frac{\partial P}{\partial x_i} \tag{3.2}$$

$$\frac{\partial \rho E}{\partial t} + \frac{\partial \rho E v_j}{\partial x_j} = -P\frac{\partial v_j}{\partial x_j} \tag{3.3}$$

The equations can also be written in a conservative form as

$$\mathbf{U}_t + \mathbf{F}\left(\mathbf{U}\right)_x + \mathbf{G}\left(\mathbf{U}\right)_y + \mathbf{H}\left(\mathbf{U}\right)_z = 0 \tag{3.4}$$

where $\mathbf{U}$ is the vector of conserved variables and $\mathbf{F}$, $\mathbf{G}$ and $\mathbf{H}$ are the fluxes in the x, y and z directions respectively. The subscript notation here indicates a partial derivative of the vector with respect to the variable in the subscript. The variable and flux vectors are given by

$$\mathbf{U} = \begin{bmatrix} \rho \\ \rho v_x \\ \rho v_y \\ \rho v_z \\ E \end{bmatrix}, \mathbf{F} = \begin{bmatrix} \rho v_x \\ \rho v_x^2 + P \\ \rho v_x v_y \\ \rho v_x v_z \\ v_x(E+P) \end{bmatrix}, \mathbf{G} = \begin{bmatrix} \rho v_y \\ \rho v_x v_y \\ \rho v_y^2 + P \\ \rho v_y v_z \\ v_y(E+P) \end{bmatrix}, \mathbf{H} = \begin{bmatrix} \rho v_z \\ \rho v_x v_z \\ \rho v_y v_z \\ \rho v_z^2 + P \\ v_z(E+P) \end{bmatrix} \tag{3.5}$$

where ρ is the density, v the velocity, P the pressure and E the total energy of the fluid. The total energy is the sum of the kinetic and internal energies,

$E = \rho(\frac{1}{2}v^2 + e)$, and the internal energy is given as a function of ρ and P by an equation of state, which closes the system of equations.

It is common to use the Euler equations rather than the full Navier-Stokes equations to model neutron stars due to the relatively negligible effects of viscosity expected from mergers. There are two main approaches to solving the equations of hydrodynamics numerically: grid-based and smoothed particle hydrodynamics (SPH). SPH codes simulate the fluid using particles whose mass is 'smeared' out in order to reconstruct the density distribution of the fluid. This is a Lagrangian approach, formulating the equations in a frame of reference that follows the flow. On the other hand, grid-based schemes solve the equations, by discretizing the fluid variables on a grid that is spatially fixed with respect to the flow. Here we shall focus on the grid based approach, as this is the basis of the methods used in our simulations.

3.1 Conservation Laws

A general system of conservation laws in 1D can be written as

$$\mathbf{U}_t + \mathbf{F}(\mathbf{U})_x = \mathbf{0} \tag{3.6}$$

By applying the chain rule, these equations can be rewritten in quasi-linear form as

$$\mathbf{U}_t + \mathbf{A}\mathbf{U}_x = \mathbf{0} \tag{3.7}$$

where the matrix $\mathbf{A} = \partial\mathbf{F}/\partial\mathbf{U}$ is the Jacobian matrix. A system is said to be hyperbolic if its Jacobian matrix has m real eigenvalues $\lambda_1, ..., \lambda_m$ and a set of corresponding linearly independent eigenvectors $\mathbf{K}^1, ..., \mathbf{K}^m$. The Euler equations are hyperbolic and we can therefore express the matrix $\mathbf{A}$ as

$$\mathbf{A} = \mathbf{K}\mathbf{\Lambda}\mathbf{K}^{-1} \tag{3.8}$$

where the diagonal elements of $\mathbf{\Lambda}$ are the eigenvalues of A and the columns of $\mathbf{K}$ are the corresponding right eigenvectors of $\mathbf{A}$. Using the matrix $\mathbf{K}^{-1}$ we can define a new set of variables, known as characteristic variables (or often primitive/physical variables), as

$$\mathbf{W} = \mathbf{K}^{-1}\mathbf{U}. \tag{3.9}$$

This transformation allows us to rewrite the system of equations as

$$\mathbf{W}_t + \mathbf{\Lambda}\mathbf{W}_x = \mathbf{0}. \tag{3.10}$$

3.1.1 The Riemann Problem

A Riemann problem consists of a system of conservation equations with piecewise constant initial conditions, where two constant states are separated by a single discontinuity (as shown in the bottom panel of Figure 3.1). We can find a solutions to Riemann problems using the method of characteristics [Toro, 2013] (Chapter

2). Characteristics are defined as curves in the $x-t$ plane along which a PDE becomes an ODE. The rate of change of the solution $u(t)$ along a path $x(t)$ is given by

$$\frac{du}{dt} = \frac{\partial u}{\partial t} + \frac{dx}{dt}\frac{\partial u}{\partial x} \tag{3.11}$$

Comparing this with a single conservation law in primitive form,

$$\frac{\partial u}{\partial t} + \lambda(u)\frac{\partial u}{\partial x} = 0 \tag{3.12}$$

and applying the condition that the solution is not time varying along a characteristic ($du/dt = 0$), we see that the characteristics are given by

$$\frac{dx}{dt} = \lambda(u) \tag{3.13}$$

Integrating this gives

$$x(t) = x_0 + \lambda(u_0(x_0))t \tag{3.14}$$

showing that the solution at a given point (x, t) can be found by tracing the characteristic it lies on back to find the initial condition it originated from:

$$u(x,t) = u_0(x - \lambda(u_0(x_0))t). \tag{3.15}$$

The diagram in Figure 3.1 shows a schematic of the result of this applying this method to a Riemann problem. The left and right initial states produce two sets of characteristics, and along the line on which they meet we get a shock front, where the solution changes discontinuously.

As the slope of the characteristics generally depends on the solution itself, how the discontinuity between the two initial constant states propagates depends on the difference in behaviour between the two sets of characteristics. If the characteristics are converging, a shock wave across which the solution changes discontinuously propagates through the domain at a speed $S \neq \lambda_{1,2}$, as shown in Figure 3.1. If the characteristics diverge, a rarefaction fan fills the gap between the two states, across which the solution changes continuously. In the case of equal characteristics, the original discontinuity simply propagates through the domain with speed λ, which is known as a contact discontinuity. Generalising this to system of m conservation equations, we find a set of m discontinuities propagating with speeds S_m separating $m+1$ regions of the solution, and so to fully determine the solution, we need to find which kind of discontinuity each wave represents.

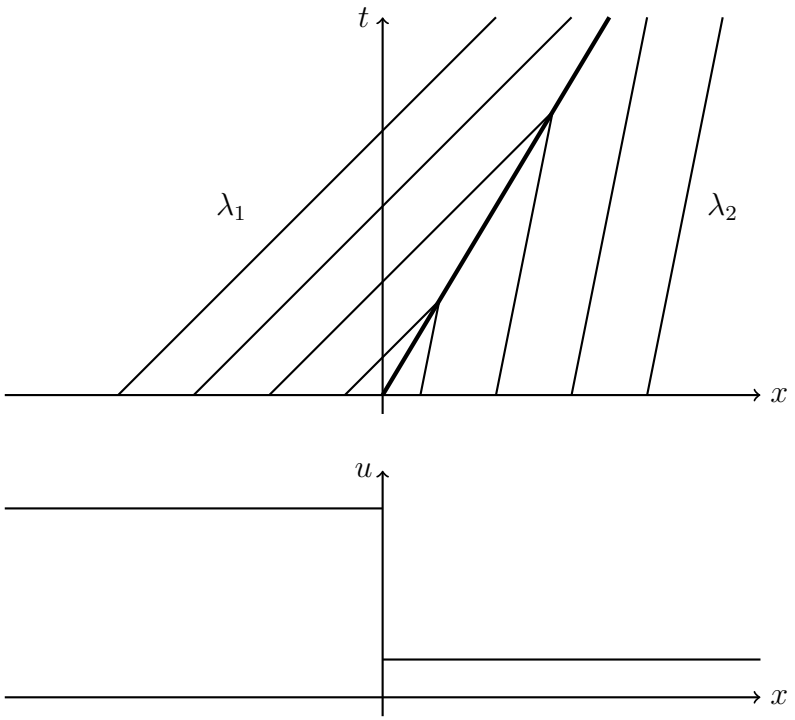

Figure 3.1: The upper panels shows an example of characteristics produced by a Riemann problem in the $x - t$ plane, while the lower shows the initial conditions for the Riemann problem. The left initial condition produces characteristics of slope λ_1 and the right λ_2. The bold line marks where the characteristics meet, and represents a shock wave travelling with speed S.

3.1.2 Numerical Conservation laws

To solve a system of conservation laws using a grid based approach we must first discretize the system, so that it describes the evolution of variables that are averaged over finite grid cells rather than continuous functions of space. A conservation law can be written in integral form as

$$\oint (udx - f(u)dt) = 0 \tag{3.16}$$

where $u(x,t)$ is the solution and $f(u)$ is the flux as seen in the differential form (Equation 3.6) [Toro, 2013] (Section 2.4.1). Here we will discuss a single equation in 1D, the extension to multiple equations in more dimensions is straightforward and is summarised in Section 3.2.1. Choosing the spatial limits of our domain to be (x_1, x_2), we can use the integral form of the conservation law to find an equation for the solution at a given time, t_2, from the value at a previous one, t_1, as

$$\int_{x_1}^{x_2} u(x,t_2)dx = \int_{x_1}^{x_2} u(x,t_1)dx + \int_{t_1}^{t_2} f(u(x_1,t))dt - \int_{t_1}^{t_2} f(u(x_2,t))dt \tag{3.17}$$

This approach can be applied to find a method for evolving the solution of the equations on a whole domain in time by discretizing the domain into cells of size Δx. We specify the cell's position using an index i, where the position x_i corresponds to the center of the ith cell and $x_{i-\frac{1}{2}}$ refers to the boundary between the x_{i-1} and x_i cells, as shown in Figure 3.2.

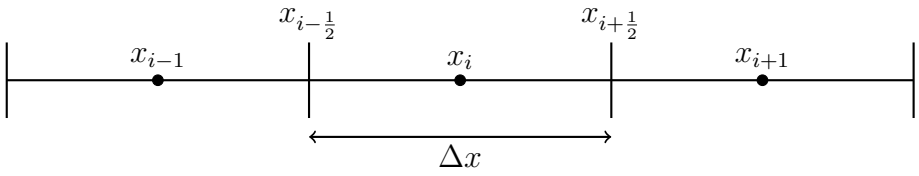

Figure 3.2: A schematic of how a 1D domain is discretized.

The solution within each cell is defined as the spatial average of the solution across the cell, $u_i^n = \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} u(x, t^n)dx$. Using this definition we can find a formula for the value of the solution a time Δt later by using Equation 3.17 and our definition of the cell-averaged solution,

$$u_i^{n+1} = u_i^n + \frac{\Delta t}{\Delta x}\left(f_{i-\frac{1}{2}} - f_{i+\frac{1}{2}}\right) \tag{3.18}$$

where u_i^n is the solution in the ith cell at a time $t^n = n\Delta t$ and $f_{i-\frac{1}{2}}$ is the time averaged flux evaluated at the left interface of the ith zone. In order to advance the solution in time, we then need a method of calculating these inter-cell fluxes.

3.2 Godunov's Method

Godunov's method [Godunov, 1959] applies the solution of Riemann problems to find the inter-cell fluxes in Equation 3.18, allowing us to numerically solve a system of conservation laws, like the Euler equations. In this method, the inter-cell fluxes are defined as

$$f_{i-\frac{1}{2}} = \frac{1}{\Delta t}\int_0^{\Delta t} f(\tilde{u}(x_{i-\frac{1}{2}}, t))dt \tag{3.19}$$

where $\tilde{u}$ is the exact solution to the Riemann problem defined at the cell interface by the cell solution values u_i and u_{i-1}. In the local coordinates of the Riemann problem $\tilde{u}$ is the solution on the t-axis, and so is often written as $u_{i-\frac{1}{2}}(0)$. By solving these Riemann problems at the cell interfaces, we can calculate the inter-cell fluxes, to be used in Equation 3.18 to advance the solution in each cell to the next timestep. This method relies on neighbouring Riemann problems not interacting, which means that the wave from a given problem should not travel more than the width of a zone in the chosen timestep, Δt. We can define a parameter, known as the Courant number,

$$C_{cfl} = \frac{\Delta t |S_{max}^n|}{\Delta x} \tag{3.20}$$

where S^n_{max} is the maximum wave speed present in the system at time t^n. Therefore, the above condition requires that Δt is chosen such that

$$0 < C_{cfl} < 1$$

in order for the method to remain stable.

3.2.1 Higher dimensions

Whilst we have presented this method in only one dimension and for a single equation, it is straightforward to extend it to a system of equations in higher dimensions. Firstly, as discussed at the end of Section 3.1.1, having multiple equations produces multiple waves, which splits the solution into several regions, therefore in solving for the inter-cell fluxes, we just need to determine which of these regions the t axis lies in and find the solution there.

In higher dimensions, we can see from the Euler equations in 3.4 that there are now flux terms corresponding to each spatial dimension. To advance the solution in time we solve Riemann problems on the cell interfaces in each direction to find the fluxes across all the faces, before summing their contributions to the total change in the cell's solution value:

$$u_i^{n+1} = u_i^n + \frac{\Delta t}{\Delta x}\left(f_{i-\frac{1}{2}} - f_{i+\frac{1}{2}}\right) + \frac{\Delta t}{\Delta y}\left(g_{j-\frac{1}{2}} - g_{j+\frac{1}{2}}\right) + \frac{\Delta t}{\Delta z}\left(h_{k-\frac{1}{2}} - h_{k+\frac{1}{2}}\right) \quad (3.21)$$

Broadly, this can be approached in two ways, referred to as dimensionally split or unsplit. In a dimensionally split algorithm, the contributions from each direction are included individually in a series of single dimension passes. This means that the x direction update is performed for the entire domain, then the y dimension update is performed using the result of that first pass and so on. This has the advantage of being simple to implement, easy to extend from a one-dimensional code and easy to parallelise. However, this approach encounters issues when applied to MHD, due to the multidimensional nature of the magnetic field terms and the requirement that the $\nabla \cdot \mathbf{B} = 0$ condition be maintained. Therefore unsplit schemes, which perform the directional updates from all directions at once, are generally favoured for solving the MHD equations. This is discussed in Section 3.3.

3.2.2 Higher order Godunov-type methods

The initial version of Godunov's method using only the cell averaged values as the left and right states for the Riemann problems gives us only first order spatial accuracy for our solution. To improve on this, there are many schemes which instead reconstruct an approximation to the function within the cell, in order to obtain more accurate values for the left and right Riemann states. For example, the Monotone Upwind-centered Scheme for Conservation Laws (MUSCL) [van Leer, 1974] is a collection of second order methods that use a linear approximation

to the function to obtain the left and right edge states for the Riemann problems. This scheme has three main steps:

1. The cell averaged values for the solution variables are replaced with linear functions of the form

$$\overline{u}_i^n(x) = u_i^n + \frac{(x - x_i)}{\Delta x}\Delta_i$$

where Δ_i controls the gradient of the solution, and is defined differently in various sub-models of the MUSCL scheme. These reconstructions are then used to obtain the cell edge values of the solution variables, $\overline{u}^n_{i-\frac{1}{2}}$ and $\overline{u}^n_{i+\frac{1}{2}}$.

2. The edge states of each cell are advanced by a half timestep using

$$\overline{u}^{n+\frac{1}{2}}_{i-\frac{1}{2}} = \overline{u}^n_{i-\frac{1}{2}} + \frac{\Delta t}{2\Delta x}\left(F(\overline{u}^n_{i-\frac{1}{2}}) - F(\overline{u}^n_{i+\frac{1}{2}})\right)$$

3. The half timestep edge states are used as the initial conditions in a Riemann problem, which is used to calculate the inter-cell fluxes at the half timestep. These are used to perform the full timestep update to the cell average values as described by the original Godunov method.

3.3 Extending to MHD

3.3.1 The equations of ideal MHD

The equations of magnetohydrodynamics (MHD) are an extension of the hydrodynamics that govern the motion of a fluid endowed with a magnetic field. More specifically the ideal MHD equations describe a non-viscous, perfectly conducting fluid [Alfvén, 1943]. The addition of the magnetic field requires us to include terms in the momentum and energy equations that account for the effect of the magnetic field on the flow and introduce a new equation for the evolution of the magnetic field itself.

To derive an equation for the evolution of the magnetic field, we consider the Maxwell equations

$$4\pi\mathbf{j} + \frac{\partial \mathbf{E}}{\partial t} = c\nabla \times \mathbf{B} \tag{3.22a}$$

$$\frac{\partial \mathbf{B}}{\partial t} = -c\nabla \times \mathbf{E} \tag{3.22b}$$

$$\nabla \cdot \mathbf{E} = 4\pi\sigma \tag{3.22c}$$

$$\nabla \cdot \mathbf{B} = 0 \tag{3.22d}$$

The Faraday equation 3.22b describes the change in the magnetic field with time, so we use this as the basis for our evolution equation. In a perfectly conducting fluid, the electric field in the fluid vanishes from the point of view of an observer

moving with the fluid, as in the presence of a non-zero electric field, arbitrarily large currents would form to neutralise any separation of charge. In frames where the fluid is moving relative to the observer however, there will be a non-zero electric field component perpendicular to the motion as described by the Lorentz transform:

$$\mathbf{E}'_{\perp} = \gamma(\mathbf{E}_{\perp} + \mathbf{v} \times \mathbf{B}/c). \tag{3.23}$$

Therefore, the only electric fields present in the fluid is that given by Lorentz transform in Equation 3.23, and so we can rewrite Equation (3.22b) as

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}). \tag{3.24}$$

This is the MHD induction equation, and governs how the magnetic field in the fluid changes with the fluid's motion. By taking the divergence of 3.24 we find

$$\frac{\partial}{\partial t} (\nabla \cdot \mathbf{B}) = 0, \tag{3.25}$$

ensuring that the condition in (3.22d) is maintained automatically, provided it is satisfied by the initial conditions. Therefore (3.22d) is a constraint on our initial conditions in MHD, rather than part of the system of equations we must solve. Further to this Equation (3.22c) becomes unnecessary, as our fluid is electrically neutral and perfectly conducting, so there are no non-zero charge densities or electric fields.

The magnetic field introduces a Lorentz force to the momentum equation of the form

$$\mathbf{F}_L = \frac{1}{c}\mathbf{j} \times \mathbf{B} \tag{3.26}$$

where $\mathbf{j}$ is the current density. As described above, in a perfectly conducting fluid there can be no separation of charge and therefore no currents, and so the only contributions to this current density are due to the relative motion of one fluid parcel with respect to another. Therefore, we can rewrite this force in terms of the magnetic field and fluid velocity only. We can rewrite the Ampere-Maxwell law (3.22a) without the electric field, using the Lorentz transform in (3.23),

$$4\pi\mathbf{j} + \frac{\partial}{\partial t} (\mathbf{v} \times \mathbf{B}) /c = c\nabla \times \mathbf{B}.$$

Restricting to non-relativistic fluid velocities ($v << c$), we consider the flow to have length scale L and timescale T and define typical velocity $V = L/T$. The displacement current term then has typical magnitude

$$|\frac{\partial}{\partial t} (\mathbf{v} \times \mathbf{B}) /c| \approx (BV)/(cT) \approx B(V/c)(V/L),$$

whereas the RHS term has magnitude Bc/L. Using the Lorentz transform in Equation (3.23), in the equivalent transform for the magnetic field,

$$\mathbf{B}_{\perp}^{'} = \gamma(\mathbf{B}_{\perp} - \mathbf{v} \times \mathbf{E}/c),$$

we see that changes in the magnetic field with frame of reference vanish to second order in v/c. Comparing the magnitude of the two terms in the Ampere-Maxwell law, we see that the displacement current is a factor of $(v/c)^2$ smaller than the RHS, and therefore we may reduce (3.22a) to

$$\mathbf{j} = \frac{c}{4\pi} \nabla \times \mathbf{B}. \tag{3.27}$$

We can use this equation to rewrite the Lorentz force term in the momentum equation as

$$\mathbf{F}_L = \frac{1}{4\pi} (\nabla \times \mathbf{B}) \times \mathbf{B} \tag{3.28}$$

Finally, the total conserved energy now includes the energy of the magnetic field $B^2/2\mu_0$, and this introduces a Poynting flux term into the energy equation,

$$\frac{1}{\mu_0} (\mathbf{E} \times \mathbf{B}) = \frac{1}{\mu_0} (\mathbf{v} \times \mathbf{B}) \times \mathbf{B}.$$

With these additions, we can write the full ideal MHD equations in conservative form as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{3.29a}$$

$$\frac{\partial \rho \mathbf{v}}{\partial t} + \nabla \cdot (\rho \mathbf{v}\mathbf{v} - \mathbf{B}\mathbf{B}) + \nabla p_{tot} = \rho \mathbf{g} \tag{3.29b}$$

$$\frac{\partial \rho E}{\partial t} + \nabla \cdot (\mathbf{v}(\rho E + p_{tot}) - \mathbf{B}(\mathbf{v} \cdot \mathbf{B})) = \rho \mathbf{g} \cdot \mathbf{v} \tag{3.29c}$$

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla \cdot (\mathbf{v}\mathbf{B} - \mathbf{B}\mathbf{v}) = 0 \tag{3.29d}$$

3.3.2 Numerical MHD

Solving the MHD equations numerically proves to be more difficult than for hydrodynamics primarily because of the constraint $\nabla \cdot \mathbf{B} = 0$. Whilst this condition is conserved analytically, it is not necessarily maintained by a finite difference approach. This can produce unphysical forces parallel to the magnetic field if not controlled for. Many different methods have been developed for dealing with this issue. An early approach [Brackbill and Barnes, 1980], showed that rewriting the momentum equation in a non-conservative form substantially reduced the errors caused by unphysical forces, but fails to correct the unphysical behaviour of the magnetic field itself. A different approach is to evolve the magnetic vector potential $\mathbf{A}$ and then calculate $\mathbf{B}$ from this, since the identity $\nabla \cdot (\nabla \times \mathbf{A}) = 0$, ensures that the divergence constraint is satisfied. However, this method requires the calculation of second derivatives of the vector potential (in the Lorentz force

term in Equation 3.29b), which can lead to a substantial loss of accuracy, as each derivative reduces the order of accuracy. Therefore, this method requires the vector potential to be calculated to very high order accuracy in order to retain a reasonable order of accuracy in the solution as a whole.

If instead we want to evolve the magnetic fields directly, there are two main approaches that have been used extensively: divergence cleaning and constrained transport. Divergence cleaning methods evolve the magnetic field like any other fluid variable, and then remove the unphysical, rotational part of the field [Ramshaw, 1983]. There many approaches to calculating and removing this rotational component, one of the simplest involves solving the following Poisson equation

$$\nabla \cdot \mathbf{B} = -\nabla^2 \phi.$$

The gradient of the potential ϕ is then used to correct the magnetic field. This can be extended to treat the potential as a vector rather than a scalar quantity [Balsara, 1998].

Constrained transport [Evans and Hawley, 1988] on the other hand, involves constructing a finite-difference scheme for the induction equation such that $\nabla \cdot \mathbf{B}$ is conserved. The key feature of this method is the use of a staggered grid, where the majority of the fluid variables are calculated at the cell center positions, but the magnetic field variables are calculated on the cell faces, as shown in Figure 3.3.

With the fields defined on cell surfaces, the magnetic flux through a face is simply given by $\Phi_{x,i-\frac{1}{2}} = B_{x,i-\frac{1}{2}} \Delta y \Delta z$ (for the x direction lower face, and similarly for all other faces). The total flux across the surface of the cell can be calculated by summing up the fluxes for each of the six faces, and it can be shown via Gauss' theorem that this total must be zero in the case that $\nabla \cdot \mathbf{B} = 0$. The magnetic flux across a face can be evolved using the integral of the electromotive force (EMF) around the edge of the face. Therefore, we define the EMF values to be centered on the cell edges, and the rate of change of the flux across the face is

$$\frac{1}{\Delta t}\left(\Phi^{n+1}_{i-\frac{1}{2}} - \Phi^{n}_{i-\frac{1}{2}}\right) = -(\epsilon_{i-\frac{1}{2},j-\frac{1}{2},k} + \epsilon_{i-\frac{1}{2},j,k+\frac{1}{2}} - \epsilon_{i-\frac{1}{2},j+\frac{1}{2},k} - \epsilon_{i-\frac{1}{2},j,k-\frac{1}{2}}) \quad (3.30)$$

where the signs of the ϵ terms are determined by the anticlockwise direction of integration. Summing up this across all six faces of the cell will give the total change in flux across the surface of the cell. The direction of integration on each face means that each edge EMF appears twice but with opposing signs, meaning the total change in flux with time is always zero. Therefore, using the relation $\Phi_{x,i-\frac{1}{2}} = B_{x,i-\frac{1}{2}} \Delta y \Delta z$ in Equation 3.30, and calculating the EMF as defined on the RHS of the induction equation 3.24 we can evolve the magnetic fields in such a way that $\nabla \cdot \mathbf{B}$ remains constant to machine precision. With this, we only need to ensure that $\nabla \cdot \mathbf{B} = 0$ in our initial conditions and it will remain so throughout the simulation.

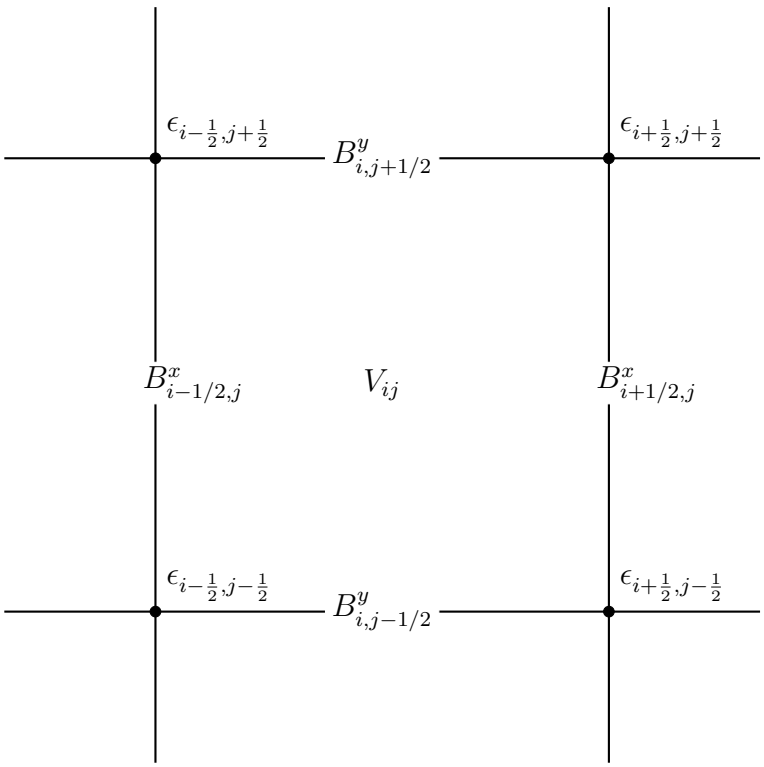

Figure 3.3: 2D schematic of the staggered grid structure. Hydrodynamic variables, V, are defined at the center of cell i, j. The magnetic field component normal to each cell face is defined on that face. The EMF, ϵ, from which the magnetic field is calculated is evaluated at the cell edges.

[Balsara and Kim, 2004] compared the performance of several divergence cleaning schemes with a staggered grid approach, finding that the divergence cleaning methods are prone to spikes in the magnetic energy which are not observed in the staggered grid approach. Further to this, the structure of the fields was different between the two types of method. These deficiencies are particularly pronounced for turbulent flows, and as such staggered mesh methods like constrained transport are preferred for their better accuracy in these situations.

3.4 Public HD/MHD codes

The early stages of this work involved finding a suitable hydrodynamics code with which to create our simulations. As part of this we surveyed many of the publicly available hydrodynamics codes, with a specific view to which were well suited for our work. The main requirement was the inclusion of an MHD solver and self-gravity. On top of this, as we expected to need to make significant additions to the model, it was important to see how easily adaptable the codes were. Here we summarize the results of this survey of codes, noting the main features of each and evaluating their suitability to our application.

The large array of available hydrodynamics simulations can be separated broadly

into three categories based on the capabilities of the central hydrodynamics solver: hydrodynamics (HD) only, MHD and general relativistic MHD (GRMHD). The first group of codes have only hydrodynamics solvers (at the time of the survey in late 2018). These include cosmology focused codes such as GANDALF [Hubber et al., 2018] and RAMSES [Teyssier, 2002], which implement other physical effects, such as coupling to N-body gravity solvers. Other codes in this category include those with more complex approaches to the basic hydrodynamical solver, for example the AREPO code [Springel, 2011], which uses a moving mesh approach to combine the advantages of Eulerian and Lagrangian methods. Ultimately, this codes were of limited applicability to this work due to the lack of magnetic fields.

The next and largest group of codes are those with MHD solvers. There are several codes in this category designed for a broad range of applications, including ATHENA/++ [White et al., 2016], FLASH [Fryxell et al., 2000] and PLUTO [Mignone et al., 2012]. They commonly have a range of different physical effects implemented, which can be easily swapped out allowing a user to easily select the required physics. Whilst it is easy to make use of physical effects already implemented, there is some variability in how difficult it is to incorporate further physics manually into these codes. Further to this, there are again codes written with more specific intended uses, including ENZO [Bryan et al., 2014] which is primarily used for cosmological simulations and PHANTOM [Price et al., 2018] which is an SPH code designed for accretion flows.

Finally, the last category includes codes with general relativistic HD/MHD solvers. The mostly widely used of these is the EinsteinToolkit [Löffler et al., 2012], which is a collection of several GRHD and GRMHD models created using the Cactus framework [Goodale et al., 2003], but others include Spritz [Cipolletta et al., 2020] and WhiskyTHC [Radice and Rezzolla, 2012]. These codes have all been used extensively to study mergers and as such they contain much of the relevant additional physics, including nuclear equations of state and neutrino emission. However, solving the equations of MHD in a fully relativistic form is substantially more computationally intensive than Newtonian calculations.

After reviewing a large range of codes, we chose to work with the FLASH code framework [Fryxell et al., 2000]. This was due to its relatively large range of included physics that was of interest to us, including MHD and self-gravity. Further to this, while we did not use them in this work, the inclusion of neutrino emission and nuclear equations of state were also of interest as these effects could then be relatively easily included future work. Another consideration was also how easy it would be to include a gravitational wave emission scheme in the code. FLASH's modular structure and relatively extensive documentation suggested that it would be somewhat easier to make the additions required than some other comparable codes, which seemed less well structured for additions by the user. Finally, whilst using a GR code would have removed the need to introduce an approximate gravitational wave solver, it also increases the computational cost of the simulations. We therefore decided to work with a Newtonian code, as this would allow us to dedicate more of our available CPU time to producing higher

resolution simulations. This is because the focus of our study is the amplification of magnetic field by turbulence, which is very sensitive to resolution.

Chapter 4

NS-NS mergers with FLASH

The FLASH code framework [Fryxell et al., 2000] is a publicly available hydrodynamics code, designed to be adaptable to many different physical applications. Its primary application was to astrophysical systems, but it has since been used in other areas, such as laser and condensed matter physics. We chose it also because it is structured such that the different physics units can be easily added to a simulation by the use, making it easier to add in the extra physics we need. It is a grid-based code, with Adaptive Mesh Refinement (AMR) provided by the Paramesh package [Olson et al., 1999] and is fully MPI parallelized.

4.1 Simulation physics

4.1.1 MHD with the USM algorithm

The FLASH code framework [Fryxell et al., 2000] contains multiple solvers for both hydro and magnetohydrodynamics, including dimensionally split and unsplit methods. The basis for all our models is the Unsplit Staggered Mesh scheme (USM), which we describe briefly here. Further details can be found in [Lee and Deane, 2009].

The algorithm follows a similar structure to most other Gondunov type methods, the fluid variables in each cell are defined at the cell centers and evolved through the solution of Riemann problems (see Chapter 3). The magnetic fields are defined on cell faces and must be treated differently in order to maintain the divergence-free condition. First we perform a reconstruction-evolution step, where the cell-centered variables are reconstructed to give the values at the cell faces and are evolved by a half timestep. For the fluid variables and the components of the magnetic field transverse to the cell face, this is done using a MUSCL-Hancock type TVD (Total Variation Diminishing) method. This is a linear (2nd order) reconstruction method designed to reduce spurious oscillations that can be generated near sharp changes in the solution variables, and the details are described in [Lee and Deane, 2009]. It is important to note that the normal component of the magnetic field is treated separately, as applying this TVD reconstruction could lead to the introduction of divergence to the field. The calculated face

states for the cell centered variables are then used to solve a Riemann problem and the resulting fluxes are used to update the face-centered normal field by a half-timestep. With the edge states for all variables now known at timestep $n + \frac{1}{2}$, a further Riemann problem is solved to obtain the fluxes used to update the cell-centered values to timestep $n + 1$. Finally, the normal fields are updated using the induction equations as described in section 3.3.2, by integrating the electric fields (equivalent to the normal flux terms in the MHD equations) around the cell edges. The required edge-centered electric fields can be calculated by averaging the face centered fluxes from the four adjoining faces or FLASH also includes a method that uses higher derivatives of the electric fields in the construction, in order to increase the accuracy of the construction.

To summarize, this scheme uses higher order Godunov-type methods to evolve the fluid and transverse field variables, whilst applying a CT-type method to normal fields in order to maintain the $\nabla \cdot \mathbf{B} = 0$ condition. Since the introduction of this method to the code, third order PPM and fifth order WENO reconstruction methods have been included in the algorithm, further improving the accuracy of the method.

4.1.2 Self-gravity

FLASH contains several different solvers for including Newtonian self-gravity in a simulation, which solve the Poisson equation,

$$\nabla^2 \phi = -4\pi G \rho, \tag{4.1}$$

for the gravitational potential ϕ, and then calculate the gravitational acceleration from this. There are three main methods: the multipole solver, the multigrid solver and the Barnes-Hut tree solver.

The multipole solver [Couch et al., 2013] calculates the gravitational potential using an expansion in the mass multipole moments of the density distribution. The angular resolution of this expansion depends on the number of moments used and the memory usage and computational cost scale as the number of multipole moments squared. Therefore, it is best-suited to spherical or approximately spherical distributions, since the number of moments required to reach acceptable resolution will be lower in these cases.

The multigrid solver [Ricker, 2008] uses Fourier transforms to solve the Poisson equation for the potential and is therefore suitable for use with general source distributions. Due to the periodic nature of Fourier expansions the boundaries of non-periodic domains require special handling, and this can sometimes lead to unphysical behaviour near such boundaries. It is therefore important to make sure that the boundaries are sufficiently far from the region of the simulation domain that is of interest when using such solvers.

Finally the Barnes-Hut tree solver [Barnes and Hut, 1986] constructs an oct-

tree by recursively dividing cells into eight zones (in 3D) of half the side length, beginning from a single zone covering the entire domain and terminating once the actual grid resolution is reached. At each level of the tree, the mass of this pseudo-cell is given by the sum of the mass in cell and is located at the center of mass of the cell. When calculating the contribution to a given cell's potential, the largest pseudo-cell in the tree for which

$$\frac{S}{D} < \theta_{BH}$$

where S is the side length of the pseudo-cell and D is the distance between the psuedo-cell and the point we are calculating the potential at. The opening angle θ_{BH} is a parameter which controls the accuracy of the calculation, a narrower angle will increase the accuracy but require us to use more highly resolved cells for each calculation. The version of this algorithm used in FLASH is detailed in [Wünsch et al., 2018] and includes more sophisticated error criteria for determining which nodes should contribute.

In our simulations we chose to use the BH tree solver for self gravity, as our matter distribution is substantially non-spherical and therefore the multipole solver would be more computationally expensive in achieving the same accuracy. Whilst we could also have used the Multigrid solver, attempts to include it resulted in technical issues and it is not more accurate the BH tree solver for similar computational costs, we did not pursue these issues further.

4.1.3 Equation of State

For all the simulations presented in this work we use a constant-Γ equation of state. This models an ideal fluid with constant heat capacity ratio $\Gamma = C_v/C_p$. In our implementation, the density and internal energy are calculated by solving the hydrodynamic equations, and in the case of internal energy subtracting the kinetic component from the total energy. These are then used to calculate the remaining thermodynamic variables: pressure and temperature. The pressure is related to the density and internal energy through

$$P = (1 - \Gamma)\rho\epsilon \tag{4.2}$$

where ϵ is the specific internal energy. The temperature is given by

$$T = \frac{1}{\Gamma - 1}\frac{N_a k}{\bar{A}}\epsilon \tag{4.3}$$

where N_a is Avogadro's number, k the Boltzmann constant and $\bar{A}$ the average atomic mass of the matter. This temperature is calculated assuming we are modelling an ideal, thermal-pressure supported gas and as such is not physically meaningful for neutron star matter, but as it does not contribute to the calculations in any way this does not cause any issues.

For our simulations we choose a value of $\Gamma = 2$, in line with many similar studies of neutron star mergers (for example [Shibata et al., 1992] [Giacomazzo et al., 2011] and [Anderson et al., 2008b]). Whilst the equation is idealised, it provides a good first approximation to a relatively stiff equation of state for neutron stars. This analytical equation of state also has a very low computational load relative to more physical equations of state, most of which use tables of pre-calculated values and therefore need to read and interpolate values from these tables each time the equation of state is used.

4.2 Gravitational waves

Gravitational waves are necessary for the simulation of compact object mergers. It is this gravitational radiation which carries away energy and angular momentum from the binary system that allows the objects to merge. As a general relativistic effect we need to manually include them in FLASH's Newtonian hydrodynamics solver. Therefore we have created a new source term module in FLASH designed to reproduce the main effects of gravitational wave emission on the dynamics of mergers.

In Section 2.2.1, we showed that the luminosity emitted in gravitational waves can be calculated from a matter distribution's mass quadrupole moment as

$$\frac{dE_{GW}}{dt} = \frac{1}{5}\frac{G}{c^5}\dddot{Q}_{ij}\dddot{Q}_{ij} \tag{4.4}$$

and that for a circular two-body orbit this becomes

$$\frac{dE_{GW}}{dt} = \frac{32G^4}{5c^5}\frac{(M_1 + M_2)M_1^2 M_2^2}{r^5} \tag{4.5}$$

where here we have replaced the total mass M and reduced mass μ with their definitions in terms of the masses of the two bodies $M_{1,2}$.

The source term method we have implemented in FLASH calculates the energy lost by the domain using, optionally, one of the above formulas, and reduces the orbital velocity of the stars such that the drop in kinetic energy is equal to this value. The first stage of the calculation is therefore to calculate the energy loss via either the quadrupole moment formula or the point mass formula. When using the quadrupole moment formula, we wish to maintain numerical accuracy by avoiding performing many numerical time derivatives of the quadrupole tensor. Therefore we can instead rewrite the third time derivative of the quadrupole moment in terms of fluid variables and first order derivatives of these. This is done by substituting for time derivatives of density and velocity using the mass and momentum conservation equations and then applying integration by parts to reduce the number of spatial derivatives further. The full derivation is detailed in

Appendix A, where we show that we can write the quadrupole moment as

$$\dddot{Q}_{ij} = STF\left[2\int dV\left(2P\frac{\partial v_i}{\partial x^j} + \frac{\partial \psi}{\partial x^j}\left(r_i\frac{\partial \rho v^k}{\partial x^k} - 3\rho v_i\right) - \rho r_i\frac{\partial \dot{\psi}}{\partial x^j}\right)\right], \quad (4.6)$$

where ψ is the Newtonian gravitational potential, P is the pressure and v_i is the ith component of the velocity [Blanchet et al., 1990]. In this way we reduce the numerical inaccuracy associated with taking multiple approximate time derivatives, only needing to calculate one for the gravitational potential rather than three. Using this in Equation 4.4, we calculate the total energy in gravitational waves leaving the binary system in a given timestep.

The energy losses must then be distributed across the grid in a way which produces approximately the correct rate of inspiral for the two objects. My approach is inspired by that of [Zhuge et al., 1994], in which the total energy lost by the system is used to define a friction force acting on the centre of mass of each of the two neutron stars. As their simulation uses smoothed particle hydrodynamics, they apply a constant acceleration to all particles inside each star. Instead of subtracting a fixed amount from velocities in zones deemed to be part of the stars, I reduce the velocities inside the stars by a fixed scale factor, avoiding any issues where the fluid's momentum could be reversed unphysically. This scale factor is calculated by equating the change in total orbital kinetic energy with the energy lost to gravitational waves in a single timestep, $\Delta E_{GW} = L_{GW}\Delta t$. In terms of the scaling factor this change is given by

$$\begin{aligned}\Delta E_{GW} &= E_{orb} - E_{orb,new} \\ &= \frac{1}{2}v_{orb}^2(1 - A_v^2) \\ &= E_{orb}(1 - A_v^2)\end{aligned}$$

where E_{orb} is the orbital kinetic energy of the two stars before reductions are made. Therefore, the scaling factor is calculated from the total orbital energy and the energy lost due to gravitational wave emission using

$$A_v = \sqrt{1 - \frac{\Delta E}{E_{orb}}} \quad (4.7)$$

The orbital kinetic energy is given by sum of the contribution to the kinetic energies in all zones within the stars (determined by a density threshold) that come from the velocity component parallel to the orbital velocity of the star (taken to be the velocity of the centre of mass of the star). This component of the velocity is then reduced by the scaling factor, slowing of the star's orbital motion which then gives rise to the required inspiral. It should be noted that in this approximation we assume the stars are of equal mass and therefore the scaling is the same for both stars, so the treatment would need to be revised in the case of an unequal mass binary.

4.3 Initial Conditions

The initial conditions for each neutron star are given by polytropes, which are density-radius relations that describe self-gravitating spheres of fluid. The relations are derived from a system of two first order differential equations,

$$\frac{dP}{dr} = -\frac{GM_r}{r^2}\rho \tag{4.8}$$

$$\frac{dM_r}{dr} = 4\pi r^2 \rho, \tag{4.9}$$

where M_r is the mass contained within a radius r and P and ρ are the pressure and density at r respectively. Equation 4.8 describes the hydrostatic equilibrium of the star, while Equation 4.9 describes how the mass contained within a radius r changes within the star. As described in Section 4.1.3, we use an ideal gas equation of state. For these initial conditions only, we then additionally assume that the fluid is adiabatic, allowing pressure and density to be related by $P = \kappa\rho^{\Gamma}$. Using this relation, the two differential equations can be combined into a single second order equation, known as the Lane-Emden equation [Lane, 1870]. This equation has a family of solutions parameterized by the polytropic index $n = 1/(1-\Gamma)$, with the only analytic solutions at $n = 0, 1, 5$. Taking $\Gamma = 2$, this gives a polytropic index of $n = 1$ and the initial density as a function of radius is therefore described by,

$$\rho(r) = \rho_c \frac{\sin \xi}{\xi} \tag{4.10}$$

where $\rho_c = \pi M_{NS}/4R_{NS}^3$ is the central density of the star (determined by the choice of neutron star mass and radius) and $\xi = \pi r/R_{NS}$ is the ratio of r to the neutron star radius, R_{NS}. The constant κ used to calculate the pressure is given by $\kappa = 2GR_{NS}^2/\pi$, and therefore determined by our choice of neutron star radius.

Outside the neutron star the fluid is set to a constant density of 10^7 g cm^{-3}, which is around 10^{-7} times lower than the maximum neutron star density initially. Between the star and this lower density exterior, a 'patching' region must be employed for numerical stability, as the steep density gradient at the outer edge of the polytrope solution otherwise leads to unphysical spikes in energy in the zones nearby which eventually cause the code to crash. This patching region is set so that the reduction in density between adjacent zones does not exceed an order of magnitude per 3 zones, and so in most simulations this corresponds a physical extension of the surface of the star by approximately 500-1000m at the resolutions we use.

Each of our binary neutron star merger simulations consist of two $1.4M_\odot$ neutron stars with radii of 14.3km. The grid for this simulation 150km wide in all directions and the stars begin at a separation of 45km. The stars are irrotational and so the velocity of each star given by the Keplerian orbital velocity, with a radial correction due to the emission of gravitational waves calculated from the point

mass formula,

$$\mathbf{v_r} = -\frac{64}{5}\frac{G^3}{c^5}\frac{M}{d^3}\hat{\mathbf{r}}. \tag{4.11}$$

4.3.1 Initial magnetic field

We endow each neutron star with a magnetic field that is confined to the interior of the star. The internal field configuration is generated from a vector potential of the form

$$A_\phi = A_b\omega^2 \max(P - P_{cut}, 0) \tag{4.12}$$

where $\omega^2 = (x - x_c)^2 - (y - y_c)^2$ is the radius in cylindrical coordinates centered on the star, A_b is a parameter which sets the maximum field strength and P_{cut} is a threshold pressure value, below which the magnetic field is set to zero. Figure 4.1 shows the resulting field applied to two stars, one of which is rotated 90° out of alignment with the other.

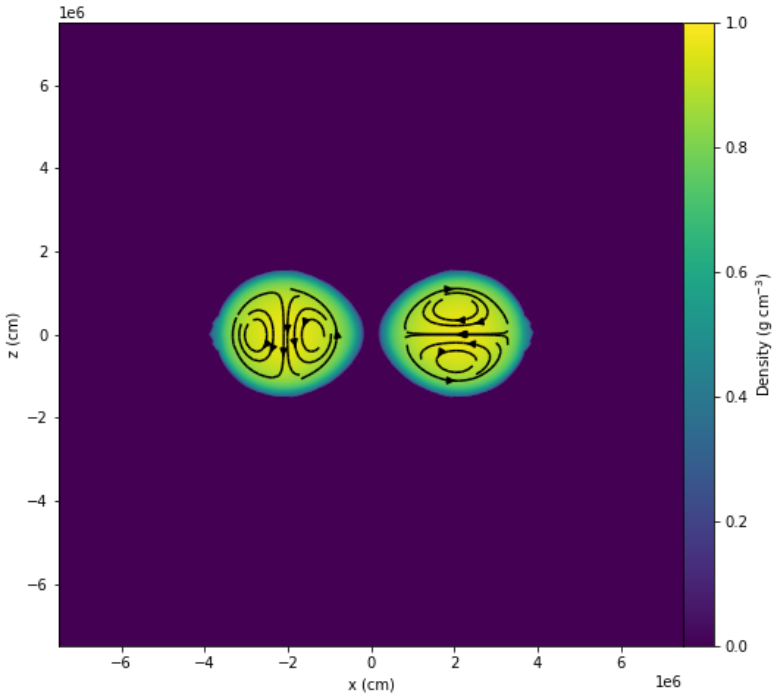

Figure 4.1: Density in the $y = 0$ plane overplotted with the magnetic field lines for model IntT at time 1.5ms when the field is initially seeded.

As the magnetic fields are not dynamically important during the inspiral of the stars, we initially evolve unmagnetized stars and then seed them with the desired magnetic fields just before significant tidal disruption occurs.

4.4 Model tests and stability

In order to maintain numerical stability in our models, we reduce the velocity of gas in the low density medium surrounding the stars to zero in regions where the gas density is below 10^8 gcm^{-3}. This cut off is low enough that it does not affect the motion of gas within the stars, but does stop the infall of gas onto the stars due to gravity, which left unchecked leads to steepening of the density gradient at the star's surface to such a degree that the simulation becomes unstable. Further to this, to avoid spurious heating of the low density atmosphere gas, which can again lead to large energy gradients that destabilise the simulation, we limit the pressure in this atmosphere to a range of $0.5P_p < P < 10P_p$, where P_p is the polytropic pressure $P_p = \kappa\rho^\Gamma$ for gas with density ρ. This limiting is applied to all atmosphere gas below $\rho = 10^{10}$ gcm^{-3}, which is again low enough that it does not affect matter within the stars, but does prevent unphysical spikes in energy from forming at the star's surface.

4.4.1 Stability of an isolated neutron star

The single, unmagnetized neutron star is constructed using the initial conditions described in Section 4.3 and evolved for many dynamical times, to assess how stationary it remained. The neutron star has mass $M_{NS} = 1.4\mathrm{M}_\odot$ and a radius of 14.3km. The grid has width of 50km, and utilises 3 levels of adaptive mesh refinement (AMR). This allows the resolution in the constant density surrounding to remain low whilst giving the finest zones a width of 390m which is comparable to the finest grid resolution in our merger simulations.

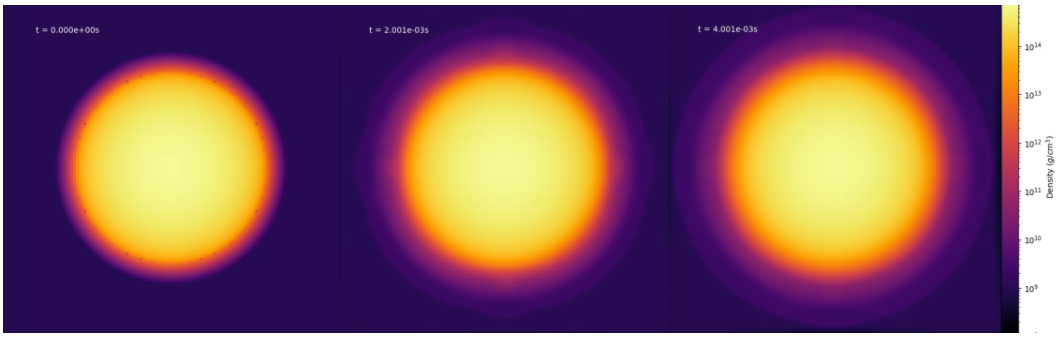

Figure 4.2: Slices density in the $x-y$ plane at several times during the simulation.

Figure 4.2 shows that the star expands slightly over the course of the simulation and that the surface oscillates slightly as it does so. Figure 4.3 shows that the ratio of kinetic to internal energy remains low and approximately constant over many dynamical times, following an initial period of increase. The dynamical time is defined as the free-fall time for the star, given by $t_{dyn} = \sqrt{R_{NS}^3/2GM_{NS}}$. This demonstrates that after an initial relaxation, the star remains in an approximately steady state. It continues to oscillate slightly as we would expect due to the discrete nature of the grid making perfect hydrostatic equilibrium impossible.

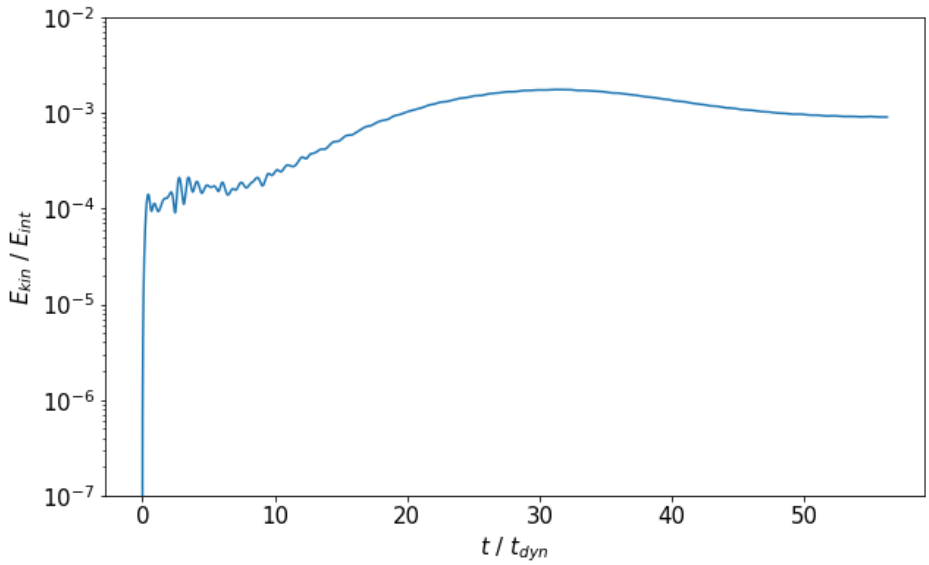

Figure 4.3: Ratio of the total kinetic and internal energy of the isolated neutron star simulation.

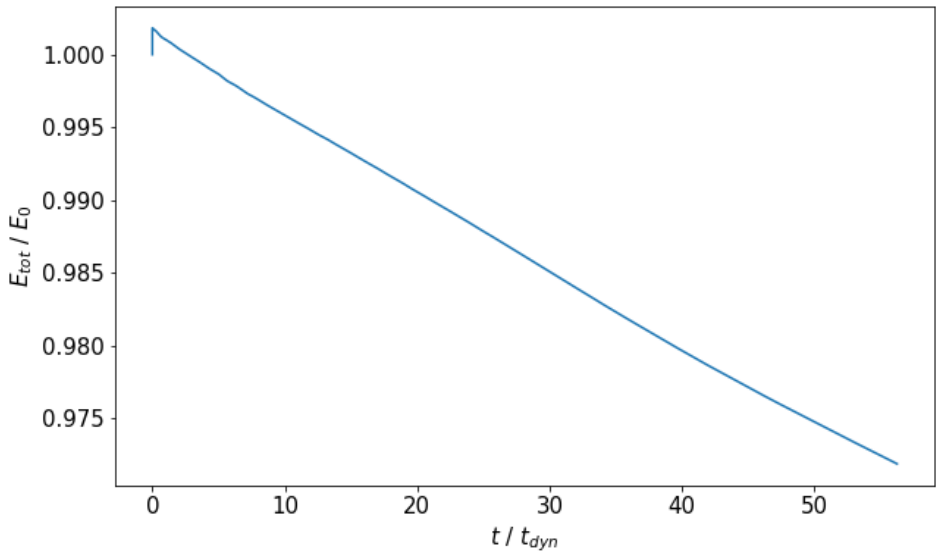

Figure 4.4: Total energy (sum of kinetic, internal and gravitational potential energies) in the isolated NS E_{tot} normalized by the initial total energy, plotted against dynamical time for the duration of the simulation.

Figure 4.4 shows that the total energy ($E_{tot} = E_{kin} + E_{int} + E_{grav}$) over 50 dynamical times falls by around 2.5%. This is due to the controls on the atmosphere, which leach energy out of the simulation over time, although the magnitude of this effect remains relatively small and only affects zones outside the stars themselves.

4.5 Unmagnetized neutron star mergers

Table 4.1 summarizes the parameters for the unmagnetized merger simulations. Our central model is Quad7, and this is the one that will be discussed in detail in the following sections. Quad6 is a lower resolution run using the same gravitational wave solver and Pmass7 uses the alternative point mass gravitational wave solver. In all cases the neutron stars have equal mass and are irrotational. The main model, Quad7 required approximately 150,000 CPUh for a 15ms simulation. The magnetized cases presented in Section 4.6 required approximately twice as much CPU time per ms of simulation time.

Model	No. refinement levels	Δx_{min} (km)	GW solver	M_{NS} ($M_\odot$)	R_{NS} (km)	D_0 (km)
Quad7	7	0.29	Quad	1.4	14.3	45
Quad6	6	0.58	Quad	1.4	14.3	45
Pmass7	7	0.29	Point Mass	1.4	14.3	45

Table 4.1: The parameters used in the unmagnetized neutron star mergers.

4.5.1 Results

Figure 4.5 shows the densities in the $x-y$ plane at several times throughout the merger, for model Quad7. At $t = 2$ms we see that the stars have become elongated by tidal forces and have made contact with one another. From the right panel of Figure 4.6, which shows that the magnitude of the vorticity (the curl of the velocity vector), we can see that the shear motion along the interface between the two stars is driving turbulent motion. This is generated by the Kelvin-Helmholtz instability, which is activated by the shearing along this contact plane. At the edges of the this interface see the curls of two larger vorticies forming at the edges of the interface between the two stars, and such vorticies appear and vanish as the two stars continue to merge.

By $t = 6$ms the stars are fully tidally disrupted, the central object appears bar-like and tidal tails trail from its ends. The shear plane between the matter of the two stars is still present, leaving a narrow band of lower density between the two stars. The bar shape then develops into a more spherical central object, although there are still two distinct density cores. A large part of the initial tidal tails visible in the 6 and 7ms plots are ejected from the grid, with what remains winding up around the remnant and beginning to form a disc around the central density core. We see that in the final plot at 20ms there are still two distinct density cores remaining, surrounded by a disc with prominent spiral waves eminating from the central object.

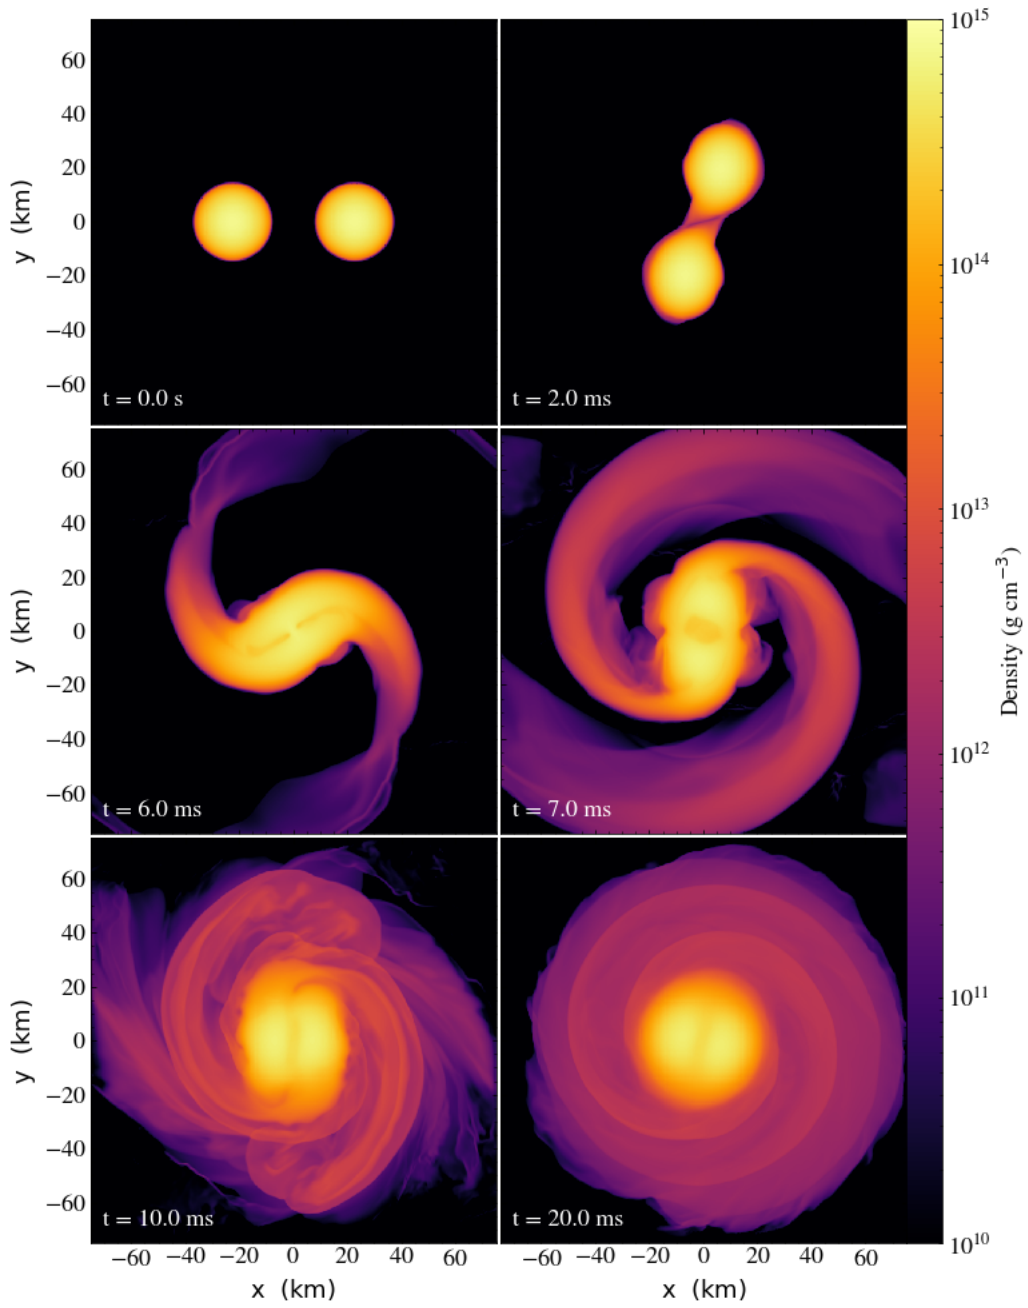

Figure 4.5: Density in the $x - y$ plane of the the binary at t =0, 2, 6, 7, 10 and 20ms. The initial orbital period of the binary is around 3ms.

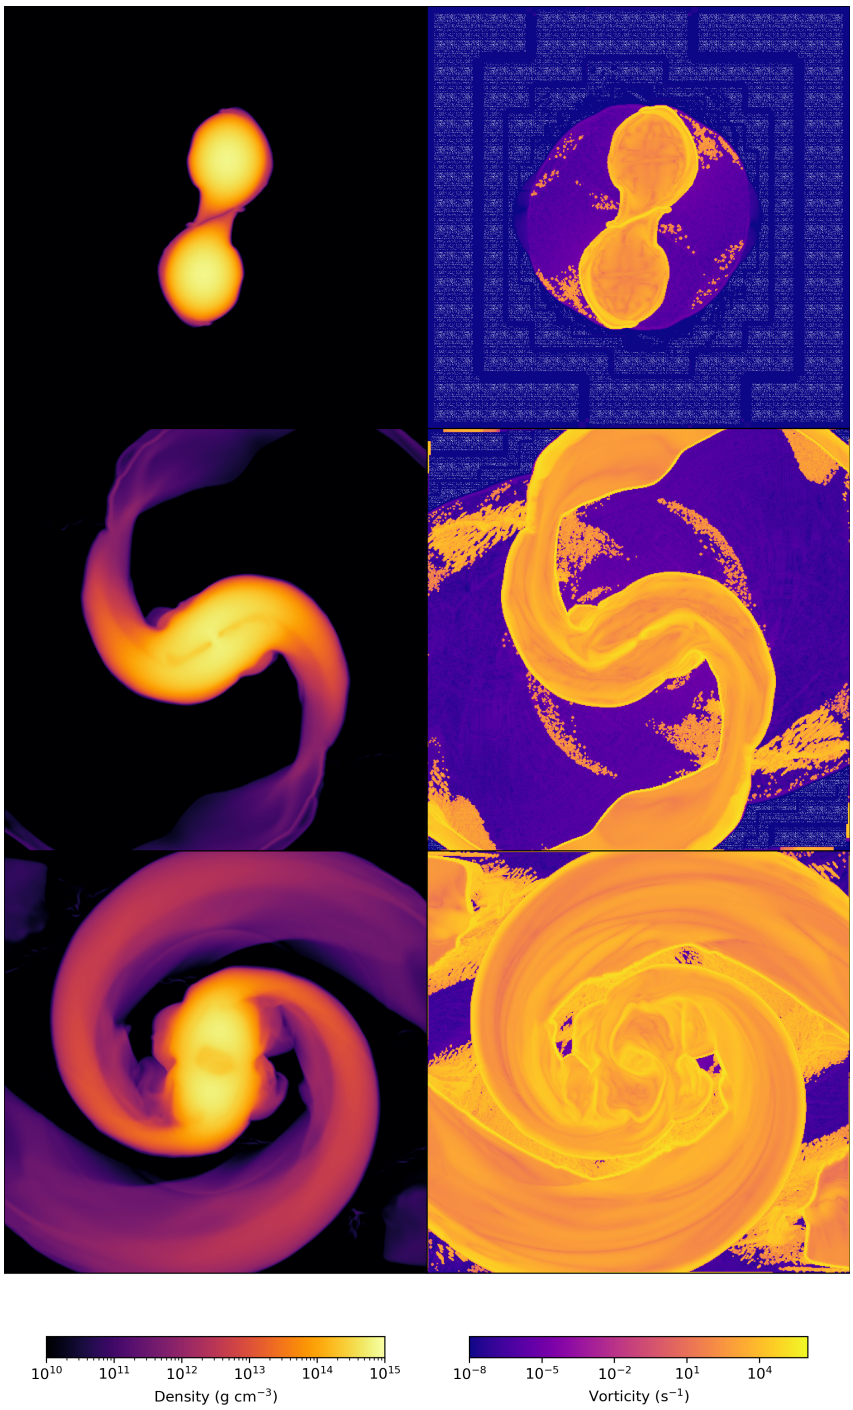

Figure 4.6: Density and vorticity magnitude (left and right respectively) in the orbital plane of the binary at $t = 2.05$ms (top) $t = 6$ms (middle) and $t = 7$ms (bottom).

Figure 4.7 shows the evolution of the total energy (including the energy emitted as gravitational waves) and mass normalised by their respective initial values throughout the simulation. We see that over the course of the merger about 5% of the original total energy is lost through the controls on low density atmosphere and the loss of mass off the grid. Between 6 and 7ms there is a plateau in the energy loss, which is due to the spreading of the tidal tails. This temporarily raises the density in a large part of the grid above the levels at which the atmosphere controls would kick in, thereby reducing the energy loss. Following this, there is a short period of more rapid energy loss that corresponds with the only significant period of mass loss, which is due to the initial tidal tails leaving the edge of the grid.

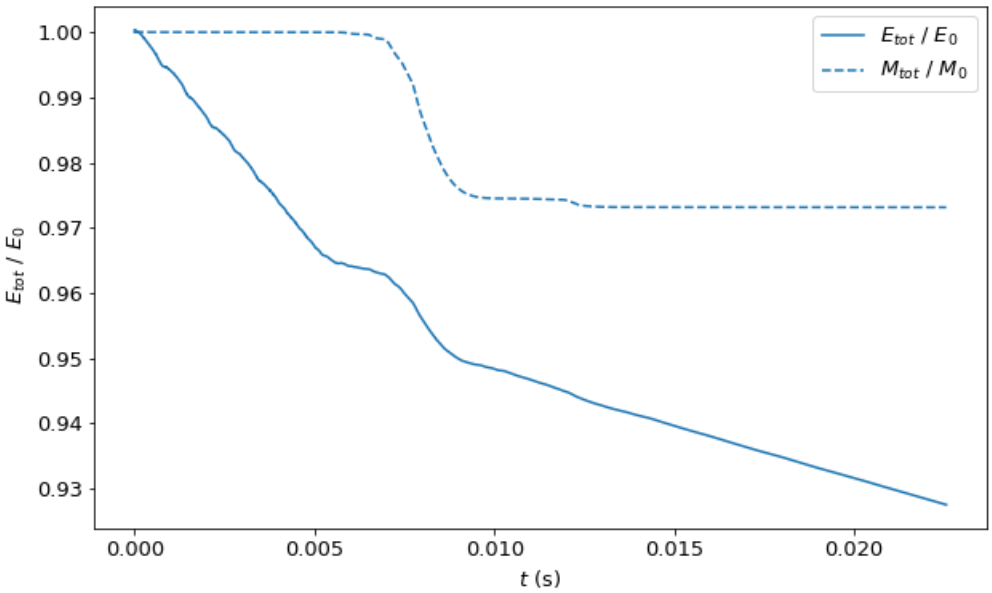

Figure 4.7: The total energy and mass on the grid plotted with time and normalised by their initial values.

Figure 4.8 shows the evolution of the totals of each of the energy components during the simulation. We see that the internal energy falls and the kinetic energy rises in the initial inspiral, reaching a minimum and maximum respectively at the point of merger. This drop in internal energy is due to the tidal stretching of the stars during the inspiral, while the rising kinetic energy simply comes from the stars inward spiral trajectory, which leads to an increase in the orbital speed as the orbit's radius falls. This is also evident in the increase gravitational potential energy (becoming more negative) up to the point of merging. Post-merger there is a much shallower rise in internal energy and fall in kinetic energy, as shock heating in the merger remnant converts kinetic to internal energy. The gravitational potential energy remains roughly constant as the mass distribution begins to settle into its final disc and central object configuration.

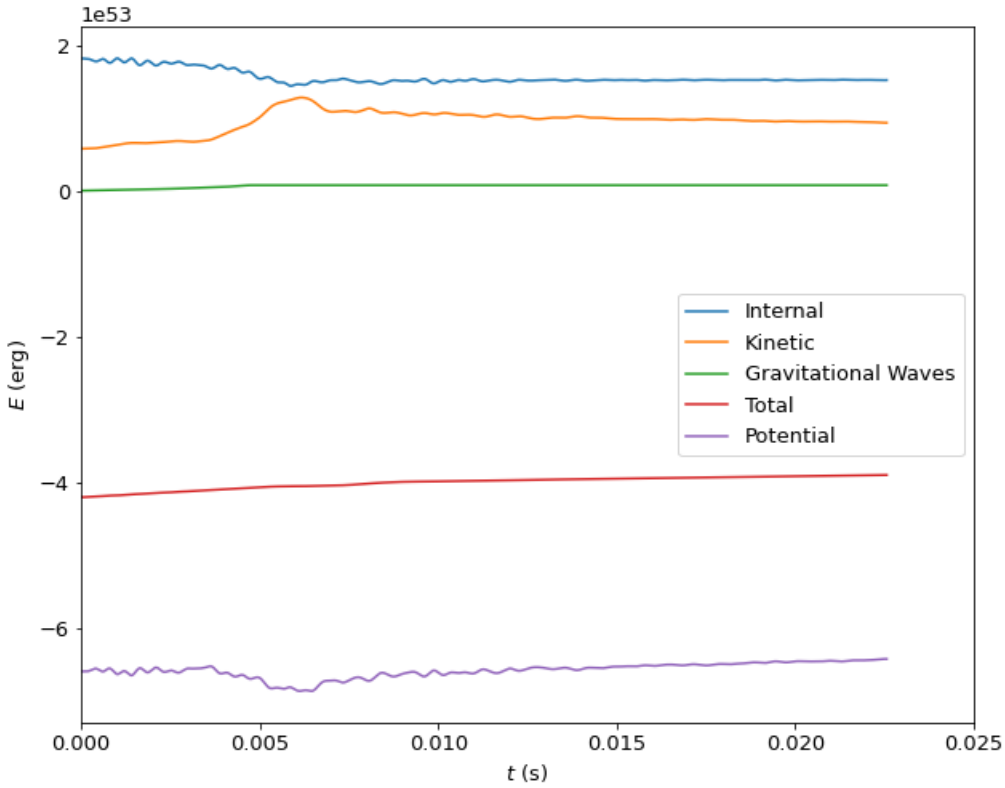

Figure 4.8: The total internal, kinetic, gravitational wave and gravitational potential energies and the overall total energy (the sum of these components) in the simulation volume for the duration of the simulation.

The evolution of the maximum density in the simulation is shown in Figure 4.9, for models Quad6 and Quad7. We see that both initially decline, reaching a minimum that occurs at the merger time, before recovering slightly. In both models, the minimum of the density maxima coincides with the point at which the stars are at their most tidally stretched, resulting in the bar-like shape we see in the lower panels of Figure 4.6. Following this, the tidal tails continue to expand whilst the inner remnants of the stars collapse to become more spherical again, causing the density to grow again. The initial decrease is more rapid in the lower resolution (Quad6) model, as the merger precedes more rapidly in this case. It also leads to a lower minimum and recovers to a lower value. Furthermore, in both cases the maximum density appears to be declining slightly in the post-merger phase. Neither case produces a higher maximum density post-merger than the initial value, although we do see a peak in the Quad7 value at 6ms which corresponds to a minimum in the separation of the density cores.

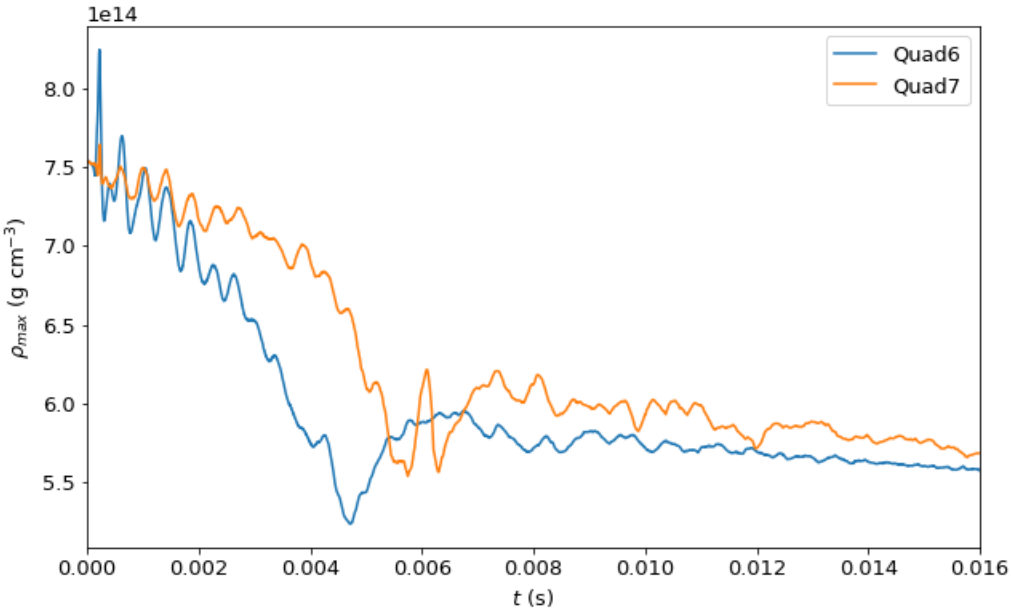

Figure 4.9: The evolution of the maximum density in the simulation volume, for the two different maximum grid resolutions.

Comparing Figure 4.9 with Figure 15 of [Baiotti et al., 2008], which corresponds to a model with the same equation of state, similar initial masses and separations but uses a fully relativistic hydrodynamics solver, we see a drop in ρ_{max} immediately before merger of similar magnitude. Their peak density at the merger time is substantially higher, becoming briefly around 10% higher than the initial maximum. The difference in the height of these peaks is likely related to simulation resolution, since the finest grid resolution in [Baiotti et al., 2008] is approximately twice ours and we see no peak at all in our lower resolution run. The lower grid resolution will smooth over the contact interface between the stars, where the fluid is compressed to produce the peak. Additionally, the post-merger density maximum in [Baiotti et al., 2008] recovers to approximately the same as the pre-merger value, whereas ours recovers slightly but then steadily declines. This is likely a result of our gravitational wave treatment, since their fully relativistic treatment includes gravitational wave emission throughout the merger whereas we neglect the gravitational back reaction after 5ms into the simulation. To assess this, Figure 4.10 shows the maximum density for simulations in which we turn off the gravitational wave backreaction at different times. We see that in the case where the backreaction remains switched on throughout the post-merger evolution, the maximum density begins to recover over similar timescales to those in [Baiotti et al., 2008]. This is because the gravitational wave backreaction extracts angular momentum from the remnant, causing it to contract and become more compact.

Whilst continuing to include the backreaction improves this aspect of the simulations in comparison with a fully relativistic model, the method by which we reduce

the velocities relies on being able to distinguish the two stars. Once the stars have merged, continuing to apply the backreaction leads to unphysical asymmetries in the remnant, since the mixing of the matter from the two stars can cause the reduction in the velocities to be applied erroneously.

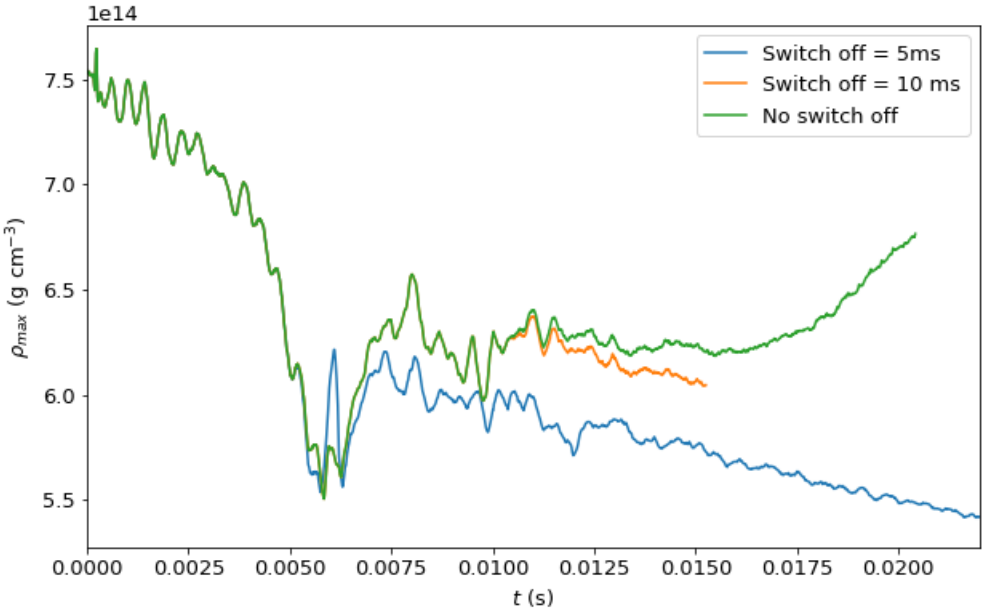

Figure 4.10: The evolution of the maximum density in the simulation volume, for models in which we switch off the gravitational wave backreaction at different times. Note that the 5ms switch off corresponds to our Quad7 model.

4.5.2 Gravitational wave solvers

To examine the effectiveness of our included gravitational wave solvers we show the gravitational wave luminosity for all three unmagnetized simulations in Figure 4.11. The luminosity in the point mass model very rapidly becomes large as a result of its limited validity once the stars begin become tidally deformed. The tidal deformation causes the stars to fall inward quicker than a pair of point masses would, which drives the point mass luminosity to grow rapidly as it scales with the star's separation as $1/d^5$. This further drives the star's infall, creating a feedback loop that very quickly drives the luminosity to unphysically high values. As such this solver is only useful in the very early stages of this simulation, as the stars become significantly tidally deformed during the first 3ms of the simulation given our initial separation.

The large-scale shape of the signal for the quadrupolar solver is broadly as expected, there is a large initial peak that coincides with the rapid decline in orbital separation and complete tidal disruption of the two stars. Following this there are several smaller peaks which in turn correspond to periods in with the

remaining cores of the stars have 'bounced' apart before pulling closer together again. The signal is however very noisy, with strong oscillations the amplitude of which is correlated with the luminosity itself. The frequency of these oscillations is around four times the orbital frequency of the binary, indicating a link with the orbital motion of the binary.

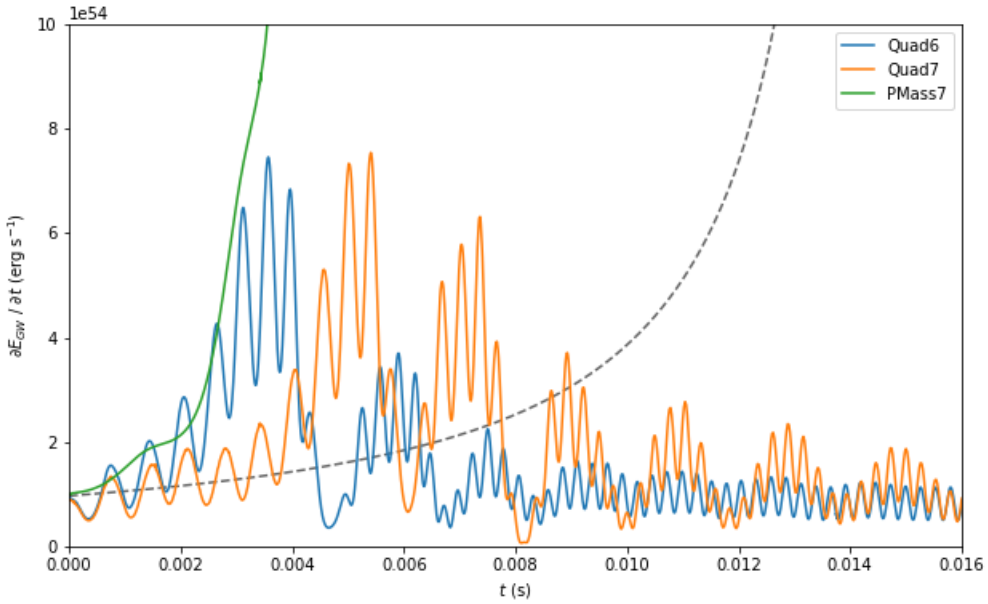

Figure 4.11: The evolution of the gravitational wave luminosities for models Quad6, Quad7 and Pmass7. The dashed line shows the analytic solution for a pair of point masses of equivalent mass and initial separation.

There is a substantial difference between the luminosities for the two quadrupolar simulations, the Quad6 model peaks around 1.6ms earlier and only has two smaller peaks following the initial maximum. The Quad7 model in contrast shows many more peaks with larger sizes relative to the main peak that are still present by the end of the simulation. This is due the stronger quadrupolar structure in the post-merger Quad7 model, which we can see by comparing the density distributions for the two simulations in Figure 4.12. The higher resolution case obviously shows more structural detail in the both remnant and surrounding spiral structure at all times. Further to this, we see that whilst both models still have a double-core structure in their centers even in the 15ms panels (bottom), the surrounding material in the Quad6 case is more uniformly distributed and spreads more widely in the orbital plane.

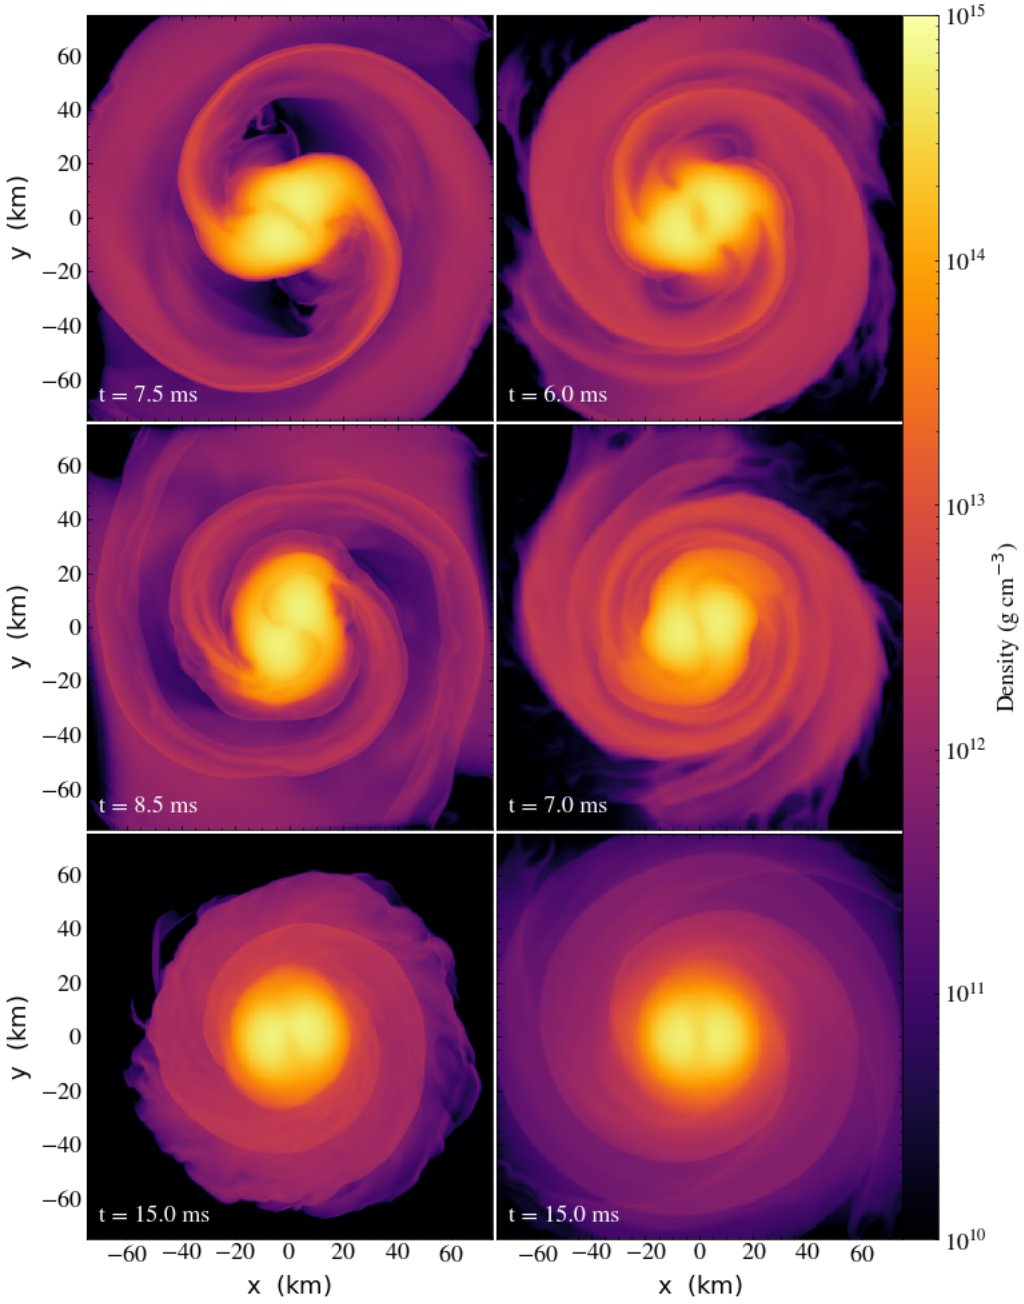

Figure 4.12: Density in the $x - y$ plane of the binary for models Quad7 (left) and Quad6 (right) at several times throughout each simulation. In the first two rows we show the Quad7 densities 1.5ms later than the Quad6, as this is the separation between the gravitational wave luminosity maxima and so the density distributions are more comparable at these shifted times.

To see more clearly the underlying signal in the gravitational wave luminosities, we apply a low-pass Butterworth filter to the data in order to remove the oscillations. This filters out frequencies in the signal above a specified threshold, which we choose to be between the frequency of the large scale structure in the signal and the oscillations we wished to remove. We also time-shift the Quad6 data by 1.62ms to align the luminosity maxima, the results of which are shown in Figure 4.13. We see that the main peaks are quite similar, they have similar width and height, although there is a slower initial growth in the Quad7 case that leads to the time offset. Following this, the subsequent peaks in the Quad7 case have a shorter period and are much larger than the corresponding features in the Quad6 signal.

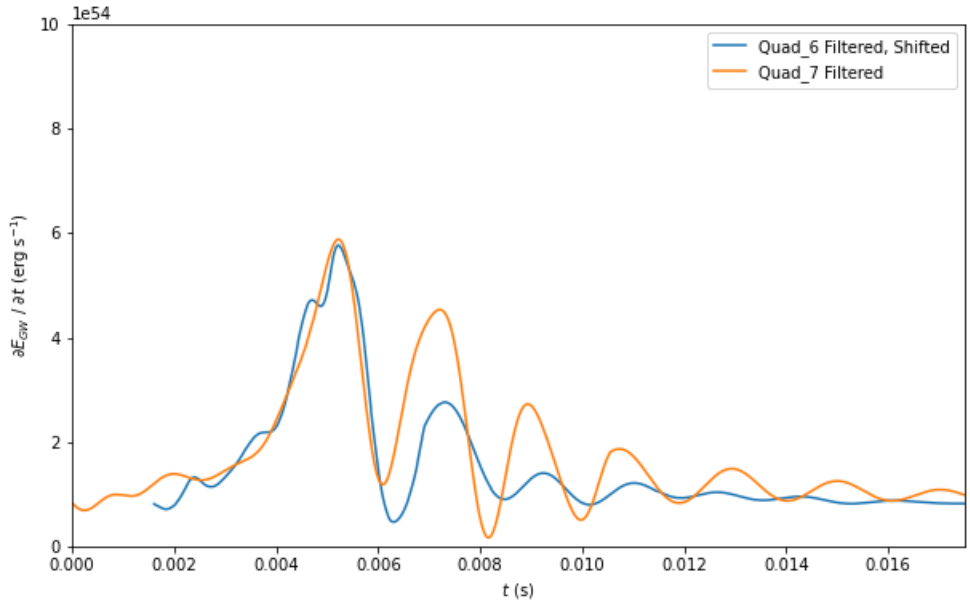

Figure 4.13: The evolution of the gravitational wave luminosities for models Quad6 and Quad7, filtered with a low-pass filter and time shifted by 1.62ms such that the maximum luminosities match.

To examine the source of the oscillatory behaviour in the quadrupolar models, we show the contribution of each term in Equation 4.6 to the components of the $\dddot{Q}_{ij}$ tensor in Figure 4.14. We see oscillatory behaviour is present in all the terms and the total for both the diagonal and off-diagonal terms. In particular, the oscillations of the total agree with what we would expect from the point mass case described by the third derivative of Equation 2.4, the diagonal terms are roughly sinusoidal and the off-diagonals are cosines. The contributions from the individual terms appear to come in pairs that almost but don't fully cancel with each other. The amplitude of the 2nd and 4th terms is much larger than that of the 1st and 3rd terms, but both pairs contribute relatively equally to the oscillations in the total, although which pair is dominant varies during the cycle.

As the point mass case suggests that the total components should still exhibit oscillations however, we don't expect that these terms should fully cancel to give a constant signal.

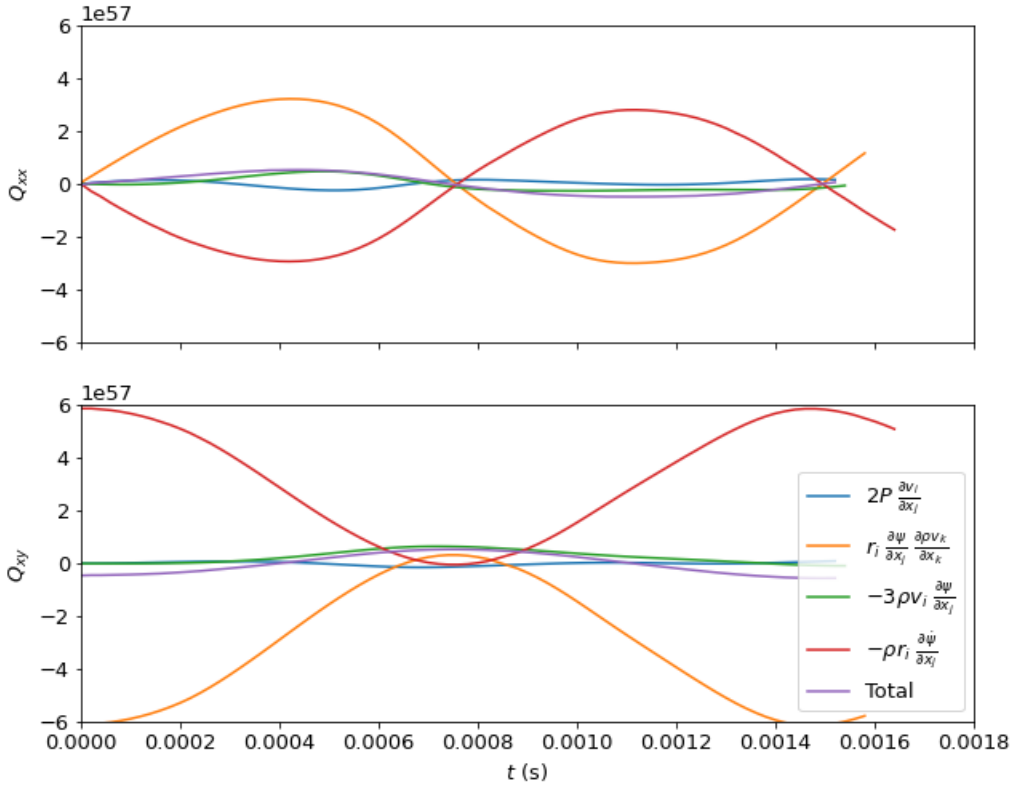

Figure 4.14: The contributions to the $\dddot{Q}_{xx}$ (top) and $\dddot{Q}_{xy}$ (bottom) components from each term in Equation 4.6.

In Figure 4.15 we show the non-zero components of the $\dddot{Q}^2_{ij}$ tensor. We see that the off-diagonal and diagonal terms oscillate with opposite phase, but that there is some variation in the diagonal terms leading to incomplete cancelling and therefore oscillations in the total. We know from the derivation in Section 2.2.1 that in the point mass limit these pairs of terms should cancel exactly. This incomplete cancelling we see is therefore the source of the oscillations in the gravitational wave signal. The causes of this appear to be twofold. Firstly, since we begin our simulations with the stars quite close together, they become tidally deformed almost immediately. The deformation is additionally quite asymmetric, resulting in a teardrop like shape which can be seen in the 2ms panel of Figure 4.6. This produces a mismatch between the $\ddot{Q}_{xx}$ and $\ddot{Q}_{yy}$ components, which leads to the oscillatory deviations from the point mass signal. In addition, there was a small error in the application of the traceless condition within our model which exacerbated this effect. By reconstructing the corrected luminosity we see that this increased the amplitude of the oscillations around a third to a half.

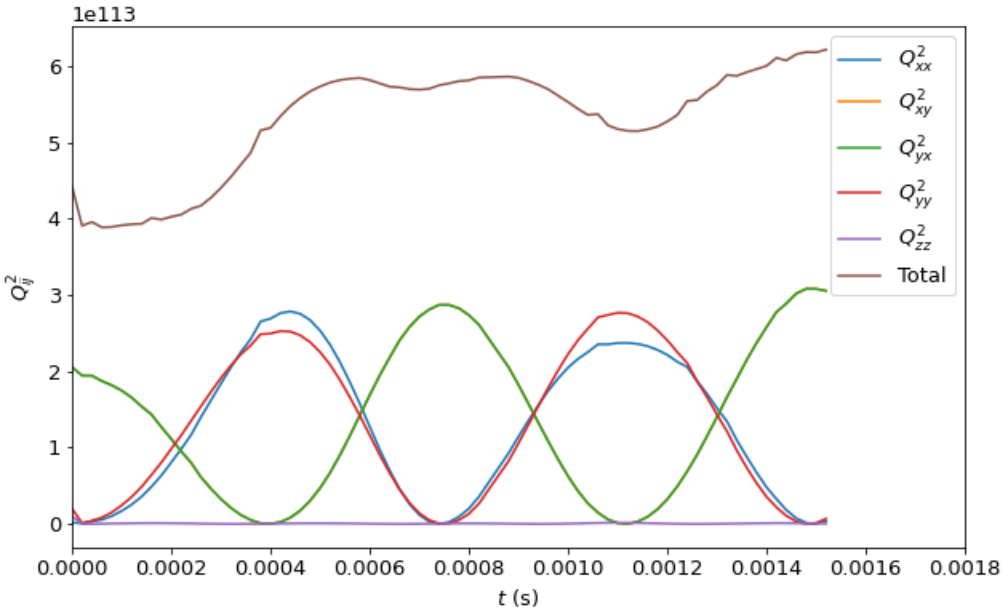

Figure 4.15: The non-zero components of the $\dddot{Q}_{ij}^2$ tensor, the sum of which gives the gravitational wave luminosity (absent the constant factors). Note that the symmetry of the tensor means that $Q_{xy} = Qyx$. The Q_{zz} term and all others not shown remain roughly constant and close to zero.

Figure 4.16 shows the separation of the density maxima for the Quad6, Quad7 and Pmass7 models. A short section of data between 2 and 4.5 ms for Quad7 and the Pmass7 separations are given instead by the centre of mass of the star, due to data loss. This is calculated using all mass above a threshold of 10^{11}g cm^{-3} which is assigned to a given star by being closer to the previous timestep's center of mass for that star than for the other star. This agrees very well with the density maximum location during this early stage of the calculation but becomes an unreliable measure of the separation between the density cores of the stars once they are tidally disrupted. We see that in all cases there is an initial period of slow decline followed by a rapid fall. For the quadrupolar models this is followed by a long period in which the separation slowly decays but with some oscillations present. The peaks of these oscillations correspond quite clearly to the peaks in the post-merger gravitational wave signal, and to the density cores in the center of the remnant appearing to 'bounce' off one another moving apart and pulling closer together again. This widening of the separation corresponds to an increase in the quadrupole moment, which is why we then see the increase in gravitational wave luminosity. Although the onset of the merger is more rapid in the Quad6 case, the Quad7 separation shows a more prominent decline in the post-merger phase, suggesting the final merging of the density cores will occur faster in the Quad7 model.

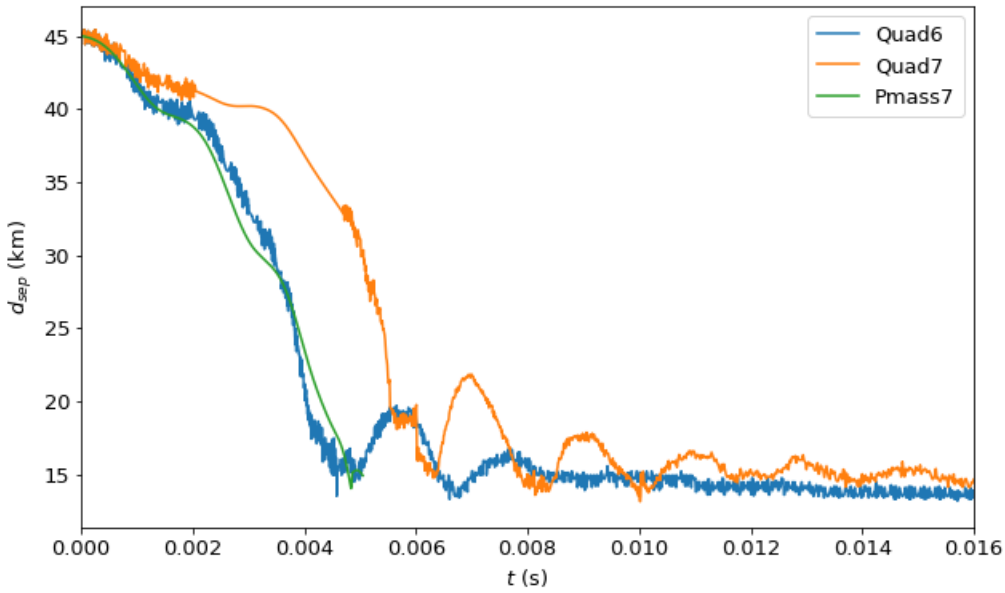

Figure 4.16: The evolution of the separation of the density maxima for models Quad6 and Quad7

4.5.3 Discussion

We performed this unmagnetized merger to examine the dynamics of the merger without magnetic fields and to assess how well our approximate treatment of gravitational wave emission recreates the inspiral and merging of the neutron stars. The gravitational wave scheme we use was inspired by the similar scheme used in [Zhuge et al., 1994], and the work of [Ruffert et al., 1997b] also ran a model with similar parameters to compare the effects of the different numerical schemes. Our results are broadly similar to those presented for model ZCM in [Ruffert et al., 1997b] and Run 2 of [Zhuge et al., 1994]. The qualitative features of the density in the orbital plane (Figure 1 of [Ruffert et al., 1997b] and Figure 4.5) are similar, showing the same tidal disruption into an elongated central object with trailing tidal tails followed by collapse to a remnant with a double-core structure. The evolution of the total energy components in Figure 4.8 are also in good agreement with those in Figure 2 of [Ruffert et al., 1997b], showing very similar trends for each component.

Further to this, the gravitational wave luminosity has a similar overall structure, however here there are some differences. Firstly, the models of [Zhuge et al., 1994] and [Ruffert et al., 1997b] are not subject to the same oscillations as our models, since [Zhuge et al., 1994] use a point mass luminosity formula and [Ruffert et al., 1997b] use a Post-Newtonian approximation to apply the gravitational wave back reaction as source terms directly in the energy and momentum equations. Models Quad6 and ZCM show the greatest similarity, with all three visible peaks of similar height in both simulations, whereas in Quad7, the later peaks in luminosity

are much larger relative to the initial maximum and more numerous. In contrast, Run 2 of [Zhuge et al., 1994] shows a similar initial peak but much lower subsequent peaks than Quad6 and ZCM. It is suggested in [Ruffert et al., 1997b] that this discrepancy could due to the lower resolution used in [Zhuge et al., 1994], since the increased numerical viscosity at lower resolution will act to damp out oscillations in the remnant, driving it towards a spherically symmetric configuration more rapidly. This is supported by our results, which show comparatively larger later peaks in the higher resolution case. Further to this, the ZCM model includes a full gravitational wave backreaction in the post-merger phase whereas our simulations do not. This gravitational wave emission will act the damp the oscillations of the remnant, which further explains why we see stronger secondary peaks in the gravitational wave luminosity.

The total energy emitted in gravitational waves over the course of the ZCM model (around 8ms) is reported as 2×10^{52}erg, which our Quad7 model agrees with. The Quad6 model shows a slightly lower overall emission of around 1.7×10^{52}erg. Finally, the models of ZCM predicts a much earlier merger than either of our models (Run 2 of [Zhuge et al., 1994] begins with a larger initial separation, meaning we do not have a direct comparison for this), the maximum of the gravitational wave luminosity comes at around 2ms, whereas in models Quad6 and Quad7 it is at 3ms and 5ms respectively. This could again be a feature of our greater resolution and therefore lower numerical viscosity (Quad6 is approximately twice the resolution of ZCM at its finest AMR level).

When compared with more recent, fully relativistic models, the dynamics of our merger are qualitatively very similar (for example Figure 11 of [Thierfelder et al., 2011], Figure 7 of [Baiotti et al., 2008], Figure 2 of [Bauswein and Stergioulas, 2015] and Figures 6 and 9 of [Espino et al., 2022]). We see similar tidal disruption and evolution of the remnant immediately post-merger, which we would expect as this stage of the evolution is driven mainly by hydrodynamic effects. In particular, the double-core structure remaining in our remnant is seen in other simulations [Bauswein and Stergioulas, 2015] [Paschalidis et al., 2015a], which also show the cores 'bouncing' apart and together several times before they eventually merge. Obviously our Newtonian code is unable to replicate the collapse of the remnant into a black hole however, for the $\Gamma = 2$ ideal fluid equation of state and initial star masses of $1.4M_{\odot}$, [Baiotti et al., 2008] do not see a collapse until 20ms into the simulation. Therefore in this early stage of the post-merger evolution our results appear to be comparable to what we expect from GR codes.

Comparing the separation of the density maxima in our simulation to those in Figure 4 in [Baiotti et al., 2008], we again see similar features in the early evolution. Both show a slight kink in the shallow part of the inspiral, followed by the rapid decline as the stars plunge inward. Following this they then see short period of oscillations as the double cores in the remnant bounce before collapsing to a black hole. This particular figure corresponds to a higher mass binary, however they report that the initial inspiral and onset of the merger are largely the same apart from the longer timescale for collapse to a black hole.

Figure 1 in [Bernuzzi, 2020] shows the gravitational wave signals for various merger remnant scenarios. We see that although the shape is very different (with our simplified emission model this is very much expected), the relative magnitude of the post-merger emission to the main peak is comparable to our results for the short-lived remnant case, which is what we would expect based on our initial NS masses and EOS.

To summarize, the results of our simulations show good agreement with those of various others, including both Newtonian and fully relativistic calculations, within the bounds of what we expect from our physical approximations and CPUh constraints. We find that resolution has a strong impact on merger times and the evolution of the remnant in the short-term, as the heightened numerical viscosity causes the merger to proceed more rapidly. This reduces the lifetime of the non-axisymmetric remnant, thereby decreasing the predicted gravitational wave emission, and has also been observed in other studies both Newtonian [Ruffert et al., 1997b] and relativistic [Yamamoto et al., 2008].

4.6 Magnetic field orientation in BNS mergers

To examine the effect magnetic fields of different orientations have on a neutron star merger, we perform two initial simulations with different magnetic field configurations. The first model, referred to as IntA, represents the standard configuration used in most magnetized merger simulations, with the two stars endowed with dipolar fields aligned with the orbital angular momentum and with maximum field strength of 10^{15}G, the construction of which is described in Section 4.3.1. The second model, IntT, has one star with it's field aligned to the orbital angular momentum while the second is rotated so that it is initially aligned with the line connecting the two stars (although since the stars are irrotational, this does not remain true as they orbit) and the same field strength, which while physically very large was chosen for comparison with other simulations. The simulations are otherwise identical to the unmagnetized Quad7 model and for computational efficiency the stars are only seeded with the magnetic field 1.5ms into the simulation, which is around a quarter orbit before they make contact.

Technical problems and a shortage of time meant that we were unable to evolve these models for as long as intended, with each model only reaching about half the runtime of the Quad7 run described in the previous section. Therefore, we can only comment so far on the field behaviour in the initial stages of the merger.

4.6.1 Results

The density in the orbital plane at several times is shown for models IntA and Quad7 in Figure 4.17. Whilst the main features of the two are mostly the same, we see that in the magnetized case the merger happens more rapidly. Further to this, the tidal tails are wider in the magnetized case. This is more easily seen in Figure 4.18, where we compare the density in the equatorial plane for IntA with those in Quad7 around 0.25ms later, so that the mergers are more closely matched in progress. We see here also that the region of lower density along the plane of contact between material from the two stars is wider in the magnetized case. This is due to stronger mixing along this shear interface where the two stars meet, the initial thin interface evolves into multiple vorticies which merge into a larger central one just before 6ms to produce the larger region of lower density seen in Figure 4.18. This is the case in for both magnetic field variants, although there is some small difference in the location of the vorticies.

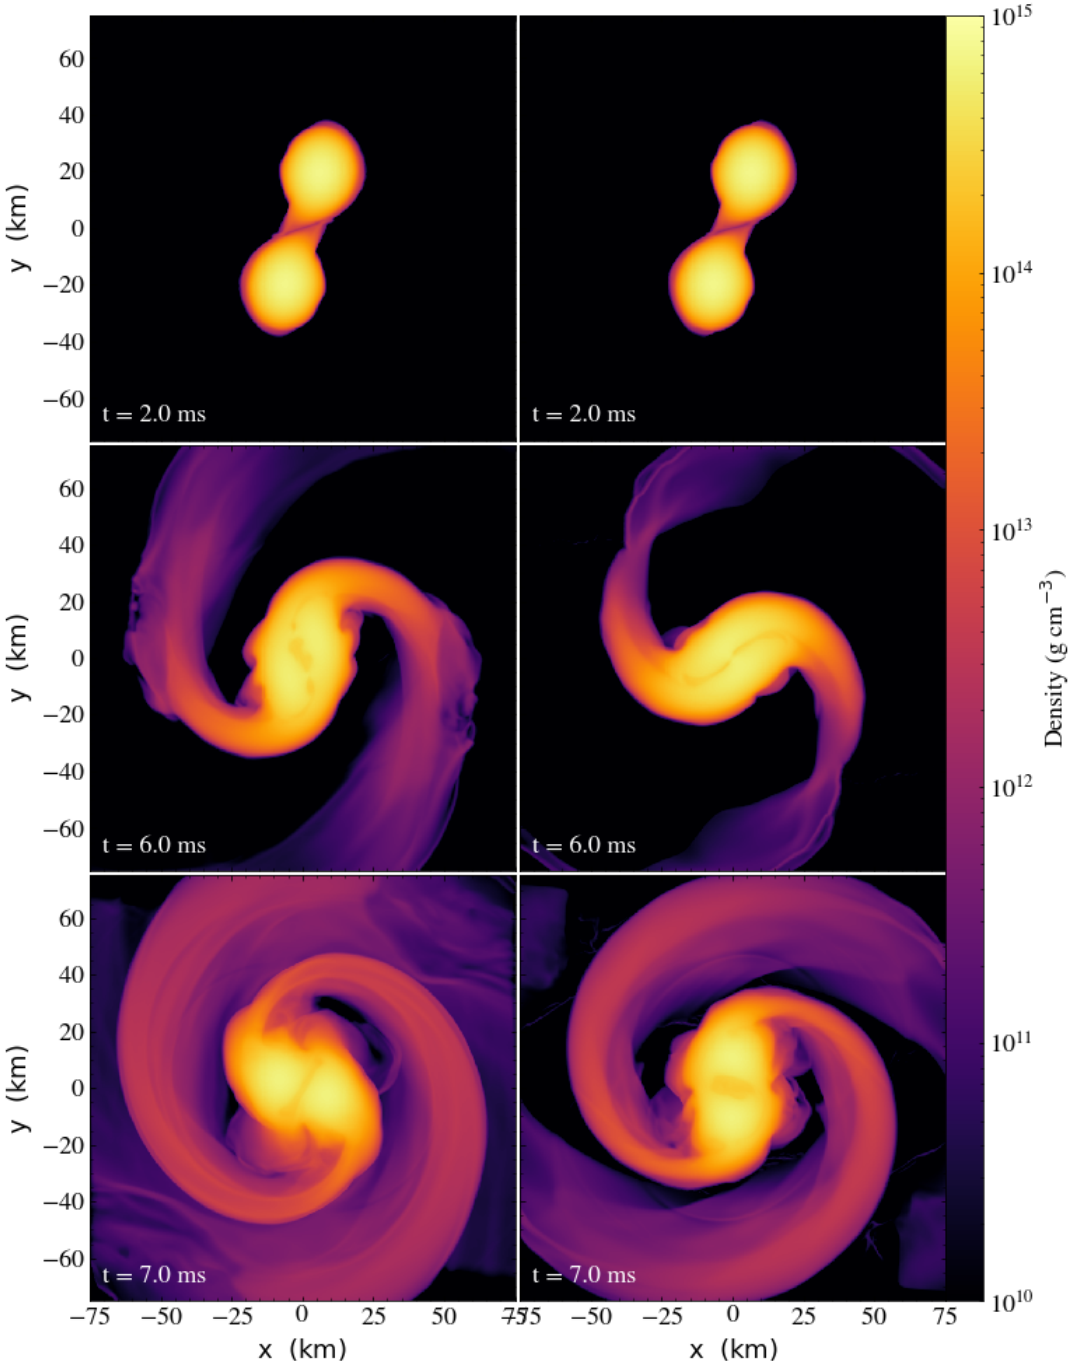

Figure 4.17: Density in the $x - y$ plane of the the binary at t =2, 6, and 7ms, for the IntA and Quad7 simulations on the left and right respectively. The initial orbital period of the binary is around 3ms.

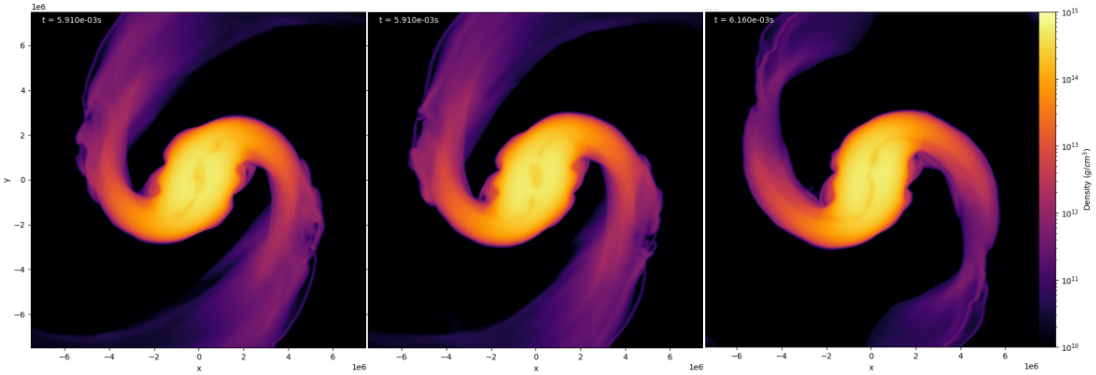

Figure 4.18: Density in the $x - y$ plane models IntT (left), IntA (middle) and Quad7(right) at 5.9ms for the magnetized models and 6.1ms for the unmagnetized case.

Figure 4.19 shows the density, magnetic field strength and vorticity magnitude in the orbital plane for model IntA. We see that along the shear interface at 5ms there is high vorticity as the Kelvin-Helmholtz instability activates, which in turn drives the magnetic field growth through the twisting of field lines. This evolves into a zone between the two density cores where there is some turbulent mixing. There is also strong field growth along the leading edge of each star visible at both 3 and 5ms, as tidal disruption causes the leading edge of the star to become stretched out compared to the trailing one. This bends the field lines along the the leading edge, causing energy to be converted from kinetic to magnetic energy by the magnetic tension force. By 7ms the field surrounding the density cores is now stronger than the field within them, whereas initially the maximum of the field was at the center of each star.

Figure 4.20 shows the same variables but for model IntT at 3, 5 and 7ms into the simulation. The only large differences are in the magnetic field strength, we see the asymmetry in the orbital plane field strength due the the different field configurations in each star. The field in the tilted star appears stronger in this plane from the outset, however this is a feature of the viewing angle, since the circular minima we see in the angular momentum aligned star is simply rotated out of the plane. Due to this different orientation, the growth in field strength along the leading edge of the star is less pronounced. It is not visible at all in the 3ms panel and is difficult to distinguish still in the 5ms panel.

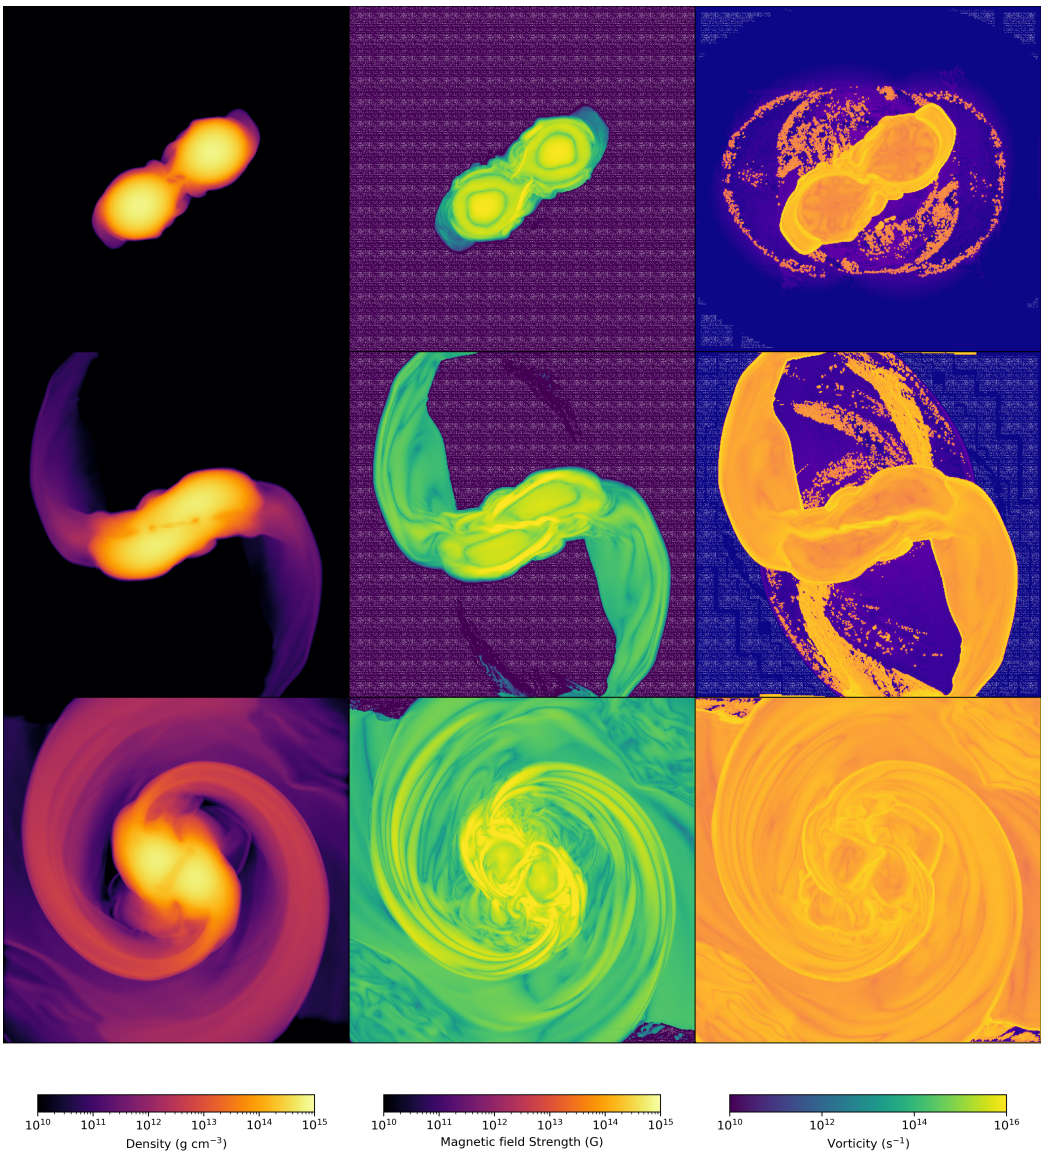

Figure 4.19: Density, magnetic field strength and vorticity magnitude (left, middle and right respectively) in the orbital plane of model IntA at $t = 3$ms (top) and $t = 5$ms (middle) and $t = 7$ms (bottom).

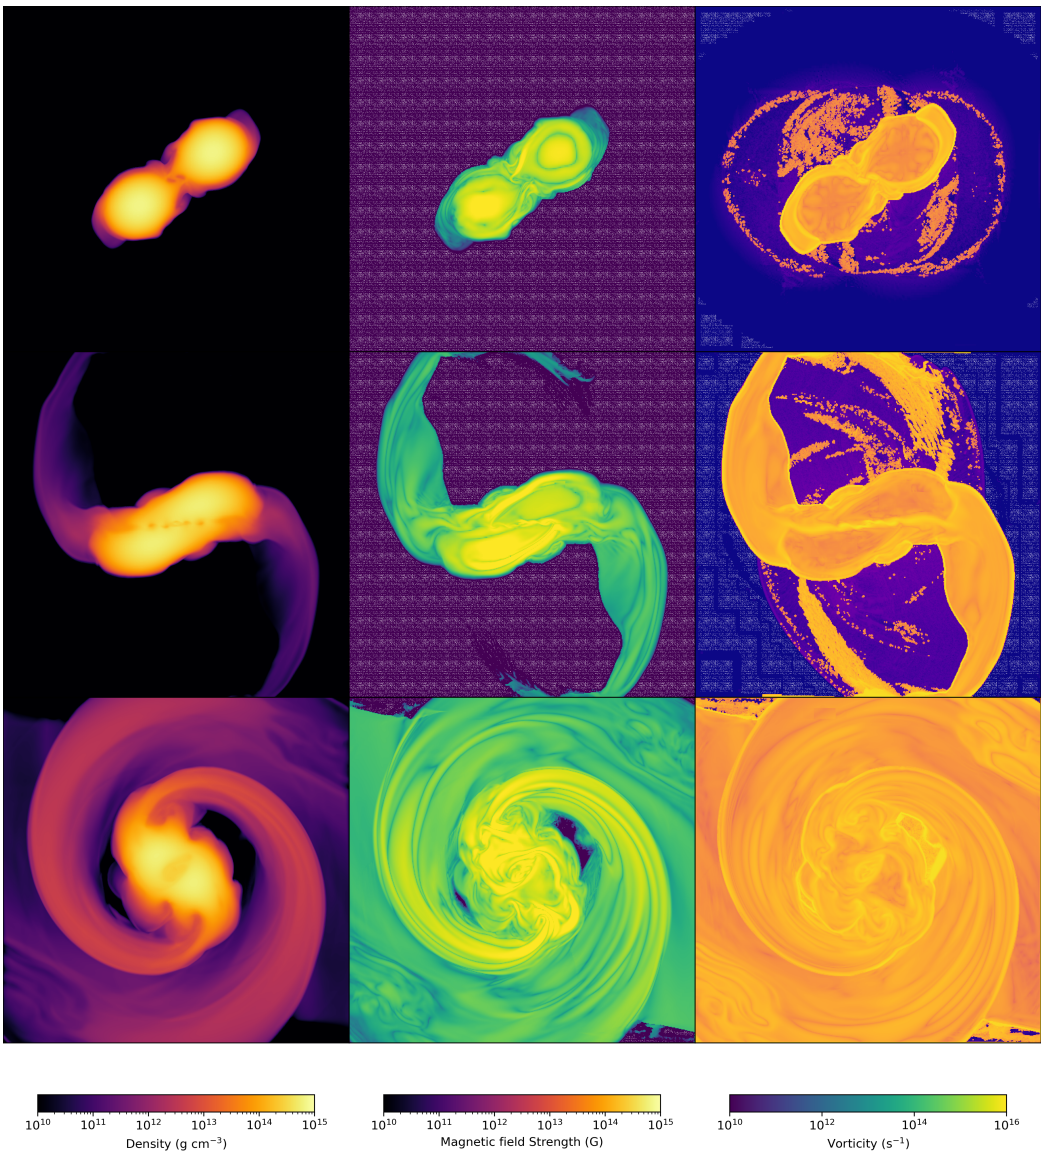

Figure 4.20: Density, magnetic field strength and vorticity magnitude (left, middle and right respectively) in the orbital plane of model IntT at $t = 3$ms (top) and $t = 5$ms (middle) and $t = 7$ms (bottom).

Figure 4.21 shows how the overall magnetic energies compare between models IntA and IntT. We see that whilst the initial energies are equal, the total magnetic energy initially grows more rapidly in model IntA. The growth is temporarily halted at around 6ms, before growing again to a peak and finally falling off. In contrast, the initial growth of the total energy is slower in model IntT, however instead of reaching a peak it continues to grow at a shallower rate, appearing to plateau towards the end of the simulation. The peak energy reached during the simulation is around 10% higher in the IntT case than the IntA.

In both models the energy in poloidal and toroidal components (defined as the azimuthal and radial plus polar components respectively in spherical coordinates

centered on the binary center of mass) grow at the same rate until around 5-6ms into the simulation (at which point the tidal tails are forming). In model IntA the toroidal component is initially weaker, but grows larger than the poloidal component and reaches its peak slightly later. In model IntT, the initial components oscillate, due to the field on one of the stars lying in the orbital plane, meaning its field contributes to different components predominantly depending on where in its orbit it is. Following the onset of tidal disruption the growth in the poloidal component becomes shallower, whilst the toroidal component's growth goes through phases of more and less rapid growth before seemingly reaching a peak at around 10ms.

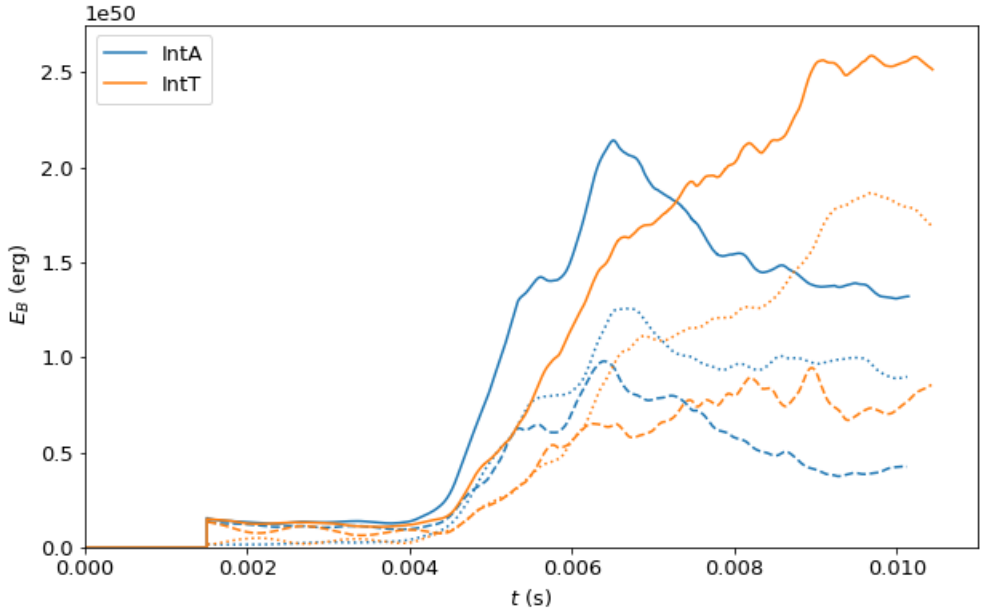

Figure 4.21: The poloidal (dashed), toroidal (dotted) and total (solid) magnetic energies for models IntA and IntT.

4.6.2 Discussion

The dynamics of both magnetic field cases are similar to those of the unmagnetized case, however for both initial field configurations the merger precedes more quickly, happening around 0.5ms earlier for our initial conditions. We believe that this is a result of the magnetic pressure causing the stars to expand slightly in the early stages of the simulation, exaggerating the effects of the tidal distortions. This can be seen in Figure 4.22, where the magnetized simulations show a slightly larger tidal elongation of the stars along with more diffuse matter around the surface of the stars. Such effects are due to the exceptionally high magnetic field strength, and simulations at similar field strengths [Ruiz et al., 2019] also see this slight decrease in the merger time. [Giacomazzo et al., 2009] further demonstrate that the magnetic field only impacts the inspiral of the stars when it is above 10^{14}G.

As a result, we would not expect this delay in merging to be a feature of observed binaries since the initial field strengths required are highly unlikely to be present initially in the merger.

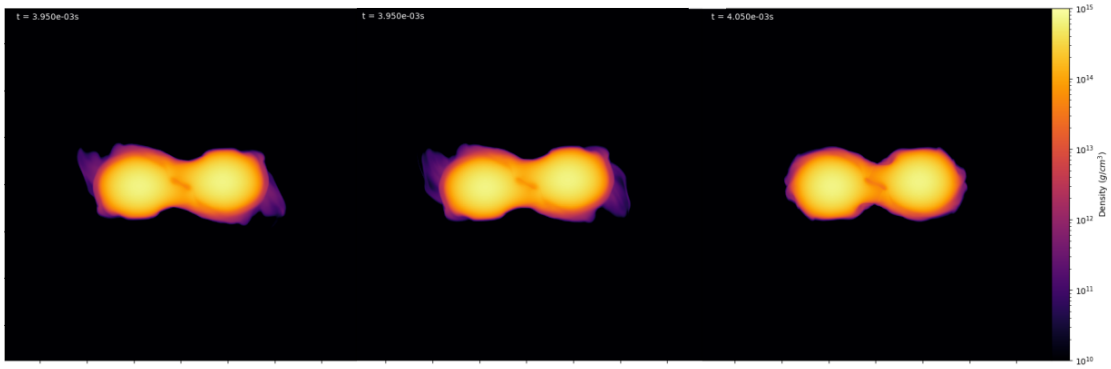

Figure 4.22: Density in the $x - y$ plane models IntT (left), IntA (middle) and Quad7(right) at 3.95ms for the magnetized models and 4.05ms for the unmagnetized case.

The growth of the magnetic field strength due to shear instabilities and tidal distortion of the stars during the merger is qualitatively similar to that observed in other simulations [Price and Rosswog, 2006] [Kiuchi et al., 2015]. For model IntA, the total magnetic energy increases by a factor of 8 during the disruption and initial merger of the stars, reaching its initial peak at around 6.5 ms, although as the merger precedes we would expect to see further amplification, particularly once a differentially rotating remnant has formed as this is when the magnetorotational instability (MRI) will activate. Starting with similar field strengths, [Ruiz et al., 2019] see an increase in the total magnetic energy by around a factor of 15 by the time the dense cores in the remnant merge, which appears within the bounds of what we may expect also from our simulations.

The amplification in our simulations so far comes primarily from turbulent motions along the contact interface between the two stars, but also from tidal elongation of the leading edge of the star. The IntT model shows slower growth than the IntA model, as a result of the tidal disruption happening parallel to the field lines in one of the stars. Since the magnetic tension force acts to resist curvature of field lines, this stretching parallel to field lines produces less field strength growth than stretching which acts perpendicular to the field lines and so acts to bend them. At this intermediate stage in the merger, the magnetic energy both models is evenly split between poloidal and toroidal components, however longer simulations are required to see whether this persists in the final remnant. Furthermore, Figure 4.23 shows the magnetic energy spectrum at several times throughout model IntT's evolution. We can see that initially the majority of the energy is at wavenumbers of around 4-5 $\times 10^{-7}$. This corresponds to scales of around 10-20km, which is approximately the size of the neutron stars and therefore the scale of the initial field as we would expect. At later times much more of the energy is recorded at smaller physical scales, and the dominance of the initial peak is greatly diminished. This increase in small scale magnetic energy is characteristic

of turbulent amplification of the field, as the vorticies twist up the field lines, increasing the strength of the field through the magnetic tension but also giving the field more small scale structure. Furthermore, the small scale turbulence at 5ms follows quite closely a Kolmogorov $-5/3$ power law [Kolmogorov, 1941], a theoretical model that describes the cascade of energy from large scales to small ones in turbulent flows. Although formulated for a highly idealized system, this power law is often observed in experimental studies of turbulence, and here it is a further indication of the turbulent mixing that occurs along the contact interface between the two stars at around 5ms. The apparent increase at large wavenumber at all times is an artifact, which results from our domain being non-periodic.

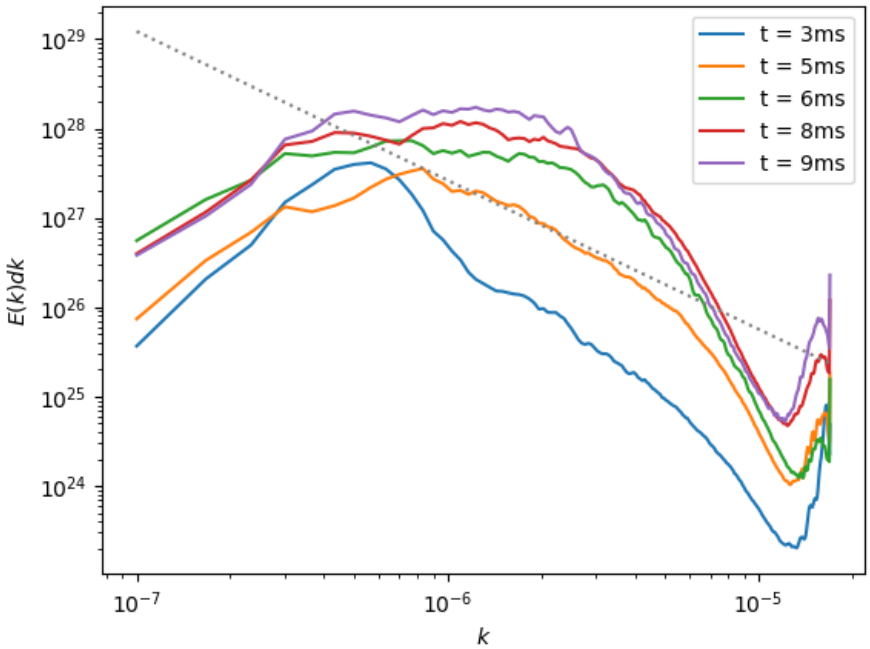

Figure 4.23: The magnetic energy spectrum for model IntT at $t = 3$, 5, 6, 8 and 9ms. The black dotted line represents a Kolmogorov spectrum, which is a $-5/3$ power law.

To summarize, our simulations show that the inclusion of a very large magnetic field causes the stars to merge earlier than in a non-magnetised case. During the initial tidal disruption and early merging stages we see close to a factor 10 increase in the total magnetic energy in the simulation for the case where both fields are aligned with the orbital angular momentum (IntA). The growth of the field in the case where one star's field is rotated so that it's axis lies in the orbital plane (IntT) is slower but grows eventually to larger strengths, as a result of the tidal amplification of the field being less effective when the field is parallel to the direction of stretching. To get a full picture of how the field configuration affects the merger however, we need to continue these simulations to examine the later stages of the merger.

Chapter 5

Black Holes in FLASH

The second part of this thesis work has been to create a model black hole (BH) in the FLASH code framework. This was achieved in large part by building upon the model accretion disc simulations [Hawley, 2000] that are included with the FLASH distribution.

5.1 Accretion discs

Accretion discs are found throughout astrophysics, in a range of environments and at various scales from protoplanetary discs to active galactic nuclei (AGN). They are also expected to be a part of NS-NS merger remnants, with the properties of such a disk determined by the properties of the original merging system. As such, they have been studied extensively, and we will here present a few key results of interest to this work.

Early studies of accretion physics were focussed on one main problem: how does accretion from the disc onto the central object actually occur? In order for matter to fall inward and be accreted it must lose angular momentum. However as the total angular momentum of the disc must be conserved, this means there must be corresponding increase in angular momentum in the outer regions of the disc. The stresses due to the gas' viscosity alone are far insufficient to provide the observed accretion rates, and so identifying a process which would produce this outward angular momentum transport took several decades to achieve. Shakura and Sunyaev [Shakura and Sunyaev, 1973] made major progress on the problem with the realization that turbulence within the disc could enhance the viscosity to the required levels. They found that for subsonic turbulence, assuming the disc's scale height z_0 to be the maximum eddy scale allows the viscosity from turbulence in the disc to be given as $\nu = \alpha c_s z_0$, where c_s is the sound speed and α is a free parameter. Equations of the disc structure in terms of the parameter α can then be constructed. The model's main advantage is that it can be easily compared with observations, allowing constraints to be placed on the value of α in different astrophysical settings [Hartmann et al., 1998] [King et al., 2007], however it provides no explanation for the source of the turbulence in the disc.

It was initially hoped that due to the high Reynold's number of the flows, the differential rotation of the disc would generate turbulence through shearing [Shakura and Sunyaev, 1973]. This is the case in planar shear flows, however in a rotating flow such as an accretion disc the Coriolis force acts as to stabilize disc, preventing the onset of turbulence. However, it was later found that the presence of even a weak magnetic field could generate turbulence in the disc very rapidly through the magnetorotational instability (MRI) [Balbus and Hawley, 1991].

5.1.1 The Magnetorotational Instability (MRI)

The MRI was first discovered in 1960 [Chandrasekhar, 1960], but its relevance to accretion discs was only realised in the early 1990s [Balbus and Hawley, 1991]. The instability can be understood in terms of a toy model, where we think of two fluid parcels at different radii in a magnetized disc as two masses connected by a spring. Due to the differential rotation of the disc, the inner mass is rotating faster and so moves ahead of the outer one. However the spring pulls back on the inner mass, slowing it down whilst speeding up the outer element. The inner mass then moves further inward, the cycle repeats and drives the process to a runaway instability.

The spring-like behaviour of the magnetic field was demonstrated by [Balbus and Hawley, 1991] in a simple case where we assume the disc is threaded by a constant vertical magnetic field. If we displace a fluid element by ξ, which has spatial dependence of e^{ikz} and lies within the plane of the disc, then the corresponding change in the magnetic field is found from the induction equation to be $\delta\mathbf{B} = ikB\xi$. The force exerted by the magnetic field in ideal MHD is given as $\mathbf{F} = (\nabla \times \mathbf{B}) \times \mathbf{B}$. By applying a vector calculus identity we can write this as

$$\mathbf{F_B} = -\nabla\left(\frac{B^2}{2}\right) + (\mathbf{B}\cdot\nabla)\mathbf{B}. \tag{5.1}$$

The first term is a magnetic pressure, which in the specific case we consider here only acts in the vertical direction and therefore does not contribute to forces in the plane of the disc. The second term however describes tension along magnetic field lines and does act in the disc plane. Substituting in our magnetic field $(\mathbf{B} + \delta\mathbf{B})$ the leading order term in δB is

$$\mathbf{F_B} = B\frac{\partial}{\partial z}\left(ikB\xi\right) = -k^2B^2\xi \tag{5.2}$$

This tension force therefore scales linearly with the displacement, acting as a spring force would. Continuing this analysis it can be shown [Balbus and Hawley, 1991] that in order for a magnetized disc to remain stable, its angular velocity, Ω, must satisfy

$$\frac{d\Omega^2}{dR} \geq 0. \tag{5.3}$$

This condition is not met in the vast majority of astrophysical systems (for example, a Keplerian rotation profile has $\Omega \propto R^{-3/2}$) and so the MRI is likely

to be active in them. Examinations of other initial field geometries beyond a simple vertical field have revealed complex fully three dimensional instabilities in the disc. For example, a toroidal field component is also unstable to the MRI, producing low wavenumber, non-axisymmetric spiral structures in a disc [Acheson and Acheson, 1978]. This means that simulations of discs in with and without axisymmetry produce quite different results [Hawley, 2000]. The destabilisation of the disc by the magnetic field in turn affects the magnetic field itself, as in the weak field regime of ideal MHD the field is 'frozen' into the fluid. This stretches an initially vertical field in both the azimuthal and radial directions, and in doing the tension force converts some of the fluid's kinetic energy into magnetic energy.

5.2 The Accretion torus simulation

The simulations here consist of an accretion torus around a central black hole. The fluid is not subject to self gravity as it is assumed that the gravitational potential of the black hole dominates. It is modelled with an ideal fluid equation of state, with $\Gamma = 5/3$. The initial condition for the simulation is an axisymmetric equilibrium torus with constant specific angular momentum, the construction of which is detailed in Section 5.2.1. Such tori are widely used in numerical studies of accretion discs because they are initially stable and have finite extent, meaning they can be contained completely inside the simulation domain which reduces the importance of boundary effects. Throughout this section we use units where $GM = 1$, used to simplify the simulation units since the model i

5.2.1 Disc initial conditions

To derive an equation for the initial density profile of our torus, we follow the methods set out in [Papaloizou and Pringle, 1984]. Here we use cylindrical coordinates (although the simulation is performed with Cartesian coordinates), with the axis of the disc's rotation aligned with the z-axis and the black hole located at the origin. A fluid flow in equilibrium is described by the steady state Euler momentum equation,

$$(\rho \mathbf{v} \cdot \nabla)\mathbf{v} = -\nabla P - \rho \nabla \psi_g \tag{5.4}$$

where $\mathbf{v} = v_\phi \hat{\phi}$ is the fluid velocity, P is pressure, ρ the density and ψ_g is the gravitational potential of the central black hole. We assume that the only non-zero velocity component is in the azimuthal direction and that the disc is axisymmetric, therefore the first term in 5.4 can be rewritten as $-\rho\Omega^2 r\hat{r}$ where $\Omega = v_\phi/r$ is the angular velocity of the fluid. For a constant angular momentum torus with an equation of state of the form $P = P(\rho)$, such as a polytrope, it can be shown that $\partial\Omega/\partial z = 0$ (Chapter 4 of [Tassoul, 1978]). Therefore we can rewrite this once more in terms of a rotational potential $\partial\psi_{rot}/\partial r = -\Omega^2 r$ giving

$$\nabla\psi_{rot} + \frac{\nabla P}{\rho} + \nabla\psi_g = 0 \tag{5.5}$$

Finally, using the polytropic relation $P = \kappa\rho^{\Gamma}$ and applying the chain rule, we can rewrite the equation as

$$\nabla\left(\psi_{rot} + \psi_g + \frac{\Gamma}{\Gamma - 1}\frac{P}{\rho}\right) = 0. \tag{5.6}$$

Integrating this gives

$$\psi_{rot} + \psi_g + \frac{\Gamma}{\Gamma - 1}\frac{P}{\rho} = C, \tag{5.7}$$

where C is an integration constant that is defined by the boundary of the torus, at which $P = \rho = 0$.

For the specific angular momentum of the torus to remain constant with radius, the angular velocity will be given by $\Omega = l/r^2$, where l is the specific angular momentum. Substituting into the definition for the rotational potential and integrating we find

$$\psi_{rot} = \frac{l^2}{2r^2} \tag{5.8}$$

where the integration constant is set to zero by assuming that the rotational potential is zero in the limit of infinite radius. Finally, we use the Pseudo-Newtonian Paczyński-Wiita potential for the black hole [Paczyński and Wiita, 1980] (described further in Section 5.2.2), which gives

$$\frac{P}{\rho} = \frac{\Gamma - 1}{\Gamma}\left(\frac{GM}{R - R_g} - \frac{l^2}{2r^2} + C\right) \tag{5.9}$$

where M is the mass of the black hole, R is the spherical coordinate radius and the parameter $R_g = 2GM/c^2$ is given by the Schwarzschild radius of the black hole.

Equation 5.9 and the polytropic relation give us a description of the initial density distribution for a general torus with constant angular momentum, and the values of the constants κ, l and C can then be used to describe the specific case we wish to simulate. As these parameters don't describe the physical extent of the torus clearly, we instead specify our system in terms of the torus' inner radius r_{in}, the radius of maximum density r_{max} and the maximum density itself, ρ_{max}. These new parameters relate to the original ones as follows:

- The specific angular momentum is chosen to be given by the Keplerian angular momentum at the maximum density radius, which is

$$l_k = \frac{(GMR_{max})^{\frac{1}{2}}R_{max}}{R_{max} - R_g} \tag{5.10}$$

  where $R_{max} = \sqrt{r_{max}^2 + z^2}$ is the spherical coordinate radius of the density maximum.

- Choosing the inner radius of the torus to be the boundary at which $P = \rho = 0$, the constant C is then

$$C = \frac{GM}{R_{in} - R_g} + \frac{l^2}{2r_{in}^2}. \tag{5.11}$$

- Using the two above conditions and the polytropic relation in Equation 5.9, we use the value of the maximum density to calculate the polytropic constant κ as

$$\kappa = \frac{\Gamma - 1}{\Gamma} \left(\frac{GM}{R_{max} - R_g} - \frac{l_k^2}{2r_{max}^2} + C \right) \frac{1}{\rho_{max}^{\Gamma - 1}} \tag{5.12}$$

With the density distribution in the torus fully defined, the pressure and temperature are calculated through the equation of state and the azimuthal velocities are given by l_k/r.

Finally, in magnetized cases, the torus is endowed with a poloidal magnetic field, whose vector potential is given by

$$A_\phi = B_0 \max(\rho - \rho_c, 0) \tag{5.13}$$

where B_0 describes the maximum field strength and ρ_c is a cut-off density below which the magnetic field is zero. The parameter B_0 is set by specifying a target plasma beta β_0 to by reached at the maximum density, which gives it a value of $B_0 = \sqrt{2\kappa\rho_{max}^{\Gamma}/\beta_0}/\rho_{max}$.

Outside the torus, the simulation domain is filled with a constant low density atmosphere that is around 2-3 orders of magnitude lower than the density within the torus. This is chosen so that it does not impact the dynamics of the torus significantly, provided the boundaries of the simulation domain are located suitably.

5.2.2 Black hole model

In these simulations, the black hole is situated at the coordinate origin. As these are Newtonian simulations, we must directly implement several features of the black hole manually, in order to obtain an approximation to the behaviour of the hole in relativity. Firstly, the gravitational potential of the BH is modelled using a Paczyński-Wiita potential [Paczyński and Wiita, 1980], which is given by

$$\psi_G = \frac{-GM_{BH}}{R - R_g} \tag{5.14}$$

where M_{BH} is the mass of the black hole and $R_g = 2GM/c^2$ is the 'gravitational' radius, which corresponds to the Schwarzchild radius of the black hole. This potential simulates some of the features of a Schwarzschild black hole that are important in the study of accretion discs, including replicating the correct posi-

tions of both the innermost stable circular orbit (ISCO), $R_{ISCO} = 3R_g$, and the marginally bound orbit, $R_{mb} = 2R_g$. It has also been shown to produce similar accretion efficiencies to a general relativistic Schwarzschild solution [Paczyński and Wiita, 1980].

On the simulation domain, the black hole is a region of radius $R_{BH} = 1.1R_g$ within which the hydrodynamic variables are not evolved. The radius is set to be slightly larger than R_g to avoid the singularity in the gravitational potential at R_g, but close enough that the large potential values drive a supersonic inflow into the hole in area around it. The black hole region is filled with a low density (5 orders of magnitude lower than the minimum density of the evolved fluid) cold fluid.

The biggest challenge in this work has been setting conditions on the boundary of the black hole which both ensure supersonic inflow into the hole and maintain numerical stability. Whilst enforcing inflow-only conditions across the cell boundaries is relatively simple it is more difficult to then ensure that zones on the edges of the hole do not develop very high velocities and very low densities, which lead to the timestep becoming unfeasibly small. Our solution uses a general boundary condition that works well for the majority of zones on the boundary and then applies small stability fixes to problem zones as they arise.

At the boundaries of the black hole region, we apply a diode condition, where the velocity in boundary cells is copied from the velocity in the final zone if it is directed into the black hole (and so out of the simulation domain) and is set to zero if the velocities are directed out of the black hole. This ensures that matter cannot flow out of the hole. In practice, the steep acceleration gradient produced by the gravitational potential ensures supersonic and super-Alfvénic inflow in the zones immediately surrounding the hole, however without the additional inflow condition we found that occasionally a single zone may develop a spurious, large outward velocity. This is a result of the surface of the the black hole region being stepped due to the Cartesian nature of the grid. The problem develops in zones where there is an interface with the black hole in a direction that was close to perpendicular to the gravitational acceleration from the black hole and for whom the inward direction across this interface was the opposite sense to the rotation of the torus (as shown in Figure 5.1). This meant that the gravitational acceleration across this particular face was not strong enough to enforce inflow and the rotation of the material was driving a flow past this face that pointed outward from the face, leading to this cell gaining an outward velocity. Left unchecked this could cause the simulation to fail therefore to maintain the stability of the simulation we set these spurious velocities to zero. As this occurred in only a very small number of zones (around 10 out of more than 5000 BH face zones) it does not affect the simulation at large.

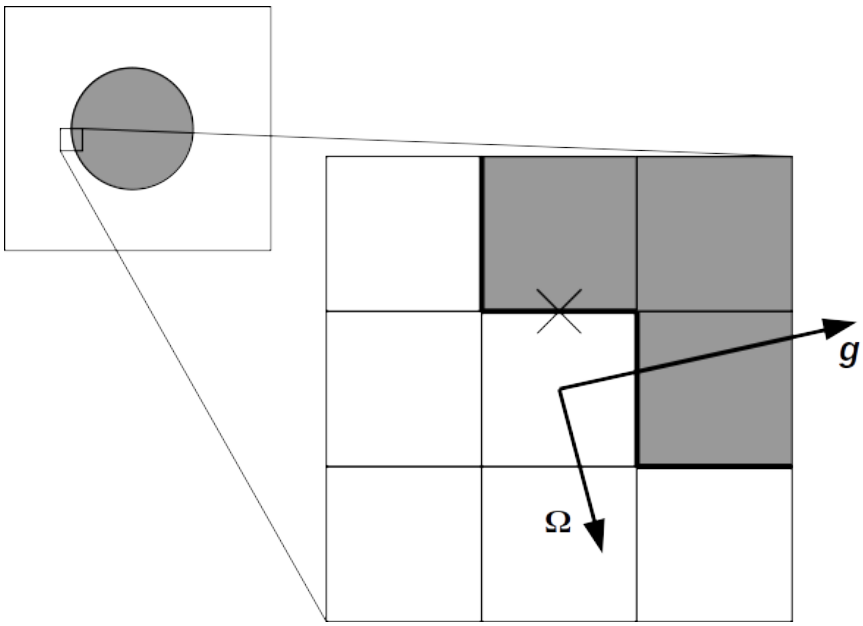

Figure 5.1: Schematic showing the location of zones in which spurious velocities develop. The shaded region is the black hole interior, we show here a slice in the x-y plane. The gravitation acceleration, $\mathbf{g}$, and torus orbital velocity, $\mathbf{\Omega}$, are oriented such that the velocity in the problem zone is likely to develop a spurious outward component in the y direction, across the face marked with an 'x'.

In magnetized simulations we also force the magnetic field to remain zero on the surface of and within the black hole. This ensures that no spurious magnetic fields develop near the hole as a result of the low densities and temperatures of the interior of the hole. We also zero fields in the small number of zones which develop spurious velocities, again to maintain the simulation's stability.

5.3 Results

Table 5.1 summarizes the parameters for the different torus simulations performed. We first performed a selection of non-magnetized cases at different grid sizes and resolution to assess the effects of changing these. With an optimal grid configuration chosen, we then performed a non-magnetized and a magnetized simulation, in order to compare the effects of adding the magnetic fields. The torus parameters are chosen for comparison with similar models used by [Hawley, 2000].

Model	Grid width	No. refinement levels	Δx_{min}	r_{in}	r_{max}	β_0
HD-SG	20	6	0.08	3.0	4.7	0.0
HD-LG	40	7	0.08	3.0	4.7	0.0
HD	30	7	0.06	3.0	4.7	0.0
MHD	30	7	0.06	3.0	4.7	350

Table 5.1: A summary of the parameters used in the accretion torus simulations. Δx_{min} is the side length of the smallest grid cell in the simulation, r_{in} is the inner radius of the torus, r_{max} is the radius at which the density is at its maximum, ρ_{max} is the maximum density and β_0 is the plasma-β ($\beta = P_{gas}/(B^2/2)$) at the density maximum. All distances are in units of R_g.

5.3.1 HD models

Figure 5.2 shows the torus density at several times throughout the HD simulation. Note that in the time units of the simulation, the orbital periods at the inner and outer edge of the torus are 21 and 112 respectively and the orbital period at the density maximum $P_{max} = 50$.

At the start of the simulation, the relaxation of the initial conditions causes matter from the inner edge of the torus falls inward towards the BH in a narrow accretion stream. The torus also oscillates slightly, which can be seen in the difference in shape between the $t = 200$ and $t = 300$ panels in Figure 5.2. The $x = 0$ slices on the left-hand sides show that the torus becomes more elongated radially between t =0 and 200, and then contracts to return to a shape more like the initial condition at $t = 300$.

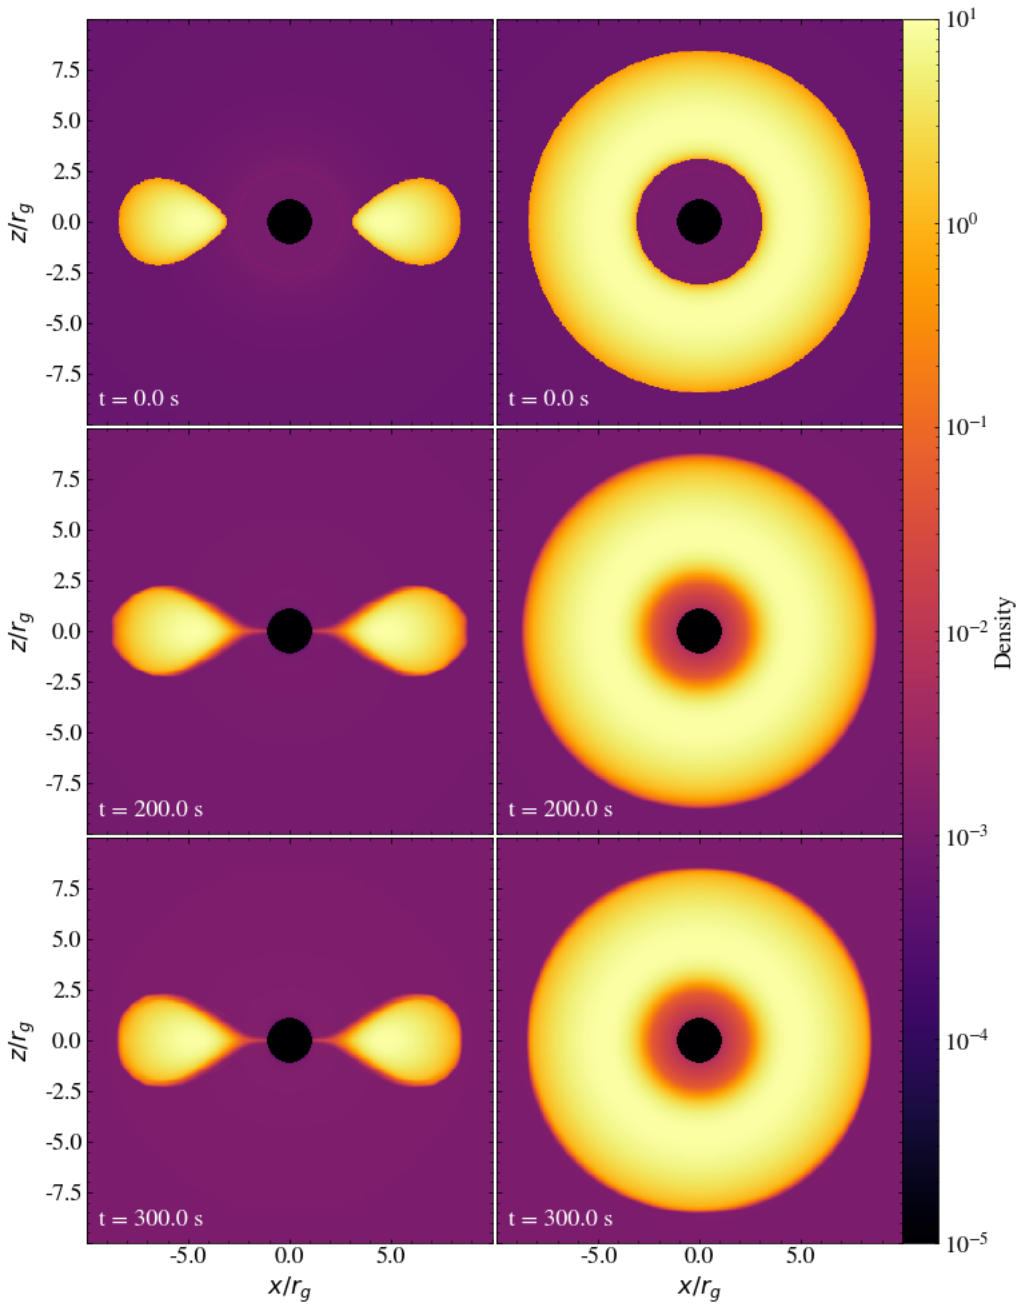

Figure 5.2: Slices of density taken along the $x = 0$ and $z = 0$ planes in the left and right panels respectively, at times 0, 200, 300. The orbital period at the inner, density maximum and outer edges of the torus are 21, 50 and 112 respectively.

The vertically and azimuthally averaged density profiles in Figure 5.3 show that after the relaxation of the torus, which raises the density slightly in the central regions, the density distribution remains constant throughout the simulation. Further to this, Figure 5.4 shows the radial angular momentum distribution of the torus at various times throughout the simulation.

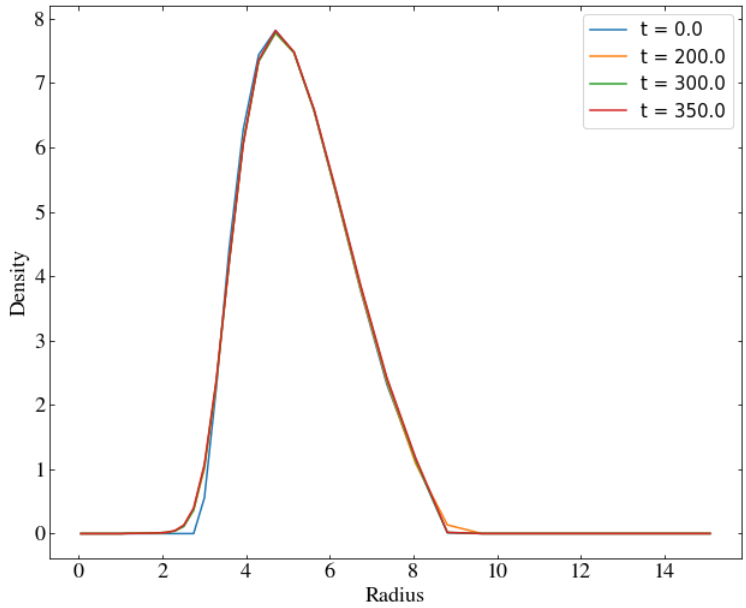

Figure 5.3: Radial density profiles for model HD at times 0, 200, 300, 350. The initial onset of accretion raises the density slightly at the inner boundary of the torus, but once this is established the profile remains the same for the rest of the simulation.

We see again that the initial relaxation raises the angular momentum slightly both within the inner edge of the torus and outside it. After this initial change however, the distributions of both density and angular momentum remain constant. This is as expected for a torus without magnetic fields, since there is no mechanism by which angular momentum can be transported.

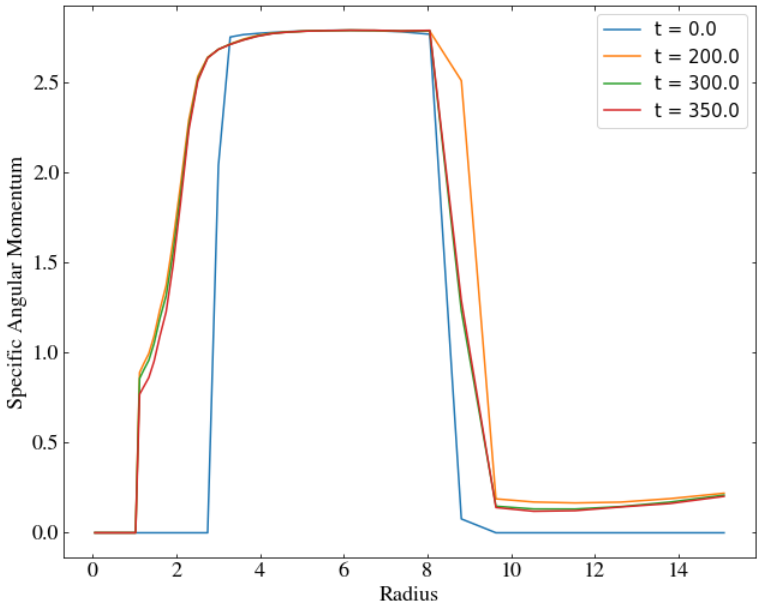

Figure 5.4: Radial specific angular momentum profiles for model HD at times 0, 200, 300, 350.

Testing the black hole model

The accretion rate onto the black hole for model HD is shown in Figure 5.5. There is an initial spike in accretion as the relaxation of the torus occurs, since the inner edge of the torus is at the r_{ISCO} and so the matter that moves inward during the relaxation is accreted. Following this the accretion rate returns to a lower value with a much smaller, steady increase across the rest of the simulation. This steady increase is as a result of the infall of the non-rotating atmosphere gas, which grows slowly over the course of the simulation under the gravitational acceleration from the black hole.

In Figure 5.6 we plot the total conserved mass, which is the total gas mass on the grid at each timestep combined with the cumulative mass accreted by the black hole and the total mass gained/lost through the grid's outer boundary. This shows that mass is conserved to within 10^{-6} of the total initial gas mass on the grid, implying that our measurements of the black hole accretion rate are sound.

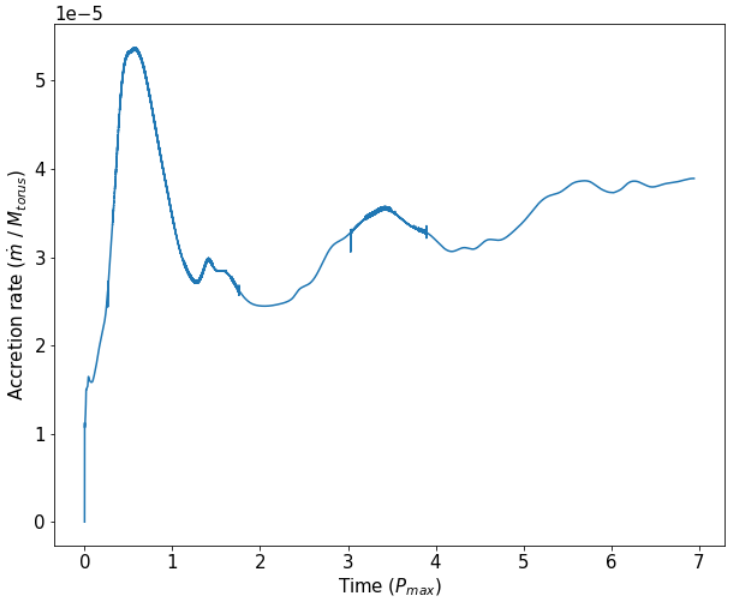

Figure 5.5: Mass accretion rate as a fraction of the total torus mass for model HD.

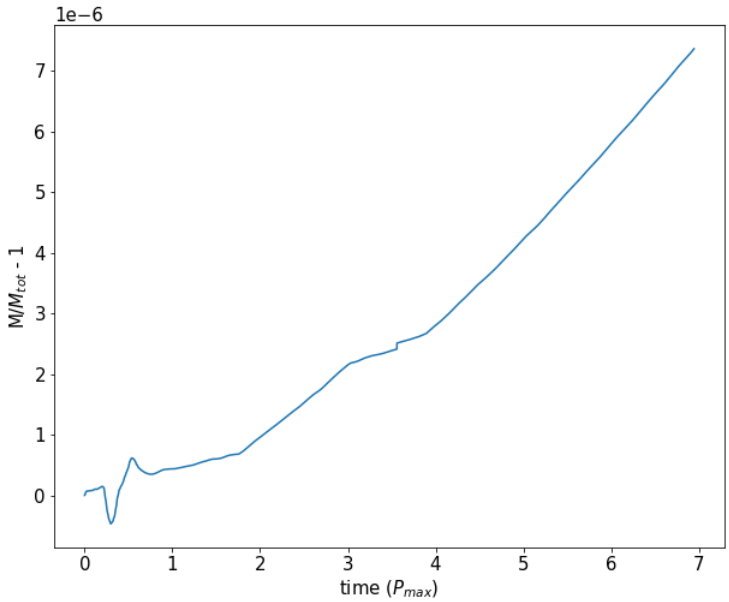

Figure 5.6: Total gas mass plus the total mass accreted by the black hole and the total mass gained/lost at the outer boundary.

Effect of boundary size

In model HD-SG, we used a smaller simulation volume, with a domain side length of $20r_g$ rather than the $30r_g$ used in model HD. The effects of using a smaller domain are quite large in this case. Figure 5.7 shows that the torus surface becomes unstable after several orbits in this model. The inner disc becomes warped bending up and down alternately out of the $z = 0$ plane, creating a many armed spiral appearance in the slices along the equatorial plane. This instability emerges after around 4 orbits and persists for several more before the inner accretion stream is completely disrupted. The velocity vectors in Figure 5.7 show the inflow of the low density background, driven by the black hole's gravitational potential, which builds over the course of the simulation to give a steady inflow along the disk's axis. Further from the axis, interactions with the disk prevent such a flow from forming, but eventually the deflection of the inflowing matter by the denser torus leads to outflows across the surface of the torus. This generates streams of matter that flow outward from the torus' surface.

The increased strength of the atmosphere inflow is a result of the boundary conditions. The gravitational acceleration from the black hole falls off as $1/r^2$, however at the boundary of the simulation domain we apply a zero gradient outflow. This assumes that the fluid variables outside the grid are the same as those in the final zone of the grid, therefore they simply grow with time like the velocity in the final zone of the grid does under the black hole's constant acceleration. As such, if the boundary is too small, it causes the atmospheric velocity to grow very large within the simulation time, and these large velocities disrupt the otherwise stable torus.

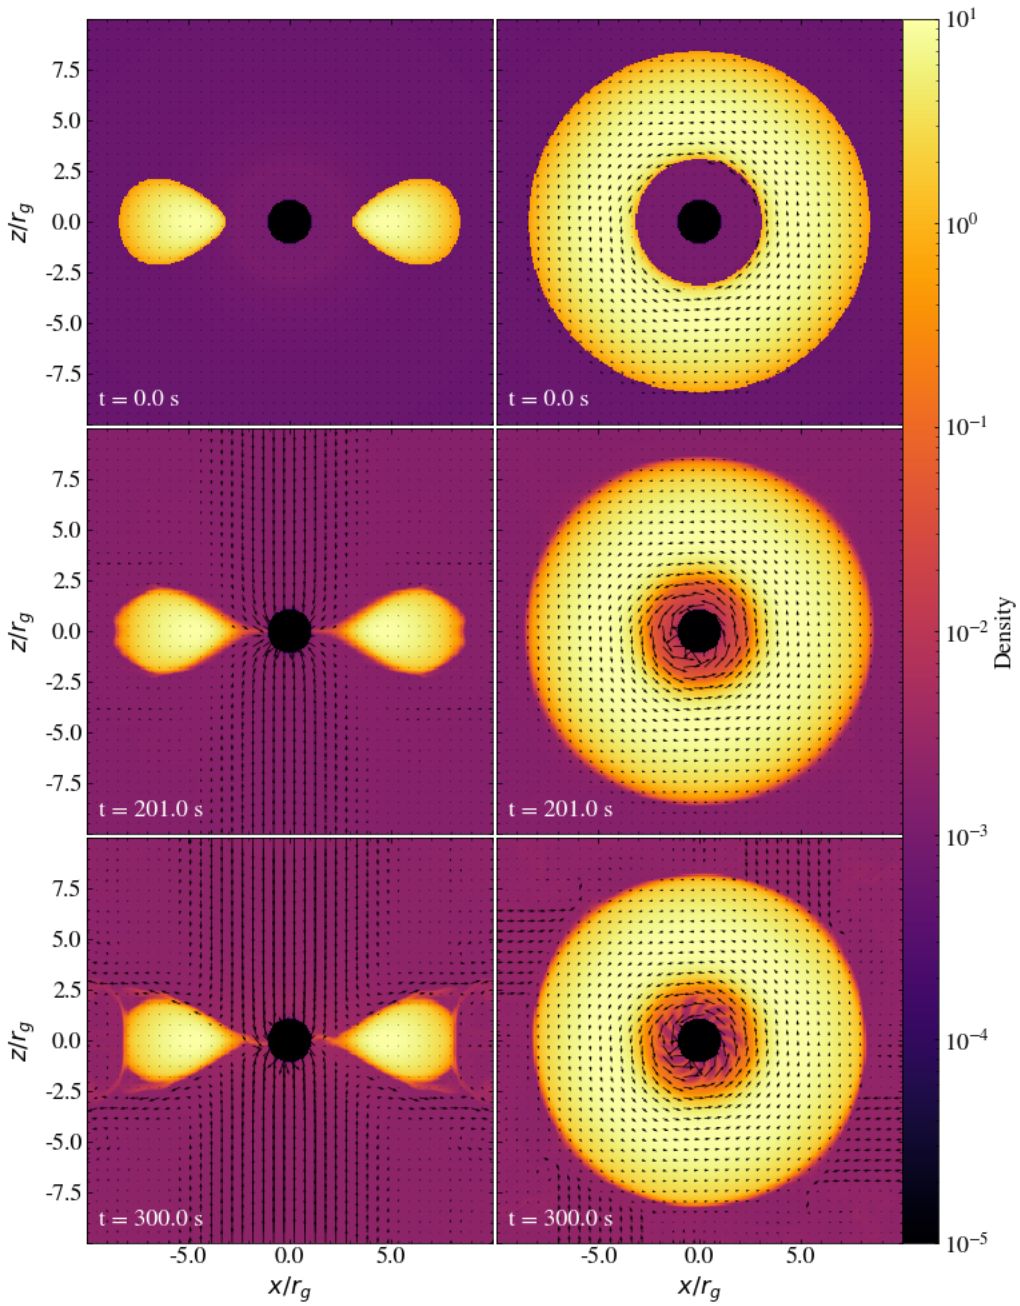

Figure 5.7: Slices of density taken along the $x = 0$ and $z = 0$ planes in the left and right panels respectively, at times 0, 200, 300. The orbital period at the inner, density maximum and outer edges of the torus are 21, 50 and 112 respectively. The velocity vectors in the plane are marked with arrows.

5.3.2 MHD model

To examine the performance of our MHD boundary conditions, we endowed the initial torus with a weak poloidal magnetic field of the form described in Equation 5.13. We again evolved the torus for around 7 orbits at the pressure maximum ($R_{max} = 4.7r_g$, $P_{max} = 50$), Figure 5.8 shows density slices at several points during the simulation. The slices at t =200 ($4P_{max}$) show turbulent 'rolls' across the surface of the torus, which begin to arise after only 1-2 orbits. By t =300, the torus has become fully turbulent, and the slice in the $x - y$ plane shows that it is also losing its axisymmetry.

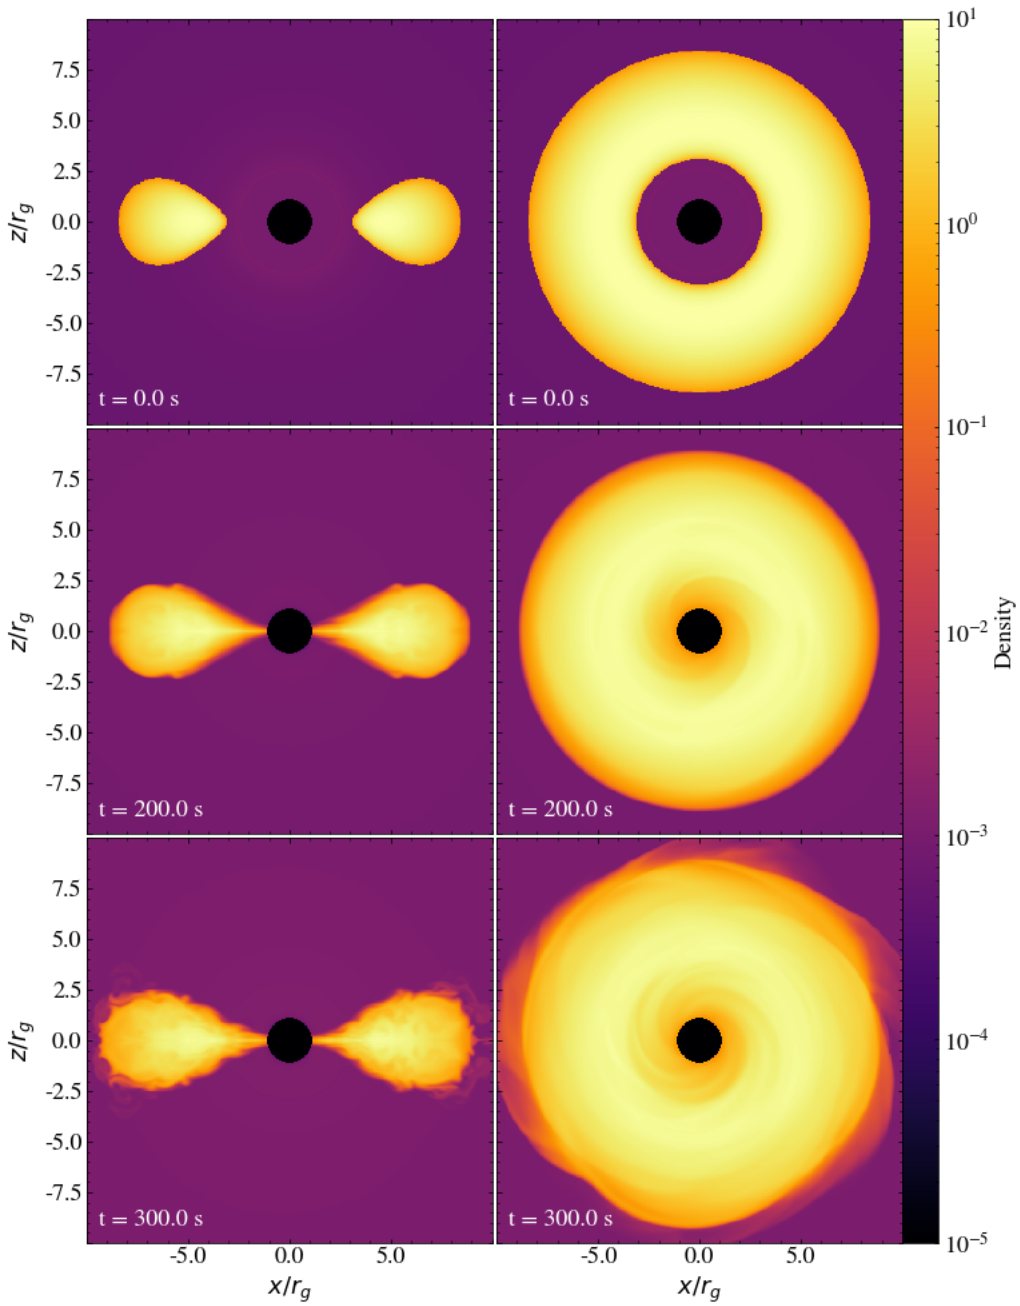

Figure 5.8: Slices of density taken along the $x = 0$ and $z = 0$ planes in the left and right panels respectively, at times 0, 200, 300. The orbital period at the inner, density maximum and outer edges of the torus are 21, 50 and 112 respectively. The velocity vectors in the plane are marked with arrows.

Figure 5.9 shows a low-wavenumber spiral wave structure visible in the inner disc. Such features appear and are disrupted periodically throughout the simulation starting from around orbit 3. The plasma-β is defined as $\beta = P_{gas}/(B^2/2)$ and is a measure of whether the flow is dominated by hydrodynamic (high β) or magnetic forces (low β). The right hand panels of Figure 5.9 show that the magnetic field is most dominant in the inner regions of the torus above and below the equatorial plane. The initial peak of the magnetic field strength was already located radially between the density maximum and the inner edge of the torus, however the initial β was still quite high. The shearing of magnetic field lines by differential rotation occurs throughout the disc, however in this region there is additional shearing in the radial direction due to accretion, which drives stronger field growth. Figure 5.9 also demonstrates how the disk turbulence has affected the vertical magnetic field structure within the torus, which is also shown by the field lines in Figure 5.10. We see that the initially poloidal field is dragged along by the rotation, gaining a growing toroidal component. The vertical structure is disrupted by the overturning fluid cells within the torus, creating distorted loops which are approximately, but not exactly, symmetric across the equatorial plane.

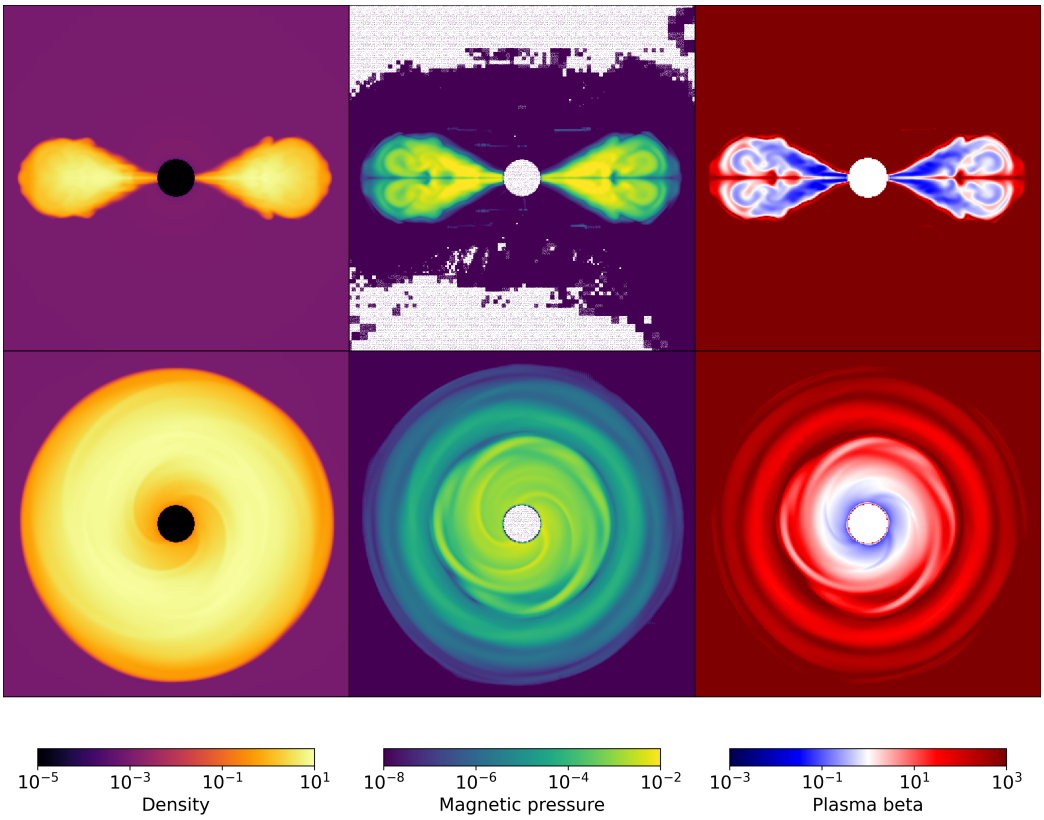

Figure 5.9: Slices of density, magnetic pressure and plasma beta taken along the $x = 0$ and $z = 0$ planes in the upper and lower panels respectively, at t =225, where $P_{max} = 50$.

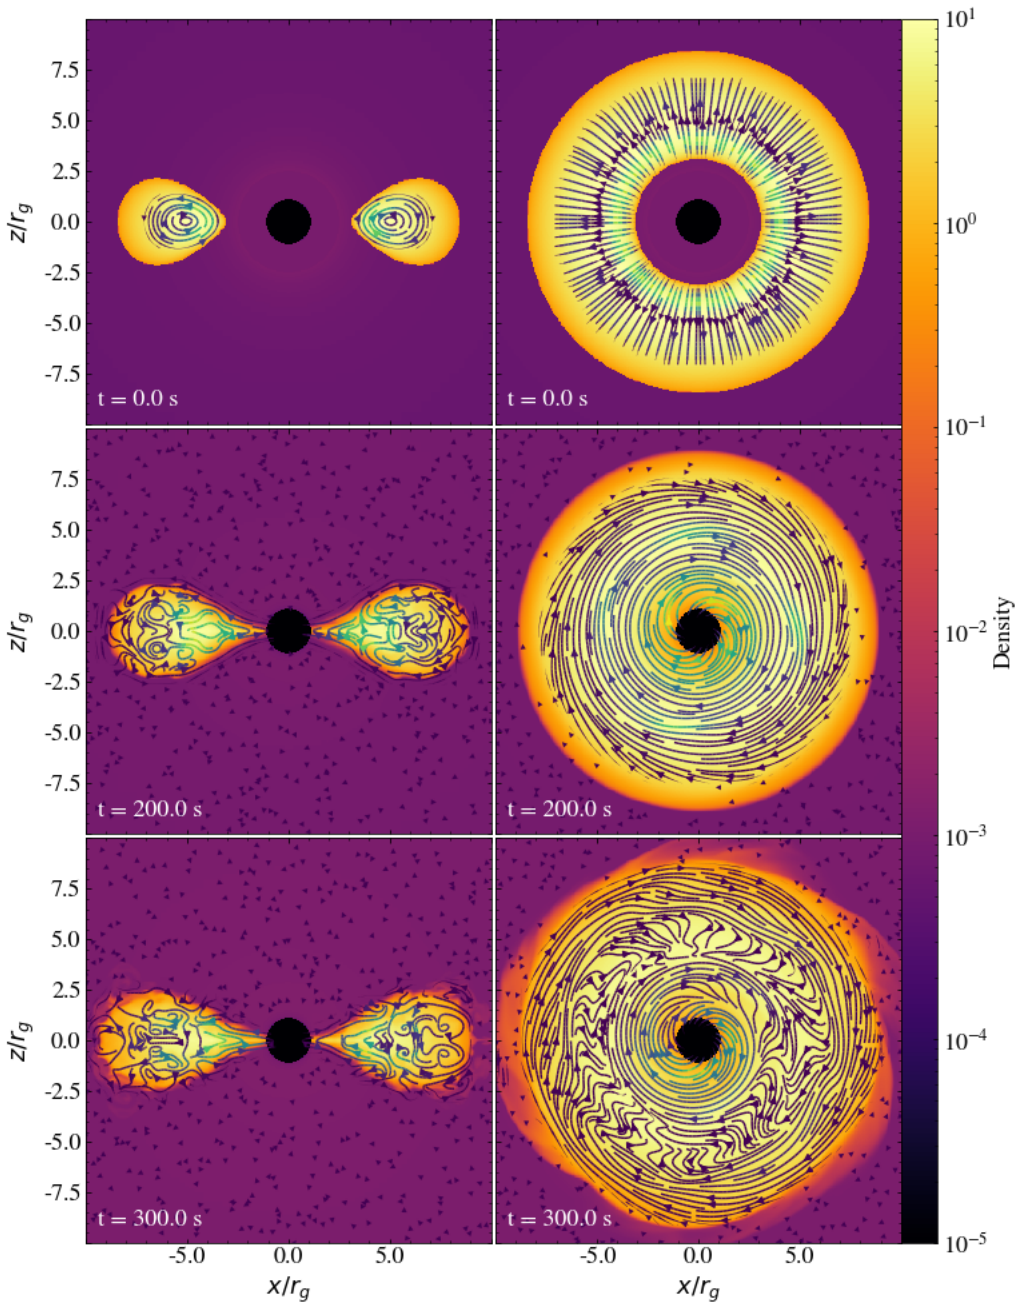

Figure 5.10: Slices of density taken along the x and z-axes in the left and right panels respectively, at times 0, 200, 300, where $P_{max} = 50$. The magnetic field lines in the plane are overplotted, coloured according to the magnetic field strength.

Figures 5.11 and 5.12 show the vertically and azimuthally averaged density and specific angular momentum as functions of radius. In contrast to the HD cases, they show significant spreading of the the torus' mass and greater angular momentum transport. The density in the inner regions of the torus is much higher, indicating that a stronger accretion flow is present. The angular momentum profile appears to be moving from the initial constant distribution to one that is increasing with radius, as the MRI induced turbulence transports momentum outward. Once the accretion stream in the inner region of the disc forms, the angular momentum profile in that region remains roughly constant and has a slope that is slightly shallower than the Keplerian profile. The profile in the outer regions of the disc meanwhile is still changing quite substantially, as the disc is spreading radially outward and the angular momentum increase with radius steepens. At later times, if this trend continues we may expect the disc to become thinner and the angular momentum distribution to steepen until it becomes close to Keplerian, as observed in [Hawley, 2000] however longer simulations need to be undertaken to probe this.

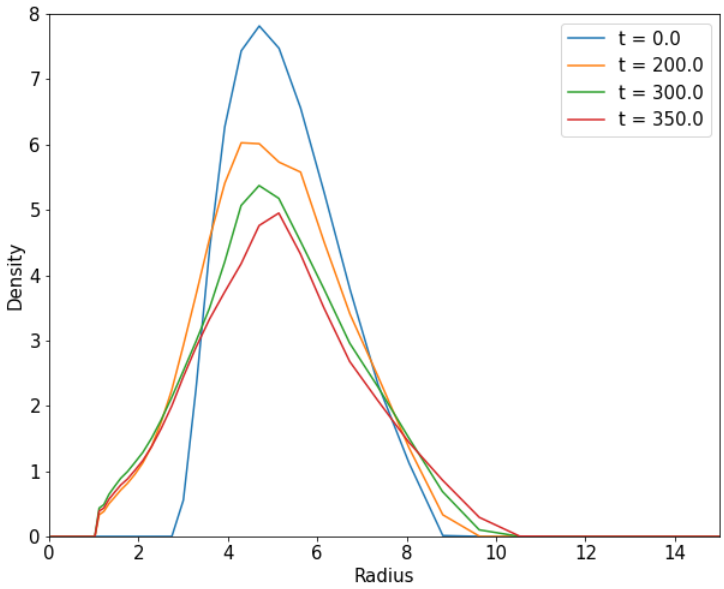

Figure 5.11: Radial density profiles for model MHD at times 0, 200, 300, 350.

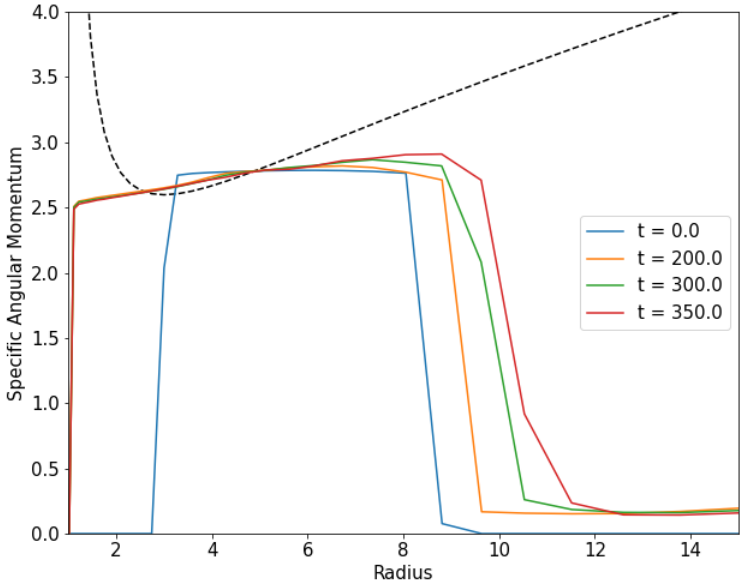

Figure 5.12: Radial specific angular momentum profiles for model MHD at times 0, 200, 300, 350. The dotted line marks a Keplerian angular momentum profile.

Figures 5.13 and 5.14 show the total poloidal and toroidal magnetic energies and the ratio of these components as functions of time. We can see that shearing due to the differential rotation in the torus causes the toroidal component to grow linearly during the first two orbits, surpassing the weak, initial poloidal field very quickly. The difference between the two components peaks at around 2.5 orbits in, after which the toroidal component falls slightly and the poloidal field begins to grow. This corresponds roughly with the first appearance of the turbulent 'rolls' visible in the vertical slices through the torus. These rolls show that circulation in the vertical direction is occurring within the torus, which is driving the growth of the poloidal field. The 3D rendering of the magnetic field lines after approximately 4 orbits in Figure 5.15 further shows the dominance of the toroidal field, with none of the original poloidal structure visible by this stage.

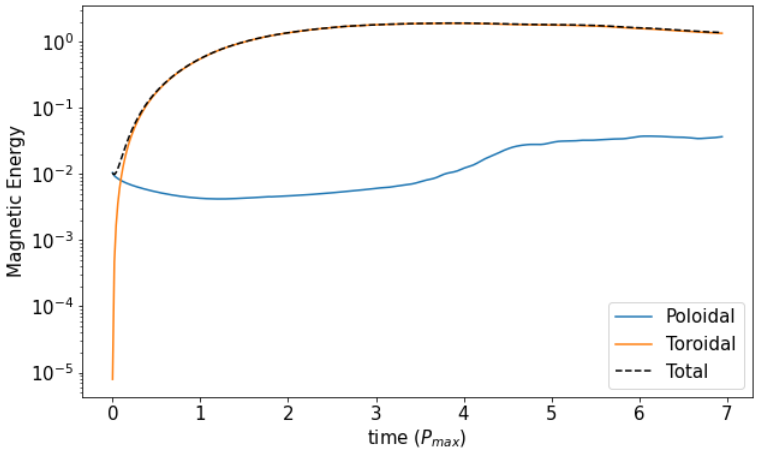

Figure 5.13: Toroidal and poloidal components of magnetic energy with time.

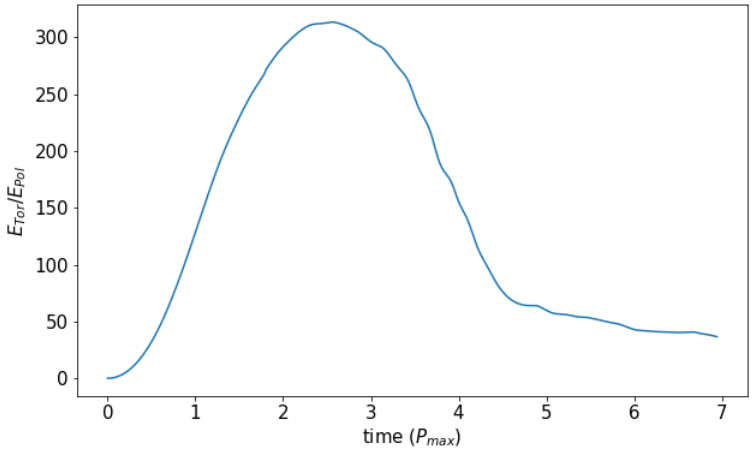

Figure 5.14: Ratio of toroidal and poloidal magnetic energies with time.

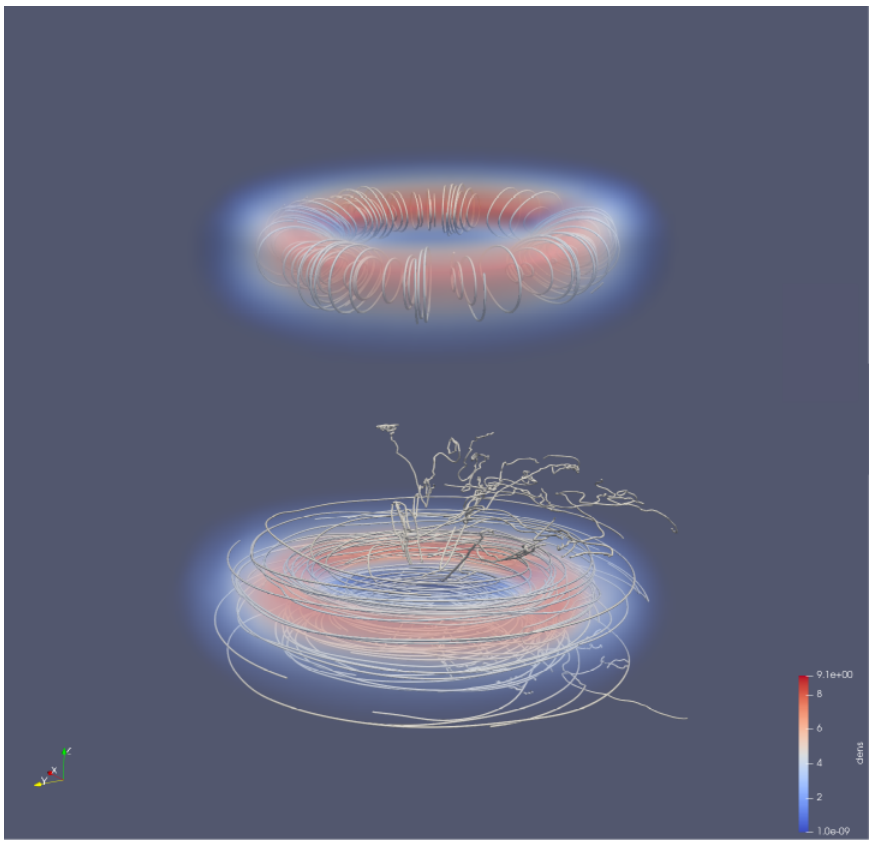

Figure 5.15: A 3D visualisation of the magnetic field lines at $t = 0$ and $t = 225$, overplotted onto a volume rendering of the density.

Testing the black hole

Figure 5.16 compares the mass accretion rate for the HD and MHD cases. We see that for the first orbit these are very similar, both showing the same initial peak in accretion. Following this, while the HD case falls to a low and relatively constant value for the remainder of the simulation, accretion in the MHD case grows rapidly over the following two orbits. The rate peaks at around 5 orbits, coinciding with the onset of major turbulence in the torus. The field lines in Figure 5.10 show that at the black hole boundary, the field maintains the same direction as in the inflowing gas around it. The $x - y$ plane plots show that the field does not end up anchored to the horizon, as in that case we would see the spiral in the field lines tightening with time. The plasma β slices in Figure 5.9 do show a peak in the regions immediately surrounding the hole, which is due to the zeroing of the magnetic field required in zones a step outside the hole. There is no evidence of any larger scale anomalous behaviour in the field near the hole.

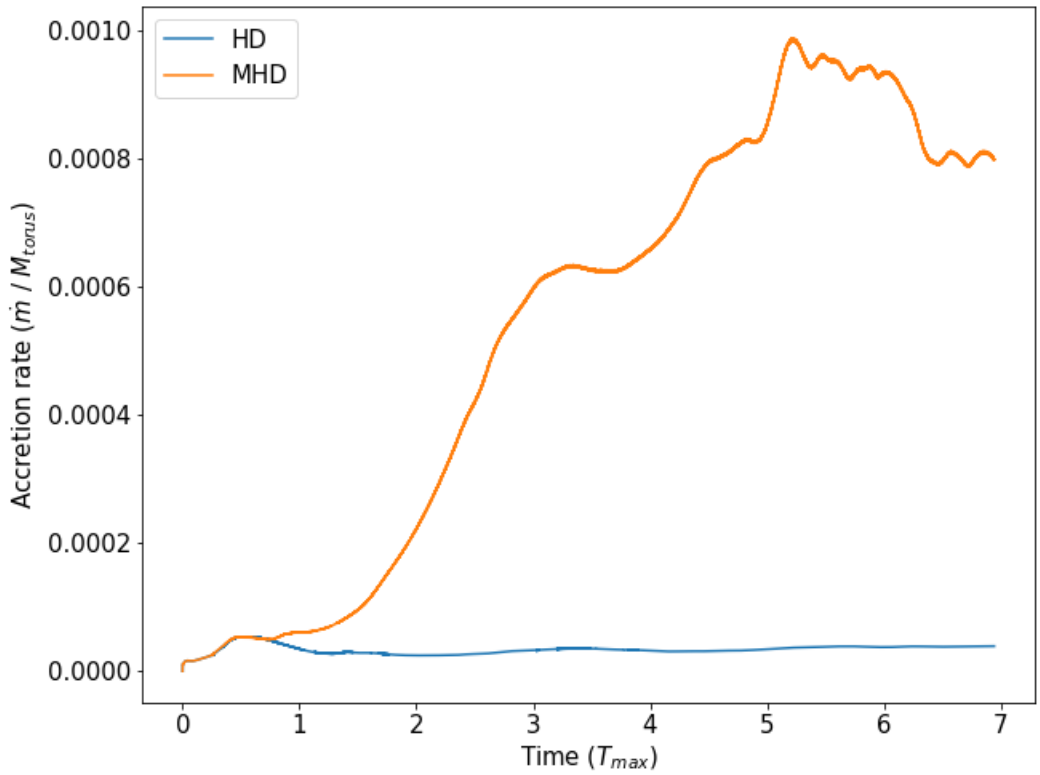

Figure 5.16: Mass accretion rate across the black hole boundary, for both the HD and MHD models.

5.4 Discussion

The primary function of these simulations was to set up and test a model black hole in FLASH. Previous Newtonian disc studies use cylindrical or spherical coordinates, where the black hole is represented via the inner radial boundary condition, which is not possible on a Cartesian grid. We developed this black hole model instead to be used in the future to simulate BH-NS mergers and the remnant of NS-NS mergers, for which a Cartesian grid is more appropriate due to the lack of spherical or axial symmetry.

The parameters of our magnetized disc closely resemble case GT1 from [Hawley, 2000], the only parameter difference being that we use a slightly weaker field (β =350 compared with $\beta = 100$). Those simulations are carried out in cylindrical coordinates, with an angular coordinate range of $0 \leq \phi \leq \pi/2$, and extend for about twice the time of our simulations. Qualitatively, the two simulations show much similarity, they both exhibit the emergence of non-axisymmetric spiral waves and eventual turbulent disruption of the disc. The evolution of the radial density and angular momentum distributions are also very similar. The onset of turbulence happens earlier in [Hawley, 2000] as the stronger magnetic field drives

more rapid growth of the instabilities. As the magnetic Maxwell stress is the dominant stress in the disc, a magnetic alpha value $\alpha_m = B_r B_\phi / 4\pi P_B$ can be considered to be a more relevant parameter to consider than the more general α of Shakura and Sunyaev. For comparison with [Hawley, 2000] we calculate a volume averaged value for α_m, which is shown in Figure 5.17 as a function of time.

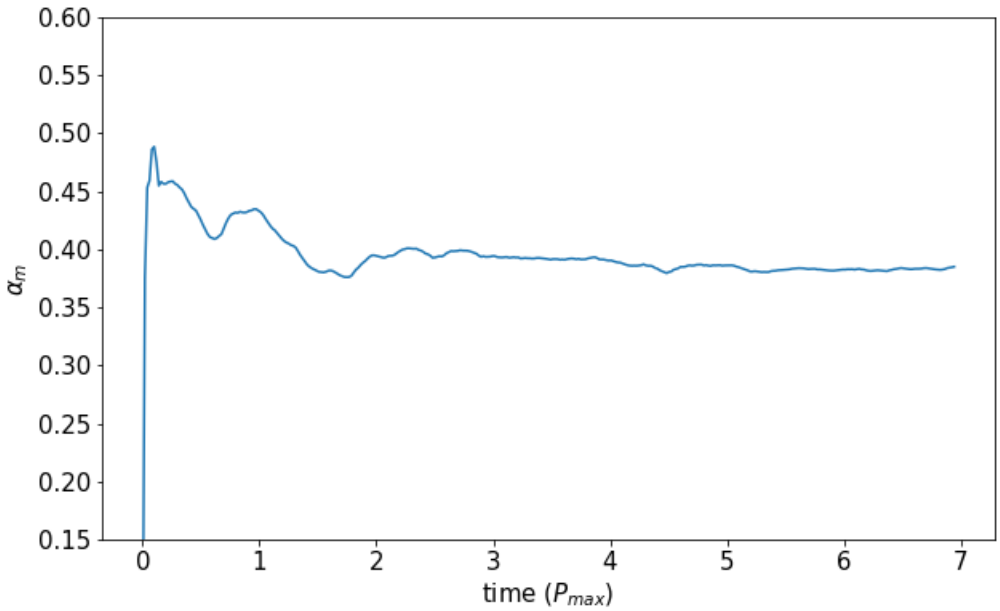

Figure 5.17: The volume averaged α_m as a function of time.

Comparing this with Figure 7 in [Hawley, 2000], we see that the two simulations converge to the same value α_m=0.4, however ours shows much lower amplitude variance in the early stages of the simulation. Their initial peak reaches around 0.7, whereas ours is only 0.5, and we do not see the drop to 0.2 following this peak.

Figure 5.16 shows that the accretion rates we record are comparable with those of [Hawley, 2000] (Figure 5), although they are slightly smaller, likely as a result of the weaker magnetic field. The time variability in our results is less pronounced and we do not see as prominent an initial spike, but the broad features, periods of growth in the accretion rate separated by short plateaus during which there is some variability in the early stages, are common to both simulations.

Examining the magnetic field structure near the hole, we find that field lines drawn into contact with the horizon do not exhibit any anomalous behaviour, giving us confidence in our approach here. The field is advected with the gas into the hole, and our boundary conditions terminate any normal fields on the horizon. This is broadly in agreement with what we may expect from the theory of magnetic field interactions with black holes described in the membrane paradigm [Thorne et al., 1986], which models the horizon as behaving as a conducting membrane. In such

a case, normal field lines are terminated on the horizon, and these fields are free to move across the surface, i.e. they do not remain tethered to point at which they make contact.

To summarize, these simulations serve a test of the black hole model we have introduced to the FLASH framework. They demonstrate that the magnetized torus simulations are stable and behave as expected with this new feature. We show our measured accretion rates to be consistent with other comparable work, and that the behaviour of the magnetic field around the horizon is broadly consistent with what we expect for a non-spinning black hole. Following this work, there are further features which we intend to include, such as including a black hole in a simulation with self-gravitating fluid, allowing the black hole mass and horizon radius to change as mass is accreted and finally allowing the black hole to move on the grid, to facilitate the simulation of BH-NS mergers.

Chapter 6

Conclusions and Outlook

In this thesis we present the creation of binary neutron star merger models using the FLASH code framework. We implement an approximate gravitational wave emission scheme, based on calculating the gravitational wave luminosity using the quadrupole moment of the mass distribution and extracting this energy by applying a friction-like force to the stars during their orbit. We find that for an unmagnetized merger, this scheme replicates the expected inspiral trajectory to give broad agreement with comparable works within the limits of our physical approximations. We then apply this model to study the effects of magnetic field configuration on the merger. We find that the amplification of the magnetic field is stronger in the case where the star's initial fields where one of the star's has its magnetic axis parallel to the orbital plane and the other is aligned with the orbital angular momentum. In both cases we observe strengthening of the magnetic field along the shear interface where the two stars first meet, which continues long after the initial contact as a region of turbulent mixing between the two stars.

Due to time constraints we were only able to study the early stages of the merger. As such we cannot yet comment on how far the field is amplified over the course of the whole merger, therefore continuing these simulations to the completion of the merger is the next stage of this work. Further to this, it is of particular interest to investigate how much of the high magnetic energy generated at small scales by turbulence can be transferred to the larger scale field. This is because in order for a merger remnant to successfully form an sGRB jet, there needs to be a very strong large scale field present in the remnant. Simulations of binary neutron star mergers which have successfully produced jets [Ruiz et al., 2016], have used very large initial neutron star magnetic fields. This is to due to current computational limits, which prevent us from resolving turbulence well enough to see amplification of magnetic fields to the large values required for jet production. As such, it remains unclear whether turbulence during the merger and in a HMNS remnant will produce the same strong, large scale fields seen in these simulations. Furthermore, other processes such as neutrino emission are expected to play important roles in jet production, for example by clearing matter from above the remnant poles which makes it easier for a highly relativistic magnetic outflow to form. Therefore detailed studies including these effects are also required to build a fuller picture of mergers which we could then test against

observations of events including sGRBs, kilonovae and gravitational wave signals.

The second part of this work was to create a model black hole in the Newtonian FLASH code framework. Doing so on a Cartesian grid proved to be challenging as the uneven surface of the resulting hole created many issues with stability, where a problem that occurred in only a handful of zones could crash the entire simulation if not dealt with. We built and tested this model using simulations of both magntized and unmagnetized accretion tori, and found that our scheme produced accretion rates for the magnetized model similar in magnitude to other comparable works, and that the schemes designed to maintain the model's stability were only active in a very small fraction (0.002% on average) of zones on the black hole boundary.

The two projects presented here both apply Newtonian hydrodynamics to relativistic compact objects, using simplified models to approximate the dynamics of the fully relativistic systems. Using such simplified models allows us to simulate these systems with much lower computational costs, which allows for the exploration of a greater range of model parameters or for focussing computational power toward increased resolution and included additional physics. The next stages for this project would be to combine the two individual components into a more complete set of tools capable of modelling black hole - neutron star merger simulations and other similar systems. This would be achieved primarily by further developing the model black hole by allowing the black hole to move on the grid and updating its Schwarzchild radius as mass is accreted during the simulation. Further refinements could also be made to the gravitational wave solver, for example by adapting the general Poisson solvers in FLASH to implement a more complete Post-Newtonian gravitational wave backreaction similar to those used in [Ruffert et al., 1996] and [Shibata et al., 1992].

Appendix A

Derivation of quadrupole moment time derivatives

The quadrupole moment of a mass distribution with density $\rho(\mathbf{r}, t)$ is given by

$$Q_{ij} = \int \rho \left(r_i r_j - \frac{1}{3}\delta_{ij} r_k r_k \right) dV \tag{A.1}$$

where ρ is the matter density and r_i is the ith component of the position vector. The gravitational wave luminosity produced by a given matter distribution depends on the third time derivative of Q_{ij}, as shown in Equation 2.6. To calculate this luminosity numerically in our simulations we seek to minimise the number of numerical derivatives we must perform as these reduce the accuracy of our calculation. Therefore, we choose to rewrite the time derivatives of the quadrupole moment in terms of fluid quantities we already have available to us in our simulation. To simplify these calculations, we will rewrite the symmetric and trace free definition of the quadrupole moment tensor in Equation A.1 as

$$Q_{ij} = STF \left[\int \rho r_i r_j dV \right] \tag{A.2}$$

where the operator STF denotes

$$STF[A_{ij}] = \frac{1}{2}A_{ij} + \frac{1}{2}A_{ji} - \frac{1}{3}\delta_{ij}A_{kk}$$

Taking the first time derivative of Equation A.2 and applying the product rule we find

$$\dot{Q}_{ij} = STF \left[\int dV \left(\dot{\rho} r_i r_j + \rho(v_i r_j + r_i v_j) \right) \right] \tag{A.3}$$

where v_i is the ith component of the velocity vector $\mathbf{v}(\mathbf{r}, t)$. We can rewrite the first term using the continuity equation (Equation 3.1) and then apply integration

by parts to give

$$\begin{aligned}\int dV\,\dot{\rho} r_i r_j &= -\int dV \frac{\partial(\rho v_k)}{\partial x_k} r_i r_j \\ &= -\int \rho r_i r_j v_k dS^k + \int dV \rho v_k \left(\frac{\partial(r_i)}{\partial x_k} r_j + \frac{\partial(r_j)}{\partial x_k} r_i\right) \\ &= \int dV \rho \left(v_i r_j + r_i v_j\right)\end{aligned}$$

Since we are integrating over all space, the surface terms must be zero on the boundary at infinity. Therefore, the first time derivative of the quadrupole moment is

$$\dot{Q}_{ij} = STF\left[\int 2\rho(v_i r_j) dV\right] \tag{A.4}$$

where we use the symmetry of the tensor enforced by the STF transform to equate the $v_i r_j$ and $v_j r_i$ terms.

Taking the second time derivative we find

$$\ddot{Q}_{ij} = STF\left[2\int \left(\dot{\rho} v_i r_j + \rho v_i v_j + \rho \dot{v}_i r_j\right) dV\right]. \tag{A.5}$$

For the first term, we can repeat the procedure of using the continuity equation to substitute for $\dot{\rho}$ and applying integration by parts to rewrite this as

$$\int \dot{\rho} v_i r_j dV = \int \rho \left(v_k \frac{\partial v_i}{\partial x_k} r_j + v_i v_j\right) dV.$$

The term containing $\dot{v}$ can be rewritten in a similar way, by substituting in the momentum conservation equation (Equation 3.2) to give

$$\int \rho \dot{v}_i r_j dV = \int \rho x_j \left(-\frac{\partial \psi}{\partial x_i} - v_k \frac{\partial v_i}{\partial x_k} - \frac{1}{\rho}\frac{\partial P}{\partial x_i}\right),$$

where ψ is the gravitational potential and P is the pressure. We can then apply integration by parts to the final term to remove the derivative of pressure giving

$$\int \rho \dot{v}_i r_j dV = \int \left(\rho x_j \left(-\frac{\partial \psi}{\partial x_i} - v_k \frac{\partial v_i}{\partial x_k}\right) + P\delta_{ij}\right) dV.$$

Considering the trace-free condition, the pressure term can now be neglected since it will vanish under the STF transform. Gathering together the remaining terms we find

$$\ddot{Q}_{ij} = STF\left[2\int \left(2\rho v_i v_j - \rho r_i \frac{\partial \psi}{\partial x_j}\right) dV\right] \tag{A.6}$$

Finally, taking the third time derivative of this, we obtain

$$\dddot{Q}_{ij} = STF\Bigg[2\int\Bigg(\underbrace{2\dot{\rho}v_iv_j}_{A} + \underbrace{2\rho(\dot{v}_iv_j}_{B} + v_i\dot{v}_j) - \underbrace{\dot{\rho}r_i\frac{\partial\psi}{\partial x_j}}_{C} - \rho\left(v_i\frac{\partial\psi}{\partial x_j} + r_i\frac{\partial\dot{\psi}}{\partial x_j}\right)\Bigg)dV\Bigg]$$

To simplify this, we again substitute for $\dot{\rho}$ and $\dot{v}$ using the continuity and momentum equations respectively. For term A, substituting the continuity equation and applying integration by parts gives

$$A = 2\int \rho v_k\left(\frac{\partial v_i}{\partial x_k}v_j + \frac{\partial v_j}{\partial x_k}v_i\right)dV$$

For B we substitute for $\dot{v}$ using the momentum equation and then integrate the pressure term by parts as we did before to get

$$B = -2\int\left(\rho v_j\left(\frac{\partial\psi}{\partial x_i} + v_k\frac{\partial v_i}{\partial x_k}\right) + P\frac{\partial v_j}{\partial x_i}\right)dV.$$

Finally, for term C we substitute for $\dot{\rho}$ to give

$$C = \int\left(\frac{\partial(\rho v_k)}{\partial x_k}r_i\frac{\partial\psi}{\partial x_j}\right)dV.$$

Putting these terms together and applying the symmetries enforced by the STF transform we find

$$\dddot{Q}_{ij} = STF\left[2\int\left(2P\frac{\partial v_i}{\partial x_j} + \frac{\partial\psi}{\partial x_i}\left(r_j\frac{\partial(\rho v_k)}{\partial x_k} - 3\rho v_j\right) - \rho r_i\frac{\partial\dot{\psi}}{\partial x_j}\right)dV\right] \quad \text{(A.7)}$$

References

[Abbott et al., 2020] Abbott, B. P., Abbott, R., Abbott, T. D., Abraham, S., et al. (2020). GW190425: Observation of a Compact Binary Coalescence with Total Mass ~ 3.4 $M_{\odot}$. *The Astrophysical Journal Letters*, 892(1):L3.

[Abbott et al., 2019] Abbott, B. P., Abbott, R., Abbott, T. D., Acernese, F., others, LIGO Scientific Collaboration, and Virgo Collaboration (2019). Properties of the Binary Neutron Star Merger GW170817. *Physical Review X*, 9(1):011001.

[Abbott et al., 2017a] Abbott, B. P., Abbott, R., Abbott, T. D., et al. (2017a). GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters*, 119(16):161101.

[Abbott et al., 2017b] Abbott, B. P., Abbott, R., Abbott, T. D., (INTEGRAL, et al. (2017b). Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *Astrophysical Journal Letters*, 848(2):L13.

[Abbott et al., 2017c] Abbott, B. P., Abbott, R., Abbott, T. D., LIGO Scientific Collaboration, Virgo Collaboration, INTEGRAL, IceCube Collaboration, AstroSat Cadmium Zinc Telluride Imager Team, IPN Collaboration, Insight-HXMT Collaboration, ANTARES Collaboration, Swift Collaboration, 1M2H Team, Dark Energy Camera GW-EM Collaboration, DES Collaboration, DLT40 Collaboration, GRAWITA: GRAvitational Wave Inaf TeAm, Fermi Large Area Telescope Collaboration, Australia Telescope Compact Array, A., Australian SKA Pathfinder, A., OzGrav, (Deeper, D., Wider, program, F., AST3, CAASTRO Collaborations, J-GEM, GROWTH, JAGWAR, Caltech-NRAO, TTU-NRAO, NuSTAR Collaborations, Pan-STARRS, Optical Telescope, N., Tech University, T., SALT Group, TOROS: Transient Robotic Observatory of South Collaboration, BOOTES Collaboration, Murchison Widefield Array, M. C., IKI-GW Follow-up Collaboration, H. E. S. S. Collaboration, Wavelength Array, L. L., HAWC Collaboration, Pierre Auger Collaboration., ALMA Collaboration, Euro VLBI Team, Pi of Sky Collaboration., Chandra Team at McGill University, Desert Fireball Network, D., ATLAS, Time Resolution Universe Survey, H., RIMAS, RATIR, and South Africa/MeerKAT, S. (2017c). Multi-messenger Observations of a Binary Neutron Star Merger. *The Astrophysical Journal Letters*, 848(2):L12.

[Abbott et al., 2018] Abbott, B. P., Abbott, R., LIGO Scientific Collaboration, Virgo Collaboration, et al. (2018). GW170817: Measurements of Neutron Star Radii and Equation of State. *Physical Review Letters*, 121(16):161101.

[Abbott et al., 2021] Abbott, R., Abbott, T. D., Abraham, S., Acernese, F., Ackley, K., Adams, A., et al. (2021). Observation of gravitational waves from two neutron star–black hole coalescences. *The Astrophysical Journal Letters*, 915(1):L5.

[Acheson and Acheson, 1978] Acheson, D. J. and Acheson, D. J. (1978). On the instability of toroidal magnetic fields and differential rotation in stars. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 289:459–500.

[Alexander et al., 2017] Alexander, K. D., Berger, E., Fong, W., Williams, P. K. G., Guidorzi, C., Margutti, R., Metzger, B. D., Annis, J., Blanchard, P. K., Brout, D., Brown, D. A., Chen, H.-Y., Chornock, R., Cowperthwaite, P. S., Drout, M., Eftekhari, T., Frieman, J., Holz, D. E., Nicholl, M., Rest, A., Sako, M., Soares-Santos, M., and Villar, V. A. (2017). The electromagnetic counterpart of the binary neutron star merger LIGO/virgo GW170817. VI. radio constraints on a relativistic jet and predictions for late-time emission from the kilonova ejecta. *The Astrophysical Journal*, 848(2):L21.

[Alexander et al., 2018] Alexander, K. D., Margutti, R., Blanchard, P. K., Fong, W., Berger, E., Hajela, A., Eftekhari, T., Chornock, R., Cowperthwaite, P. S., Giannios, D., Guidorzi, C., Kathirgamaraju, A., MacFadyen, A., Metzger, B. D., Nicholl, M., Sironi, L., Villar, V. A., Williams, P. K. G., Xie, X., and Zrake, J. (2018). A Decline in the X-Ray through Radio Emission from GW170817 Continues to Support an Off-axis Structured Jet. *The Astrophysical Journal Letters*, 863(2):L18.

[Alfvén, 1943] Alfvén, H. (1943). On the Existence of Electromagnetic-Hydrodynamic Waves. *Arkiv for Matematik, Astronomi och Fysik*, 29B:1–7.

[Anderson et al., 2008a] Anderson, M., Hirschmann, E. W., Lehner, L., Liebling, S. L., Motl, P. M., Neilsen, D., Palenzuela, C., and Tohline, J. E. (2008a). Magnetized Neutron-Star Mergers and Gravitational-Wave Signals. *Physical Review Letters*, 100(19):191101.

[Anderson et al., 2008b] Anderson, M., Hirschmann, E. W., Lehner, L., Liebling, S. L., Motl, P. M., Neilsen, D., Palenzuela, C., and Tohline, J. E. (2008b). Simulating binary neutron stars: Dynamics and gravitational waves. *Physical Review D*, 77(2):024006.

[Arcavi et al., 2017] Arcavi, I., Hosseinzadeh, G., Howell, D. A., McCully, C., Poznanski, D., Kasen, D., Barnes, J., Zaltzman, M., Vasylyev, S., Maoz, D., and Valenti, S. (2017). Optical emission from a kilonova following a gravitational-wave-detected neutron-star merger. *Nature*, 551(7678):64–66.

[Arcones et al., 2006] Arcones, A., Scheck, L., and Janka, H.-T. (2006). Nucleosynthesis relevant conditions in neutrino-driven supernova outflows. In Mengoni, A., Cederkall, J., D'Auria, J., Fraile, L., Fynbo, H., Hass, M., Jeanne, M. M., Hjorth Jensen, M., Kaeppeler, F., Kratz, K. L., Lindroos, M., Menna, M., Meynet, G., and Riisager, K., editors, *International Symposium on Nuclear Astrophysics - Nuclei in the Cosmos*, page 231.1.

[Baiotti et al., 2008] Baiotti, L., Giacomazzo, B., and Rezzolla, L. (2008). Accurate evolutions of inspiralling neutron-star binaries: Prompt and delayed collapse to a black hole. *Physical Review D*, 78(8):084033.

[Baker et al., 2006] Baker, J. G., Centrella, J., Choi, D.-I., Koppitz, M., and van Meter, J. (2006). Gravitational-Wave Extraction from an Inspiraling Configuration of Merging Black Holes. *Physical Review Letters*, 96(11):111102.

[Balbus and Hawley, 1991] Balbus, S. A. and Hawley, J. F. (1991). A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. *The Astrophysical Journal*, 376:214.

[Balsara, 1998] Balsara, D. S. (1998). Total Variation Diminishing Scheme for Adiabatic and Isothermal Magnetohydrodynamics. *The Astrophysical Journal Supplement Series*, 116(1):133–153.

[Balsara and Kim, 2004] Balsara, D. S. and Kim, J. (2004). A Comparison between Divergence-Cleaning and Staggered-Mesh Formulations for Numerical Magnetohydrodynamics. *The Astrophysical Journal*, 602(2):1079–1090.

[Band et al., 1993] Band, D., Matteson, J., Ford, L., Schaefer, B., Palmer, D., Teegarden, B., Cline, T., Briggs, M., Paciesas, W., Pendleton, G., Fishman, G., Kouveliotou, C., Meegan, C., Wilson, R., and Lestrade, P. (1993). BATSE Observations of Gamma-Ray Burst Spectra. I. Spectral Diversity. *The Astrophysical Journal*, 413:281.

[Barnes and Hut, 1986] Barnes, J. and Hut, P. (1986). A hierarchical O(N log N) force-calculation algorithm. *Nature*, 324(6096):446–449.

[Bauswein et al., 2020] Bauswein, A., Blacker, S., Vijayan, V., Stergioulas, N., Chatziioannou, K., Clark, J. A., Bastian, N.-U. F., Blaschke, D. B., Cierniak, M., and Fischer, T. (2020). Equation of state constraints from the threshold binary mass for prompt collapse of neutron star mergers. *Phys. Rev. Lett.*, 125:141103.

[Bauswein et al., 2012] Bauswein, A., Janka, H. T., Hebeler, K., and Schwenk, A. (2012). Equation-of-state dependence of the gravitational-wave signal from the ring-down phase of neutron-star mergers. *Physical Review D*, 86(6):063001.

[Bauswein and Stergioulas, 2015] Bauswein, A. and Stergioulas, N. (2015). Unified picture of the post-merger dynamics and gravitational wave emission in neutron star mergers. *Physical Review D*, 91(12):124056.

[Berger et al., 2013] Berger, E., Fong, W., and Chornock, R. (2013). An r-process Kilonova Associated with the Short-hard GRB 130603B. *The Astrophysical Journal Letters*, 774(2):L23.

[Bernuzzi, 2020] Bernuzzi, S. (2020). Neutron star merger remnants. *General Relativity and Gravitation*, 52(11):108.

[Bhattacharya and van den Heuvel, 1991] Bhattacharya, D. and van den Heuvel, E. P. J. (1991). Formation and evolution of binary and millisecond radio pulsars. *Physics Reports*, 203(1-2):1–124.

[Blanchet, 2006] Blanchet, L. (2006). Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Reviews in Relativity*, 9(1):4.

[Blanchet et al., 1990] Blanchet, L., Damour, T., and Schaefer, G. (1990). Post-Newtonian hydrodynamics and post-Newtonian gravitational wave generation for numerical relativity. *Monthly Notices of the Royal Astronomical Society*, 242:289–305.

[Blandford and Znajek, 1977] Blandford, R. D. and Znajek, R. L. (1977). Electromagnetic extraction of energy from Kerr black holes. *Monthly Notices of the Royal Astronomical Society*, 179:433–456.

[Bloom et al., 2002] Bloom, J. S., Kulkarni, S. R., and Djorgovski, S. G. (2002). The Observed Offset Distribution of Gamma-Ray Bursts from Their Host Galaxies: A Robust Clue to the Nature of the Progenitors. *The Astrophysical Journal*, 123(3):1111–1148.

[Bonetti et al., 2019] Bonetti, M., Perego, A., Dotti, M., and Cescutti, G. (2019). Neutron star binary orbits in their host potential: effect on early r-process enrichment. *Monthly Notices of the Royal Astronomical Society*, 490(1):296–311.

[Brackbill and Barnes, 1980] Brackbill, J. U. and Barnes, D. C. (1980). The Effect of Nonzero $\nabla \cdot$ B on the numerical solution of the magnetohydrodynamic equations. *Journal of Computational Physics*, 35(3):426–430.

[Bromberg et al., 2017] Bromberg, O., Tchekhovskoy, A., Gottlieb, O., Nakar, E., and Piran, T. (2017). The gamma-rays that accompanied GW170817 and the observational signature of a magnetic jet breaking out of NS merger ejecta. *Monthly Notices of the Royal Astronomical Society*, 475(3):2971–2977.

[Bryan et al., 2014] Bryan, G. L., Norman, M. L., O'Shea, B. W., Abel, T., Wise, J. H., Turk, M. J., Reynolds, D. R., Collins, D. C., Wang, P., Skillman, S. W., Smith, B., Harkness, R. P., Bordner, J., Kim, J.-h., Kuhlen, M., Xu, H., Goldbaum, N., Hummels, C., Kritsuk, A. G., Tasker, E., Skory, S., Simpson, C. M., Hahn, O., Oishi, J. S., So, G. C., Zhao, F., Cen, R., Li, Y., and Enzo Collaboration (2014). ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *The Astrophysical Journal Supplement Series*, 211(2):19.

[Cavallo and Rees, 1978] Cavallo, G. and Rees, M. J. (1978). A qualitative study of cosmic fireballs and gamma -ray bursts. *Monthly Notices of the Royal Astronomical Society*, 183:359–365.

[Chandrasekhar, 1960] Chandrasekhar, S. (1960). The Stability of Non-Dissipative Couette Flow in Hydromagnetics. *Proceedings of the National Academy of Science*, 46:253–257.

[Cipolletta et al., 2020] Cipolletta, F., Kalinani, J. V., Giacomazzo, B., and Ciolfi, R. (2020). Spritz: a new fully general-relativistic magnetohydrodynamic code. *Classical and Quantum Gravity*, 37(13):135010.

[Cipolletta et al., 2021] Cipolletta, F., Kalinani, J. V., Giangrandi, E., Giacomazzo, B., Ciolfi, R., Sala, L., and Giudici, B. (2021). Spritz: general relativistic magnetohydrodynamics with neutrinos. *Classical and Quantum Gravity*, 38(8):085021.

[Côté et al., 2017] Côté, B., Belczynski, K., Fryer, C. L., Ritter, C., Paul, A., Wehmeyer, B., and O'Shea, B. W. (2017). Advanced LIGO Constraints on Neutron Star Mergers and r-process Sites. *The Astrophysical Journal*, 836(2):230.

[Couch et al., 2013] Couch, S. M., Graziani, C., and Flocke, N. (2013). An Improved Multipole Approximation for Self-gravity and Its Importance for Core-collapse Supernova Simulations. *The Astrophysical Journal*, 778(2):181.

[Coulter et al., 2017] Coulter, D. A., Foley, R. J., Kilpatrick, C. D., Drout, M. R., Piro, A. L., Shappee, B. J., Siebert, M. R., Simon, J. D., Ulloa, N., Kasen, D., Madore, B. F., Murguia-Berthier, A., Pan, Y. C., Prochaska, J. X., Ramirez-Ruiz, E., Rest, A., and Rojas-Bravo, C. (2017). Swope Supernova Survey 2017a (SSS17a), the optical counterpart to a gravitational wave source. *Science*, 358(6370):1556–1558.

[Cromartie et al., 2020] Cromartie, H. T., Fonseca, E., Ransom, S. M., Demorest, P. B., Arzoumanian, Z., Blumer, H., Brook, P. R., DeCesar, M. E., Dolch, T., Ellis, J. A., Ferdman, R. D., Ferrara, E. C., Garver-Daniels, N., Gentile, P. A., Jones, M. L., Lam, M. T., Lorimer, D. R., Lynch, R. S., McLaughlin, M. A., Ng, C., Nice, D. J., Pennucci, T. T., Spiewak, R., Stairs, I. H., Stovall, K., Swiggum, J. K., and Zhu, W. W. (2020). Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar. *Nature Astronomy*, 4:72–76.

[Daigne and Mochkovitch, 1998] Daigne, F. and Mochkovitch, R. (1998). Gamma-ray bursts from internal shocks in a relativistic wind: temporal and spectral properties. *Monthly Notices of the Royal Astronomical Society*, 296(2):275–286.

[Deaton et al., 2013] Deaton, M. B., Duez, M. D., Foucart, F., O'Connor, E., Ott, C. D., Kidder, L. E., Muhlberger, C. D., Scheel, M. A., and Szilagyi, B. (2013). Black Hole-Neutron Star Mergers with a Hot Nuclear Equation of State: Outflow and Neutrino-cooled Disk for a Low-mass, High-spin Case. *The Astrophysical Journal*, 776(1):47.

[East et al., 2016] East, W. E., Paschalidis, V., and Pretorius, F. (2016). Equation of state effects and one-arm spiral instability in hypermassive neutron stars formed in eccentric neutron star mergers. *Classical and Quantum Gravity*, 33(24):244004.

[Einstein, 1918] Einstein, A. (1918). Über Gravitationswellen. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pages 154–167.

[Espino et al., 2022] Espino, P. L., Bozzola, G., and Paschalidis, V. (2022). Quantifying uncertainties in general relativistic magnetohydrodynamic codes. *arXiv e-prints*, page arXiv:2210.13481.

[Etienne et al., 2012] Etienne, Z. B., Paschalidis, V., and Shapiro, S. L. (2012). General-relativistic simulations of black-hole-neutron-star mergers: Effects of tilted magnetic fields. *Physical Review D*, 86(8):084026.

[Evans and Hawley, 1988] Evans, C. R. and Hawley, J. F. (1988). Simulation of Magnetohydrodynamic Flows: A Constrained Transport Model. *The Astrophysical Journal*, 332:659.

[Fong and Berger, 2013] Fong, W. and Berger, E. (2013). The Locations of Short Gamma-Ray Bursts as Evidence for Compact Object Binary Progenitors. *The Astrophysical Journal*, 776(1):18.

[Fruchter et al., 2006] Fruchter, A. S., Levan, A. J., Strolger, L., Vreeswijk, P. M., Thorsett, S. E., Bersier, D., Burud, I., Castro Cerón, J. M., Castro-Tirado, A. J., Conselice, C., Dahlen, T., Ferguson, H. C., Fynbo, J. P. U., Garnavich, P. M., Gibbons, R. A., Gorosabel, J., Gull, T. R., Hjorth, J., Holland, S. T., Kouveliotou, C., Levay, Z., Livio, M., Metzger, M. R., Nugent, P. E., Petro, L., Pian, E., Rhoads, J. E., Riess, A. G., Sahu, K. C., Smette, A., Tanvir, N. R., Wijers, R. A. M. J., and Woosley, S. E. (2006). Long γ-ray bursts and core-collapse supernovae have different environments. *Nature*, 441(7092):463–468.

[Fryxell et al., 2000] Fryxell, B., Olson, K., Ricker, P., Timmes, F. X., Zingale, M., Lamb, D. Q., MacNeice, P., Rosner, R., Truran, J. W., and Tufo, H. (2000). FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *The Astrophysical Journal Supplement Series*, 131:273–334.

[Giacomazzo et al., 2009] Giacomazzo, B., Rezzolla, L., and Baiotti, L. (2009). Can magnetic fields be detected during the inspiral of binary neutron stars? *Monthly Notices of the Royal Astronomical Society*, 399(1):L164–L168.

[Giacomazzo et al., 2011] Giacomazzo, B., Rezzolla, L., and Baiotti, L. (2011). Accurate evolutions of inspiralling and magnetized neutron stars: Equal-mass binaries. *Physical Review D*, 83(4):044014.

[Godunov, 1959] Godunov, S. K. (1959). A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Mat. Sb. (N.S.)*, 47 (89):271–306.

[Goodale et al., 2003] Goodale, T., Allen, G., Lanfermann, G., Massó, J., Radke, T., Seidel, E., and Shalf, J. (2003). The Cactus framework and toolkit: Design and applications. In *Vector and Parallel Processing – VECPAR'2002, 5th International Conference, Lecture Notes in Computer Science*, Berlin. Springer.

[Grindlay et al., 2006] Grindlay, J., Portegies Zwart, S., and McMillan, S. (2006). Short gamma-ray bursts from binary neutron star mergers in globular clusters. *Nature Physics*, 2(2):116–119.

[Hansen and Lyutikov, 2001] Hansen, B. M. S. and Lyutikov, M. (2001). Radio and X-ray signatures of merging neutron stars. *Monthly Notices of the Royal Astronomical Society*, 322(4):695–701.

[Hartmann et al., 1998] Hartmann, L., Calvet, N., Gullbring, E., and D'Alessio, P. (1998). Accretion and the Evolution of T Tauri Disks. *The Astrophysical Journal*, 495(1):385–400.

[Hawley, 2000] Hawley, J. F. (2000). Global magnetohydrodynamical simulations of accretion tori. *The Astrophysical Journal*, 528(1):462–479.

[Hjorth et al., 2003] Hjorth, J., Sollerman, J., Møller, P., Fynbo, J. P. U., Woosley, S. E., Kouveliotou, C., Tanvir, N. R., Greiner, J., Andersen, M. I., Castro-Tirado, A. J., Castro Cerón, J. M., Fruchter, A. S., Gorosabel, J., Jakobsson, P., Kaper, L., Klose, S., Masetti, N., Pedersen, H., Pedersen, K., Pian, E., Palazzi, E., Rhoads, J. E., Rol, E., van den Heuvel, E. P. J., Vreeswijk, P. M., Watson, D., and Wijers, R. A. M. J. (2003). A very energetic supernova associated with the γ-ray burst of 29 March 2003. *Nature*, 423(6942):847–850.

[Hotokezaka et al., 2011] Hotokezaka, K., Kyutoku, K., Okawa, H., Shibata, M., and Kiuchi, K. (2011). Binary neutron star mergers: Dependence on the nuclear equation of state. *Physical Review D*, 83(12):124008.

[Hubber et al., 2018] Hubber, D. A., Rosotti, G. P., and Booth, R. A. (2018). GANDALF - Graphical Astrophysics code for N-body Dynamics And Lagrangian Fluids. *Monthly Notices of the Royal Astronomical Society*, 473(2):1603–1632.

[Hulse and Taylor, 1975] Hulse, R. A. and Taylor, J. H. (1975). Discovery of a pulsar in a binary system. *The Astrophysical Journal Letters*, 195:L51–L53.

[Ji et al., 2016] Ji, A. P., Frebel, A., Chiti, A., and Simon, J. D. (2016). R-process enrichment from a single event in an ancient dwarf galaxy. *Nature*, 531(7596):610–613.

[Kasen et al., 2015] Kasen, D., Fernández, R., and Metzger, B. D. (2015). Kilonova light curves from the disc wind outflows of compact object mergers. *Monthly Notices of the Royal Astronomical Society*, 450(2):1777–1786.

[Kasen et al., 2017] Kasen, D., Metzger, B., Barnes, J., Quataert, E., and Ramirez-Ruiz, E. (2017). Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event. *Nature*, 551(7678):80–84.

[King et al., 2007] King, A. R., Pringle, J. E., and Livio, M. (2007). Accretion disc viscosity: how big is alpha? *Monthly Notices of the Royal Astronomical Society*, 376(4):1740–1746.

[Kiuchi et al., 2015] Kiuchi, K., Cerdá-Durán, P., Kyutoku, K., Sekiguchi, Y., and Shibata, M. (2015). Efficient magnetic-field amplification due to the Kelvin-Helmholtz instability in binary neutron star mergers. *Physical Review D*, 92(12):124034.

[Kiuchi et al., 2009] Kiuchi, K., Sekiguchi, Y., Shibata, M., and Taniguchi, K. (2009). Long-term general relativistic simulation of binary neutron stars collapsing to a black hole. *Physical Review D*, 80(6):064037.

[Kolmogorov, 1941] Kolmogorov, A. (1941). The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds' Numbers. *Akademiia Nauk SSSR Doklady*, 30:301–305.

[Kouveliotou et al., 1993] Kouveliotou, C., Meegan, C. A., Fishman, G. J., Bhat, N. P., Briggs, M. S., Koshut, T. M., Paciesas, W. S., and Pendleton, G. N. (1993). Identification of Two Classes of Gamma-Ray Bursts. *The Astrophysical Journal Letters*, 413:L101.

[Lane, 1870] Lane, H. J. (1870). On the theoretical temperature of the Sun, under the hypothesis of a gaseous mass maintaining its volume by its internal heat, and depending on the laws of gases as known to terrestrial experiment. *American Journal of Science*, 50(148):57–74.

[Lattimer and Schramm, 1974] Lattimer, J. M. and Schramm, D. N. (1974). Black-Hole-Neutron-Star Collisions. *The Astrophysical Journal Letters*, 192:L145.

[Lazzati, 2005] Lazzati, D. (2005). Precursor activity in bright, long BATSE gamma-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 357(2):722–731.

[Lazzati et al., 2001] Lazzati, D., Ramirez-Ruiz, E., and Ghisellini, G. (2001). Possible detection of hard X-ray afterglows of short gamma -ray bursts. *Astronomy and Astrophysics*, 379:L39–L43.

[Lee and Deane, 2009] Lee, D. and Deane, A. E. (2009). An unsplit staggered mesh scheme for multidimensional magnetohydrodynamics. *Journal of Computational Physics*, 228(4):952–975.

[Li, 2007] Li, L.-X. (2007). Gamma-ray burst precursors as the remnant of the thermal radiation initially trapped in the fireball. *Monthly Notices of the Royal Astronomical Society*, 380(2):621–636.

[Li and Paczyński, 1998] Li, L.-X. and Paczyński, B. (1998). Transient Events from Neutron Star Mergers. *The Astrophysical Journal Letters*, 507(1):L59–L62.

[Löffler et al., 2012] Löffler, F., Faber, J., Bentivegna, E., Bode, T., Diener, P., Haas, R., Hinder, I., Mundim, B. C., Ott, C. D., Schnetter, E., Allen, G., Campanelli, M., and Laguna, P. (2012). The Einstein Toolkit: a community computational infrastructure for relativistic astrophysics. *Classical and Quantum Gravity*, 29:115001.

[MAGIC Collaboration et al., 2019] MAGIC Collaboration, Acciari, V. A., Ansoldi, S., et al. (2019). Teraelectronvolt emission from the γ-ray burst GRB 190114C. *Nature*, 575(7783):455–458.

[Margalit and Metzger, 2017] Margalit, B. and Metzger, B. D. (2017). Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817. *The Astrophysical Journal Letters*, 850(2):L19.

[Margutti et al., 2017] Margutti, R., Berger, E., Fong, W., Guidorzi, C., Alexander, K. D., Metzger, B. D., Blanchard, P. K., Cowperthwaite, P. S., Chornock, R., Eftekhari, T., Nicholl, M., Villar, V. A., Williams, P. K. G., Annis, J., Brown, D. A., Chen, H., Doctor, Z., Frieman, J. A., Holz, D. E., Sako, M., and Soares-Santos, M. (2017). The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. V. Rising X-Ray Emission from an Off-axis Jet. *The Astrophysical Journal Letters*, 848(2):L20.

[Margutti et al., 2012] Margutti, R., Zaninoni, E., Bernardini, M. G., Chincarini, G., Pasotti, F., Guidorzi, C., Angelini, L., Burrows, D. N., Capalbi, M., Evans, P. A., Gehrels, N., Kennea, J., Mangano, V., Moretti, A., Nousek, J., Osborne, J. P., Page, K. L., Perri, M., Racusin, J., Romano, P., Sbarufatti, B., Stafford, S., and Stamatikos, M. (2012). The prompt-afterglow connection in gamma-ray bursts: a comprehensive statistical analysis of Swift X-ray light curves. *Monthly Notices of the Royal Astronomical Society*, 428(1):729–742.

[Martínez-Pinedo et al., 2012] Martínez-Pinedo, G., Fischer, T., Lohs, A., and Huther, L. (2012). Charged-Current Weak Interaction Processes in Hot and Dense Matter and its Impact on the Spectra of Neutrinos Emitted from Protoneutron Star Cooling. *Physical Review Letters*, 109(25):251104.

[Maselli et al., 2013] Maselli, A., Cardoso, V., Ferrari, V., Gualtieri, L., and Pani, P. (2013). Equation-of-state-independent relations in neutron stars. *Physical Review D*, 88(2):023007.

[Mathews et al., 1992] Mathews, G. J., Bazan, G., and Cowan, J. J. (1992). Evolution of Heavy-Element Abundances as a Constraint on Sites for Neutron-Capture Nucleosynthesis. *The Astrophysical Journal*, 391:719.

[Metzger and Fernández, 2014] Metzger, B. D. and Fernández, R. (2014). Red or blue? A potential kilonova imprint of the delay until black hole formation following a neutron star merger. *Monthly Notices of the Royal Astronomical Society*, 441(4):3444–3453.

[Metzger et al., 2010] Metzger, B. D., Martínez-Pinedo, G., Darbha, S., Quataert, E., Arcones, A., Kasen, D., Thomas, R., Nugent, P., Panov, I. V., and Zinner,

N. T. (2010). Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *Monthly Notices of the Royal Astronomical Society*, 406(4):2650–2662.

[Metzger et al., 2008] Metzger, B. D., Quataert, E., and Thompson, T. A. (2008). Short-duration gamma-ray bursts with extended emission from protomagnetar spin-down. *Monthly Notices of the Royal Astronomical Society*, 385(3):1455–1460.

[Mignone et al., 2012] Mignone, A., Zanni, C., Tzeferacos, P., van Straalen, B., Colella, P., and Bodo, G. (2012). The PLUTO Code for Adaptive Mesh Computations in Astrophysical Fluid Dynamics. *The Astrophysical Journal Supplement Series*, 198:7.

[Mösta et al., 2020] Mösta, P., Radice, D., Haas, R., Schnetter, E., and Bernuzzi, S. (2020). A magnetar engine for short GRBs and kilonovae. *The Astrophysical Journal*, 901(2):L37.

[Nicholl et al., 2017] Nicholl, M., Berger, E., Kasen, D., Metzger, B. D., Elias, J., Briceño, C., Alexander, K. D., Blanchard, P. K., Chornock, R., Cowperthwaite, P. S., Eftekhari, T., Fong, W., Margutti, R., Villar, V. A., Williams, P. K. G., Brown, W., Annis, J., Bahramian, A., Brout, D., Brown, D. A., Chen, H. Y., Clemens, J. C., Dennihy, E., Dunlap, B., Holz, D. E., Marchesini, E., Massaro, F., Moskowitz, N., Pelisoli, I., Rest, A., Ricci, F., Sako, M., Soares-Santos, M., and Strader, J. (2017). 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(2):L18.

[Nousek et al., 2006] Nousek, J. A., Kouveliotou, C., Grupe, D., et al. (2006). Evidence for a Canonical Gamma-Ray Burst Afterglow Light Curve in the Swift XRT Data. *The Astrophysical Journal*, 642(1):389–400.

[Olson et al., 1999] Olson, K. M., MacNeice, P., Fryxell, B., Ricker, P., Timmes, F. X., and Zingale, M. (1999). PARAMESH: A Parallel, Adaptive Mesh Refinement Toolkit and Performance of the ASCI/FLASH code. In *American Astronomical Society Meeting Abstracts*, volume 195 of *American Astronomical Society Meeting Abstracts*, page 42.03.

[Oohara and Nakamura, 1989] Oohara, K. and Nakamura, T. (1989). Gravitational Radiation from Coalescing Binary Neutron Stars. I. *Progress of Theoretical Physics*, 82(3):535–554.

[Paczyński and Wiita, 1980] Paczyński, B. and Wiita, P. J. (1980). Thick Accretion Disks and Supercritical Luminosities. *Astronomy and Astrophysics*, 88:23.

[Papaloizou and Pringle, 1984] Papaloizou, J. C. B. and Pringle, J. E. (1984). The dynamical stability of differentially rotating discs with constant specific angular momentum. *Monthly Notices of the Royal Astronomical Society*, 208:721–750.

[Paschalidis et al., 2015a] Paschalidis, V., East, W. E., Pretorius, F., and Shapiro, S. L. (2015a). One-arm spiral instability in hypermassive neutron stars formed by dynamical-capture binary neutron star mergers. *Physical Review D*, 92(12):121502.

[Paschalidis et al., 2015b] Paschalidis, V., Ruiz, M., and Shapiro, S. L. (2015b). Relativistic Simulations of Black Hole-Neutron Star Coalescence: The Jet Emerges. *The Astrophysical Journal Letters*, 806(1):L14.

[Perego et al., 2017] Perego, A., Radice, D., and Bernuzzi, S. (2017). AT 2017gfo: An Anisotropic and Three-component Kilonova Counterpart of GW170817. *The Astrophysical Journal Letters*, 850(2):L37.

[Perego et al., 2014] Perego, A., Rosswog, S., Cabezón, R. M., Korobkin, O., Käppeli, R., Arcones, A., and Liebendörfer, M. (2014). Neutrino-driven winds from neutron star merger remnants. *Monthly Notices of the Royal Astronomical Society*, 443(4):3134–3156.

[Portegies Zwart et al., 1999] Portegies Zwart, S., , and McMillan, S. (1999). Black hole mergers in the universe. *The Astrophysical Journal*, 528(1):L17.

[Pretorius, 2005] Pretorius, F. (2005). Evolution of Binary Black-Hole Spacetimes. *Physical Review Letters*, 95(12):121101.

[Price and Rosswog, 2006] Price, D. J. and Rosswog, S. (2006). Producing Ultra-strong Magnetic Fields in Neutron Star Mergers. *Science*, 312(5774):719–722.

[Price et al., 2018] Price, D. J., Wurster, J., Tricco, T. S., Nixon, C., Toupin, S., Pettitt, A., Chan, C., Mentiplay, D., Laibe, G., Glover, S., Dobbs, C., Nealon, R., Liptai, D., Worpel, H., Bonnerot, C., Dipierro, G., Ballabio, G., Ragusa, E., Federrath, C., Iaconi, R., Reichardt, T., Forgan, D., Hutchison, M., Constantino, T., Ayliffe, B., Hirsh, K., and Lodato, G. (2018). Phantom: A Smoothed Particle Hydrodynamics and Magnetohydrodynamics Code for Astrophysics. *Publications of the Astronomical Society of Australia*, 35:e031.

[Radice et al., 2022] Radice, D., Bernuzzi, S., Perego, A., and Haas, R. (2022). 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(1):1499–1521.

[Radice and Rezzolla, 2012] Radice, D. and Rezzolla, L. (2012). THC: a new high-order finite-difference high-resolution shock-capturing code for special-relativistic hydrodynamics. *Astronomy and Astrophysics*, 547:A26.

[Ramshaw, 1983] Ramshaw, J. D. (1983). A Method for Enforcing the Solenoidal Condition on Magnetic Field in Numerical Calculations. *Journal of Computational Physics*, 52(3):592–596.

[Rees and Meszaros, 1994] Rees, M. J. and Meszaros, P. (1994). Unsteady Outflow Models for Cosmological Gamma-Ray Bursts. *The Astrophysical Journal Letters*, 430:L93.

[Rezzolla et al., 2011] Rezzolla, L., Giacomazzo, B., Baiotti, L., Granot, J., Kouveliotou, C., and Aloy, M. A. (2011). The Missing Link: Merging Neutron Stars Naturally Produce Jet-like Structures and Can Power Short Gamma-ray Bursts. *The Astrophysical Journal Letters*, 732(1):L6.

[Ricker, 2008] Ricker, P. M. (2008). A Direct Multigrid Poisson Solver for Oct-Tree Adaptive Meshes. *The Astrophysical Journal Supplement Series*, 176(1):293–300.

[Ruffert et al., 1996] Ruffert, M., Janka, H.-T., and Schaefer, G. (1996). Coalescing neutron stars - a step towards physical models. I. Hydrodynamic evolution and gravitational-wave emission. *Astronomy and Astrophysics*, 311:532–566.

[Ruffert et al., 1997a] Ruffert, M., Janka, H.-T., Takahashi, K., and Schaefer, G. (1997a). Coalescing neutron stars - a step towards physical models. II. Neutrino emission, neutron tori, and gamma-ray bursts. *Astronomy and Astrophysics*, 319:122–153.

[Ruffert et al., 1997b] Ruffert, M., Rampp, M., and Janka, H. T. (1997b). Coalescing neutron stars - gravitational waves from polytropic models. *Astronomy and Astrophysics*, 321:991–1006.

[Ruiz et al., 2016] Ruiz, M., Lang, R. N., Paschalidis, V., and Shapiro, S. L. (2016). Binary Neutron Star Mergers: A Jet Engine for Short Gamma-Ray Bursts. *The Astrophysical Journal Letters*, 824(1):L6.

[Ruiz et al., 2019] Ruiz, M., Tsokaros, A., Paschalidis, V., and Shapiro, S. L. (2019). Effects of spin on magnetized binary neutron star mergers and jet launching. *Physical Review D*, 99(8):084032.

[Ruiz et al., 2020] Ruiz, M., Tsokaros, A., and Shapiro, S. L. (2020). Magnetohydrodynamic simulations of binary neutron star mergers in general relativity: Effects of magnetic field orientation on jet launching. *Physical Review D*, 101(6):064042.

[Ryde, 2005] Ryde, F. (2005). Is Thermal Emission in Gamma-Ray Bursts Ubiquitous? *The Astrophysical Journal Letters*, 625(2):L95–L98.

[Sadowski et al., 2008] Sadowski, A., Belczynski, K., Bulik, T., Ivanova, N., Rasio, F. A., and O'Shaughnessy, R. (2008). The Total Merger Rate of Compact Object Binaries in the Local Universe. *The Astrophysical Journal*, 676(2):1162–1169.

[Shakura and Sunyaev, 1973] Shakura, N. I. and Sunyaev, R. A. (1973). Black holes in binary systems. Observational appearance. *Astronomy and Astrophysics*, 24:337–355.

[Shibata et al., 2017] Shibata, M., Fujibayashi, S., Hotokezaka, K., Kiuchi, K., Kyutoku, K., Sekiguchi, Y., and Tanaka, M. (2017). Modeling GW170817 based on numerical relativity and its implications. *Physical Review D*, 96(12):123012.

[Shibata et al., 1992] Shibata, M., Nakamura, T., and Oohara, K.-i. (1992). Coalescence of Spinning Binary Neutron Stars of Equal Mass — 3D Numerical Simulations —. *Progress of Theoretical Physics*, 88(6):1079–1095.

[Shibata and Uryū, 2000] Shibata, M. and Uryū, K. ō. (2000). Simulation of merging binary neutron stars in full general relativity: Γ=2 case. *Physical Review D*, 61(6):064001.

[Springel, 2011] Springel, V. (2011). Moving-mesh hydrodynamics with the AREPO code. In Alves, J., Elmegreen, B. G., Girart, J. M., and Trimble, V., editors, *Computational Star Formation*, volume 270, pages 203–206.

[Sun et al., 2022] Sun, L., Ruiz, M., Shapiro, S. L., and Tsokaros, A. (2022). Jet launching from binary neutron star mergers: Incorporating neutrino transport and magnetic fields. *Physical Review D*, 105(10):104028.

[Symbalisty and Schramm, 1982] Symbalisty, E. and Schramm, D. N. (1982). Neutron Star Collisions and the r-Process. *Astrophysical Letters*, 22:143.

[Tanvir et al., 2013] Tanvir, N. R., Levan, A. J., Fruchter, A. S., Hjorth, J., Hounsell, R. A., Wiersema, K., and Tunnicliffe, R. L. (2013). A 'kilonova' associated with the short-duration γ-ray burst GRB 130603B. *Nature*, 500(7464):547–549.

[Tassoul, 1978] Tassoul, J. L. (1978). *Theory of rotating stars / by Jean-Louis Tassoul.* Princeton series in astrophysics. Princeton University Press, Princeton, N.J.

[Taylor et al., 1979] Taylor, J. H., Fowler, L. A., and McCulloch, P. M. (1979). Measurements of general relativistic effects in the binary pulsar PSR1913 + 16. *Nature*, 277(5696):437–440.

[Teyssier, 2002] Teyssier, R. (2002). Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *Astronomy and Astrophysics*, 385:337–364.

[Thierfelder et al., 2011] Thierfelder, M., Bernuzzi, S., and Brügmann, B. (2011). Numerical relativity simulations of binary neutron stars. *Physical Review D*, 84(4):044012.

[Thompson and Duncan, 1995] Thompson, C. and Duncan, R. C. (1995). The soft gamma repeaters as very strongly magnetized neutron stars - I. Radiative mechanism for outbursts. *Monthly Notices of the Royal Astronomical Society*, 275(2):255–300.

[Thompson et al., 2001] Thompson, T. A., Burrows, A., and Meyer, B. S. (2001). The Physics of Proto-Neutron Star Winds: Implications for r-Process Nucleosynthesis. *The Astrophysical Journal*, 562(2):887–908.

[Thorne et al., 1986] Thorne, K. S., Price, R. H., and MacDonald, D. A. (1986). *Black holes: The membrane paradigm.* Yale University Press.

[Toro, 2013] Toro, E. (2013). *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer Berlin Heidelberg.

[Troja et al., 2008] Troja, E., King, A. R., O'Brien, P. T., Lyons, N., and Cusumano, G. (2008). Different progenitors of short hard gamma-ray bursts. *Monthly Notices of the Royal Astronomical Society: Letters*, 385(1):L10–L14.

[Troja et al., 2017] Troja, E., Piro, L., van Eerten, H., Wollaeger, R. T., Im, M., Fox, O. D., Butler, N. R., Cenko, S. B., Sakamoto, T., Fryer, C. L., Ricci, R., Lien, A., Ryan, R. E., Korobkin, O., Lee, S. K., Burgess, J. M., Lee, W. H., Watson, A. M., Choi, C., Covino, S., D'Avanzo, P., Fontes, C. J., González, J. B., Khandrika, H. G., Kim, J., Kim, S. L., Lee, C. U., Lee, H. M., Kutyrev, A., Lim, G., Sánchez-Ramírez, R., Veilleux, S., Wieringa, M. H., and Yoon, Y. (2017). The X-ray counterpart to the gravitational-wave event GW170817. *Nature*, 551(7678):71–74.

[Troja et al., 2010] Troja, E., Rosswog, S., and Gehrels, N. (2010). Precursors of short gamma-ray bursts. *The Astrophysical Journal*, 723(2):1711–1717.

[Turk et al., 2011] Turk, M. J., Smith, B. D., Oishi, J. S., Skory, S., Skillman, S. W., Abel, T., and Norman, M. L. (2011). yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *The Astrophysical Journal Supplement*, 192(1):9.

[van Leer, 1974] van Leer, B. (1974). Towards the ultimate conservative difference scheme. ii. monotonicity and conservation combined in a second-order scheme. *Journal of Computational Physics*, 14(4):361–370.

[Wallner et al., 2015] Wallner, A., Faestermann, T., Feige, J., Feldstein, C., Knie, K., Korschinek, G., Kutschera, W., Ofan, A., Paul, M., Quinto, F., Rugel, G., and Steier, P. (2015). Abundance of live ^{244}Pu in deep-sea reservoirs on Earth points to rarity of actinide nucleosynthesis. *Nature Communications*, 6:5956.

[Wang et al., 2020] Wang, J.-S., Peng, Z.-K., Zou, J.-H., Zhang, B.-B., and Zhang, B. (2020). Stringent Search for Precursor Emission in Short GRBs from Fermi/GBM Data and Physical Implications. *The Astrophysical Journal Letters*, 902(2):L42.

[White et al., 2016] White, C. J., Stone, J. M., and Gammie, C. F. (2016). An Extension of the Athena++ Code Framework for GRMHD Based on Advanced Riemann Solvers and Staggered-mesh Constrained Transport. *The Astrophysical Journal Supplement Series*, 225:22.

[Wünsch et al., 2018] Wünsch, R., Walch, S., Dinnbier, F., and Whitworth, A. (2018). Tree-based solvers for adaptive mesh refinement code FLASH - I: gravity and optical depths. *Monthly Notices of the Royal Astronomical Society*, 475(3):3393–3418.

[Yamamoto et al., 2008] Yamamoto, T., Shibata, M., and Taniguchi, K. (2008). Simulating coalescing compact binaries by a new code (SACRA). *Physical Review D*, 78(6):064054.

[Zhuge et al., 1994] Zhuge, X., Centrella, J. M., and McMillan, S. L. W. (1994). Gravitational radiation from coalescing binary neutron stars. *Physical Review D*, 50(10):6247–6261.