To the Graduate Council:

I am submitting herewith a dissertation written by Reuben Donald Budiardja entitled "Towards Simulations of Binary Neutron Star Mergers and Core-Collapse Supernovae with GenASiS." I have examined the final paper copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Physics.

Michael Guidry, Major Professor

We have read this dissertation
and recommend its acceptance:

Christian Y. Cardall

Soren Sorensen

George Siopsis

Jim Chambers

Accepted for the Council:

Carolyn R. Hodges

Vice Provost and Dean of the Graduate School

To the Graduate Council:

I am submitting herewith a dissertation written by Reuben Donald Budiardja entitled "Towards Simulations of Binary Neutron Star Mergers and Core-Collapse Supernovae with GenASiS." I have examined the final electronic copy of this dissertation for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy, with a major in Physics.

Michael Guidry, Major Professor

We have read this dissertation
and recommend its acceptance:

Christian Y. Cardall

Soren Sorensen

George Siopsis

Jim Chambers

Accepted for the Council:

Carolyn R. Hodges

Vice Provost and Dean of the Graduate School

(Original signatures are on file with official student records.)

Towards Simulations of Binary Neutron Star Mergers and Core-Collapse Supernovae with GenASiS

A Dissertation
Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Reuben Donald Budiardja
August 2010

Abstract

This manuscript describes the current version of GenASiS and reports recent progress in its development. GenASiS is a new computational astrophysics code built for large-scale and multi-dimensional computer simulations of astrophysical phenomena, with primary emphasis on the simulations of neutron star mergers and core-collapse supernovae. Neutron star mergers are of high interest to the astrophysics community because they should be the prodigious source of gravitation waves and the most promising candidates for gravitational wave detection. Neutron star mergers are also thought to be associated with the production of short-duration, hard-spectral gamma-ray bursts, though the mechanism is not well understood. In contrast, core-collapse supernovae with massive progenitors are associated with long-duration, soft-spectral gamma-ray bursts, with the 'collapsar' hypothesis as the favored mechanism. Of equal interest is the mechanism of core-collapse supernovae themselves, which has been in the forefront of many research efforts for the better half of a century but remains a partially-solved mystery. In addition supernovae, and possibly neutron star mergers, are thought to be sites for the *r*-process nucleosynthesis responsible for producing many of the heavy elements. Until we have a proper understanding of these events, we will have only a limited understanding of the origin of the elements. These questions provide some of the scientific motivations and guidelines for the development of GenASiS. In this document the equations and numerical scheme for Newtonian and relativistic magnetohydrodynamics are presented. A new FFT-based parallel solver for Poisson's equation in GenASiS are described. Adaptive mesh refinement in GenASiS, and a novel way to solve Poisson's equation on a mesh with refinement based on a multigrid algorithm, are also presented. Following these descriptions, results of simulations of neutron star mergers with GenASiS such as their evolution and the gravitational wave signals and spectra that they generate are shown. In the context of core-collapse supernovae, we explore the capacity of the stationary shock instability to generate magnetic fields starting from a weak, stationary, and radial magnetic field in an initially spherically symmetric fluid configuration that models the stalled shock in the post-bounce supernova environment. Our results show that the magnetic energy can be amplified by almost 4 orders of magnitude. The amplification mechanisms for the magnetic fields are then explained.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

We are developing a new computational astrophysics code, GenASiS, that is suitable for large-scale computer simulations of astrophysical phenomena, with primary emphasis on the simulations of neutron star mergers and core-collapse supernovae. This manuscript describes the current implementation of GenASiS and reports recent progress in the development. In this chapter, we provide the scientific motivations guiding GenASiS' development, introduce the computational challenges we are trying to meet, and give an outline of the rest of the document.

1.1 Scientific Background

Neutron star mergers and core-collapse supernovae are of fundamental interest on their own merits, and because they are likely sources of gamma-ray bursts, strong gravitational-wave emission, and nucleosynthesis of many heavy elements. In the rest of this section, these observables and phenomena are discussed more fully, providing the scientific motivations for studies involving computer simulations of neutron star mergers and core-collapse supernovae.

1.1.1 Observables

1.1.1.1 Gamma-ray Burst

In the late 1960s, United States Vela military satellites began registering gamma radiation pulses, radiations that are commonly emitted by nuclear weapons. The Vela satellites were originally built to monitor compliance with the nuclear test ban treaty. The detected radiation signatures, however, were unlike any other known nuclear weapon signatures. Several more bursts were observed by additional Vela satellites before Klebesadel et al. (1973) ruled out the terrestrial and solar origin of these bursts by estimating the sky positions of sixteen bursts using data from multiple satellites. This was quickly confirmed by data from the Soviet Konus satellites (Mazets et al., 1982).

This initial discovery of Gamma-ray bursts (GRB), followed by even more events of GRB recorded by various satellites, spurred interests in their nature and origins. However, it was not until about two decades later, with the launch of the Compton Gamma-Ray Observatory (CGRO) and its all-sky survey instrument Burst and Transient Source Experiment (BATSE), that more definitive origins of GRB were obtained. Over 2700 bursts that were recorded by BATSE showed essentially isotropic distribution in the sky, suggesting a cosmological origin (Meegan et al., 1992). These were confirmed and complemented by data from Oriented Scintillation Spectrometer Experiment (OSSE), Compton Telescope (CompTel), and Energetic Gamma Ray Experiment Telescope (EGRET) on the CGRO (Fishman and Meegan, 1995). The launch of CGRO marked a new era of GRB observations, obtaining spectral data (Band et al., 1993), light curves (Fishman and Meegan, 1995), energy ranges up to GeV (Schneid et al., 1995; Hurley et al., 1994), and durations of GRB (Kouveliotou et al., 1993).

GRB are the most luminous electromagnetic events in the entire universe. Yet for decades after the initial discovery of GRB, astronomers unsuccessfully searched for the optical counterparts of GRB to pinpoint the sources of the bursts (Fishman and Meegan, 1995). A debate on whether the GRB were of galactic or extragalactic origin arose (Hurley, 1992; Paczynski, 1995; Lamb, 1995). These issues were finally resolved with the successes of the Beppo-SAX satellite in obtaining the X-ray images of the fading afterglow of GRB in 1997 (Costa et al., 1997). The afterglows of GRB had been expected on theoretical grounds (Paczynski and Rhoads, 1993), but earlier satellites were unable to find these due to the difficulties in observing at longer wavelengths after the initial bursts. The detection of X-ray images finally made possible the optical detection because of the arc-minute position accuracy (Frail and Taylor, 1997; van Paradijs et al., 1997). Measurement of redshift distances and identification of candidate host galaxies were finally possible with these new data (Metzger et al., 1997; Frail et al., 1999; Kulkarni et al., 1999). More recent observations of GRB and their afterglows were reviewed by van Paradijs et al. (2000) and Weiler et al. (2002).

Kouveliotou et al. (1993) studied the duration distribution of GRB and showed that there is a well-defined bimodal distribution corresponding to bursts longer or shorter than about 2 seconds. Therefore GRB are generally classified into short bursts that last only up to one second or less and long bursts that last anywhere from 2 seconds to several minutes. Furthermore, Fishman and Meegan (1995) show that there is an anticorrelation between the spectral hardness and duration, the short ones being harder.

The detection of the afterglows, X-ray, and optical, of long-duration soft-spectrum GRB shows that they occur in star-forming regions with massive stars (van Paradijs et al., 2000). The kinetic energies of ejecta and electromagnetic emission associated with these GRB are comparable to those of core-collapse supernovae (see §1.1.2.1 for description), in particular for a rare class of supernovae with high-velocity ejecta. This suggests a link between the two phenomena (Gehrels et al., 2005; Galama et al., 1998; Matheson et al., 2003; Woosley and Bloom, 2006), and provides indirect evidence for the ‘collapsar’ hypothesis proposed as the origin of the bursts (MacFadyen and Woosley, 1999). In the collapsar hypothesis, the end result of the collapse of a massive star with significant rotation is a black hole with an accretion torus surrounding it. The bursts are the result of shocks in electron/positron pair plasma moving near the speed of light in an environment nearly devoid of neutrons, protons, and nuclei. The long time scale corresponds to the infall and fallback time of material from the stellar collapse (Woosley, 1993; MacFadyen and Woosley, 1999).

In contrast to the long-soft GRB, a growing body of evidence shows that the short-duration hard-spectrum GRB occur in regions where no star formation nor massive stars are present (Hjorth et al., 2005; Gehrels et al., 2005; Fox et al., 2005; Soderberg et al., 2006; Prochaska et al., 2006). The favored progenitors for these short-hard GRB are mergers of compact objects such as neutron stars and black holes (Rosswog, 2003; Belczynski et al., 2006; Lee and Ramirez-Ruiz, 2007). Neutrino-antineutrino pair annihilation (Meszaros and Rees, 1992; Woosley, 1993) and/or the generation of strong magnetic fields during the mergers (Mészáros and Rees, 1997; Meszaros, 2002) resulting in an energetic relativistic plasma are the possible candidates for the central engines of short-hard GRB.

Despite advances in the theories and models for the sources and central engines of GRB (see for examples Meszaros (2002); Piran (2005); Lee and Ramirez-Ruiz (2007) for reviews), the detailed mechanisms for both classes of GRB are still not properly understood. Numerical simulations of core-collapse supernovae and neutron star mergers (described further in §1.1.2.1 and §1.1.2.2) are needed to test the many assumptions of these models and give insight of the central engines that power the GRB.

1.1.1.2 Gravitational Waves

The existence of gravitational waves were first predicted by Albert Einstein in 1916 as a consequence of the general theory of relativity. In this theory, concentration of mass or energy produces curvature in spacetime, and accelerating masses produce ripples of spacetime that propagate with the speed of light called gravitational waves.

Although their existence have been indirectly confirmed (Taylor and Weisberg, 1989; Taylor, 1994), at the time of writing gravitational waves have not been directly detected because the weak coupling of gravitational waves to matter makes them very difficult to detect. However, a new generation of gravitational waves detectors such as LIGO (Abbott et al., 2004b), VIRGO (Acernese et al., 2005), TAMA300 (Takahashi and The TAMA Collaboration, 2004), and GEO-600 (Willke et al., 2004) are functioning and coming online. These are ground-based laser interferometers detectors. It is the expectation of researchers in this field that these waves will be detected in the near future (Frey, 2007; Willke, 2007).

One of the most promising gravitational waves sources for these ground-based detectors is neutron star mergers (Abramovici et al., 1992; Bradaschia, 1990; Luck, 1997). There are two reasons for this. Mergers of neutron stars involve two extremely dense objects in rapid orbit, producing strong gravitational waves in the frequency range 10 – 10000 Hz, the range covered by these broadband detectors. The second reason is the event rate of detectable neutron star mergers, which is predicted to be between 10 to 100 per year for advanced LIGO by population synthetic calculations (Belczynski et al., 2002; Nakar et al., 2006; Sadowski et al., 2008). Combined, these two reasons make neutron star mergers the likeliest candidates for gravitational waves detection.

Extracting gravitational wave signals from noisy background relies on a 'matched filtering' procedure where theoretical waveforms are compared against the interferometer data sets (Baumgarte et al., 2008). Several detectors have performed this search for signals of compact remnant mergers in their initial science runs (Tagoshi et al., 2001; Abbott et al., 2004a, 2005, 2006, 2010). The gravitational wave signals from merging neutron stars can be derived analytically and are well understood in the regime where the separation is large compared to the radii of the neutron stars. In this regime, the point-mass formula is a good approximation. However, at late times in the inspiral, when the separation becomes comparable to the radii of the stars, hydrodynamical effects become dominant and the gravitational wave signals deviate significantly from the point-mass approximation. Therefore detailed modelings from numerical simulations are required to produce the theoretical waveforms. Knowing the expected waveforms would also provide valuable information in guiding the design and tuning of future detectors (Harry et al., 2002; Mandel et al., 2008).

The weakly-interacting nature of gravitational waves makes them very difficult to detect, but also means that they are a very unique tool to study our universe, complementing the already long list of electromagnetic signals we employ. Because gravitational waves interact extremely weakly with matter, once produced they propagate without being absorbed or scattered. In cases where other (electromagnetic) signals cannot escape to carry information from the source, gravitational waves can. For example, combined with neutrinos, gravitational waves would be the only messengers carrying direct information from the core in a core-collapse supernova event. Combined studies of theoretical modelings and detection of these signals may then give us hints of the mechanisms of these events (for example, see Ott (2009); Yakunin et al. (2010) for recent studies of this nature). In another example, in a case of a neutron star merger followed by a formation of a black hole, gravitational waves could be the only messenger providing direct information from the new remnant. Furthermore, gravitational radiation observations will be from frequencies below 10 kHz, while observations from electromagnetic signals come from frequencies above 10 Mhz. This by itself could give new insights and is a very different way to study the universe.

1.1.2 Phenomena

1.1.2.1 Core-collapse Supernovae

Core-collapse supernovae are the violent death of massive stars and are among the most powerful explosions in universe, releasing about 10^{53} erg of energy on timescales of a few tens of seconds. This rivals the instantaneous power of all the rest of the luminous visible universe combined. They mark the birth of the most exotic states of matter known: neutron stars and black holes, while at the same time producing and disseminating most of the elements heavier than helium in nature, making life as we know it possible. These events occur about twice per century in a typical galaxy like our own and have been in the forefront research in the field of astronomy and astrophysics for almost half a century. Yet how exactly they work is still shrouded in mystery.

Stars burn hydrogen into helium for most of their existence. For stars more massive than ~ 10 solar masses ($M_\odot$), temperatures and densities are sufficiently high for burning to continue through carbon to oxygen, neon, magnesium, silicon, and to iron group elements. The star ends up in an onion-like configuration, with an iron core surrounded by layers of silicon, oxygen, carbon, helium, and hydrogen. Since the iron group elements are the most tightly bound, burning in the core ceases. At this point, the pressure in the core is dominated by electron degeneracy pressure (a consequence of the Pauli exclusion principle), which supports it against the inward pull of gravity. This balance between the gravitational pull and the electron degeneracy pressure in the core is only marginally stable.

Two processes occur in the core that result in the reduction of the degeneracy pressure support: electron capture on the free protons and nuclei, and nuclear dissociation under extreme densities and temperatures. The pressure support in the core is reduced enough that the core eventually becomes unstable and collapses. As the core collapses, the inner and outer regions behave differently. The inner core undergoes homologous collapse—velocity increases linearly with radius—as expected of fluid with relativistic, degenerate electron pressure. With increasing radius, the density decreases, and so thus the local sound speed. Thus, there is a radius where the speed of the infalling matter is the same as the local sound speed, demarcating the inner and outer core. Beyond this radius—the outer core—matter collapses supersonically.

The inner core collapses until it exceeds nuclear matter density ($\sim 1-3\times10^{14}$ g/cm^3). At this extreme density, the pressure of the inner core increases dramatically as a result of the repulsive component of the short-range nuclear force. The inner core becomes incompressible and bounces, and a shock wave forms at the boundary of the inner and outer core and begins to move out. Ultimately this shock wave will be responsible in the disruption of the star, producing the observable explosion.

It was once thought that as the shock wave propagated outward, the velocity of the bounce would grow as it moved into the outer layers of the core; the bounce is therefore the origin of the supernova's energy (Colgate and Johnson, 1960; Baron et al., 1985). From all the realistic models completed to date, we now know that this is not the case, and therein lies the core-collapse supernova problem.

As the shock propagates out, it has to move through infalling materials in the outer core, during which nuclear dissociation happens. This costs the shock energy. Additional energy losses occur when the core electrons capture on the newly available free protons because of the nuclei dissociation, which eventually results in electron neutrino burst in a core-collapse supernova. As a result of these energy losses, the shock stalls.

If this were the end of the story, no supernova would ever explode nor observed. The shock has to be reenergized so that it may continue to propagate outward and eventually produces the explosion. The details of how the stalled shock is revived is the central question in core-collapse supernova theory.

Out of the 10^{53} erg energy released during the explosion, the visible explosion energy is only 1%. The rest is released as neutrinos. Because neutrinos dominate the energetics of a supernova event, it is natural to consider neutrino heating as a mechanism for the revival of the stalled shock. This delayed neutrino-heating has been proposed as one of the mechanisms that leads to explosion (for example, see Mezzacappa (2005) and Janka et al. (2007) for reviews).

Core-collapse supernovae are asymmetric events. Observational evidence that has accumulated to support this includes spectropolarimetry (Wang et al., 2001a; Leonard et al., 2006), the large average pulsar velocities (Hansen and Phinney, 1997; Zou et al., 2005; Chatterjee et al., 2005), and the morphology of highly resolved images of supernova such as SN 1987A (Arnett et al., 1989; McCray, 1993). On the theoretical side, simulations have shown that a variety of asymmetric fluid instabilities are present. These instabilities develop convective overturn and help transport hot gas from neutrino-heating region directly to the shock, therefore enhancing the neutrino energy deposition to the stalled shock (Janka and Müller, 1996; Burrows et al., 1995; Herant et al., 1994; Mezzacappa et al., 1998; Buras et al., 2006). These multidimensional effects therefore may be important for the neutrino-heating mechanism to revive the stalled shock. Recent simulations have also revealed the existence of standing accretion shock instability (SASI), which given enough time, may also grow via the propagation of sound waves (Blondin et al., 2003; Blondin and Mezzacappa, 2006). All these multidimensional effects may play essential roles in the mechanisms of core-collapse supernovae.

Stars have both rotation and magnetic fields. It has been suggested that in more massive progenitors rotation and magnetic fields may play a more significant role (Thompson et al., 2005; Fryer and Heger, 2000; Wheeler and Akiyama, 2006), producing jet-like hypernovae, and perhaps giving birth to 'magnetars', a type of neutron star with unusually large magnetic field. Observational evidence seems to support this (Gaensler et al., 2005; Figer et al., 2005). Even in normal supernovae there seems to be observational evidence that rotation and magnetic field play some roles (Burrows et al., 2004). During collapse magnetic fields may also be amplified enough to have important dynamical effects. Recently we discovered that an amplification of magnetic fields can happen in SASI (as discussed in chapter 6), therefore extending the range of progenitors in which magnetic field may play a significant dynamical role.

Because of the complexity of the supernova mechanisms, a purely analytical investigation is not possible. Instead, supernova modeling requires sophisticated and realistic numerical simulations. All the input physics required to model core-collapse supernovae present daunting challenges that are both algorithmic and computational in nature, and tax state-of-the-art supercomputers for years to come. We describe some of these challenges in §1.2 and our plans to meet them.

1.1.2.2 Binary Neutron Star Mergers

A neutron star is one possible outcome of stellar evolution, the very compact and dense object left behind by core-collapse supernova (see also §1.1.2.1) (Fryer, 2004; Woosley and Janka, 2006). A neutron star, as the name implies, is composed largely of neutrons. It is supported by degeneracy pressure due to Pauli exclusion principle against further collapse. As it ages, it cools via neutrino and photon emissions (Prakash et al., 2001).

A typical neutron star has roughly $1 - 1.5\ \mathrm{M}_\odot$ compressed into an object of only about 20 to 10 km in radius (Lattimer and Prakash, 2004). Most neutron stars were initially discovered as solitary objects known as pulsars (rotating neutron stars emitting periodic electromagnetic pulses). It was not until 1974 that a neutron star binary was discovered (two neutron stars in mutual orbit) by Hulse and Taylor (1975), a discovery which led to a

Nobel Prize. Since then, additional discoveries of binary neutron stars in our galaxy have followed (Dewey et al., 1985; Anderson et al., 1990; Wolszczan, 1991).

Widely separated neutron stars inspiral, driven primarily by the loss of energy and momentum due to gravitational wave emissions, as predicted by Einstein's general theory of relativity. Close neutron stars binaries, for which the merger timescale is smaller than a Hubble time, typically start out as a high-mass binary systems, with both stars having masses greater than $8 - 10\ \mathrm{M}_{\odot}$ to ensure that both stars eventually become core-collapse supernovae. In such systems, the more massive (primary) star evolves faster. After leaving the main sequence and passing through the giant phase, it becomes a core-collapse supernova and eventually forms the heavier compact object (i.e. neutron star). The less massive (secondary) star follows the same evolutionary path, but when it reaches the giant phase, a common envelope is formed around the primary star. This causes dynamical friction that shrink the orbital separation dramatically (Belczynski et al., 2008). Sufficient potential energy is eventually converted into thermal energy to evaporate the envelope. This step is necessary for binaries to merge within a Hubble timescale, since loss of energy via emissions of gravitational wave alone is too slow to drive the inspiral. The secondary object eventually becomes a core-collapse supernova. Depending on the orientation and magnitude of the supernova kick, either a tight binary is left behind or a complete unbinding of the system occurs (Taam and Sandquist, 2000; van Den Heuvel, 2006; Voss and Tauris, 2003; Hobbs et al., 2005).

Binary neutron star mergers are interesting to astrophysicists because of their association with two observables already mentioned above. They are the candidates for the short-hard gamma-ray burst (GRB) progenitors, and the prodigious sources of gravitational waves that are most promising for direct detection. Neutron star merger simulations are therefore critical to improve our understanding of these aspects of observable phenomenology. Simulations would shed lights on the viabilities of the proposed GRB central engines (see §1.1.1.1). Simulations would also produce gravitational wave templates necessary for the matched filtering procedure—comparation of noisy data against theoretical waveforms—in the search for gravitational wave with detectors (see also §1.1.1.2).

Due to the complex and large-scale nature of the problem, computer simulations of neutron star mergers have been performed with various simplifications. For many years the study of neutron star mergers has largely proceeded along two lines: concentration on improving the microphysics while retaining Newtonian gravity, and concentration on improving the treatment of gravity while using simplified microphysics.

In the former approach, simulations have been done using Newtonian gravitation with a 'realistic' nuclear equation of state, with the back-reaction from gravitational wave emission usually taken into account (Ruffert et al., 1996; Ruffert and Janka, 2001a,b; Rosswog and Davies, 2002; Rosswog and Liebendörfer, 2003). Hydrodynamics codes based on the 'piecewise parabolic method' (Ruffert et al., 1996; Ruffert and Janka, 2001a,b) or 'smoothed particle hydrodynamics' (SPH) (Rosswog and Davies, 2002; Rosswog and Liebendörfer, 2003) have been used, with the energy loss and composition changes due to neutrino emission treated using a 'leakage' scheme, rather than full neutrino transport. Recent simulations using SPH have taken magnetic field into account and reported its amplification due to the Kelvin-Hemholtz instability in the shear layer between the neutron stars as they begin to merge (Price and Rosswog, 2006). This may have implications for the production of short-hard GRB.

The latter approach has focused on fully general relativistic simulations with simplified microphysics (Miller et al., 2004; Shibata and Uryu, 2002; Shibata et al., 2003). A simpler polytropic equation of state is used in these codes rather than a realistic nuclear equation of state. In this case energy loss and composition changes due to neutrino emission are largely ignored. In similar vein, an approximation to the full general relativity called the 'conformally flat approximation' (CFA) (Wilson et al., 1996) is used with SPH by Oechslin et al. (2002) and Faber et al. (2004)

Recently efforts have been made to bridge the chasm between these two approaches. Simulations with CFA have been done using realistic equation of state (Oechslin and Janka, 2006; Oechslin et al., 2007). Another group have also modeled mergers in full general relativistic simulations using zero-temperature equation of state (Shibata and Taniguchi, 2006; Shibata et al., 2005). The evolution of magnetic fields has also been considered in general relativistic simulations (Anderson et al., 2008; Liu et al., 2008; Giacomazzo et al., 2009), with efforts made to study the effects of magnetohydrodynamics instabilities (Shibata et al., 2006a; Duez et al., 2006) and neutrino energy transport (Setiawan et al., 2006; Lee et al., 2005).

Beyond these recent studies in the modeling of neutron star mergers, much work still needs to be done to improve our understanding of the merger events. Questions regarding the central engines and mechanisms of short-hard gamma-ray bursts remain unanswered, and the viabilities of the proposed mechanisms need to be explored. More accurate microphysics for GRB modeling is still needed, including better approximations, if not the full solution, to the neutrino transport problem. For the purpose of providing templates for the search of gravitational waves, more simulations need to be done covering larger parameter space to provide waveform catalogs for detectors data analysis. To answer these challenges, we intend to complement and advance the recent studies of neutron star mergers with our own. Toward that goal, we have taken the first steps in building a code suitable to perform simulations of neutron star mergers with realistic physics input.

1.2 Computational Challenges

From the discussions above, it is clear that both core-collapse supernovae and neutron star mergers are incredibly complex events where input from all major fields of modern physics are at play. Therefore, proper modelings of these events require a code with the input of multi-physics such as hydrodynamics, magnetic fields, self-gravity, relativity, radiation transport, and nuclear physics. The complexity of building such codes is challenging and therefore requires modern software engineering practices to be applied thoughtfully. Both classes of these phenomena however share the same challenges and involve related physics. Therefore, a versatile simulation system would be able to explore both those problems, maximizing the return of investment in building such system.

To meet these challenges, we have started the development of a code named *GenASiS*. GenASiS, which stands for *Gen*eral *A*strophysical *Si*mulation System, is a multidimensional self-gravitating radiation magnetohydrodynamics code with adaptive mesh refinement. GenASiS is currently written in Fortran 95 standard, with an upgrade planned to use the Fortran 2003 standard as it matures. The world "General" in GenASiS refers to the design modularity, using the object-oriented principle facility in the Fortran 95/2003 standard to allow for function overloading. In object-oriented programming, this is called *polymorphism*, which allows us to use a generic name with several different implementations, thus providing extensibility of the physics. New implementations of different equations of state, hydrodynamics scheme, coordinate system, gravity theory, and so forth can therefore be done without having to go back and change the basic parts of the code. GenASiS is also designed for scalability: the ability to use large number of processes with distributed-memory parallel computers. These two philosophies guide the development of and design choices in GenASiS, as will be described in the rest of this document.

Magnetic fields likely play significant roles in both core-collapse supernovae and neutron star mergers events. In neutron star mergers, magnetic fields may be amplified due to various instabilities and act as a central engine for gamma-ray bursts. In the context of core-collapse supernovae, magnetic fields may be generated, amplified, and affect the dynamic of the supernova explosion following the collapse. Therefore, instead of a mere hydrodynamics, both events require *magneto*hydrodynamics for proper modelings. In chapter 2 we describe the implementation of magnetohydrodynamics in GenASiS. The description of the numerical schemes is followed by test problems to verify the correctness of the implementation. We also show a scalability test for the MHD module, essential for practical execution of large scale simulations such as the ones described in chapter 5 and 6.

In chapter 3 we describe an implementation to solve Poisson's equation using FFT for an isolated system. The method solves the equation globally on mesh blocks distributed across multiple processes on a parallel computer. Test results to demonstrate the correctness of the method are presented. We also show the scaling properties of the method. As in

MHD, weak scaling for the solver is essential for large-scale simulations. The modularity in GenaSiS architecture allows us to decouple this solver module from the rest of the code with ease. We plan to release this solver as a freely-available software library as our contribution to the community (Budiardja et al., 2010).

In GenASiS, we have chosen to work primarily with multidimensional Cartesian coordinate systems in grid-based Eulerian formulation (e.g. LeVeque et al. (1998); Laney (1998)). This choice is motivated by several reasons. Many of the previous simulations of neutron star mergers have been done in Lagrangian formulation of gas particle using smoothed particle hydrodynamics (SPH) Monaghan (1992). SPH has been used extensively because of its adaptivity in resolution, where particles can be easily concentrated and adjusted with respect to variables such as the density. This seems ideal for problems such as neutron star mergers where densities are concentrated in the body of the stars surrounded by largely empty space. We have chosen to use grid-based Eulerian formulation instead to complement the previous studies and help build the confidence of the community in the results of previous simulations. It has been shown that grid-based Eulerian method are more accurate in resolving resolving dynamical instabilities such Kevin-Helmholtz or Rayleigh-Taylor (Agertz et al., 2007) and capturing shocks (Quilis, 2000) as compared to SPH.

Our choice of a Cartesian coordinate system avoids complications associated with using spherical coordinate system—commonly used in core-collapse supernova simulations—because of coordinate singularities. Coordinate singularities in curvilinear coordinates enforce very small simulation time steps due to Courant-Friedrichs-Lewy condition (see 2.1.1). Furthermore, it makes more sense in terms of code development to use Cartesian coordinate systems since there is no spherical symmetry to be exploited in the neutron star merger problem. Thus neutron star merger simulations do not lend themselves well with curvilinear coordinate systems. Out choice of a Cartesian coordinate system positions us to have a code well-suited to explore both core-collapse supernovae and neutron star mergers problems.

Both core-collapse supernovae and neutron star mergers problems are inherently multidimensional, covering large dynamical range in length scales. In the core-collapse supernovae, density increases of six orders of magnitude occur during the collapse, which must be properly resolved. In particular, adequate resolution of features of the flow (shocks, for example) is a necessity. In the context of neutron star mergers, there are different length scales that must be properly resolved, such as the individual stars, possibly including their internal dynamics; the orbital length scale; and the location of outer boundaries that must be large enough so that they do not have dynamical effects on the mergers. To best meet these needs, we incorporate the development of 'adaptive mesh refinement' (AMR) in GenASiS. The basic idea of AMR is to employ high resolution only where needed. This allows conservation of memory and computational effort. Although AMR implementation in our code is still in its infancy, in chapter 4 we describe and show the progress that has been made with its development. A novel multigrid algorithm to solve Poisson's equation in mesh with refinement is also described there, laying groundwork for future simulations with AMR.

The code description in chapters 2 – 4 is followed by two scientific applications. In chapter 5 we reports a merger simulation that we have done with GenASiS as a milestone of the code development. We show the tools we have developed to extract and analyze gravitational wave signatures from the merger, which are suitable as templates for gravitational wave detectors.

In chapter 6 we show results from simulations of SASI with magnetic fields in the context of core-collapse supernovas. We found new results showing that initially modest magnetic field strength in supernova progenitors may be amplified by up to several orders of magnitude as the non-linear mode of SASI develops. Our analysis explains the amplification mechanisms for the generations magnetic fields in this post-bounce supernova environment (Endeve et al., 2010). These simulations show that GenASiS, although still very much in development, is already capable of performing scientific simulations.

We finish this manuscript with some concluding remarks and an outlook for future work in chapter 7.

Chapter 2

Fluid and Magnetic Field Evolution

In GenASiS, we have implemented a second-order, semi-discrete central-upwind scheme for the solution of the equations of Newtonian and relativistic magnetohydrodynamics (MHD) in their conservative forms (Kurganov et al., 2001; Londrillo and Del Zanna, 2000; Del Zanna and Bucciantini, 2002; Del Zanna et al., 2003). The time-dependent evolution of the equations is integrated with a time-explicit total variation diminishing (TVD) two-step Runge-kutta scheme (Shu and Osher, 1988; Kurganov and Tadmor, 2000) to achieve second-order temporal accuracy.

The conservative form of the fluid part (coupled with magnetic fields) of the MHD equations implies that the rate of change of conserved quantities in an infitesimal volume is equal to the flux through the surface of the volume. Accurate computation of fluxes at every cell surface is therefore a key to this *finite-volume* method. We use a variation of the so-called “central schemes” (Einfeldt, 1988; Harten et al., 1983; Kurganov and Tadmor, 2000; Kurganov et al., 2001) to calculate the fluxes with “HLL” formulae given by Del Zanna and Bucciantini (2002). The central schemes have been noted for their ability to achieve comparable accuracy to other Goudonov-type Riemann solvers but with much greater simplicity and less costly computation (Lucas-Serrano et al., 2004). This finite-volume approach with the HLL central scheme handles shocks and smooth flows of the hydrodynamic evolution.

The evolution of the magnetic fields in the multidimensional (relativistic) MHD system is irreducible to a conservative form since the magnetic field, being a vector, is advanced in time by the curl operator rather than the divergence operator like the other scalar variables. This leads to the induction equation for the magnetic field $\mathbf{B}$ and the non-evolutionary divergence-free condition: $\nabla \cdot \mathbf{B} = 0$. The method of *constrained transport* (CT) by Evans and Hawley (1988) preserves the divergence-free condition by evaluating the rate of change of the magnetic flux on a cell face via the discretization of the curl of the electric field. By Faraday's law, the line integral of the electric field around an infinitesimal surface's boundary is equal to the rate of change of the flux through the surface and therefore, applied to a closed surface, the magnetic flux mush vanish, maintaining the divergence-free condition. Evaluating the electric field on the cell edges accurately is therefore a key to the CT method. To that end, we also use an HLL central scheme to evaluate the electric field on the edges that define the cell surface (Del Zanna et al., 2003; Londrillo and Del Zanna, 2004), which has been shown to handle both shocks and smooth flows of the magnetic field evolution.

In this chapter describe we describe the numerical schemes to solve the equations of ideal magnetohydrodynamics and our implementations in GenASiS. In §2.1 the HLL central scheme for Newtonian and relativistic magnetohydrodynamics is described briefly (further details on the HLL central scheme may be found on the references previously mentioned). This is followed by the description of equations of state in §2.2. In §2.3 we describe the parallelization of the algorithms and show performance and scaling with an increasing number of processes. The numerical tests of our implementations are described in chapter 2.4. In GenASiS, Newtonian and relativistic MHD are implemented as two different modules since the hydrodynamics equations could be recovered simply by setting the magnetic fields to zero. However, for the purpose of the presentation we start by describing the scheme in Newtonian hydrodynamics and its relativistic counterpart. The Newtonian and relativistic magnetohydrodynamics then follows as an extension to the hydrodynamics equations.

2.1 Numerical Schemes

2.1.1 Newtonian Hydrodynamics

The ideal Newtonian hydrodynamics (HD) equations, without sources, in conservative Eulerian form are (Landau and Lifshitz, 1959):

$$\begin{aligned}
\frac{\partial D}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^i\right) &= 0, && (2.1)\\
\frac{\partial S^j}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^j v^i + p\delta^{ij}\right) &= 0, && (2.2)\\
\frac{\partial E}{\partial t} + \frac{\partial}{\partial x^i}\left(\left[e + p + \frac{1}{2}\rho v^j v^j\right] v^i\right) &= 0. && (2.3)
\end{aligned}$$

These are respectively mass conservation, momentum conservation, and energy conservation equations. Einstein's summation convention is used here and throughout, with Greek indices to indicate four-vectors and Latin indices to denote spatial three-dimensional vectors. Note that equation 2.2 collapses three equations in three separate spatial dimensions to one. In equations 2.1 - 2.3, ρ, v^i, e, and p represent rest mass density, fluid velocity, internal energy density, and fluid pressure, respectively. The time-evolved quantities are the conserved density D, conserved energy density E, and conserved momentum density S^j. In Newtonian HD, these variables are defined as:

$$\begin{aligned}
D &= \rho, && (2.4)\\
S^j &= \rho v^j, && (2.5)\\
E &= e + \frac{1}{2}\rho v^j v^j. && (2.6)
\end{aligned}$$

The system of equations has to be closed with an equation of state specifying the pressure in terms of density (with temperature and composition in more sophisticated equations of state). For an ideal gas, we use a polytropic equation of state of the form

$$p = \kappa \rho^{\Gamma}, \tag{2.7}$$

where Γ is the adiabatic index and κ is a to-be-determined polytropic constant. This equation of state implies the relation

$$e = \frac{p}{\Gamma - 1}. \tag{2.8}$$

The conservation equations 2.1 - 2.3 may be cast as hyperbolic conservation laws of the form

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{f} = 0, \tag{2.9}$$

where $\mathbf{u} = \left[D, S^j, E\right]$ is the vector of *conserved* variables, with corresponding fluxes $\mathbf{f}^i$, defined as

$$\begin{aligned} \mathbf{u} &= \left[D, S^j, E\right], && (2.10) \\ \mathbf{f}^i &= \left[\rho v^i, \rho v^j v^i + p\delta^{ij}, \left(e + p + \frac{1}{2}\rho v^j v^j\right) v^i\right], && (2.11) \end{aligned}$$

where we call the variables $\left[\rho, v^i, e\right]$ *primitive* variables. At each time step during the evolution, the primitive variables have to be recovered from the conserved ones by inverting equations 2.4 - 2.6, which is trivial in Newtonian HD.

In integral form over the finite volume V of a particular cell, equation 2.9 becomes

$$\int \left(\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{f}\right) dV = 0. \tag{2.12}$$

By the divergence theorem $\int_V dV\,(\nabla\cdot\mathbf{f}) = \oint \mathbf{f}\cdot d\mathbf{A}$ on the cell of volume V and surface area A, this becomes

$$\begin{aligned} \frac{\partial\langle u\rangle}{\partial t} &= -\frac{1}{V}\oint \mathbf{f}\cdot d\mathbf{A} && (2.13)\\ &= -\frac{1}{V}\left[\left(A_i f^i\right)_{i\rightarrow} - \left(A_i f^i\right)_{i\leftarrow}\right], && (2.14) \end{aligned}$$

where $\langle u\rangle$ is a cell-centered, volume-average value of the conserved variable and A is the area of the cell face. Here the left arrow $(\leftarrow)$ and right arrow $(\rightarrow)$ indicate evaluation on the cell inner and outer faces in the i direction, respectively.

The primitive and conserved variables are defined as cell-centered values. During the flux computation their values are needed on the cell interfaces and must be reconstructed from the the cell-centered values. We use slope-limited linear interpolation inside our computational cells to reconstruct the primitive variables on the appropriate cell faces. The slopes are second-order in spatial accuracy for regions with smooth flows, but may reduce to first-order accuracy while maintaining non-oscillatory behavior near shocks and discontinuities. Specifically, we use a one-parameter family of *generalized MinMod* limiters (Kurganov and Tadmor, 2000). For an arbitrary variable χ inside a cell whose center is indexed by (i,j,k), the slope in (for example) the x-direction is defined by

$$\left.\frac{\partial\chi}{\partial x}\right|_{ijk} = \text{minmod}\left[\theta\left(\frac{\chi_{ijk}-\chi_{i-1jk}}{x_i - x_{i-1}}\right), \left(\frac{\chi_{i+1jk}-\chi_{i-1jk}}{x_{i+1}-x_{i-1}}\right), \theta\left(\frac{\chi_{i+1jk}-\chi_{ijk}}{x_{i+1}-x_i}\right)\right], \quad (2.15)$$

where the multi-variable minmod function is defined as

$$\text{minmod}\left(\chi_1,\chi_2,\ldots\right) = \begin{cases} \min_j\left(\chi_j\right), & \text{if } \chi_j > 0\,\forall j,\\ \max_j\left(\chi_j\right), & \text{if } \chi_j < 0\,\forall j,\\ 0 & \text{otherwise}, \end{cases} \quad (2.16)$$

where θ is the slope weighting parameter and $\theta\in[1,2]$. $\theta = 1$ is the most dissipative and reduces the limiter to the original *MinMod* (MM) limiter (Roe, 1986), while $\theta = 2$ is the least dissipative and reduces the limiter to the *Monotonized Central* (MC) limiter (van

Leer, 1977). Note than in our implementation only primitive variables are interpolated to the cell interfaces. All other quantities needed on the cell interfaces are recomputed from the reconstructed primitive variables. In our experience, this technique gives better results and is more robust than interpolating all of the conserved and primitive variables to the cell interfaces.

Once primitive variables are reconstructed on the cell interfaces as face-centered values, the *raw* fluxes and conserved variables on all cell faces are computed using equations 2.4 - 2.6 and 2.11. Each cell face then has raw fluxes and conserved variables to the left and to the right of the face, respectively, labeled $\mathbf{f}^{\mathrm{L}}$, $\mathbf{f}^{\mathrm{R}}$, $\mathbf{u}^{\mathrm{L}}$, and $\mathbf{u}^{\mathrm{R}}$. The flux through a particular cell face is given by the HLL central scheme formula (Del Zanna and Bucciantini, 2002):

$$\mathbf{f}^{\mathrm{HLL}} = \frac{\alpha^{+}\mathbf{f}^{\mathrm{L}} + \alpha^{-}\mathbf{f}^{\mathrm{R}} - \alpha^{+}\alpha^{-}\left(\mathbf{u}^{\mathrm{R}} - \mathbf{u}^{\mathrm{L}}\right)}{\alpha^{+} + \alpha^{-}}. \tag{2.17}$$

The coefficients $\alpha^{\pm}$ take into account the highest speeds of the Riemann fan at the cell interface, and can be estimated from the maximum and minimum characteristic speeds (eigenspeeds) $\lambda^{\pm}$ at the two reconstructed states as

$$\alpha^{\pm} = \max\left\{0, \pm\lambda^{\pm}\left(\mathbf{v}^{\mathrm{L}}\right), \pm\lambda^{\pm}\left(\mathbf{v}^{\mathrm{R}}\right)\right\}, \tag{2.18}$$

where, in Newtonian HD

$$\lambda^{\pm} = v \pm c_s. \tag{2.19}$$

The sound speed is given by $c_s = \sqrt{\Gamma p/\rho}$ for a polytropic equation of state. Using only the two fastest characteristic speeds ensures that shocks are handled correctly, although contact discontinuities and shear waves, which corresponds to intermediate eigenspeeds, can be more smeared compared to results from more exact Riemann solvers.

Having obtained the right-hand side of equation 2.14, we can finally integrate in time. To achieve a high-accuracy time-stepping integration, we use a second-order TVD Runge-Kutta scheme (Shu and Osher, 1988):

$$u^{(1)} = u^n + \Delta t L(u^n) \tag{2.20}$$

$$u^{(n+1)} = \frac{1}{2}u^n + \frac{1}{2}u^{(1)} + \frac{1}{2}\Delta t L(u^{(1)}) \tag{2.21}$$

where the operator $L(u^n)$ denotes the spatial differencing at the right-hand side of equation 2.14. The time interval Δt must obey the Courant-Friedrichs-Lewy (CFL) condition given by

$$v_c \cdot \Delta t < C\Delta x, \tag{2.22}$$

where v_c is the maximum of the characteristic speed on the grid

$$v_c = \max\left\{+\lambda_i^+, -\lambda_i^-\right\}, \tag{2.23}$$

Δx is the cell spacing, and C is the Courant parameter. In most cases, we use $C < 0.5$. The CFL condition can be loosely interpreted that information can only propagate a certain fraction of a cell size in a time step for the integration algorithm to be stable.

2.1.2 Relativistic Hydrodynamics

Relativistic flows and shocks are common in modern simulations of high energy astrophysics. Phenomena such as gamma-ray bursts, X-ray bursts, and jets in active galactic nuclei, are all examples where relativistic flows and shocks play essential roles. Therefore, the extension of Newtonian hydrodynamics to its relativistic counterpart, and its implementation in GenASiS, is well motivated. The same overall scheme described in §2.1.1 can be used to obtain the solution to the equations of relativistic hydrodynamics (RHD). Here we note some only differences that appear in relativistic case.

Landau and Lifshitz (1959) gives the equations of RHD in covariant form:

$$\partial_\alpha (\rho u^\alpha) = 0, \tag{2.24}$$

$$\partial_\alpha \left(w u^\alpha u^\beta + p g^{\alpha\beta} \right) = 0. \tag{2.25}$$

We have chosen to use geometrized unit throughout in which $G = c = 1$, where G is Newton's gravitational constant and c is the speed of light. For simplicity, Minkowskian flat space is assumed with $g^{\alpha\beta} = \mathrm{diag}\{-1,1,1,1\}$, with the coordinates $x^\alpha = (t, x^j)$. The variables have the same meaning as before, with the addition of the relativistic enthalpy $w = \rho + e + p$ and the four-velocity $u^\alpha = (\gamma, \gamma v^j)$, where $\gamma = (1 - v^2)^{-1/2}$ is the Lorentz factor. An equation of state is necessary to close the system of equations.

The RHD equations 2.24 and 2.25 may be cast into hyperbolic conservation laws in the form of equation 2.9 by defining the vector of conserved variables with their corresponding fluxes as:

$$\mathbf{u} = \left[D, S^j, E \right], \tag{2.26}$$

$$\mathbf{f}^i = \left[\rho\gamma v^i, w\gamma^2 v^i v^j + p\delta^{ij}, w\gamma^2 v^i - \gamma\rho v^i \right], \tag{2.27}$$

where the conserved variables are

$$D = \rho\gamma, \tag{2.28}$$

$$S^j = w\gamma^2 v^j, \tag{2.29}$$

$$E = w\gamma^2 - p - \rho\gamma, \tag{2.30}$$

with ρ, v^j, and e as the primitive variables.

The primitive variables must be recovered from the conserved variables at lease once in per time step. This is trivial in the case of Newtonian HD, but not in the RHD case, where the equations for the conserved variables form a non-linear system of equations to be inverted for the primitive variables. Another complicating factor is that in the general case the equation of state may not have an analytic form such as the ideal gas (polytropic) equation of state, but may be given in tabular form instead (see §2.2). Therefore, to invert the conserved variables to the primitives we employ an iteration procedure that does not depend on the form of the equation of state as follows.

Equations 2.28 - 2.30 may be combined to form a non-linear equation

$$(E + p' + D)(1 - \gamma^{-2}) - S^i S^i = 0, \tag{2.31}$$

where E, D, and S^i are known. Here p' as the initial guess for the pressure. It is set to the value of p from the previous known state. Equation 2.31 is solved using a numerical non-linear root-finder for γ. We use the Newton-Raphson method *rtsafe* found in *Numerical Recipes* by Press et al. (1986). Given γ, we can solve for the other primitive variable ρ, e, and v^i, in that order. A new pressure $p(\rho, e)$ is then computed using the equation of state and compared to p'. If the relative error $\varepsilon = (p' - p)/p$ is above a certain threshold, set $p' = p$ and reiterate, otherwise, quit. We set this threshold to 1×10^{-7}.

For relativistic flows, the eigenspeeds required for the flux computation are given according to the relativistic rule for velocity vector addition, after splitting the velocity to its components that are parallel and perpendicular to the direction of the spatial dimension (Del Zanna and Bucciantini, 2002):

$$\lambda^{\pm} = \frac{v_{\parallel}(1 - c_s^2) \pm c_s \sqrt{(1 - v^2)(1 - v_{\parallel}^2 - v_{\perp}^2 c_s^2)}}{1 - v^2 c_s^2}. \tag{2.32}$$

This concludes the modifications needed for the scheme described in §2.1.1 to solve the RHD equations.

2.1.3 Newtonian Magnetohydrodynamics

Equations 2.1 - 2.3 may be extended to include the magnetic fields in order to describe the evolution of the magnetized fluid. The equations of Newtonian ideal (zero resistivity) MHD, without sources, are Landau et al. (1984):

$$\begin{aligned} \frac{\partial D}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^i\right) &= 0, \qquad (2.33) \\ \frac{\partial S^j}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^j v^i + p\delta^{ij} - B^j B^i\right) &= 0, \qquad (2.34) \\ \frac{\partial E}{\partial t} + \frac{\partial}{\partial x^i}\left(\left[e + p + \frac{1}{2}\rho v^j v^j + \frac{1}{2}B^j B^j\right] v^i - B^i(\mathbf{B}\cdot\mathbf{v})\right) &= 0, \qquad (2.35) \\ \frac{\partial \mathbf{B}}{\partial t} + \nabla \times \mathbf{E} &= 0, \qquad (2.36) \\ \nabla \cdot \mathbf{B} &= 0. \qquad (2.37) \end{aligned}$$

Equations 2.33 - 2.35 are the conservation equations for mass, momentum, and energy for the magnetized fluid. Equation 2.36 is the induction equation governing the evolution of the magnetic field $\mathbf{B}$, where the electric field is given by $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$. The magnetic divergence-free equation 2.37 has to be satisfied at all times by the numerical method.

The conservation equations for the fluid part of the MHD equations retain the form of equation 2.9 if we define the conserved variables and corresponding fluxes as

$$\begin{aligned} \mathbf{u} &= \left[D, S^j, E\right], \qquad (2.38) \\ \mathbf{f}^i &= \left[\rho v^i, \rho v^j v^i + p\delta^{ij} - B^j B^i, \left(e + p + \frac{1}{2}\rho v^j v^j + \frac{1}{2}B^j B^j\right) v^i - B^i(\mathbf{B}\cdot\mathbf{v})\right], (2.39) \end{aligned}$$

where the conserved variables are defined as

$$\begin{aligned} D &= \rho, \qquad (2.40) \\ S^j &= \rho v^j, \qquad (2.41) \\ E &= e + \frac{1}{2}, \left(\rho v^j v^j + B^j B^j\right), \qquad (2.42) \end{aligned}$$

with ρ, v^j, and e as the primitive variables.

The magnetic part of the MHD equations requires a different treatment. Applying Stokes' theorem on the magnetic induction equation 2.36 gives us the integral form

$$\frac{\partial \mathbf{B}}{\partial t} = -\frac{1}{A}\oint_c \mathbf{E}\cdot dl, \tag{2.43}$$

where l is the cell-edge boundary for the cell surface of area A. Summed over all faces of a cell, two line integrals in opposite directions cancel on every cell edge, therefore satisfying the divergence-free constraint of equation 2.37. In the method of constrained transport, we evaluate the rate of change of magnetic flux on a cell face with the discrete correspondence of equation 2.43 as

$$\frac{\partial \langle B_i \rangle}{\partial t} = -\frac{1}{A_i}\sum_{j\neq i}\left[(l_k E^k)_{j\rightarrowtail} - (l_k E^k)_{\leftarrowtail j}\right]. \tag{2.44}$$

Here $\langle B_i \rangle$ is the face-centered surface average located at cell interfaces in the i direction. This is the fundamental property of the CT method, which relies on the definition of staggered field components $\langle B_i \rangle$, $\langle B_j \rangle$ and $\langle B_j \rangle$ as primary data located on the cell interfaces in the direction i, j, k respectively. The left and right arrows with tails, $(\leftarrowtail j)$ and $(j\rightarrowtail)$, denote evaluation along the inner and outer edge of the face in j direction, respectively. The index j runs over two spatial dimension orthogonal to i, while k indicates the direction orthogonal to both i and j. This discretization satisfies the divergence-free condition by construction.

The electric field component E_k on the edge along the k direction is computed using the HLL formula (Del Zanna et al., 2003)

$$\begin{aligned} E_k^{\mathrm{HLL}} &= \frac{\alpha_i^+\alpha_j^+E_k^{\mathrm{LL}} + \alpha_i^+\alpha_j^-E_k^{\mathrm{LR}} + \alpha_i^-\alpha_j^+E_k^{\mathrm{RL}} + \alpha_i^-\alpha_j^-E_k^{\mathrm{RR}}}{\left(\alpha_i^+ + \alpha_i^-\right)\left(\alpha_j^+ + \alpha_j^-\right)} \\ &\quad + \frac{\alpha_i^+\alpha_i^-}{\alpha_i^+ + \alpha_i^-}\left(B_j^{\mathrm{R}} - B_j^{\mathrm{L}}\right) - \frac{\alpha_j^+\alpha_j^-}{\alpha_j^+ + \alpha_j^-}\left(B_i^{\mathrm{R}} - B_i^{\mathrm{L}}\right), \end{aligned} \tag{2.45}$$

where i, j, k are orthogonal to each other and $i \neq j \neq k$. The values for E_k^{LL}, E_k^{LR}, E_k^{RL}, and E_k^{RR} are computed using the formula for the electric field $\mathbf{E} = -\mathbf{v}\times\mathbf{B}$, with the values for

v and **B** defined in the four cells surrounding the edge, as indicated by the two subscripts to denote the left (L) or right (R) face with respect to the edge in the first and second direction. The values for **v** and **B** have to be reconstructed from their cell-centered and face-centered values to the edge using linear interpolation with a slope limiter. Note that for cell-centered values the reconstruction to an edge takes two independent linear interpolations in two different directions. The eigenspeeds ($\alpha_i^{\pm}$) are also computed from the characteristic speeds of the reconstructed states on the edge, which in the MHD system is given by

$$\lambda_i^{\pm} = v_i \pm \left[\frac{1}{2} \left(c_s^2 + c_a^2 + \sqrt{(c_s^2 + c_a^2)^2 - \frac{4c_s^2 B_i^2}{\rho}} \right) \right]^{1/2}, \tag{2.46}$$

where c_a is the speed of Alfvén wave, defined as:

$$c_a = \sqrt{\frac{B_i B^i}{\rho}}. \tag{2.47}$$

The calculation of the fluid fluxes requires all components of the magnetic field on the cell interfaces. To accomplish this, we define the cell-center value of the magnetic field components as

$$(B_i)_{\leftrightarrow} = \frac{1}{2}\left[(B_i)_{\leftarrow i} + (B_i)_{i\rightarrow}\right], \tag{2.48}$$

$$(B_j)_{\leftrightarrow} = \frac{1}{2}\left[(B_j)_{\leftarrow j} + (B_j)_{j\rightarrow}\right], \tag{2.49}$$

$$(B_k)_{\leftrightarrow} = \frac{1}{2}\left[(B_k)_{\leftarrow k} + (B_k)_{k\rightarrow}\right], \tag{2.50}$$

where B_i, B_j and B_j are the primary data defined on the cell inner and outer faces in the i, j, and k direction. The slope-limited gradients are evaluated for each component of the magnetic fields in each direction using equation 2.15. The rest of the magnetic field components on the cell faces may then be reconstructed using the slopes.

2.1.4 Relativistic Magnetohydrodynamics

The covariant equations for relativistic MHD may be written as (Del Zanna et al., 2003)

$$\partial_\alpha(\rho u^\alpha) = 0, \tag{2.51}$$

$$\partial_\alpha\left[(w+|b|^2)u^\alpha u^\beta - b^\alpha b^\beta + (p+|b|^2/2)g^{\alpha\beta}\right] = 0, \tag{2.52}$$

$$\partial_\alpha(u^\alpha b^\beta - u^\beta b^\alpha) = 0, \tag{2.53}$$

where b^α is a magnetic induction four-vector with components

$$b^\alpha = [\gamma(\mathbf{v}\cdot\mathbf{B}), \mathbf{B}/\gamma + \gamma(\mathbf{v}\cdot\mathbf{B})\mathbf{v}], \tag{2.54}$$

and $b^\alpha = [0, \mathbf{B}]$ in the fluid comoving local rest frame, so that

$$|b|^2 \equiv b_\alpha b^\alpha = \frac{B^2}{\gamma^2} + (\mathbf{v}\cdot\mathbf{B})^2. \tag{2.55}$$

Other variables retain the same definitions as in §2.1.2 and §2.1.3. Equation 2.51 gives mass conservation and equation 2.52 gives momentum-energy conservation. The spatial part of equation 2.53 gives the classical induction equation 2.36, while the temporal component becomes the divergence-free constraint of 2.37. The Newtonian MHD equations can be obtained by letting $v^2 \ll 1$, and RHD equations can be recovered by setting $b^\alpha = 0$.

The scheme previously described in §2.1.3 may be implemented by defining the conserved variables and their corresponding fluxes as

$$\mathbf{u} = \left[D, S^j, E\right], \tag{2.56}$$

$$\mathbf{f}^i = \left[\rho u^i, w_{\text{tot}}u^i u^j - b^i b^j + p_{\text{tot}}\delta^{ij}, w_{\text{tot}}u^0 u^i - \rho u^i - b^0 b^i\right], \tag{2.57}$$

where $w_{\rm tot} = w + |b|^2$ and $p_{\rm tot} = p + |b|^2/2$. The conserved variables are defined as

$$\begin{aligned} D &= \rho\gamma, && (2.58) \\ S^j &= w_{\rm tot}u^0u^j - b^0b^j, && (2.59) \\ E &= w_{\rm tot}u^0u^0 - p_{\rm tot} - b^0b^0 - \rho\gamma, && (2.60) \end{aligned}$$

with ρ, v^j, and e as the primitive variables.

Like the case of RHD, the recovery of primitive variables from the conserved ones must be carried out using some iterative root-finding method because of the non-linearity of the equations. Despite the fact that the magnetic field **B** is both a primitive and conserved variable, the inversion is even more difficult in the relativistic MHD case. The inversion scheme needs to be agnostic to the form of the equation of state in general. In the code, we have implemented a slightly modified version of the $1D_W$ scheme found in Noble et al. (2006) to allow for the convergence of the pressure as a function of other primitive variables.

To obtain the primitive variables, first let us define

$$Q_\mu = \gamma(w + b^2)u_\mu - (p' + b^2/2)n_\mu + (n_\nu b^\nu)b_\mu, \tag{2.61}$$

where we set p' to the value of p from the last step initially, and $n_\nu = [-1, 0, 0, 0]$. A non-linear equation for $W = w\gamma^2$ may then be constructed as

$$Q_\mu n^\mu = -\frac{\mathscr{B}}{2}(1 + v^2) + \frac{\left(Q_\mu \mathscr{B}^\mu\right)^2}{2W^2} - W + p', \tag{2.62}$$

where we have used the magnetic field four-vector $\mathscr{B} = [0, \mathbf{B}]$. v^2 in equation 2.62 may be eliminated using the equation

$$v^2 = \frac{\tilde{Q}^2W^2 + \left(Q_\mu \mathscr{B}^\mu\right)^2\left(\mathscr{B}^2 + 2W\right)}{\left(\mathscr{B}^2 + W\right)^2 W^2}, \tag{2.63}$$

where $\tilde{Q}^2 = Q_\mu Q^\mu + \left(Q_\mu n^\mu\right)^2$. This yields an equation that is a function of W only, and may be solved using a numerical root-finding method. After v is recovered through equation 2.63, the rest of the primitive variables can be obtained from equations 2.58 - 2.60. A new value for pressure $p(\rho,e)$ may then be computed using the newly-recovered internal energy e. We check for convergence by comparing the relative error $\varepsilon = (p' - p)/p$ to a preset threshold of 1×10^{-7}. If ε is above this threshold, we set $p' = p$ and re-iterate.

The complete structure of the characteristic wave speeds in a relativistic MHD system was first investigated by Anile and Pennisi (1987), and reviewed later in Anile (1989). In the HLL scheme however, complete knowledge of this structure is not required since we use only the fastest waves to compute the fluid flux in equation 2.17 and electric field in equation 2.45. Here only the necessary expressions for the eigenspeeds are reported, as shown in Del Zanna et al. (2003). These are the four magneto-sonic wave speeds given by the nonlinear quartic equation

$$\left(1-\varepsilon^2\right)\left(u^0\lambda - u^i\right)^4 + \left(1-\lambda^2\right)\left[c_s^2\left(\tilde{b}^0\lambda - \tilde{b}^i\right)^2 - \varepsilon^2\left(u^0\lambda - u^i\right)^2\right] = 0, \tag{2.64}$$

where

$$\begin{aligned} \tilde{b}^\alpha &= b^\alpha/\sqrt{w_{\text{tot}}}, && (2.65)\\ |\tilde{b}|^2 &= \tilde{b}_\alpha \tilde{b}^\alpha = |b|^2/w_{\text{tot}}, && (2.66)\\ \varepsilon^2 &= c_s^2 + |\tilde{b}|^2 - c_s^2|\tilde{b}|^2. && (2.67) \end{aligned}$$

The nonlinear equation 2.64 may be solved using numerical methods suitable for finding roots of polynomials. In GenASiS, we have used the Laguerre's method with the routine *zroots* from *Numerical Recipes* (Press et al., 1986).

2.2 Equation of State

An equation of state describes the relationships between the thermodynamic variables of a system. Typically it describes the dependence of pressure and internal energy on density, temperature, and composition. Therefore, an equation of state closes the equations of (relativistic) MHD systems. We have implemented several equation of state in GenASiS: the polytropic equation of state for an ideal gas, the Lattimer-Swesty (LS) equation of state (Lattimer and Swesty, 1991), and the equation of state by Shen et al. (1998) for dense nuclear matter.

A polytrope is a simple equation of state that is often used in simplified modeling of astrophysical structure where the microphysics is not the primary interest. It has been used in the modelings of objects such as neutron stars and the pre-collapse progenitor and proto-neutron star in the context of core-collapse supernovae. The assumption of this equation of state implies that pressure is independent of temperature, depending only on density and composition. In its most familiar form, the polytropic equation of state is:

$$p = \kappa\rho^{\Gamma}, \tag{2.68}$$

$$e = \frac{p}{(\Gamma-1)}, \tag{2.69}$$

where the adiabatic index Γ is a specified parameter and the polytropic constant κ is updated in response to changes in the internal energy e. A completely ionized, fully convective star, or non-relativistic and completely degenerate gas of fermions, may be modeled with $\Gamma = 5/3$, while one can use $\Gamma = 4/3$ for a degenerate gas of ultrarelativistic fermions. Neutron star environment with nuclear matter density are often model with $\Gamma \geq 2$.

We have also implemented the Lattimer-Swesty (LS) equation of state in GenASiS (Lattimer and Swesty, 1991). LS is an equation of state for dense nuclear matter. This equation of state was formulated for the study of supernovae and neutron stars under the conditions which we are interested in. The matter is taken to consist of free protons, free neutrons, leptons, photons, alpha particles, and nuclei of a single representative heavy species. The leptons and photons are treated as ideal Fermi and Bose gases respectively, since they interact weakly. This equation of state is based on the compressible liquid drop model in a non-relativistic framework. It is designed to be fast enough for use in hydrodynamics codes. It also allows consistent alteration of the major nuclear force parameters. This equation of state has been used in many astrophysics simulations and is a ‘standard’ realistic nuclear equation of state for supernova and neutron star merger simulations.

Another dense nuclear matter equation of state was developed by Shen et al. (1998). This equation of state was developed using relativistic mean field theory for high density nuclear matter. It is also designed for use in numerical simulations under the conditions encountered in neutron stars and neutron star merger. This equation of state incorporates special relativistic effects in the nuclear structure that play a role in describing nuclear saturation. It is therefore important to study neutron star mergers and core-collapse supernovae using this relativistic equation of state. This equation of state has also been incorporated in GenASiS for usage in future simulations.

2.3 Parallel Implementation

A common architecture in modern supercomputing is distributed-memory parallel computers in a computer cluster. Machines of this type allow large problems to be decomposed—for example, into multiple spatial subdomains—and distributed across different ‘processes’ to be solved in parallel. Each ‘process’ contains its own copy of the program, can only access memory locations allocated either statically or dynamically by the program, and can communicate with other processes only through a specific protocol, with the Message Passing Interface (MPI) (Gropp and Lusk, 2010; MPI-Forum, 2010; Gropp et al., 1999) presently being the most widely used.

In many physical simulations the problem size is large enough that the computational domain is conveniently spatially decomposed into multiple subdomains, each assigned to a different process. Communications (via MPI) are then required to synchronize the

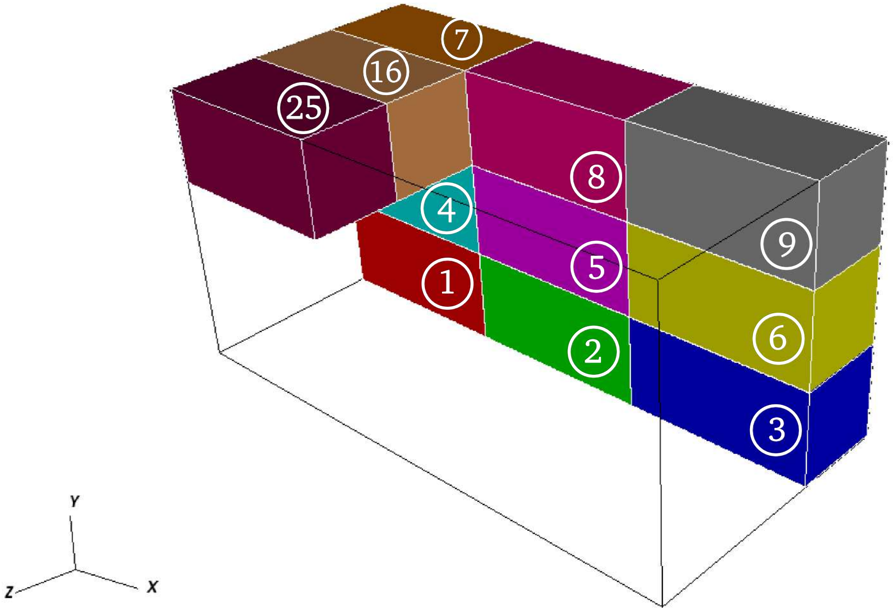

Figure 2.1: An illustrative brick decomposition in three dimensions for a computational domain assigned to twenty-seven processes. Only eleven bricks are shown to simplify the illustration. The bricks are labeled with the rank of the process that 'owns' them. (Process rank numbering here and in the following two figures begins with 1, rather than 0 as in MPI and internally in the code.)

time integration and solve the problem in parallel on the whole computational domain. In the general case the extent of the domain—and/or the number of mesh points—need not be the same in all dimensions. In the case of magnetohydrodynamics, in which only nearest-neighbor information is required (for example, gradient and flux computation), decompositions yielding subdomains with low surface-to-volume ratio are favorable to minimize communications between processes. To that end, GenASiS uses a simple 'brick' decomposition: in three dimensions, the computational domain is divided in each dimension by $n_b = \sqrt[3]{n_p}$, the cube root of the number of processes n_p. Figure 2.1 illutrates the brick decomposition.

To facilitate the computations that require nearest-neighbor cell values crossing process boundaries, each process keeps 'ghost' cells in addition to the 'proper' cells (the cells that make up the subdomain owned by the process). The value of ghost cells of a process need to be updated to reflect the values of the proper cells owned by other processes corresponding to that ghost cells at least once after every Runge-Kutta step. This is done by exchanges of values between processes. The exchanges are point-to-point communications, with each process sending to and receiving from other (known) processes that are neighbors to that process. Figure 2.2 shows this construction. For illustrative purpose, the figure is a two-dimensional slice of figure 2.1, showing only an *xy*-plane. Proper cells owned by process 1, 3, 5, 7, 9 are shown as solid lines. The colors denote the process number. The ghost cells for each process are indicated by dotted lines. In this example (in the two-dimensional slice), the ghost cells of process 1 correspond to the proper cells owned by process 2, its right neighbor (not shown), and process 4, its top neighbor (not shown). Similarly, the ghost cells of process 4 correspond to the proper cells of process 2 (bottom neighbor), 4 (left neighbor), 6 (right neighbor), and 8 (top neighbor). Conversely, some of the ghost cells of process 2, 4, 6, and 8 correspond to the proper cell of process 5. Thus in this case, process 5 needs to have communication exchanges with process 2, 4, 6, and 8 to send the values of its proper cells that make up the ghost cells of those processes, and to receive values of the proper cells of those processes that make up its ghost cells.

Notice that only a subset of proper cells needs to be exchanged between processes. For our second-order spatial scheme, this is one layer of cells that border the process boundary. (Higher-order spatial scheme may need more layers of cells). Since communications take times away from computation, in order to achieve good performance and scalability as the number of processes grows, communications are overlapped with computations. This is done by collecting the cells that need to be communicated into a set called 'sent cells', with their interfaces to 'sent faces'. Evaluation of fluxes involving the sent cells and sent faces are done first (note that the computation of fluxes and slopes on the domain can be done in arbitrary order). Once the changes to the sent cells are applied, they are communicated to

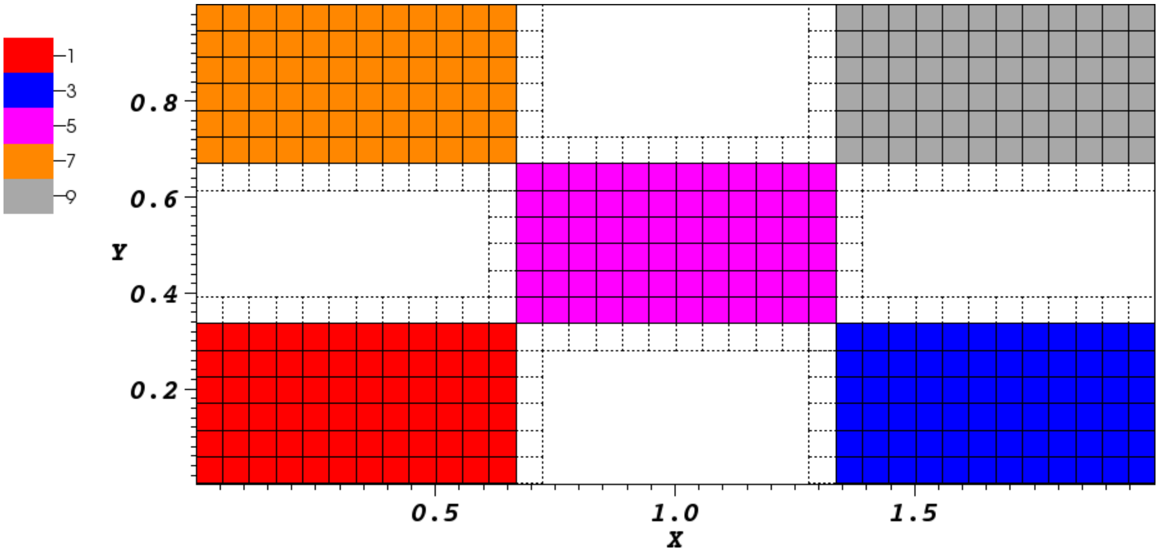

Figure 2.2: A slice of brick decomposition in the two-dimensional *xy*-plane. Process 1, 3, 5, 7, 9, with their proper cells, are shown as solid lines. The ghost cells for these processes are shown as dotted lines.

other processes with non-blocking communications. The fluxes for the rest of the cells are evaluated while waiting for the communication to finish.

This technique of overlapping communication with computation has allowed our code to maintain weak scaling, i.e. maintaining the time needed to solve the equations of MHD as we increase the number of processes for a fixed problem size (e.g. the size of subdomain) per process. Figure 2.3 shows the weak-scaling plot of the MHD implementation. Good weak-scalability is essential to the practical execution of large-scale simulations such as the ones we have in chapter 6.

2.4 Numerical Test Problems

In this section we present numerical test problems that have been done in GenASiS for the MHD and relativistic MHD implementations. Many of these test problems are well-known in the literature. Test problems serve as methods to validate the correctness of implementation and scheme. They also serve as a way to check strengths

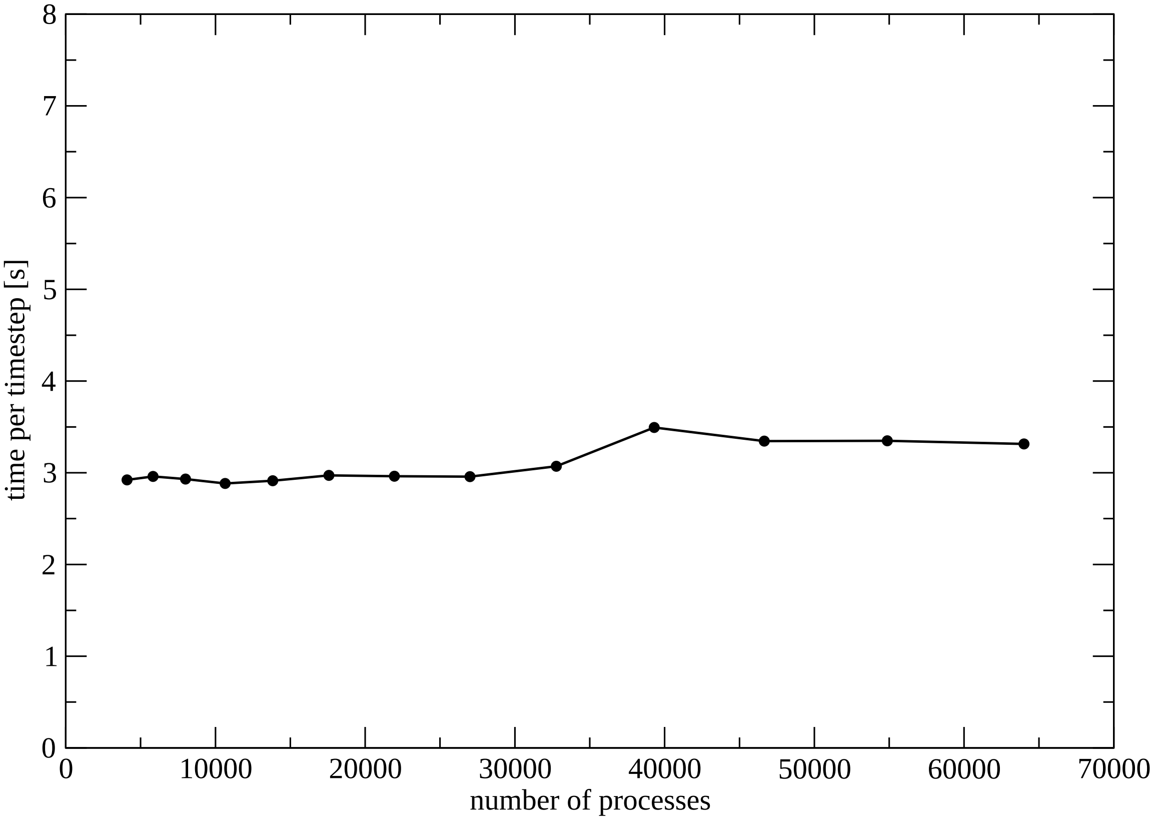

Figure 2.3: Weak scaling of magnetohydrodynamics implementation in GenASiS.

and weaknesses of different schemes to solve the same physics. Since in the future GenASiS may implement different schemes (e.g. new coordinate systems, different magnetohydrodynamics formulations, different Poisson's equation solvers), these test problems will also serve as benchmarks with which we measure the maturity of the code with regards to specific scheme's implementation. In its current state, these test problems are exercised daily to guard against unintentional introduction of new bugs as code development proceeds Yet another important purpose of the test problems is to test the ability of GenASiS to handle different physical situations and regimes that it may encounter in real physical simulation in a controlled setting. The successes of the code in this area gives us confidence in the correctness of results when we use it to do physical simulations to explore scientific problems.

Some of the test problems here have analytical solution that we can compare against. Others have semi-analytical or numerical solutions computed by a different scheme (for example, a more computationally-costly but more 'exact' scheme), or different codes that are available in the literature. Some test problems have periodicity in which we can compare the numerical solution after being evolved for some period to its initial condition (i.e. the known solution). For each of the test problems, we benchmark GenASiS' numerical result by quantifying differences of some variable χ on the mesh as the L_1 norm relative error:

$$L_1(\chi) = \frac{\sum\limits_{i,j,k} \left|\chi\left(x_i, y_j, z_k\right) - \chi_0\left(x_i, y_j, z_k\right)\right|}{\sum\limits_{i,j,k} \left|\chi_0\left(x_i, y_j, z_k\right)\right|}, \tag{2.70}$$

where χ_0 is the 'known' value being used as the standard for comparison. Here the summation is done over all cells on the computational domain. Thus the L_1 norm gives a single number as a quantitative measure of error for a certain mesh resolution. Plotting this error measurement as function of mesh resolution then gives us the order of convergence of our numerical solution. In general, because we have second-order scheme in space for MHD and second-order Runge-Kutta integration in time evolution, we expect second-order convergence for smooth flows. For flows with shocks, this reduces to first-order convergence.

2.4.1 Newtonian Riemann Shock Tube Problem

2.4.1.1 Analytical Solution

The 1D Riemann shock tube problem is a well-known staple test problem for any shock-capturing hydrodynamics code. It was first introduced by Sod (1978) to benchmark hydrodynamics algorithms. The problem is initialized by setting up two discontinuous states: high density and high pressure gas on the left, and low density with low pressure gas on the right, separated by a membrane. At time t_0 the membrane is removed, allowing the two states of the gas to interact, and the evolution of the gas interaction is followed. Nonlinear waves are generated at the discontinuities, with a shock wave propagating to the right and a rarefaction wave to the left.

The anatomy of the shock tube profile is illustrated in figure 2.4. The density, pressure, and velocity profiles of the gas at $t = 0.25$ are plotted. Five distinct regions can be identified, as indicated in figure 2.4 by thin dotted vertical lines to delineate the regions:

- region 1: the undisturbed state initially to the left of the membrane
- region 2: the rarefaction region
- region 3: the pre-shock region to the left of the contact discontinuity
- region 4: the post-shock region to the right of the contact discontinuity
- region 5: the undisturbed state initially to the right of the membrane

There is a constant pressure and velocity across the contact discontinuity in region 3 and 4. This is accompanied by a discontinuity in the specific entropy and internal energy of the fluid in region 4 just behind the shock because of heating through the shock.

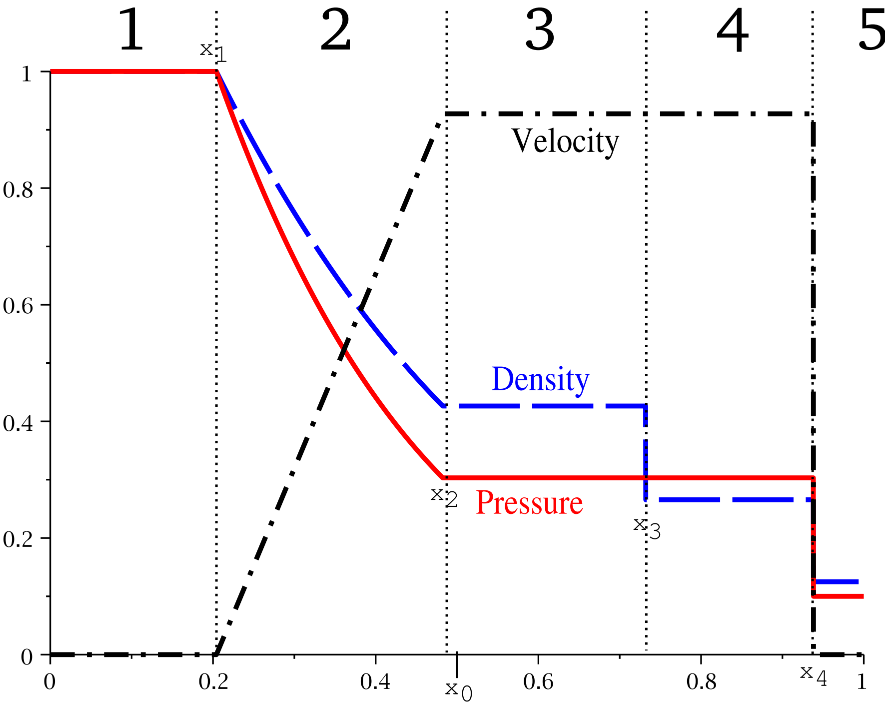

Figure 2.4: Analytical solution for shock tube problem at $t = 2.5$ with polytropic equation of state $\Gamma = 1.4$. The figure illustrates the various regions and key positions that exist in the problem. Density, pressure, and velocity profiles are shown. The vertical dotted lines delineate various regions labeled by numbers. Various key positions are identified as $x_0 \ldots x_4$, with their particular values as: $x_0 = 0.5$, $x_1 = 0.2$, $x_2 = 0.48$, $x_3 = 0.73$, $x_4 = 0.94$.

Analytical solutions can be derived to find the values of density, pressure, and velocity in each region. The following key positions, as indicated in figure 2.4, help describes the complete solution to the shock tube problem:

- x_0: the initial position of the membrane
- x_1: the position of the rarefaction head moving to the left
- x_2: the position of the rarefaction tail
- x_3: contact discontinuity that separates left fluid from the right fluid
- x_4: the shock front moving to the right of the membrane

The initial conditions of the shock tube problem gives us the values of pressure, density, and velocity in region 1 and 5, denoted as $[p_1, \rho_1, v_1]$ and $[p_5, \rho_5, v_5]$, respectively. Here number subscripts are used to indicate the region where the quantity belongs.

The states on either side of the shock can be described by the Rankine-Hugoniot jump conditions (Courant and Friedrichs, 1977):

$$\begin{aligned}
\rho_4\left(v_4 - v_S\right) &= \rho_5\left(v_5 - v_S\right), && (2.71)\\
\rho_4\left(v_4 - v_S\right)^2 + p_4 &= \rho_5\left(v_5 - v_S\right)^2 + p_5, && (2.72)\\
\rho\left(\frac{1}{2}v_4^2 + \varepsilon\right)\left(v_4 - v_S\right) + p_4 v_4 &= \rho\left(\frac{1}{2}v_5^2 + \varepsilon\right)\left(v_5 - v_S\right) + p_5 v_5, && (2.73)
\end{aligned}$$

where v_s is the velocity of the shock moving to the right. Let us define:

$$\Omega = \frac{\Gamma - 1}{\Gamma + 1},$$

$$\beta = \frac{\Gamma - 1}{2\Gamma},$$

where Γ is the adiabatic index of the polytropic equation of state of the form $p = \kappa\rho^{\Gamma}$, with sound speed is given by $c_s = \sqrt{\Gamma p/\rho}$. Using the shock jump conditions and the initial state of the fluid to the right of the membrane, we find the equation describing the possible post-shock values for P and v:

$$v = v_5 + (p - p_5)\sqrt{\frac{1-\Omega}{\rho_5\,(p+\Omega p_5)}}. \tag{2.74}$$

This is a curve on the p-v plane. The possible values of p and v on the rarefaction wave moving to the left are described by the equation

$$v = v_1 + \left({p_1}^{\beta} - p^{\beta}\right)\sqrt{\frac{(1-\Omega^2)\,{p_1}^{1/\Gamma}}{\Omega^2\rho_1}}. \tag{2.75}$$

This is also a curve on the p-v plane. The intersection of these two curves then gives us the value of pressure and velocity in region 3 and 4. By solving the following equation for p, we get the value of p_3 (and therefore also p_4):

$$v_1 + \left({p_1}^{\beta} - p^{\beta}\right)\sqrt{\frac{(1-\Omega^2)\,{p_1}^{1/\Gamma}}{\Omega^2\rho_1}} = v_5 + (p - p_5)\sqrt{\frac{1-\Omega}{\rho_5\,(p+\Omega p_5)}}. \tag{2.76}$$

We can then get the value of v_3 and v_4 from equation 2.74 by substituting $p = p_3$.

The jump conditions also imply the following relationship

$$\rho_4 = \rho_5 \frac{p_4 + \Omega p_5}{p_5 + \Omega p_4}. \tag{2.77}$$

We use this to find the density in in region 4 (ρ_4). The density in region 3 (ρ_3) can be calculated from polytropic relation:

$$\rho_3 = \left(\frac{p_3}{p_1}\right)^{\frac{1}{\Gamma}} \rho_1. \tag{2.78}$$

The rarefaction head moves to the left with the speed of sound in the initial region left to the membrane, thus giving us

$$x_1 = x_0 - c_{s1} t \tag{2.79}$$

at any time $t > 0$. The rarefaction tail location is described by

$$x_2 = x_0 - \left(c_{s1} - \frac{\Gamma + 1}{2} v_3 \right) t. \tag{2.80}$$

The position of the contact discontinuity and shock front are described by:

$$x_3 = x_0 + v_3 t, \tag{2.81}$$

$$x_4 = x_0 + v_s t, \tag{2.82}$$

where v_s is the shock front velocity. The shock front velocity can be derived from the jump conditions, yielding:

$$v_s = \frac{v_3 \rho_4}{\rho_5 \left(\frac{\rho_4}{\rho_5} - 1 \right)}. \tag{2.83}$$

This completes the analytical solution of the shock tube problem.

2.4.1.2 Numerical Solution

At $t = 0$, we initialize the states with the following conditions:

$$\begin{aligned} [v, \rho, p]_L &= [0.0, 1.0, 1.0], \\ [v, \rho, p]_R &= [0.0, 0.1, 0.125], \end{aligned} \tag{2.84}$$

where the subscript L and R indicate the left and right states of the discontinuity, respectively. Initially the discontinuity between the left and right states is placed at $x = 0.5$.

Figure 2.5 shows numerical and analytical results for this shock tube problem at $t = 0.245$. We can see from the figure that the code is able to capture all the features of the shock tube problems with high accuracy. Also notice that there is no post-shock

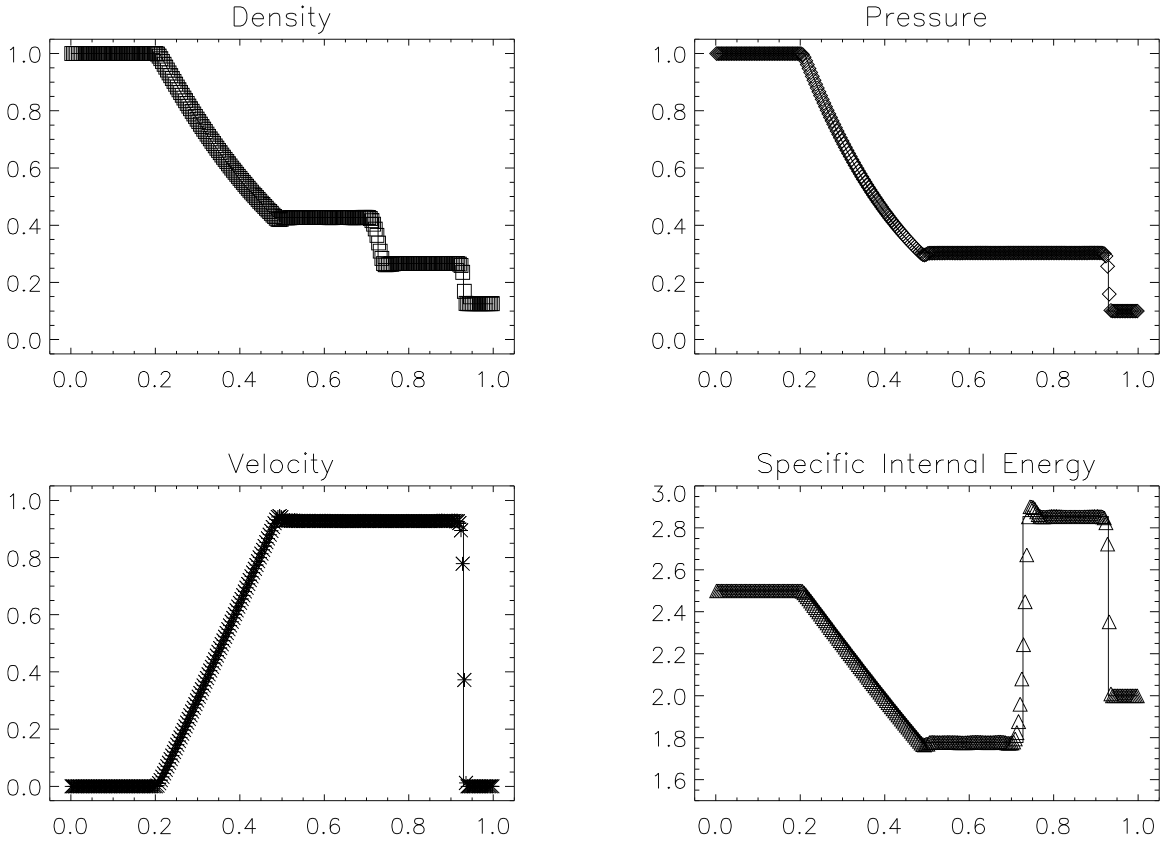

Figure 2.5: Comparison of analytical and numerical result for 1D Sod shock tube problem at $t = 0.245$. The solid lines indicate the analytical result. The numerical solution is computed with 256 cells, using a Courant number $C = 0.5$ and slope-limiter parameter $\theta = 2.0$.

oscillation visible in the numerical solution, since it is quickly damped. This is a feature of the hydrodynamics scheme we employ. The correctness of our implementation is illustrated by the ability of the code to reproduce the analytical solution numerically.

2.4.2 Newtonian Magnetized Shock Tube Problem

Brio and Wu (1988) introduced the magnetized version of Sod's shock tube problem. The hydrodynamical initial condition are identical to Sod's shock tube problem, with a magnetic field added as initial condition:

$$
\begin{aligned}
[\rho, p, v_x, B_x, B_y, B_z]_L &= [1.0, 1.0, 0.0, 0.75, 1.0, 0.0]\,, \\
[\rho, p, v_x, B_x, B_y, B_z]_R &= [0.125, 0.1, 0.0, 0.75, -1.0, 0.0]\,. \qquad (2.85)
\end{aligned}
$$

Figure 2.6 shows our numerical result for this test problem computed with 800 cells. The MHD equations give rise to additional characteristic waves, resulting in more structures in this problem compared to the non-magnetized shock tube problem. Similarly, higher resolution is needed to resolve the compound wave structure between the rarefaction wave and contact discontinuity, which exists due to the non-convexity of the equations. Our results are consistent with those of Brio and Wu (1988).

2.4.3 Relativistic Riemann Shock Tube Problems

Now we present several relativistic Riemann problems to test the implementation of relativistic MHD in GenASiS when magnetic fields are absent, therefore reducing the implementation to the relativistic hydrodynamics scheme (cf. §2.1.2). In the relativistic regime, the structure of the solution remains the same qualitatively as in the Newtonian shock tube, except that the rarefaction wave does not yield a linear profile due to the non-linearity in Lorentz transformation.

2.4.3.1 1D Relativistic Blast Waves

The first two test problems were introduced by Donat (1998), and were also used by Del Zanna and Bucciantini (2002) to test the central-scheme method. We start with a fairly easy *blast wave* exploring only a mildly relativistic regime as the first problem. The second one is similar, but more severe relativistically with a stronger shock. For both problems, initially two discontinuous states on the left and right side of the membrane are set as

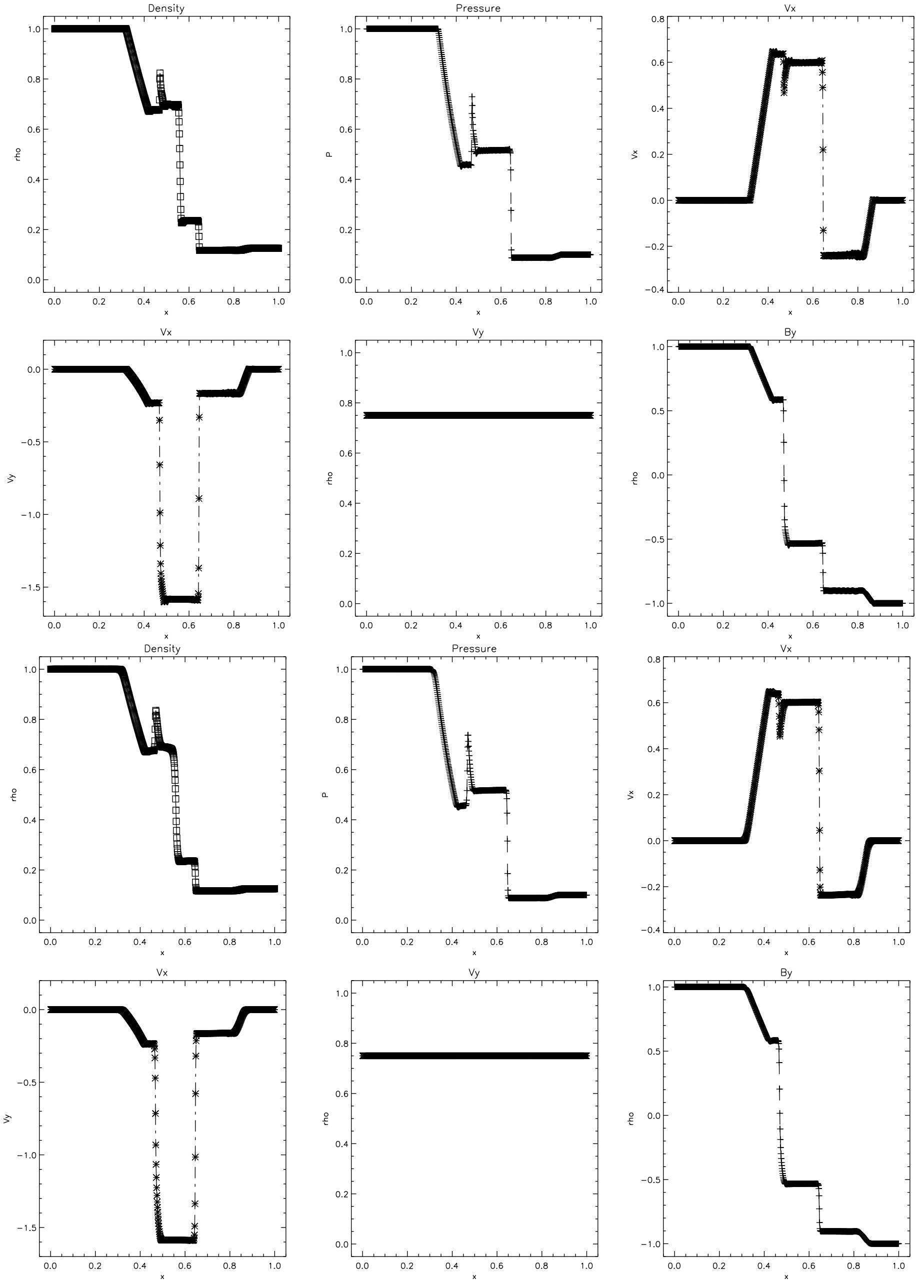

Figure 2.6: 1D magnetized shock tube problem at $t = 0.1$. The numerical solution is computed with 800 zones, with Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top panel and bottom panel, respectively.

follows on region $x = [0, 1]$:

$$\text{blast wave 1} \quad : \quad \begin{cases} [v, \rho, p]_L = [0.0, 10.0, 13.3], \\ [v, \rho, p]_R = \left[0.0, 1.0, 1.0 \times 10^{-6}\right], \end{cases} \tag{2.86}$$

$$\text{blast wave 2} \quad : \quad \begin{cases} [v, \rho, p]_L = [0.0, 1000, 1.0], \\ [v, \rho, p]_R = [0.0, 0.01, 1.0], \end{cases} \tag{2.87}$$

where the discontinuity is located at $x = 0.5$. As before, the fluid is assumed to be an ideal gas with adiabatic index $\Gamma = 5/3$. As in §2.4.1, the initial discontinuities generate a shock wave, a rarefaction wave, and a contact discontinuity during the evolution of the fluid.

The exact solution to relativistic shock tube problem was first considered by Martí and Müller (1994). A Fortran program to compute this exact solution may be found in Martí and Müller (2003). We use this program as benchmark for our numerical solution for the two blast wave problems.

Figure 2.7 shows our numerical results for the first blast wave problem with 400 cells. We show results from using MM and MC slope limiter. The figure indicates that GenASiS is able to resolve all features of the solution and accurately capture the definition of shocks. With a Minmod slope limiter the discontinuities are less sharp and more smeared out due to the diffusive nature of the limiter. Note also the lack of post-shock oscillation, which is a feature of the scheme we use.

The simulated results of the second blast wave problem are shown on figure 2.8. There is a very thin shell in the density behind the shock. The exact solution for this density peak is around 10.5, while ours is around 6.5 with the MC limiter and 5.5 with the MinMod limiter, as shown in the density plots in the upper panel (for MC limiter) and lower panel (for MinMod limiter) in figure 2.8, respectively. This provides a measure of the numerical viscosity of the scheme. Our results of these blast wave problems reproduces other published results (for example, by Del Zanna and Bucciantini (2002) and Zhang and MacFadyen (2006)).

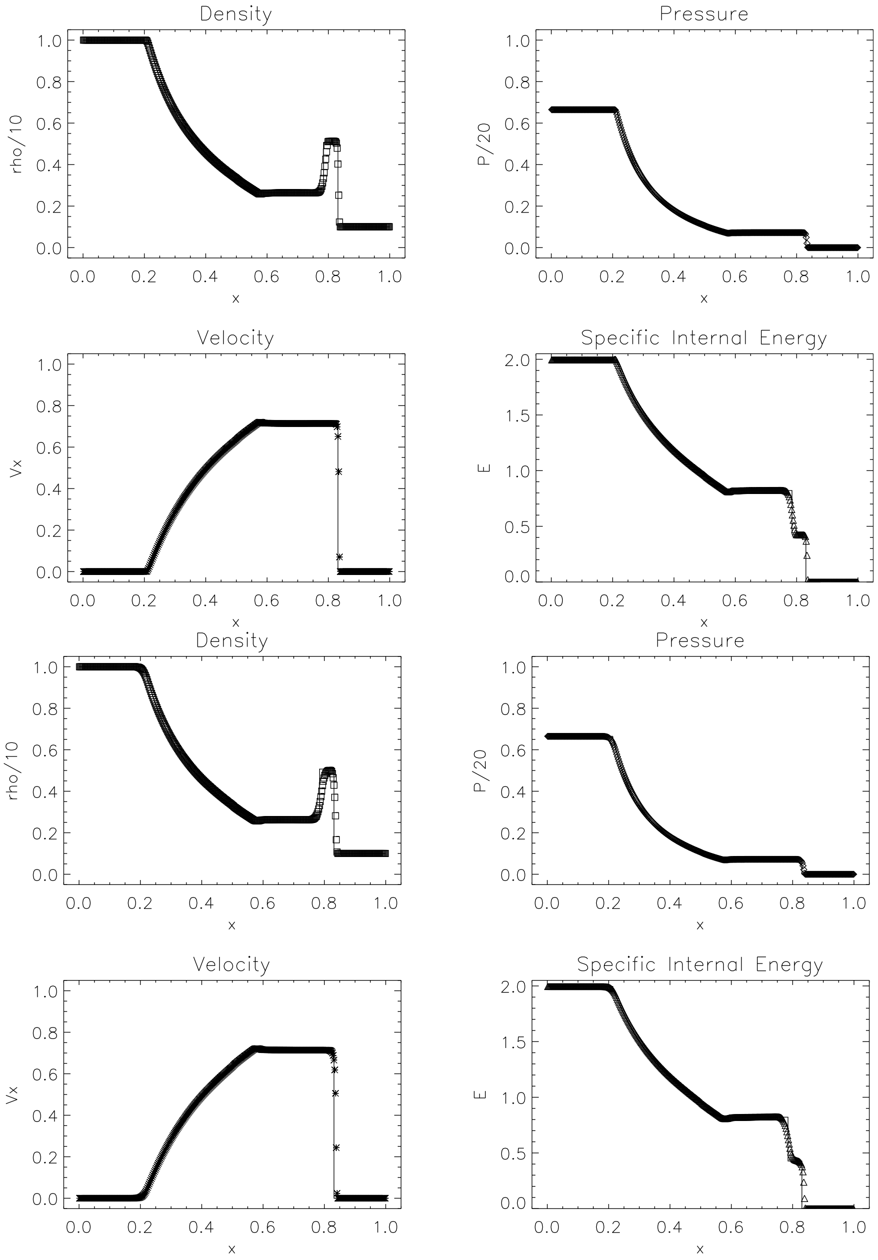

Figure 2.7: Comparison of numerical result and exact solution for blast wave 1 test problem at $t = 0.4$. The solid lines indicate the exact solution generated by program in Martí and Müller (2003). The numerical solution is computed with 400 cells, using a Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top panel and bottom panel, respectively.

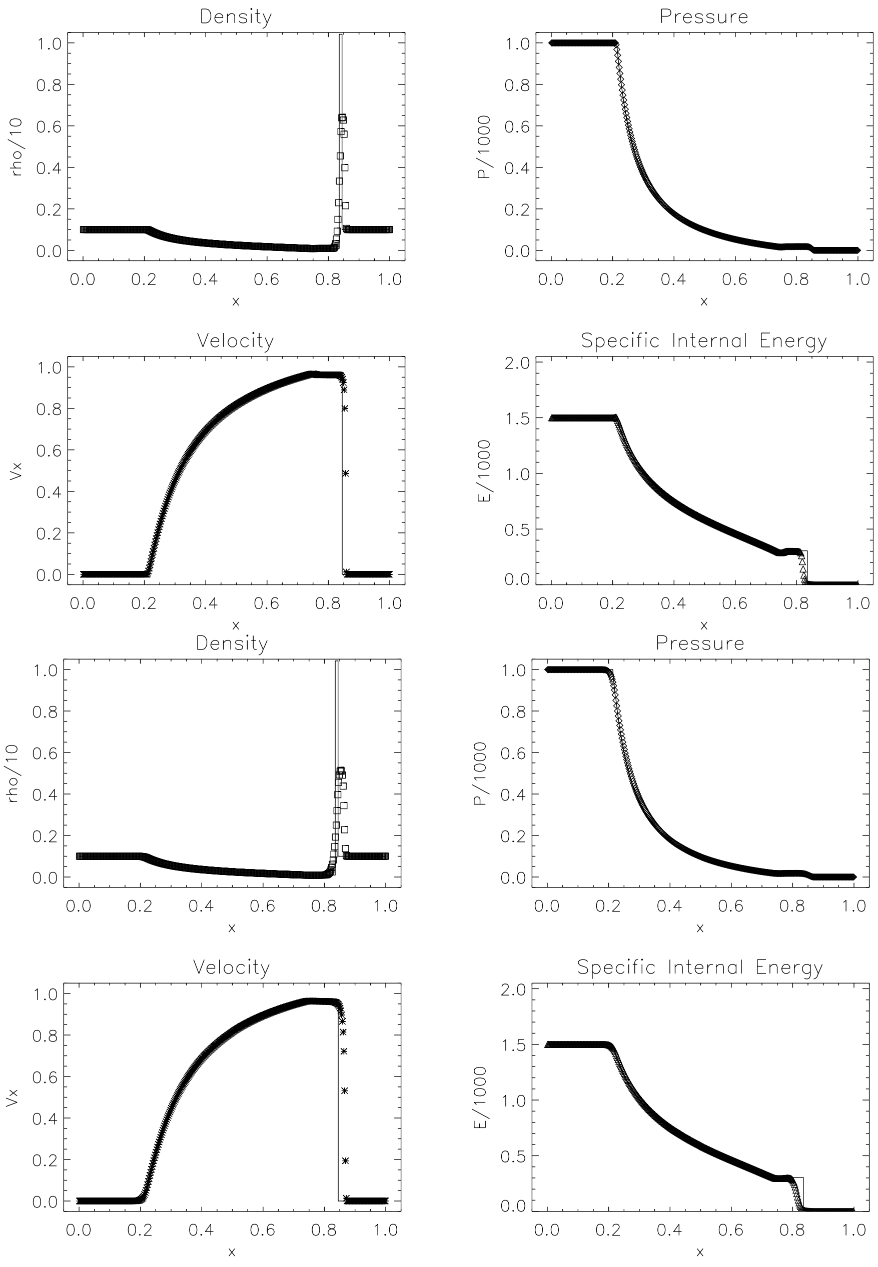

Figure 2.8: Comparison of numerical result and exact solution for blast wave 2 test problem at $t = 0.35$. The solid lines indicate the exact solution. The numerical solution is computed with 400 zones, using a Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top panel and bottom panel, respectively.

2.4.3.2 1D Shock Tube with Non-zero Velocity

Next we consider a relativistic shock tube problem with non-zero initial velocity. We set the initial condition of the left and right states as follows:

$$\begin{aligned} [v, \rho, p]_L &= [0.9, 1.0, 1.0]\,, \\ [v, \rho, p]_R &= [0.0, 10, 1.0]\,. \end{aligned} \tag{2.88}$$

In this test problem we set the adiabatic index to $\Gamma = 4/3$.

The numerical results and exact solutions are shown in figure 2.9. As before, the exact solution is obtained using the Riemann solver found in Martí and Müller (2003). In this problem a strong reverse shock is observed in the results. The numerical results shows small but visible post-shock oscillations for the MC limiter, which are quickly damped. This can be reduced by lowering the Courant number or by increasing the numerical viscosity. In the Minmod case, this oscillation is less noticeable because the limiter is more diffusive. The trade-off is less sharp definition of shocks and contact discontinuities.

We also consider a shock tube problem with with non-zero transverse velocity. In relativistic flow, transverse velocity is coupled to the dynamics along all directions by the Lorentz factor, which makes this much more difficult to solve correctly compared to the Newtonian counterpart. This test is relevant in cases where the hydrodynamics involve strong shear flows, such as in astrophysical jets.

We set up the initial condition as follows:

$$\begin{aligned} [v_x, v_y, \rho, p]_L &= [0.0, 0.0, 1.0, 1000.0]\,, \\ [v_x, v_y, \rho, p]_R &= [0.0, 0.99, 1.0, 0.01]\,. \end{aligned} \tag{2.89}$$

An adiabatic index $\Gamma = 5/3$ is used for this test problem.

Pons et al. (2000) obtained the general solution for the relativistic Riemann problem with tangential velocity. A Fortran program to compute this exact solution was provided

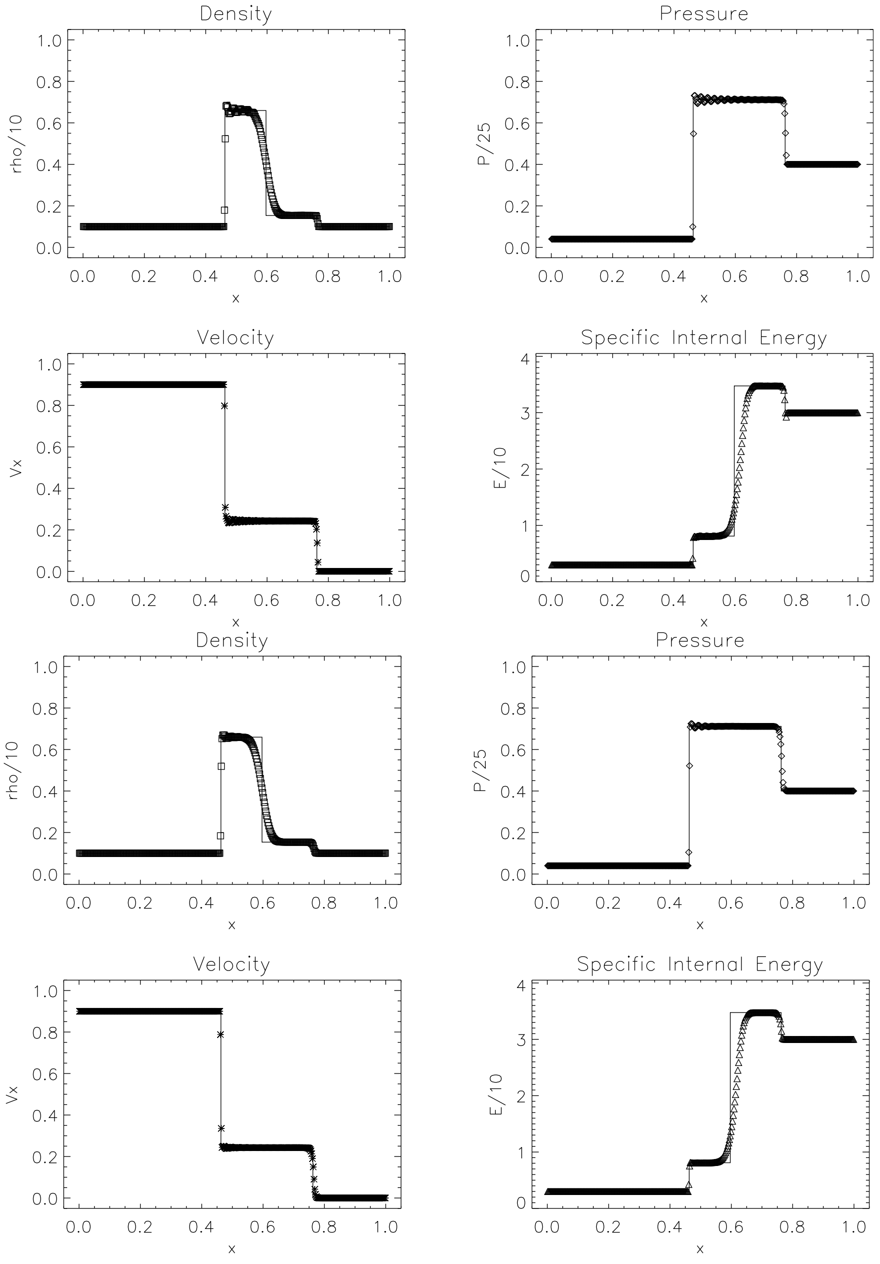

Figure 2.9: Comparison of numerical result and exact solution for shock tube problem with non-zero initial velocity at $t = 0.4$. The solid lines indicate the exact solution. The numerical solution is computed with 400 cells, using a Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top panel and bottom panel, respectively.

by Martí and Müller (2003). We use this program to compare our numerical solution for the this problem.

Figure 2.10 shows our numerical results. The exact solutions are also plotted for comparison. Our results are comparable to other published results by Zhang and MacFadyen (2006).

2.4.3.3 2D Relativistic Shock Tube

Here we consider the two-dimensional counterparts of the shock tube problems. Multidimensional relativistic simulations are harder than one-dimensional for the following reason. In our scheme for multidimensional hydrodynamics, the velocity components are interpolated spatially to the cell faces. In highly relativistic cases, the interpolation may cause a non-physical condition where $v^2 > 1$. To avoid this, we have implemented a fall-back mechanism by returning to the beginning of the failed time step and using first-order reconstruction when the second-order one produces a non-physical condition. This does not mean that the code cannot handle the highly-relativistic regime. It simply means that if we use insufficient resolution, which may cause the non-physical condition, the reconstruction may be less accurate than second-order for certain time steps.

This 2D test problem starts with a square domain divided into four quadrants of constant values at the initial time. The four boundaries defines the contact discontinuities and two 1D shocks, which are symmetric with respect to the main diagonal. The initial conditions for the four quadrants are:

$$\begin{aligned}
[v_x, v_y, \rho, p]_{NE} &= [0.0, 0.0, 0.1, 0.01]\,, \\
[v_x, v_y, \rho, p]_{NW} &= [0.99, 0.0, 0.1, 1.0]\,, \\
[v_x, v_y, \rho, p]_{SW} &= [0.0, 0.0, 0.5, 1.0]\,, \\
[v_x, v_y, \rho, p]_{SE} &= [0.0, 0.99, 0.1, 1.0]\,, \qquad (2.90)
\end{aligned}$$

with adiabatic index $\Gamma = 4/3$.

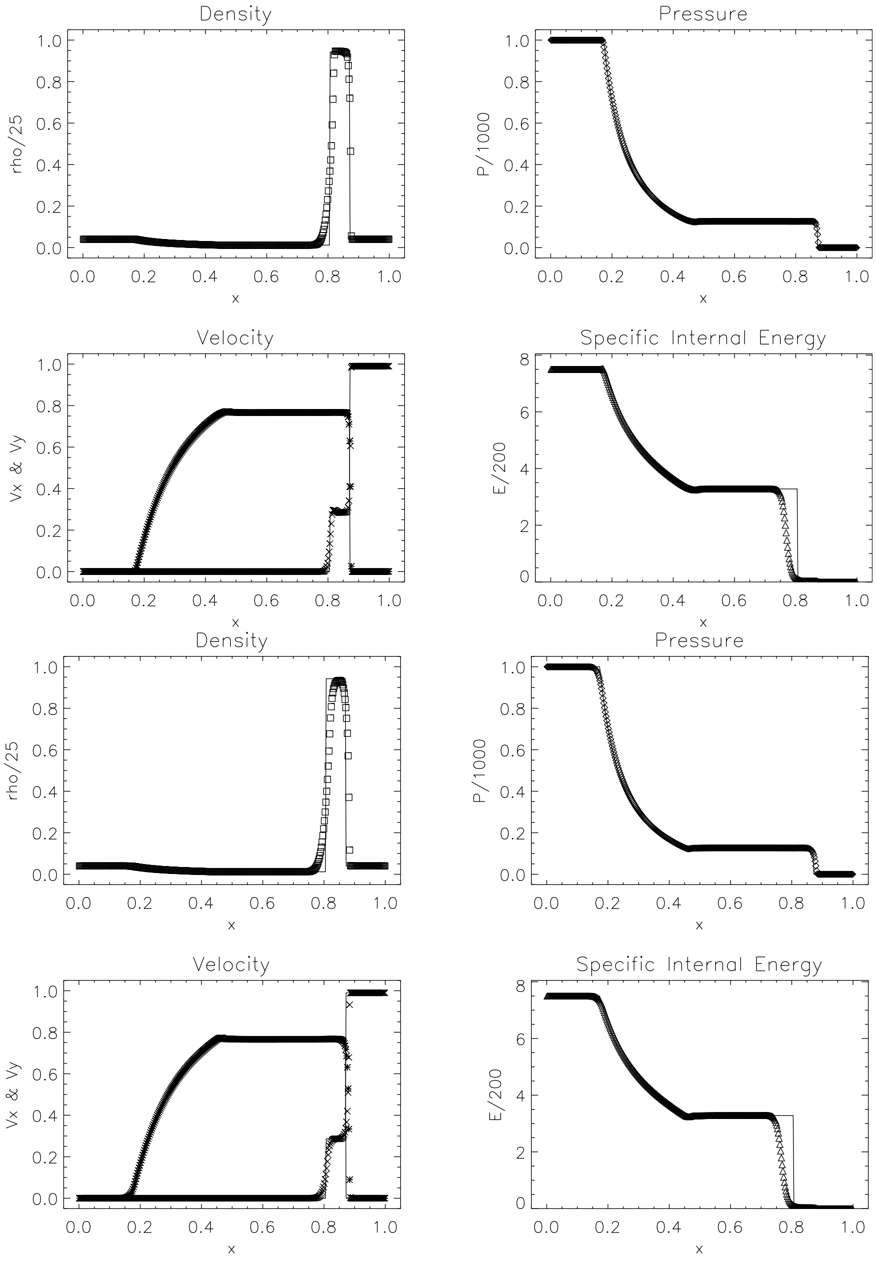

Figure 2.10: Comparison of numerical result and exact solutions for relativistic shock tube with transverse velocity at $t = 0.4$. The solid lines indicate the exact solution. Both the longitudinal and transverse velocity are plotted. The numerical solutions are computed with 400 cells, using a Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top and bottom half of the panel, respectively.

Figure 2.11 shows our numerical result as density contour in logarithmic scale. While Riemann problems can be solved exactly in 1D, this is not the case with 2D problem. Therefore we may only compare the results to other published numerical results in the literature. Our results reproduce other results published by Del Zanna and Bucciantini (2002).

2.4.4 Relativistic Magnetized Shock Tube Problems

Here we present several relativistic magnetized shock tube problems. In all of the test problems here, we set the adiabatic index $\Gamma = 5/3$, and $v_y = v_z = 0$.

2.4.4.1 1D Compound Wave

A relativistic extension of §2.4.2 may be constructed and solved with the equations of relativistic MHD. As before, the initial conditions involve two separate states in the numerical region, with discontinuities at $x = 0.5$. For this test, we choose the left and right states as

$$\begin{aligned} [\rho, p, v_x, B_x, B_y, B_z]_L &= [1.0, 1.0, 0.0, 0.5, 1.0, 0.0]\,, \\ [\rho, p, v_x, B_x, B_y, B_z]_R &= [0.125, 0.1, 0.0, 0.5, -1.0, 0.0]\,. \end{aligned} \tag{2.91}$$

Figure 2.12 shows our numerical results. As in §2.4.2, the solution gives rise to a compound wave. Analytical solution for this test problem is still under investigation. In fact, the reality of the compound wave structures is still debatable because they are not predicted by analytic calculation, yet are found in any shock capturing code (Myong and Roe, 1998; Barmin et al., 1996). However, the results presented here are consistent with other published results (Balsara, 2001; Del Zanna et al., 2003).

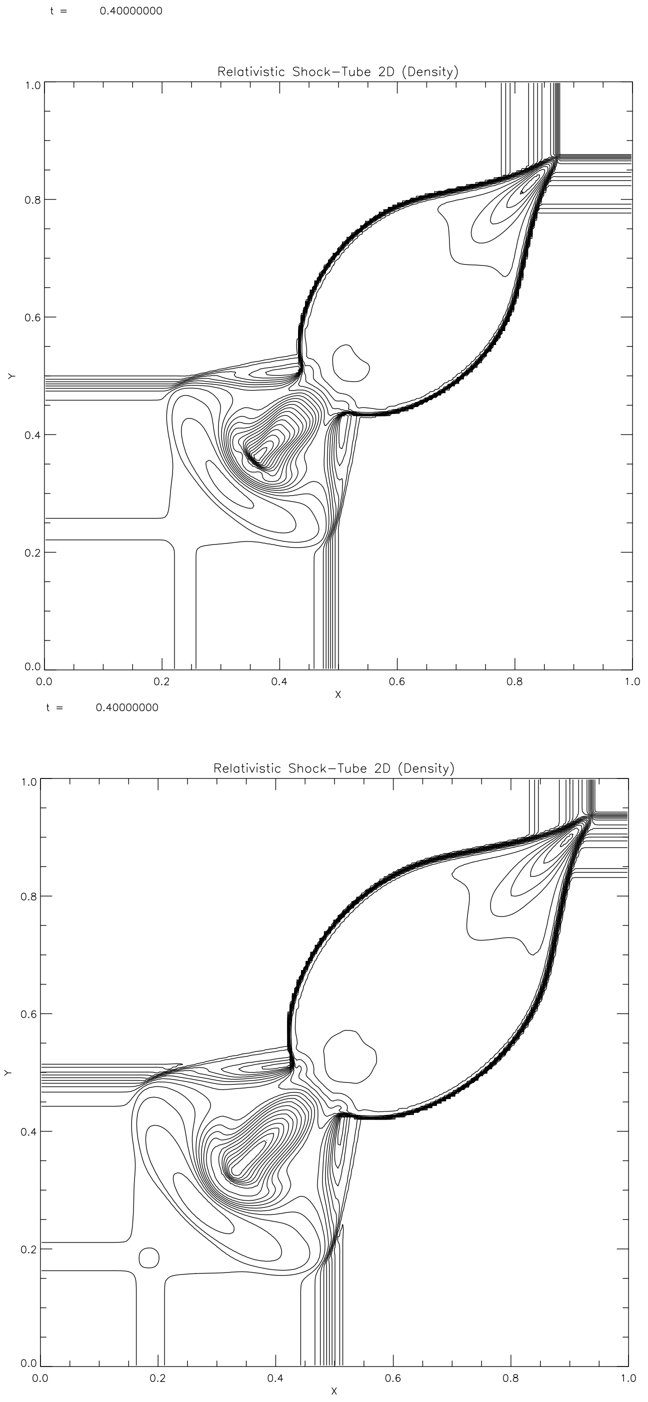

Figure 2.11: Density contours in logarithmic scale for 2D relativistic shock tube problem at time $t = 0.4$ with 400 cells per dimension and a Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top and bottom half of the panels, respectively.

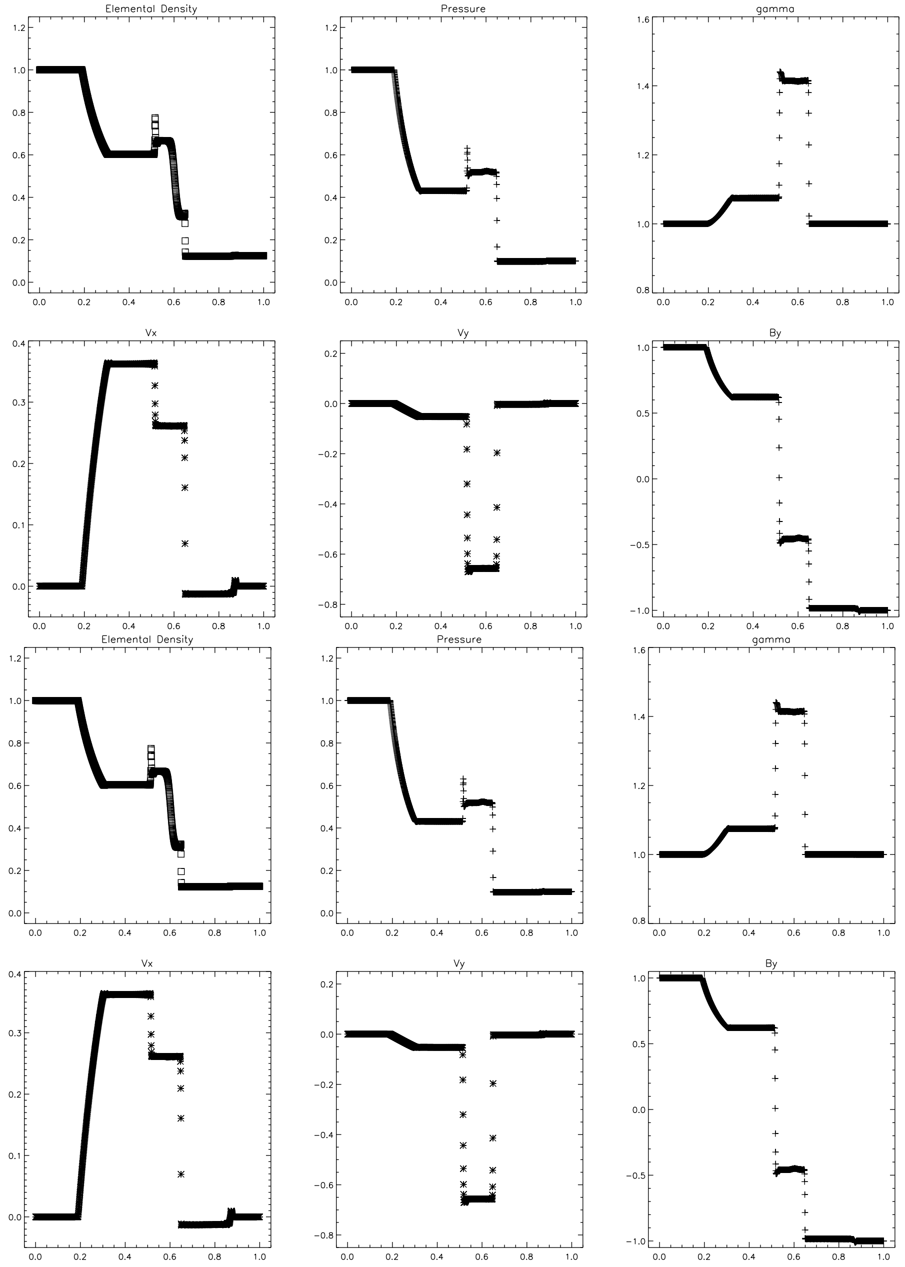

Figure 2.12: 1D relativistic compound wave test problem at $t = 0.4$. From left to right, top to bottom panels, the plot for density, pressure, Lorentz factor γ, v_x, v_y, and B_y are shown. The numerical solutions are computed with 1600 cells, using a Courant number $C = 0.5$. MC and Minmod slope limiters are used for the top and bottom half of the panels, respectively.

2.4.4.2 1D Magnetized Relativistic Blast Waves

A couple of magnetized versions of relativistic blast waves are presented here. As in 2.4.3.1, the first one is a mildly relativistic blast wave with moderate pressure jump. The second one has much stronger shocks with about a factor of 10^4 pressure jump, producing relativistic flow with Lorentz factors as large as $\gamma \cong 3.4$. The left and right states of the two blast waves problems are as follows:

$$\text{blast wave 1} \quad : \quad \begin{cases} [\rho, p, v_x, B_x, B_y, B_z]_L = [1.0, 30.0, 0.0, 5.0, 6.0, 6.0]\,, \\ [\rho, p, v_x, B_x, B_y, B_z]_R = [1.0, 1.0, 0.0, 5.0, 0.7, 0.7]\,, \end{cases} \tag{2.92}$$

$$\text{blast wave 2} \quad : \quad \begin{cases} [\rho, p, v_x, B_x, B_y, B_z]_L = [1.0, 1000.0, 0.0, 10.0, 7.0, 7.0]\,, \\ [\rho, p, v_x, B_x, B_y, B_z]_R = [1.0, 0.1, 0.0, 10.0, 0.7, 0.7]\,, \end{cases} \tag{2.93}$$

The results for these blast waves problems are shown in figures 2.13 and 2.14. Both results were computed using 1600 cells. The first blast wave problem is well resolved in figure 2.13. We can clearly see all the structures that develop due the various characteristic waves of the relativistic MHD equations. For example, on the density profile, from left to right we can identify the fast rarefaction wave moving to the left, the left rarefaction wave moving to the left, a contact discontinuity, the slow shock moving to the right, and the fast shock moving to the right. This blast wave produces a maximum Lorentz factor of almost $\gamma = 1.6$. Not all of these structures can be seen in figure 2.14 however. We see the two left-going rarefaction waves and contact discontinuity, but cannot easily identify the two right-going shocks. The more relativistic flow with higher Lorentz factor produced by this problem has the consequence of a more severe length-contraction effect to the structures moving to the right when viewed in laboratory frame, which is the also the computational grid. Therefore the same resolution that resolved the first blast wave problem is not high enough to resolve all the structures in this problem. However, this also demonstrate a feature of our scheme. Even though some structures are under-resolved, the ones that are well resolved may achieve the correct value. We see this in the well-defined structures of the rarefaction waves and contact discontinuity. Notice also that, as before, post-shock oscillations are absent. We have compared these results to other published results in the literature (for example by Balsara (2001) and Del Zanna et al. (2003)) and found them to be consistent.

2.4.5 Circularly-Polarized Alfvèn Wave

A propagating circularly polarized (CP) Alfvèn wave is a well-known analytical, nonlinear solution of the multidimensional MHD system. Since the solution is smooth (contains no shocks), it is often used to measure the convergence of numerical scheme.

We consider a propagating CP Alfvèn wave on the Cartesian plane. Our setup for this test problem is similar to Tóth (2000) and Londrillo and Del Zanna (2004). The wave propagates at an angle α relatives to the x-axis. Periodic boundary conditions are used on the computational domain with inner and outer boundary at:

$$\begin{aligned} \left[0 \leq x \leq \frac{1}{\cos\alpha}, 0 \leq y \leq \frac{1}{\sin\alpha}\right] \quad & \text{for } \alpha > 0, \\ [0 \leq x \leq 1, 0 \leq y \leq 1] \quad & \text{otherwise.} \end{aligned} \tag{2.94}$$

Let the coordinate along the direction of propagation be

$$\xi = x\cos\alpha + y\sin\alpha, \tag{2.95}$$

and the coordinate along the transverse direction be

$$\eta = y\cos\alpha - x\sin\alpha, \tag{2.96}$$

The initial values for the fluid velocity and magnetic field are given as

$$v_\eta = B_\eta = A\sin(2\pi\xi), \tag{2.97}$$

$$v_z = B_z = A\cos(2\pi\xi), \tag{2.98}$$

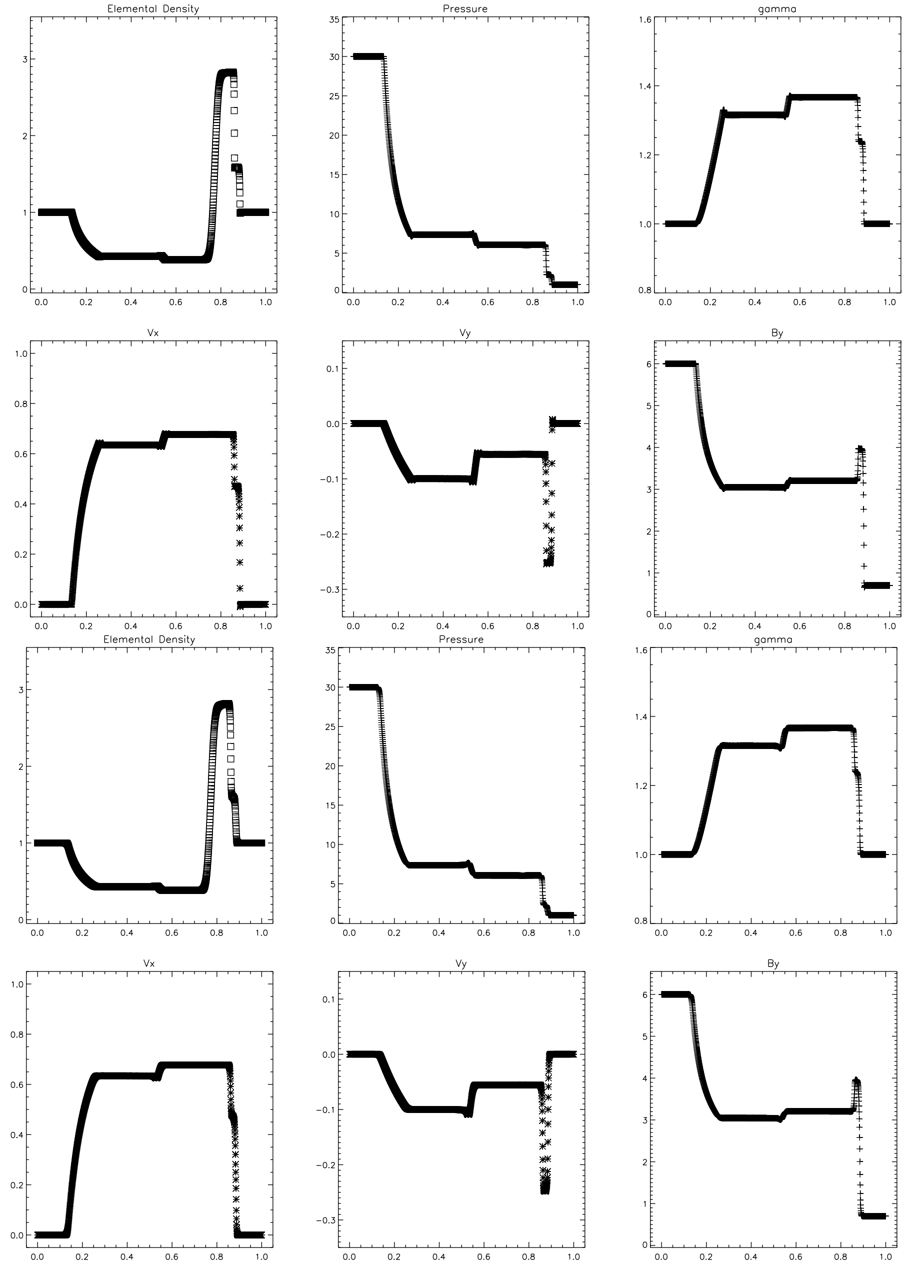

Elemental Density Pressure gamma

Vx Vy By

Elemental Density Pressure gamma

Vx Vy By

Figure 2.13: 1D mildly relativistic magnetized blast wave problem at $t = 0.4$ with 1600 cells and a Courant number $C = 0.5$. The panels shows, from left to right, top to bottom, the density, pressure, Lorentz factor γ, v_x, v_y, B_y. On the top half panels we show results computed with MC slope limiter. The results with Minmod slope limiter are shown on the bottom half panels.

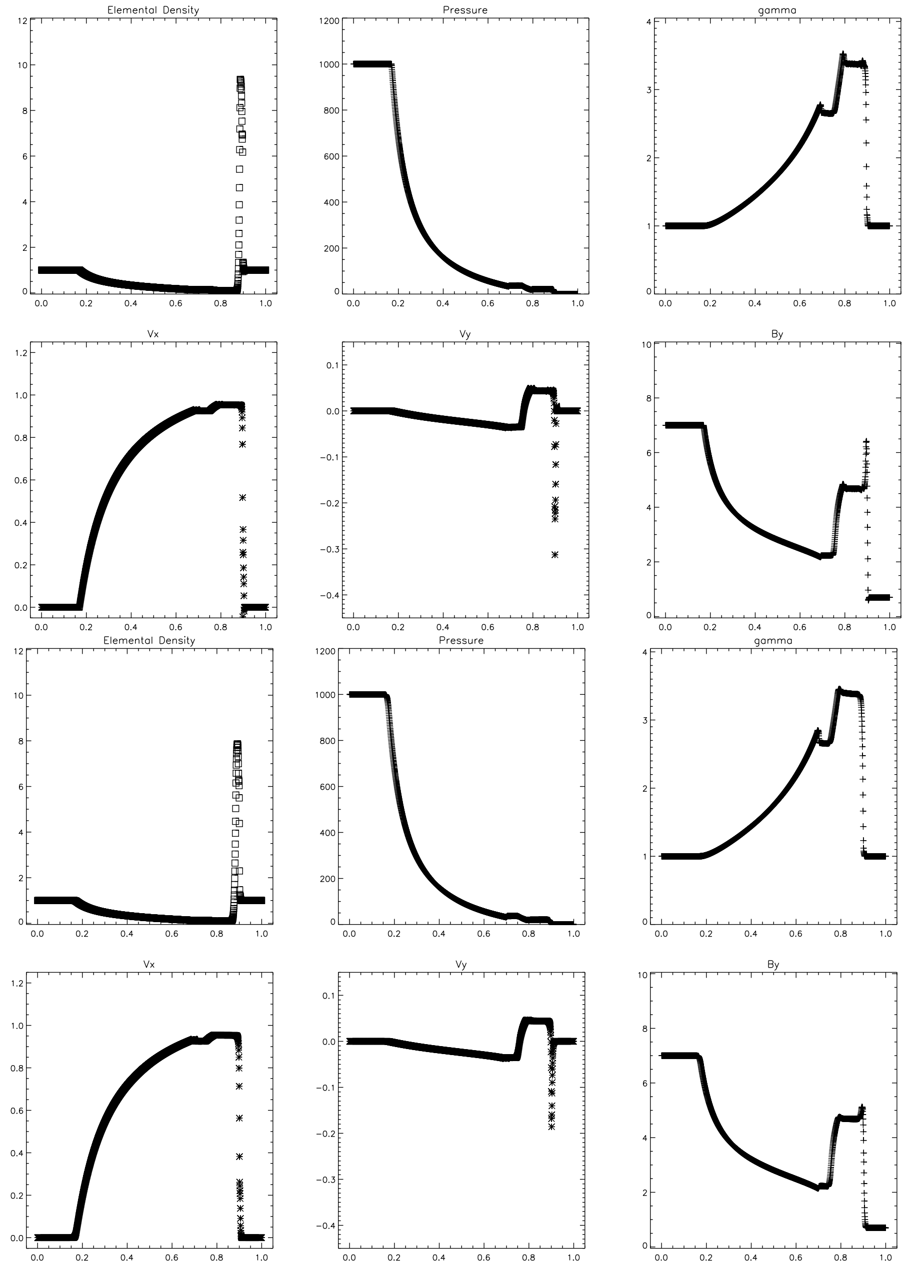

Figure 2.14: 1D relativistic magneticed blast wave 2 at $t = 0.4$ with 1600 cells using a Courant number $C = 0.5$. Panels arrangement are as described before in figure 2.13.

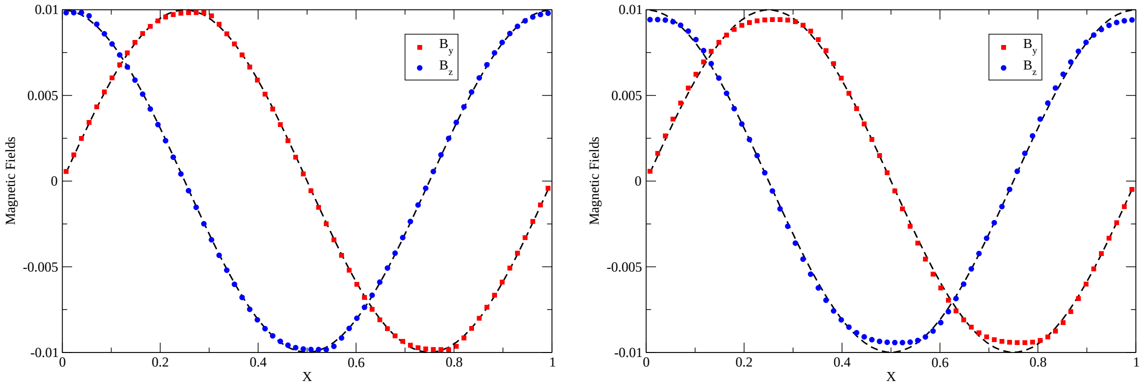

Figure 2.15: Magnetic fields of a circularly-polarized Alfvén wave after one period of evolution with 64 cells. The black dotted lines shows the initial magnetic field at $t = 0$; the symbols show the magnetic field at $t = T = 1$ after being evolved with MC (left) and Minmod (right) slope limiter. The wave's dissipation is due to numerical inaccuracies, which should decrease as resolution increases. The MC limiter is less dissipative then Minmod, but we can see its steepening in the direction of wave's propagation, causing a slight overshoot after the wave's optima.

where A measures the wave's amplitude. The parallel component for the fluid velocity and magnetic field are set to

$$v_\xi = 0, \tag{2.99}$$

$$B_\xi = 1. \tag{2.100}$$

We set the amplitude to $A = 0.01$ with uniform density $\rho = 1$, pressure $p = 0.1$, and adiabatic index $\Gamma = 5/3$. These values correspond to a wave of period $T = 1$, with the propagation Alfvènic speed $\lambda_A = 1$.

To check for numerical accuracy, we compare the initial conditions to the numerical solutions after some arbitrary number of periods n. We measure the convergence rate by measuring the relative error using equation 2.70 where χ is $\left\{B_y, B_z, v_y, v_z\right\}$ for some time $t = nT$ and χ_0 is their initial values at $t = 0$.

Figure 2.15 shows the profile of magnetic fields for the 1D version of this test problem, obtained by setting the angle of propagation to $\alpha = 0°$. The second-order convergence rate for this configuration is shown in figure 2.16.

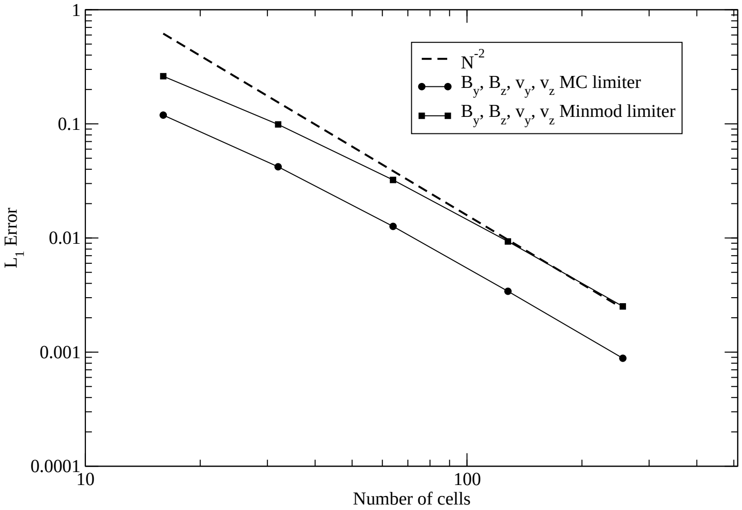

Figure 2.16: Convergence plot for 1D CP Alfvén wave with MC and Minmod slope limiter

For the 2D version, we set the propagation angle to $\alpha = 30^\circ$. This is a truly multidimensional problem since the x- and y-fluxes are different because that the wave's propagation is not along the diagonal of the computational boundary. Figure 2.17 shows the magnetic fields at initial conditions and at some selected later time. The second-order convergence rate for this 2D version is shown in 2.18.

2.4.6 Relativistic Circularly Polarized Alvèn Wave

In the limit of small amplitude where transverse relativistic effects may be neglected, the circularly polarized Alvèn Wave in previous §2.4.5 is still a valid solution to the relativistic MHD system. Here, we use it to investigate the convergence property of the relativistic MHD module in GenASiS.

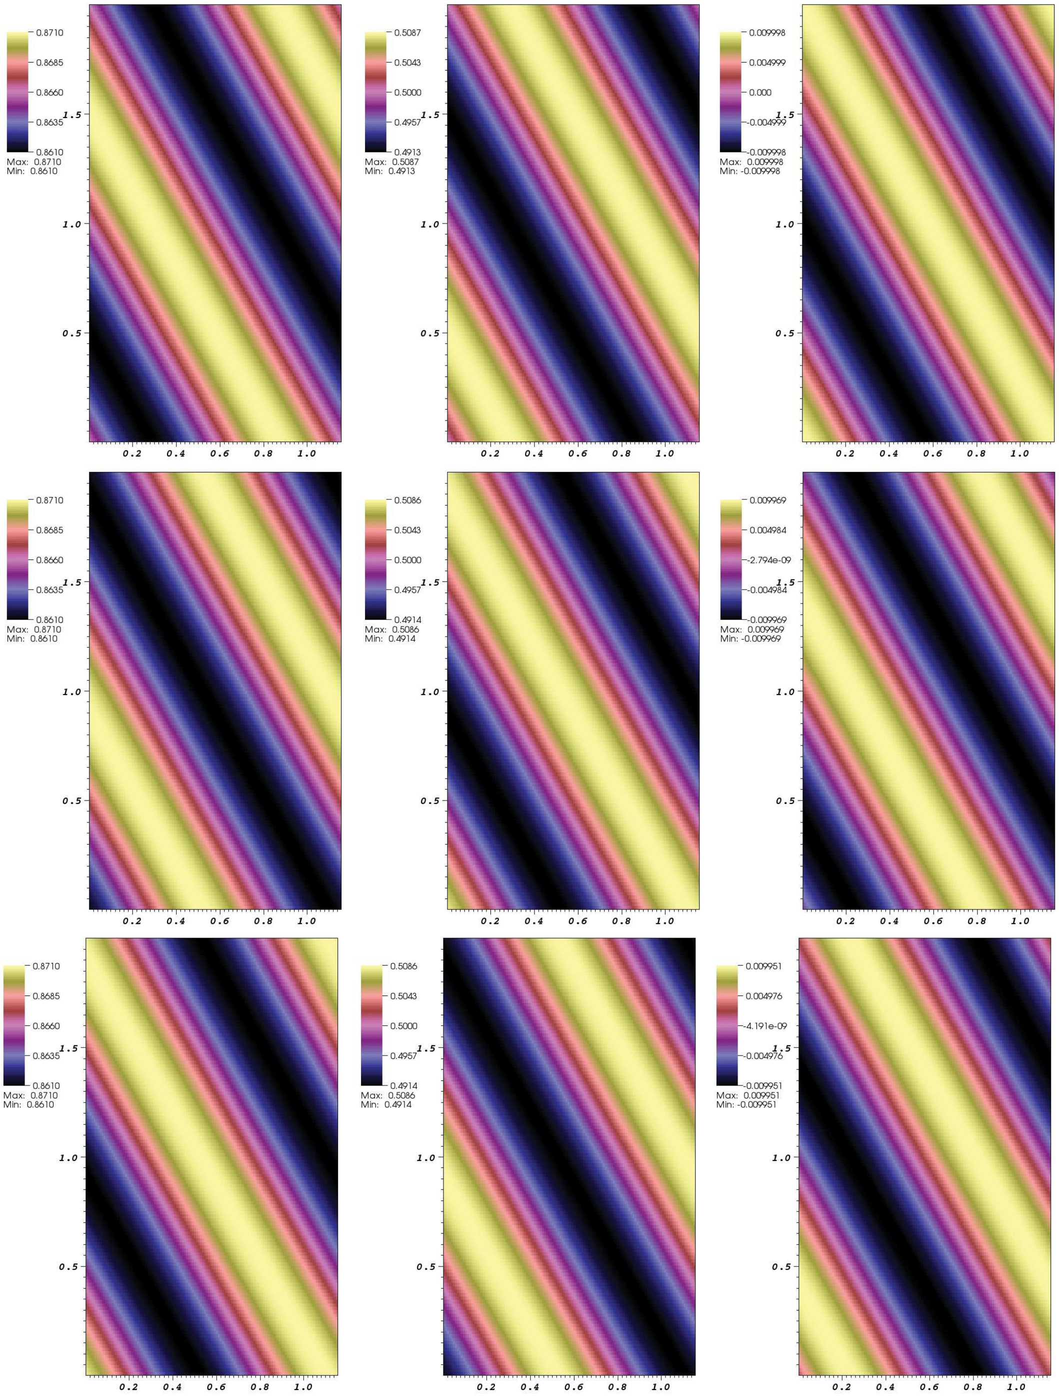

Figure 2.17: Pseudocolor plot of magnetic fields for 2D CP Alfvèn wave at time $t = 0$, $t = 0.3$, and $t = 0.8$ from top to bottom computed with 128×128 cells.

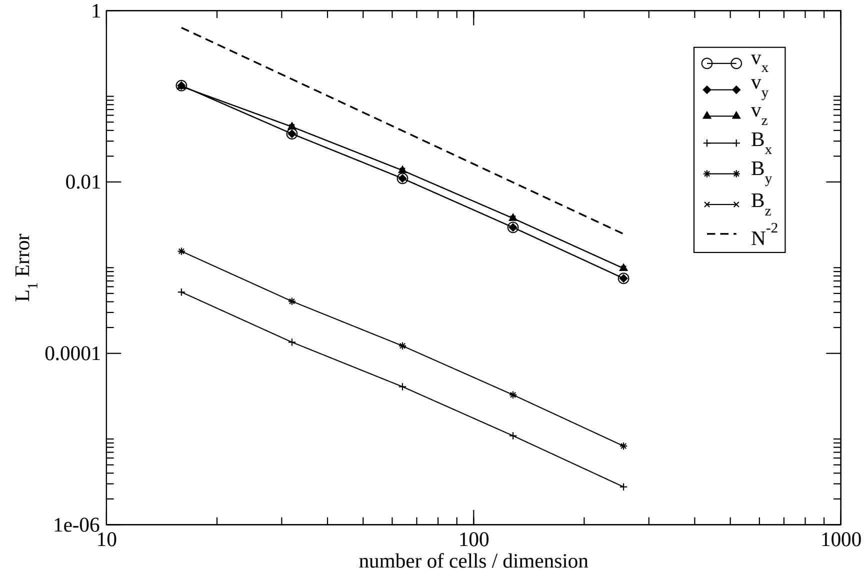

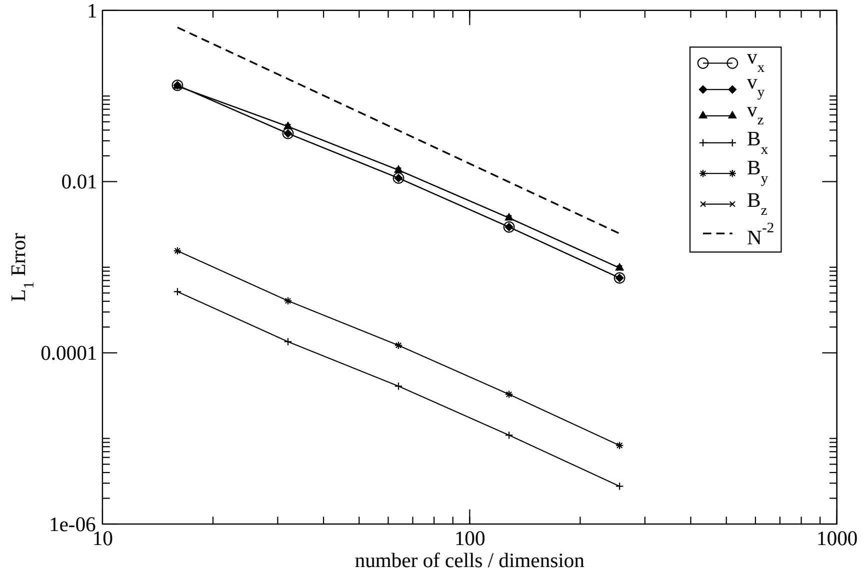

Figure 2.18: Convergence plot for 2D CP Alfvèn wave with MC (top) and Minmod (bottom) limiter for various quantities.

Following Del Zanna et al. (2003), we use initial conditions slightly modified from §2.4.5. Define a generic Cartesian reference frame: (ξ, η, ζ) and set:

$$B_\xi = B_0 = 1, \tag{2.101}$$

$$v_\xi = v_0 = 0, \tag{2.102}$$

$$v_\eta = -B_\eta = A\cos(2\pi\xi), \tag{2.103}$$

$$v_\zeta = -B_\zeta = A\sin(2\pi\xi), \tag{2.104}$$

where as before $A = 0.01$ is the amplitude of the wave. We also set the uniform density $\rho = 1$ and pressure $p = 1$. For the 1D case we have $(\xi, \eta, \zeta) = (x, y, z)$, which yields the following initial conditions for the fluid velocity and magnetic fields:

$$[v_x, v_y, v_z] = [0, A\cos(2\pi x), A\sin(2\pi x)], \tag{2.105}$$

$$[B_x, B_y, B_z] = [1, -A\cos(2\pi x), -A\sin(2\pi x)], \tag{2.106}$$

on the computational domain $[0, 1]$. For the 2D case, we consider the propagation along the $x = y$ direction; thus, we have

$$(\xi, \eta, \zeta) = \left((x+y)/\sqrt{2}, (-x+y)/\sqrt{2}, z\right). \tag{2.107}$$

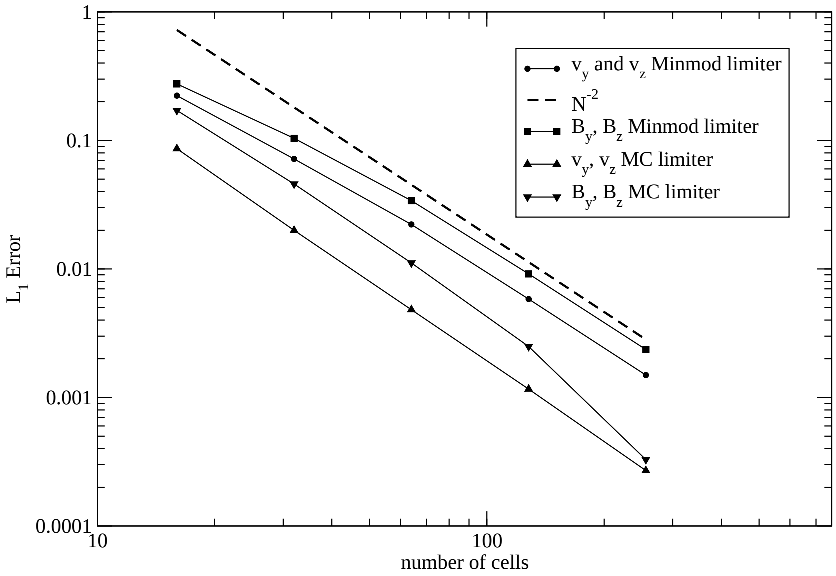

Figure 2.19: Convergence plot for 1D CP Alfvén wave with MC and Minmod slope limiter in relativistic MHD.

This gives us the following values for the fluid velocity and magnetic fields:

$$
\begin{aligned}
v_x &= -A\cos\left(2\pi(x+y)/\sqrt{2}\right)/\sqrt{2}, && (2.108)\\
v_y &= A\cos\left(2\pi(x+y)/\sqrt{2}\right)/\sqrt{2}, && (2.109)\\
v_z &= A\sin\left(2\pi\left(x+y\right)/\sqrt{2}\right), && (2.110)\\
B_x &= \frac{B_\xi - A\cos\left(2\pi(x+y)/\sqrt{2}\right)}{\sqrt{2}}, && (2.111)\\
B_y &= \frac{B_\xi + A\cos\left(2\pi(x+y)/\sqrt{2}\right)}{\sqrt{2}}, && (2.112)\\
B_z &= -A\sin(2\pi(x+y)/\sqrt{2}), && (2.113)
\end{aligned}
$$

on the computational domain $x = [0, 1/\cos(\pi/4)]$ and $y = [0, 1/\sin(\pi/4)]$.

As before, we measure the L_1 norm relative error of the magnetic fields after one period as compared to the initial condition. Figures 2.19 and 2.20 plot this for the 1D and 2D cases as a function of mesh resolution, which gives us the convergence of the scheme.

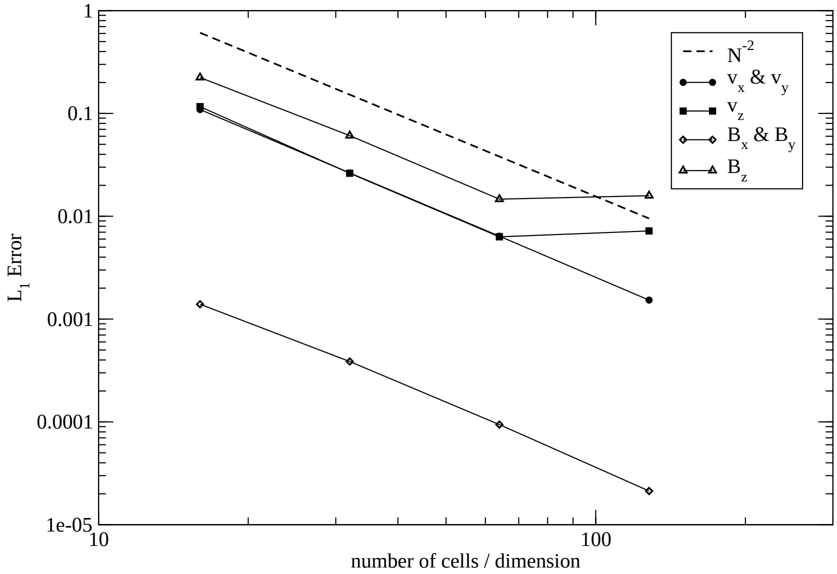

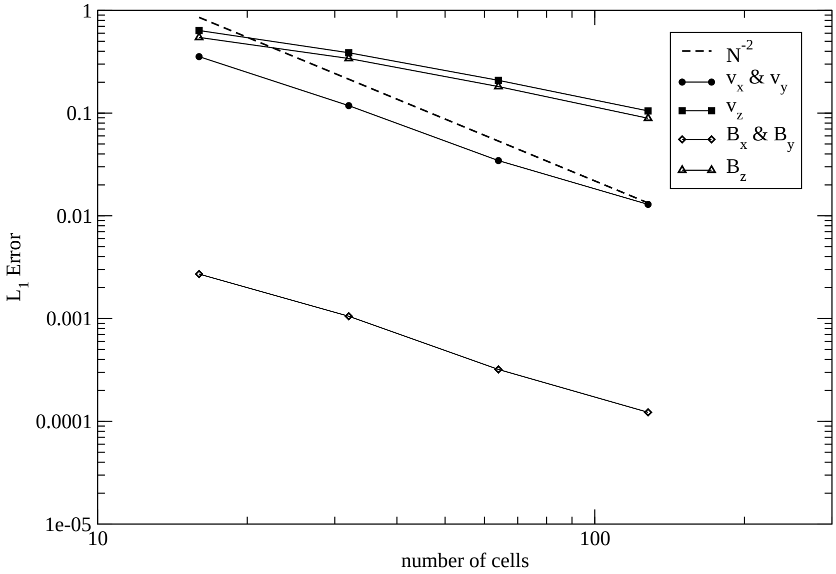

Figure 2.20: Convergence plot for 2D CP Alfvèn wave with MC (top) and Minmod (bottom) limiter for various quantities in relativistic MHD.

Chapter 3

Poisson's Equation Solver

3.1 Introduction

Some physics simulations require the solution of Poisson's equation with an isolated source distribution and a vanishing boundary condition at infinity. A common example is the calculation of the Newtonian gravitational potential of a self-gravitating astrophysical system. Poisson's equation is

$$\nabla^2 \Phi(\mathbf{x}) = S(\mathbf{x}), \tag{3.1}$$

where $S(\mathbf{x})$ describes the known distribution of the source that generates the potential $\Phi(\mathbf{x})$. For instance, $S(\mathbf{x})$ is proportional to the mass density in the context of Newtonian gravity, and to the charge density in electrostatics.

Several methods exist to solve the discretized Poisson's equation on a uniform grid. These include, for example, multigrid methods, iterative / relaxation methods, several matrix methods, and methods that employ Fourier transforms (for discussion of some these, see for instance Hockney and Eastwood (1989); Swarztrauber (490); Dorr (1970)). Here we implement, and extend to three dimensions, a particular method of the latter class by Hockney (1970) (see also Hockney and Eastwood (1989) and Eastwood and Brownrigg (1979)). (Another well-known approach for isolated systems based on Fourier transforms by James (1977), also discussed in Hockney and Eastwood (1989), would not be as straightforward to parallelize and is not discussed here.) An advantage of this approach is that discrete Fourier transform algorithms have been well-studied, with the Fast Fourier transform (FFT) being the most commonly employed; it requires $O(n \log n)$ operations, where n is the number of elements to transform. Several FFT implementations, some freely available, also exist as libraries suitable for end-users.

The key issue addressed by the implementation described here is the parallelization of an FFT-based algorithm for solving Poisson's equation for an isolated system. Obtaining such solutions in three dimensions requires resources that at present are most commonly available in distributed-memory parallel computers, where large problems are decomposed into multiple spatial sub-domains and distributed across different processes. While Poisson's equation must be solved globally on the computational domain and across multiple processes, most FFT implementations require that all data be accessible on a single process; therefore data movement and coordination across multiple processes are key ingredients of our FFT-based approach.

We use the FFTW library (Frigo and Johnson, 2005) to compute FFT, but our use of it is abstracted in such a way that other FFT libraries could be used without having to make widespread changes throughout the code. We use MPI to manage data movement and communication across processes, but our calls to message passing routines are abstracted as well.

This chapter is organized as follows. Section 3.2 outlines the algorithm as well as implementation details specific to our code. Test problems illustrating the convergence properties and performance of our implementation are presented in Section 3.4. Section 3.5 contains concluding remarks.

3.2 Solution Method

3.2.1 Formulation

Our problem is to solve equation (3.1) with the boundary condition $\Phi(\mathbf{x}) \to 0$ as $|\mathbf{x}| \to \infty$ (vanishing boundary condition). Use of the Green's function

$$G(\mathbf{x}) = -\frac{1}{4\pi\,|\mathbf{x}|}, \tag{3.2}$$

which satisfies

$$\nabla^2 G(\mathbf{x}) = \delta(\mathbf{x}) \tag{3.3}$$

and obeys the vanishing boundary condition, yields the formal solution

$$\Phi(\mathbf{x}) = \int d\mathbf{x}'\, G(\mathbf{x}-\mathbf{x}')\, S(\mathbf{x}'). \tag{3.4}$$

This integral may be evaluated with the help of the convolution theorem. Given the Fourier transforms $\tilde{G}(\mathbf{k})$ and $\tilde{S}(\mathbf{k})$ of $G(\mathbf{x})$ and $S(\mathbf{x})$, the Fourier transform of the potential is

$$\tilde{\Phi}(\mathbf{k}) = \tilde{G}(\mathbf{k})\,\tilde{S}(\mathbf{k}). \tag{3.5}$$

The potential $\Phi(\mathbf{x})$ is then obtained by an inverse Fourier transform.

When the Fourier transforms are to be done with FFTs, use of a regular mesh with the same spacings in each dimension is most natural; but in principle it should be possible to use any mesh for which a coordinate transformation can bring the mesh point positions to triplets of integers. For instance, to allow for a regular mesh with numbers of mesh points n_x, n_y, n_z and unequal mesh point spacings h_x, h_y, h_z in the three dimensions, Eqs.

(3.2)-(3.4) become

$$
\begin{aligned}
G(\bar{\mathbf{x}}) &= -\frac{h_x h_y h_z}{4\pi\sqrt{h_x{}^2\bar{x}^2 + h_y{}^2\bar{y}^2 + h_z{}^2\bar{z}^2}}, && (3.6)\\
\left(\frac{1}{h_x{}^2}\frac{\partial^2}{\partial\bar{x}^2} + \frac{1}{h_y{}^2}\frac{\partial^2}{\partial\bar{y}^2} + \frac{1}{h_z{}^2}\frac{\partial^2}{\partial\bar{z}^2}\right) G(\bar{\mathbf{x}}) &= \delta(\bar{\mathbf{x}}), && (3.7)\\
\Phi(\bar{\mathbf{x}}) &= \int d\bar{\mathbf{x}}'\, G(\bar{\mathbf{x}} - \bar{\mathbf{x}}')\, S(\bar{\mathbf{x}}'), && (3.8)
\end{aligned}
$$

where the values of the transformed coordinates $\bar{\mathbf{x}}$ corresponding to the mesh points are triplets of integers ranging from 0 to $n_x - 1,\ n_y - 1,\ n_z - 1$ in the three dimensions respectively. (Note that the Jacobian of the coordinate transformation has been absorbed into the numerator of equation (3.6), with the denominator still being 4π times the distance from the origin.)

The implementation of boundary conditions at infinity on a necessarily finite computational domain can be handled 'exactly', that is to say, with only discretization error, via a mesh doubling procedure and use of the standard periodic form of the discrete Fourier transform (Hockney, 1970) (see also Hockney and Eastwood (1989) and Eastwood and Brownrigg (1979)). Figure 3.1 illustrates this in two dimensions. The 'active' portion of the mesh corresponds to the original computational domain, while the 'inactive' portions are those arising from doubling the extent of the mesh in each dimension. The source distribution is set to its known physical values in the active region, and to zero in the inactive regions. Indexing cells by integer triplets p, q, r (the position of the mesh points in transformed coordinates $\bar{\mathbf{x}}$), the Green's function in the active and inactive regions is

$$
\left.
\begin{aligned}
G_{p,q,r} &= -h_x h_y h_z (4\pi)^{-1}\left(h_x{}^2 p^2 + h_y{}^2 q^2 + h_z{}^2 r^2\right)^{-1/2}\\
G_{2n_x-p,q,r} &= G_{p,2n_y-q,r} = G_{p,q,2n_z-r}\\
&= G_{2n_x-p,2n_y-q,2n_z-r} = G_{p,q,r}\\
G_{2n_x-p,2n_y-q,r} &= G_{2n_x-p,q,2n_z-r} = G_{p,2n_y-q,2n_z-r} = G_{p,q,r}
\end{aligned}
\right\}
\quad
\begin{aligned}
0 \le p,q,r \le n_x,n_y,n_z\\
p+q+r \ne 0
\end{aligned}
$$

$$
G_{0,0,0} = -h_x h_y h_z \left[4\pi \min(h_x, h_y, h_z)\right]^{-1}. \qquad (3.9)
$$

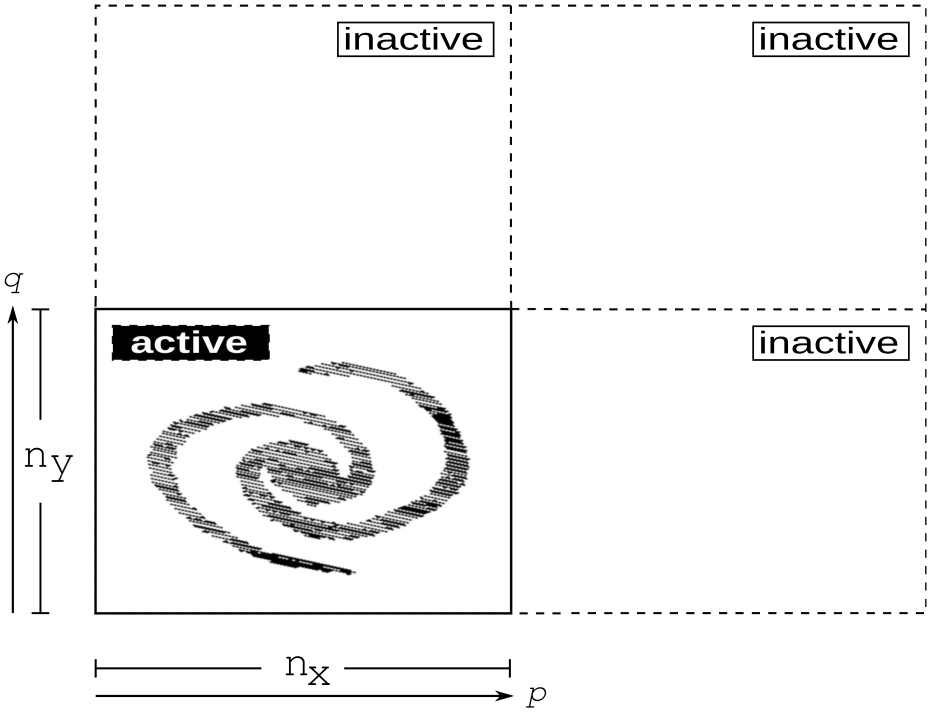

Figure 3.1: A two-dimensional illustration (redrawn after Hockney (1970)) of the mesh doubling used to solve for the potential due to an isolated source. The 'active' mesh on the lower left is the original computational domain with $n_x \times n_y$ cells.

That is, the Green's function in the active region follows equation (3.6), and is set up in the inactive regions in such a way that a periodic replication of the doubled mesh in all dimensions yields equation (3.6) in the entire region $-n_x, -n_y, -n_z \leq p, q, r \leq n_x, n_y, n_z$ surrounding the origin. The value for $G_{0,0,0}$ regularizes the singularity at the origin by assigning it to be equal to the largest off-origin value on the mesh; for $h_x = h_y = h_z = 1$, this reduces (up to conventions for sign and the 4π factor) to the prescription given in Hockney (1970).

3.2.2 Program Implementation

As with the rest of GenASiS, the implementation of the Poisson solver is also written in Fortran 95, using an object-oriented programming style (to the degree convenient and possible within that language). The library FFTW (Frigo and Johnson, 2005) provides our basic FFT functionality. We use MPI (MPI-Forum, 2010; Gropp and Lusk, 2010; Gropp et al., 1999) to implement parallelization across multiple processes.

3.2.2.1 Domain Decomposition and Storage

Our code assumes that the source $S(\mathbf{x})$ of Eq. (3.1) is available, and that the potential $\Phi(\mathbf{x})$ is desired, in a simple 'brick' decomposition: in three dimensions, the computational domain is divided in each dimension by $n_b = \sqrt[3]{n_p}$, the cube root of the number of processes n_p (see also §2.3 for more discussion of the decomposition). For simplicity we require $n_p = {n_b}^3$ to be a perfect cube (in three dimensions), and the number of mesh points in each dimension to be evenly divisible by n_b.

The brick decomposition is not convenient for FFTs, however, because a single transform is most naturally and efficiently performed on data accessible to a single process; therefore our solver has its working storage arranged in 'pillars' rather than bricks. The 'length' of what we term 'x pillars' spans the full extent of the computational domain in the x direction. The 'width' of the x pillars is their extent in the y direction, which is $1/n_b$ times the y extent of the bricks. This implies another constraint imposed by our solver: the number of mesh points n_y/n_b spanned by the y extent of a brick must itself be evenly divisible by n_b. The 'height' of the x pillars, which is their extent in the z direction, is the same as the extent of the bricks in the z direction. By similar construction (and with similar constraints on n_z and n_x), we have y pillars and and z pillars whose (width, height) are taken to be their extents in (z,x) and (x,y) respectively. These 'pillar decompositions' cover the same total volume and contain the same total number of mesh points as the brick decomposition, as illustrated in Fig. 3.2. Finally, we note that the lengths of the pillars are doubled as necessary to accommodate the mesh doubling procedure, so that the pillars span both the active and inactive portions of the mesh.

Because of the row-major nature of Fortran array storage, a pillar's length, width, and height correspond in our code to the first, second, and third dimensions of a rank-three array. This allows a (width × height)-number of one-dimensional FFTs to be performed

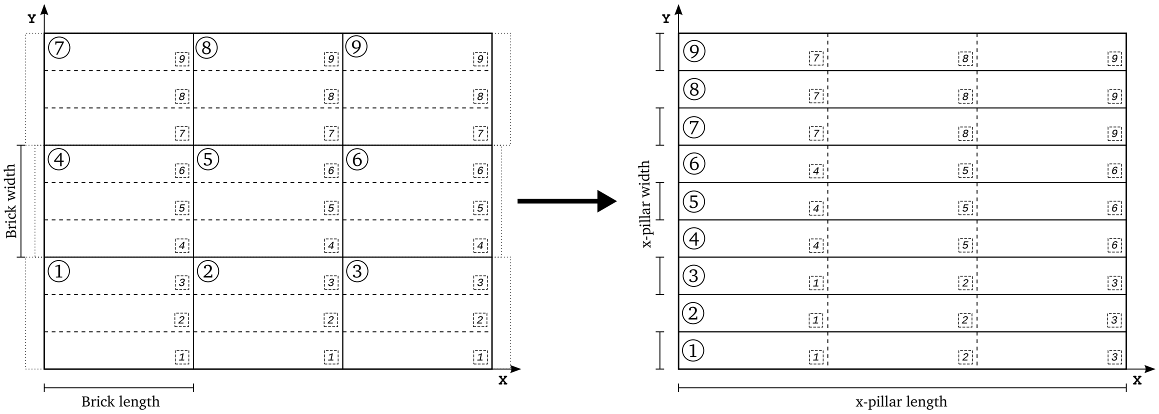

Figure 3.2: An illustration of the transformation from brick decomposition to x pillars from a three-dimensional mesh assigned to twenty-seven processes. Here only the first (lowest in the z direction) xy slab of bricks is shown; other slabs independently follow the same transformation. The left panel shows the whole computational domain decomposed into bricks, demarcated by solid lines and assigned to processes labeled by the numbers in solid circles. Dashed lines in the left panel mark the chunks of data that need to be sent to the processes labeled by the numbers in dashed squares in order to build the pillars. As indicated by the dotted boundaries, processes $[1,2,3]$, $[4,5,6]$, and $[7,8,9]$ form three separate groups, each with its own subcommunicator within which chunks of data are exchanged during the construction of the x pillars. In the right panel, we see that each process (again, labeled by numbers in solid circles) also owns a pillar. The boundaries between pillars are now marked by solid lines, and the dashed lines indicate the chunks of data that the process received from other processes labeled by numbers in dashed squares.

efficiently on contiguous data, specifically on lines containing a number of data points equal to the pillar length. The construction of pillars from bricks and vice-versa requires data movement across different processes. Using MPI, this is accomplished by creating a subcommunicator for each group of processes that need to communicate data among themselves, as illustrated in Fig. 3.2. For each group, a call to `MPI_AllToAll` can then be made with the group's subcommunicator in order to achieve the construction of pillars. These subcommunicators are saved to be used in the reverse process of deconstructing pillars back into bricks.

3.2.2.2 Multidimensional Transforms

A multidimensional FFT can be accomplished as a sequence of sets of one-dimensional transforms. The number of required operations is still of $O(n \log n)$, where $n = n_x n_y n_z$ is the total number of mesh points. One possibility for achieving computational efficiency is to transpose data between transforms in subsequent dimensions in order to achieve contiguity of memory access in each dimension. In any case, such transpose operations become a necessity in a distributed memory environment if parallelization of individual one-dimensional transforms is to be avoided.

The sequence of transforms and transposes is as follows. Data are initially loaded into the solver's x pillars: during initialization the Green's function is set up directly in the x pillars according to Eq. (3.9), while the source is transferred from the brick decomposition to the x pillars at the beginning of every solve. With data loaded in x pillars, multiple one-dimensional transforms in the x dimension are simultaneously performed by all processes. The y pillars are then populated, independently in separate xy 'slabs', as illustrated in figure 3.3. For each slab a separate MPI group with its own subcommunicator is created; thus there are $n_b = \sqrt[3]{n_p}$ subcommunicators, each of which has ${n_b}^2$ processes. Within each subcommunicator a call to `MPI_AllToAll` transposes the data from x pillars to y pillars so that FFTs can be performed in the y direction. Similar transposes in yz slabs allow FFTs to be performed in z pillars. Here the multiplication of the transforms of the Green's

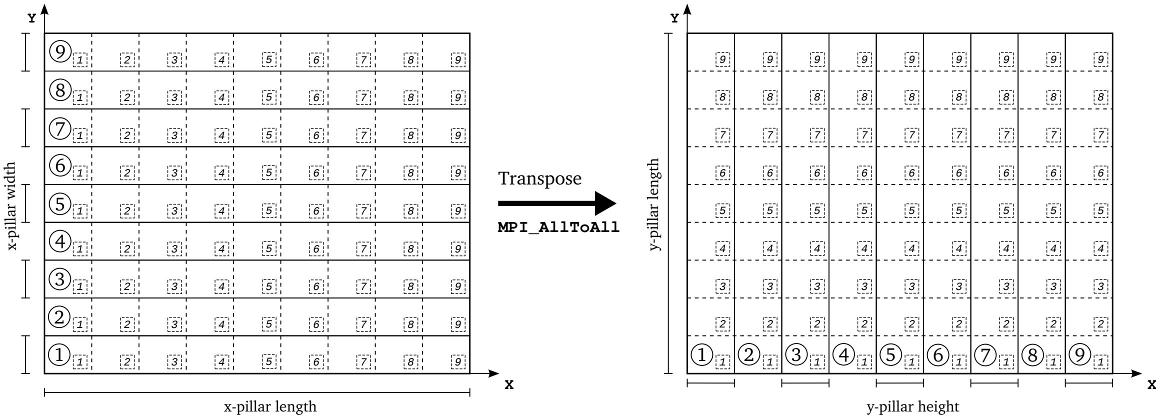

Figure 3.3: An illustration of the transpose of x pillars to y pillars on a three-dimensional mesh assigned to twenty-seven processes. As before, only the first (lowest in the z direction) xy slab is shown. The solid rectangles demarcate data owned by different processes, labeled by numbers in solid circles. Dashed lines mark chunks of data that need to be sent to (left panel) and received from (right panel) processes labeled by numbers in dashed boxes. In this example, a slab with 9 processes forms a single MPI group with its own subcommunicator, and the transpose is accomplished with a call to `MPI_AlltoAll`.

function and the source takes place as well, with the resulting Fourier-space solution of the Poisson equation overwriting the transformed source distribution. A reverse sequence of backward transforms and transposes gets the solution (modulo a normalization factor due to the multiple transforms) back into real space, stored in the x pillars. Finally the solution is redistributed from the active portion of the x pillars to the brick decomposition, overwriting the storage in which the source was delivered.

This sequence of transforms and transposes makes use of permanent storage for the source distribution in x, y, and z pillars, which at the end of the solve is reused to store the potential. This same storage is then updated with a new source distribution on the next call to the Poisson solver. The transform of the Green's function is computed only once, and stored permanently in z pillars, when the solver is initialized. Computation of the transformed Green's function requires x pillars and y pillars, but these are deallocated at the end of the initialization.

3.3 Coupling to Hydrodynamics

Gravitational potential is coupled to the hydrodynamics as momentum and energy sources in the conservation equations. To include these sources, we modify equations 2.1 – 2.3 as the following:

$$
\begin{aligned}
\frac{\partial D}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^i\right) &= 0, && (3.10)\\
\frac{\partial S^j}{\partial t} + \frac{\partial}{\partial x^i}\left(\rho v^j v^i + p\delta^{ij}\right) &= -\rho\frac{\partial \Phi}{\partial x^j}, && (3.11)\\
\frac{\partial E}{\partial t} + \frac{\partial}{\partial x^i}\left(\left[e+p+\frac{1}{2}\rho v^j v^j\right] v^i\right) &= -\rho v^i \frac{\partial \Phi}{\partial x^i}. && (3.12)
\end{aligned}
$$

The mass conservation equation 3.10 has not been modified from the original equation 2.1 and is only included here for clarity and completeness. The right-hand side of the momentum conservation equation 3.11 is the gravitational force, while we term the the right-hand side of the energy conservation equation 3.12 the ‘gravitational power’. Equations 3.10 – 3.12 describe self-gravitating fluid. The source terms retain the same forms in the MHD equations.

Because gravitational potential is smooth, its gradient is computed using second-order central difference without slope limiter on the grid. As in the case of the hydrodynamics variables, after the potential is obtained, the values in the cells that border process boundaries are exchanged to their nearest-neighbor cells that are owned by other processes via point-to-point communications (see §2.3) to facilitate the gradient computations of the gravitational potential in these cells. For the cells that border the computational domain, we use first-order forward or backward differences to compute the gradient of the gravitational potential for inner and outer boundaries, respectively. This results in a less accurate gravitational force and gravitational power near the computational boundaries, but we expect the domain boundaries to have negligible effect on the dynamics inside the computational domain.

The gravitational potential is computed in ‘operator splitting’ fashion. After initial conditions are set, the potential due to the mass distribution is computed. This is used as the source terms for the integration of the hydrodynamics equations. At the end of the timestep, new mass distribution is obtained, which is then used to compute new gravitational potential. Since we use second-order two-step Runge-Kutta scheme for the time-stepping algorithm (see §2.1.1), gravitational potential is computed after each Runge-Kutta step. In our experience this gives better and more accurate results than updating the potential only after the full Runge-Kutta integration.

Self-gravity introduces a new dynamical timescale to be considered. The free-fall time for every mass shell to reach the center ($r = 0$) of a spherical gas with uniform mass density ρ is (see also 3.4.4)

$$t_{\mathrm{ff}} = \sqrt{\frac{3Pi}{32G\rho}}. \tag{3.13}$$

In addition to the CFL condition (see 2.1.1), this gives a constraint for the size of timestep that can be taken for the integration of the conservation equations with gravitational sources to guarantee a stable algorithm. Therefore we take the timestep for the integration to be

$$\Delta t = \min\left\{\alpha t_{\mathrm{ff}}, \Delta t_{\mathrm{CFL}}\right\}, \tag{3.14}$$

where Δt_{CFL} is the timestep given by the CFL condition, and α is an arbitrary factor which we usually set to 1×10^{-2}.

3.4 Numerical Results

In this section we present some numerical results illustrated by several test problems. The test problems chosen were of similar nature with the ones commonly encountered in astrophysical simulations. The test problems described here have analytical solution, which we can use to verify the correctness of our implementation. We also investigate the numerical convergence of the implementation with respect to mesh resolution. We run each problem on a domain decomposed to large number of processes to illustrate the scaling behavior of the algorithm on parallel computers.

3.4.1 Gravitational potential of spherical uniform mass

We consider the potential of sphere with radius R that has uniform mass density ρ. The gravitational potential as a function of radius r from the center of the sphere has a simple analytical solution:

$$\Phi(r) = \begin{cases} G\pi\rho\left(\frac{2}{3}r^2 - 2R^2\right) & \text{for } r \leq R \\ -G\pi\rho\frac{4R^3}{3r} & \text{for } r > R. \end{cases} \tag{3.15}$$

We compute the potential for a sphere of radius $R = 1$ and mass density $\rho = 1$ in a Cartesian computational box with its inner and outer boundary at $[-1.2, +1.2]$ respectively in all dimensions. The sphere is centered on the origin of the coordinate system. For each mesh resolution, we calculate the potential in two ways: first, by using the analytical solution above with $r = \sqrt{x^2 + y^2 + z^2}$, and second, by using our implementation of the Poisson solver. By varying the mesh resolutions, we can check for the convergence properties of our solver with respect to spatial resolution. The potential for this test problem is shown in Figure 3.4.

We use the usual definition of L_1 norm to measure the relative error of potential computed by our solver compared to the analytical solution, as given in equation 2.70. Figure 3.5 illustrates the distribution of the relative error on the grid for a certain resolution, which is representative (by different constant factor) of the error distribution for other resolution.

Figure 3.6 shows the result of convergence test of our solver. We see that the convergence of the error trend is better than first order.

3.4.2 Gravitational potential of an homogeneous spheroid

A more general case of the previous test problem is the potential of a spheroid with a uniform mass-density distribution. Here we consider a spheroid with uniform density ρ. The spheroid was formed by an ellipse centered at the origin and rotated about the z axis,

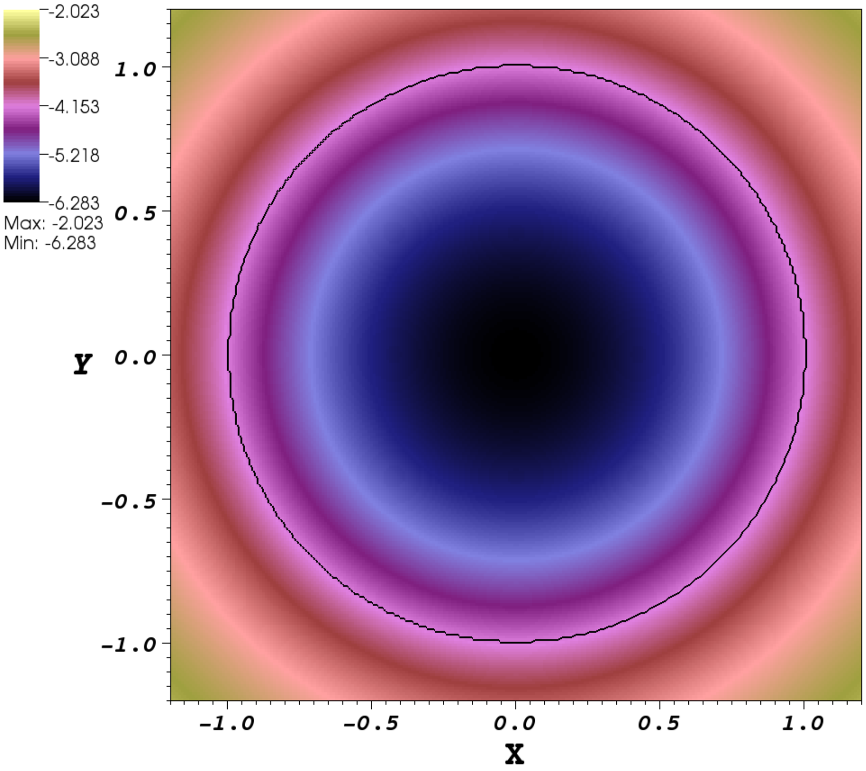

Figure 3.4: The potential of a spherical uniform mass with mass density $\rho = 1$. The figure is a slice through the three-dimensional mesh crossing the origin to show the xy-plane. The solid black line indicates the surface of the sphere at radius $R = 1$. The mesh resolution is 256 cells in each dimension.

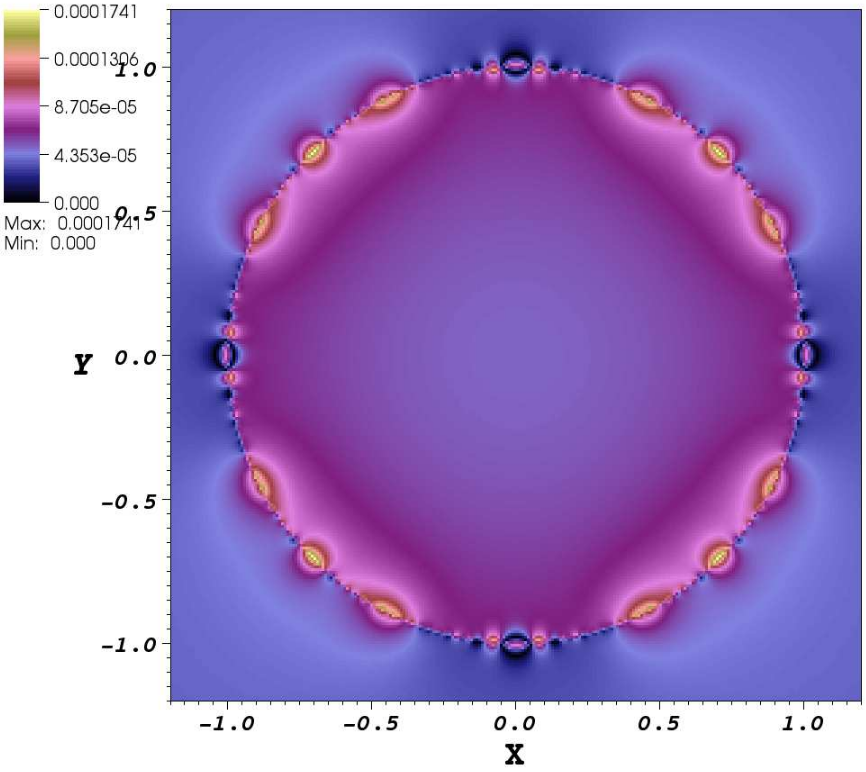

Figure 3.5: The relative error of the potential for spherical uniform mass as described by equation 2.70, but without the summation over all cells. The mesh resolution is 256 cells in each dimension. The figure is a slice through the mesh showing the *xy*-plane. We see that the largest errors are on the surface of the sphere because of the nature of the Cartesian grid.

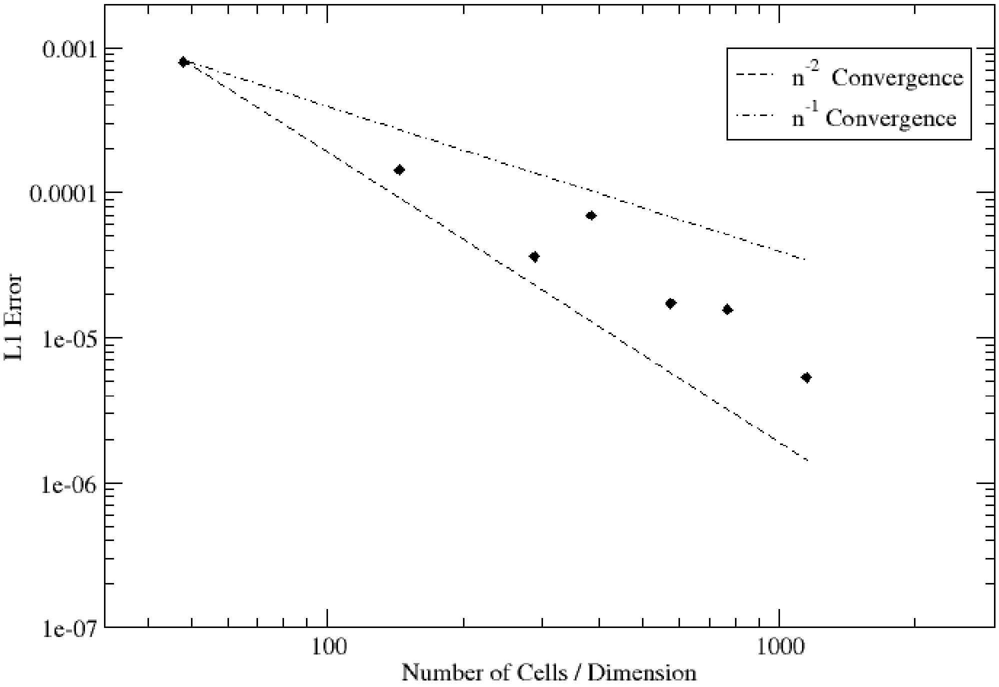

Figure 3.6: Convergence test of potential for spherical uniform mass. The L_1-norm relative error of the computed potential as compared to the analytical solution is plotted as function of the following mesh resolution: $[48^3, 144^3, 288^3, 384^3, 576^3, 768^3, 1152^3]$. The dashed and dot-dashed lines give reference for the second and first order error convergence respectively

and described by the equation

$$\frac{x^2+y^2}{a^2}+\frac{z^2}{b^2}=1, \tag{3.16}$$

where a and b are the semi-diameters of the spheroid. The spheroid is oblate when $a > b$, with its eccentricity defined as

$$e=\sqrt{1-\left(\frac{b}{a}\right)^2}. \tag{3.17}$$

The gravitational potential of a homogeneous spheroid is a simpler case of the potential of a homogeneous ellipsoid given in Chandrasekhar (1987). Inside the spheroid, it may be expressed as (Ricker, 2008):

$$\Phi(x,y,z)=-\pi G\rho\left[A\left(2a^2-x^2-y^2\right)+B\left(b^2-z^2\right)\right], \tag{3.18}$$

where

$$A=\frac{\sqrt{1-e^2}}{e^3}\sin^{-1}(e)-\frac{1-e^2}{e^2}, \tag{3.19}$$

$$B=\frac{2}{e^2}-\frac{2\sqrt{1-e^2}}{e^3}\sin^{-1}(e). \tag{3.20}$$

Outside the spheroid, the potential is given by

$$\begin{aligned}\Phi(x,y,z)&=-2\pi\rho G\frac{ab}{e}\\ &\times\left[\tan^{-1}(h)-\frac{1}{2a^2e^2}\left\{\left(x^2+y^2\right)\left(\tan^{-1}(h)-\frac{h}{1+h^2}\right)+2z^2\left(h-\tan^{-1}(h)\right)\right\}\right],\end{aligned} \tag{3.21}$$

with

$$h\equiv\frac{a\,e}{\sqrt{b^2+\lambda}}, \tag{3.22}$$

where λ is the positive root of the equation

$$\frac{x^2+y^2}{a^2+\lambda}+\frac{z^2}{b^2+\lambda}=1. \tag{3.23}$$

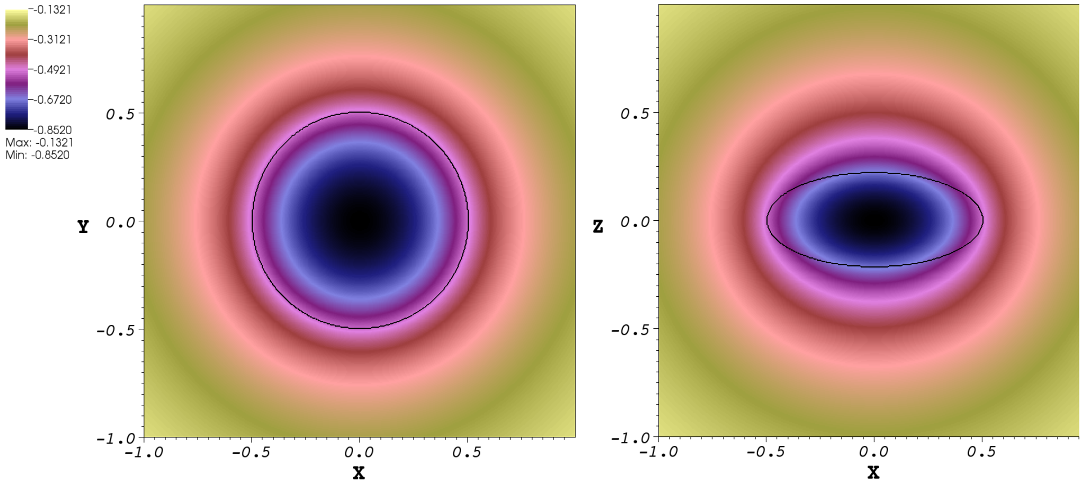

Figure 3.7: The potential of homogeneous spheroid with eccentricity $e = 0.9$ and semi-major axis $a = 0.5$ on a mesh with spatial resolution of 384 cells in each dimension. The figures are slices of the mesh through the origin showing both the xy- and xz-plane. The solid black line indicates the surface of the spheroid.

We compute the potential for a spheroid with eccentricity $e = 0.9$ and semi-major axis $a = 0.5$ on a Cartesian computational box of size two in each dimension. As before, we use a fixed uniform density $\rho = 1$. Figure 3.7 shows the computed potential for a certain mesh resolution.

As in the previous test problem, we vary the mesh resolutions to test for the convergence of the relative error L_1-norm of the numerical solution as compared to the analytical solution. This is shown in figure 3.9. Figure 3.8 illustrate the relative error distribution of the numerical solution as compared to the analytical solution. From figure 3.9, we see that our solver has higher than first order convergence, but less than second order in general.

3.4.3 Gravitational potential of homogeneous binary spheroid

In this test problem, we simply put two homogeneous spheroids separated by some distance in our computational domain. We extend our computational domain in the x

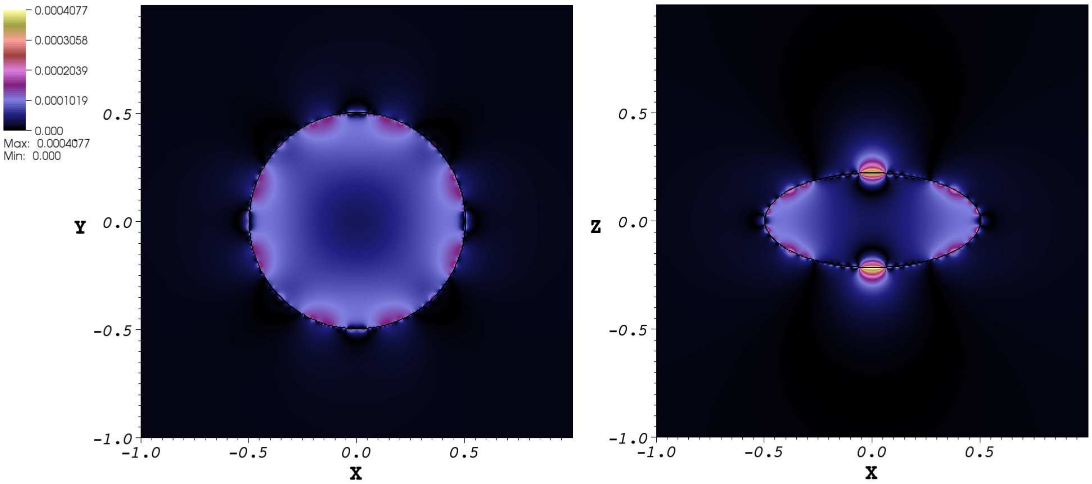

Figure 3.8: The relative error distribution of the homogeneous spheroid on a mesh with resolution of 384 cells in each dimension. Slices of *xy*- and *xz*-planes are shown. As before, the solid black line indicates the surface of the spheroid.

direction to be twice that used in the previous test problem so that the x dimension has inner and outer boundaries at coordinate $[-2,+2]$. The spheroids are then centered on coordinates $[\pm 1,0,0]$. In order to maintain the same effective resolutions as our previous test problem, we double the number of cells in the x dimension only, resulting in rectangular computational box. Figure 3.10 shows the potential for this test problem.

The gravitational potential for this configuration is the sum of the potentials due to both spheroid. Therefore we can generate the analytical solution for this test problem by modifying the analytical solution found in Section 3.4.2 to account for the shift of the spheroids' center from the coordinate origin. This is done by substituting $x-c$ for x in equations 3.18, 3.21, and 3.23, where c is the x coordinate of the center of spheroid.

As before, we vary the mesh resolution for this test problem to do convergence tests of our solver. This is shown in figure 3.11. We see a similar convergence trend of this test problem as we have seen in previous test problem: our solver converges better than first order, but not up to second order convergence.

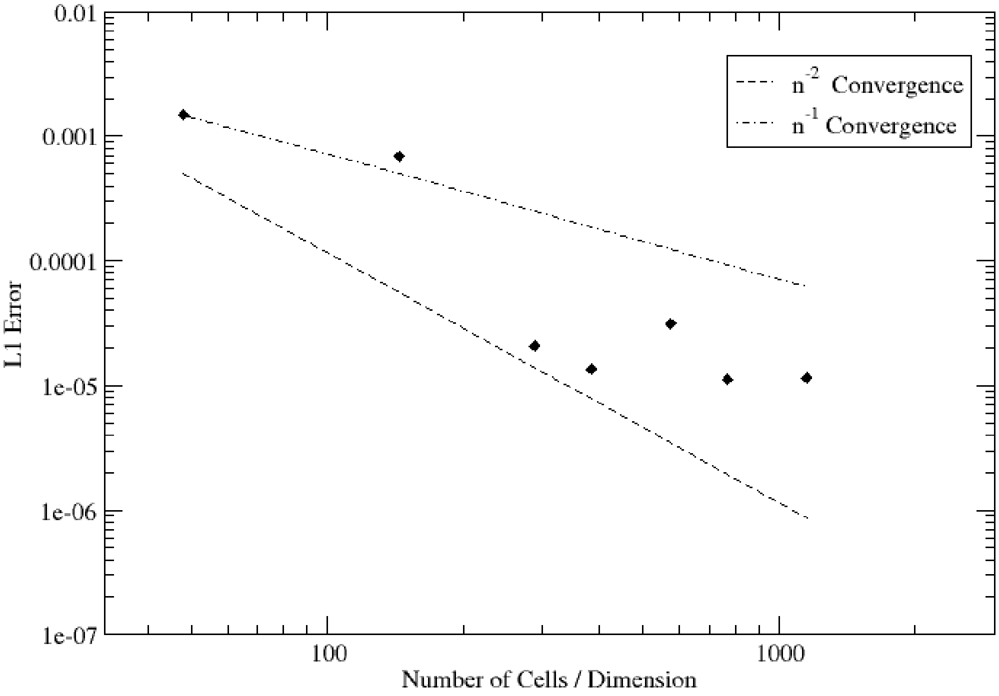

Figure 3.9: Convergence test of potential for a homogeneous spheroid with mesh resolutions $[48^3, 144^3, 288^3, 384^3, 576^3, 768^3, 1152^3]$.

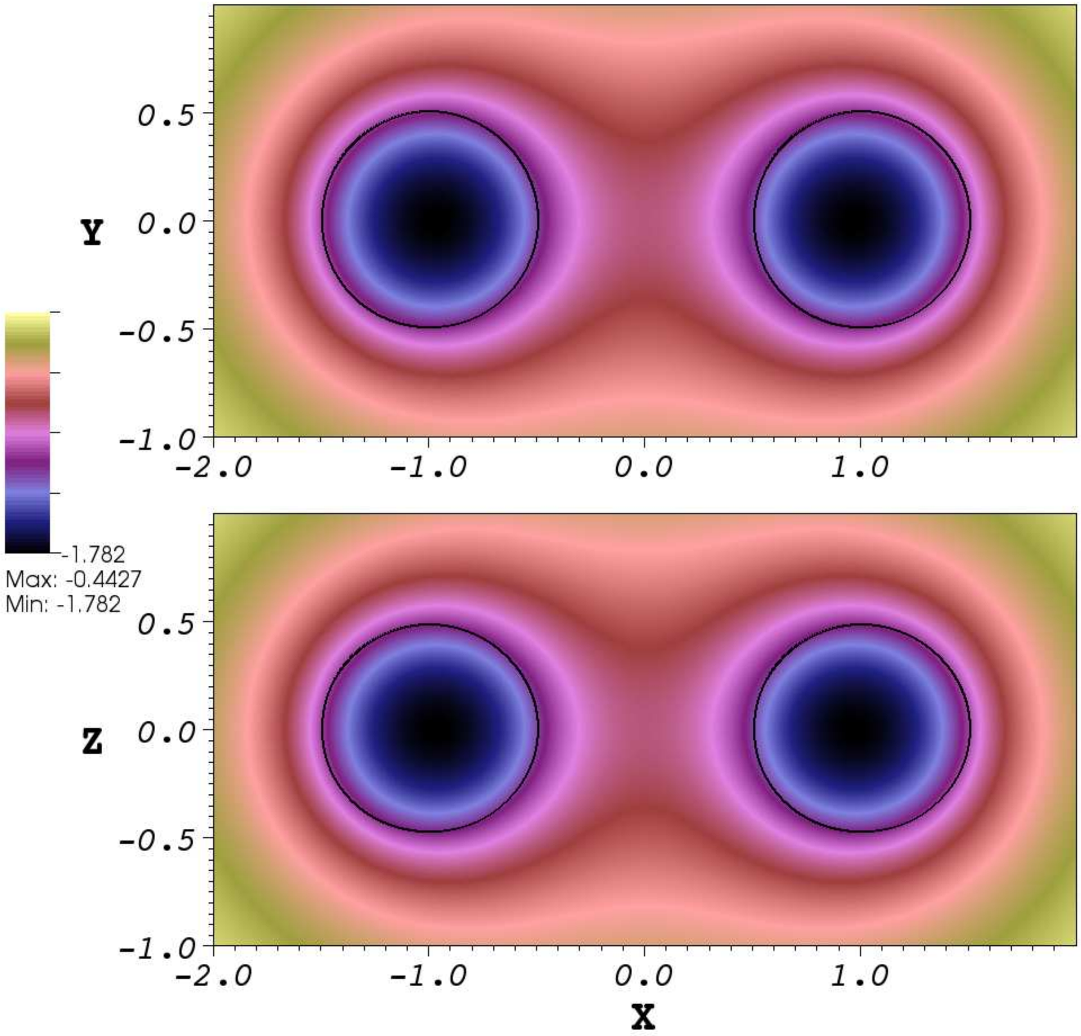

Figure 3.10: The potential for homogeneous binary spheroids. Each spheroid has mass density $\rho = 1$. The spheroid has very low eccentricity $e = 0.28$ with semi-major axis $a = 0.5$. Solid black lines indicate the surfaces of the spheroids. The mesh resolution is $768 \times 384 \times 384$ cells.

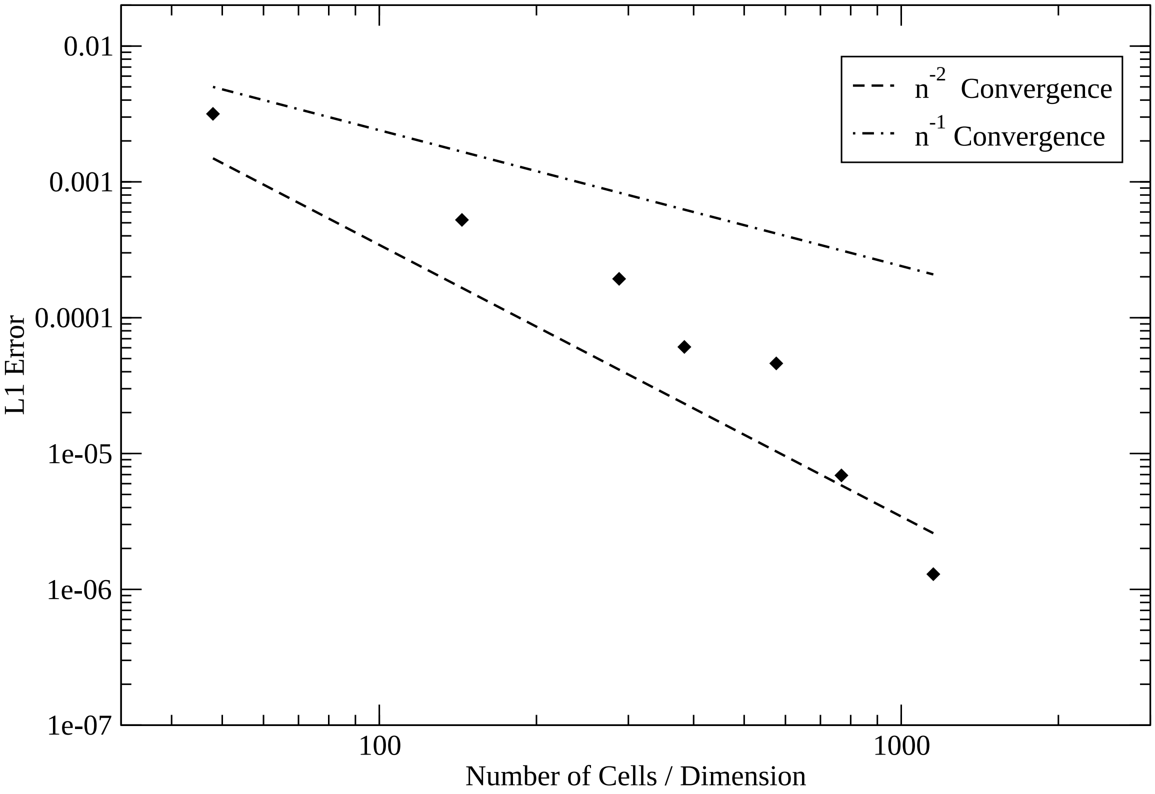

Figure 3.11: Convergence test of the potential for a binary spheroid with uniform mass with the same effective mesh resolutions as the previous test problems.

3.4.4 Pressureless (Dust) Collapse

In the previous test problems we have considered only the evolution of the magnetohydrodynamics conservation equations with the absent of source term. We have also shown test problems in which we test the accuracy and correctness of our Poisson's equation solver to solve for the gravitational potential due to a mass distribution in the static case. This test problem requires the evolution of the hydrodynamics conservation equations coupled with the gravitational source term due to the mass distribution. Therefore, it allows us to test both the Poisson's equation solver in dynamical fashion and our implementation of operator-splitting for the source term of the conservation equations.

We start with an initial condition of a homogeneous pressureless gas of initial density ρ_0 with radius r_0. Because there is no pressure support, the sphere collapses due to gravitational forces, starting with free-fall of the gas which increases the density of the sphere. Hunter (1962) and Stone and Norman (1992) give an analytic solution for the radius and density as a function of time t:

$$\begin{aligned} r &= r_0 \cos^2 \beta, & (3.24) \\ \rho &= \rho_0 \cos^{-6} \beta, & (3.25) \end{aligned}$$

where

$$\beta + \frac{1}{2} \sin 2\beta = t \sqrt{\frac{8\pi}{3} G \rho_0}. \tag{3.26}$$

We can then define the free-fall time, i.e. the time for every mass shell to reach $r = 0$, as

$$t_{\rm ff} = \sqrt{\frac{3\pi}{32 G \rho_0}}. \tag{3.27}$$

We initialize this test problem by setting a uniform sphere of gas with radius $r_0 = 1$ and density $\rho_0 = 1$ on 3D Cartesian mesh with domain boundary ± 1.2 in all dimension. This setup also allows us to check that the code does multidimensional evolution correctly. Since this problem has an inherent spherical symmetry, a 3D numerical evolution should also gives us a spherically symmetric result. For this setup, we find the free-fall time to be $t_{\mathrm{ff}} = 0.543$. We set the simulation to stop at $t = 0.535$, which gives us density increase of nearly three orders of magnitude from the initial value.

Figure 3.12 shows the results of numerical simulations as evolved by GenASiS for several mesh resolutions. With a unigrid mesh in this problem, as the sphere collapses, fewer cells are available to resolved the sphere with increasing density. This situation is illustrated clearly in figure 3.12a where only several cells are available to resolved the sphere due to the low resolution of the simulation, producing an inaccurate result as compared to the analytical result. This is improved by putting more cells on the computation domain to increase the mesh resolution. The central density of the sphere is better resolved and more accurate with higher resolution simulation, as shown on other panels in figure 3.12.

We see that there is an anomalous density spike near the density cliff. The most likely cause of this is the inaccuracy of the pressure gradient and gravitational force balance that manifests itself more in this region.

We check the convergence of our numerical relative results to the analytical solution by computing the L_1 norm relative error (equation 2.70) for each resolution. Figure 3.13 shows the L_1 norm relative error of the numerical to the analytical solution as function of mesh resolutions. This shows that the numerical solution has less than first order convergence. The bad convergence is due to the anomalous density spike previously mentioned.

3.4.5 Hydrostatic Polytrope

In this test problem, we investigate the code's ability to maintain a hydrostatic equilibrium solution numerically during dynamical evolution. As in the previous test problem, this test requires the evolution of hydrodynamics equations with the gravitational source term. A polytropic equation of state is used in this test problem; therefore, unlike the dust collapse problem, the energy conservation equation is also evolved in this test problem.

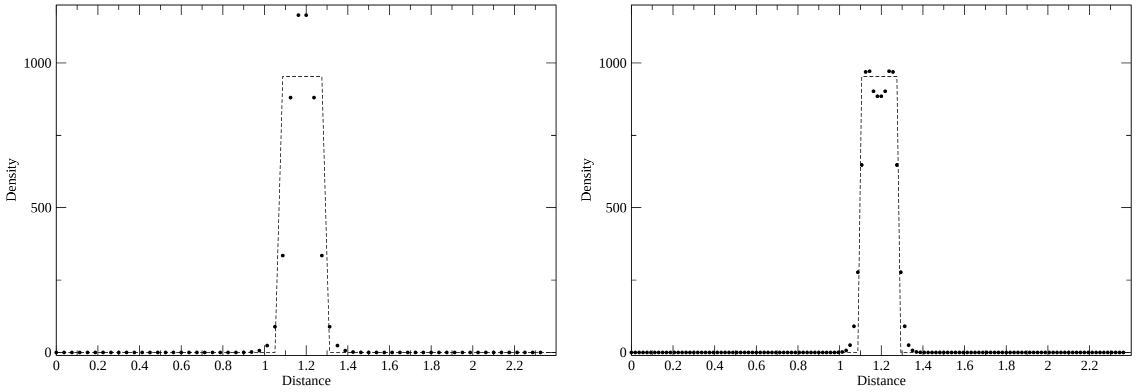

(a) 64^3 cells (b) 128^3 cells

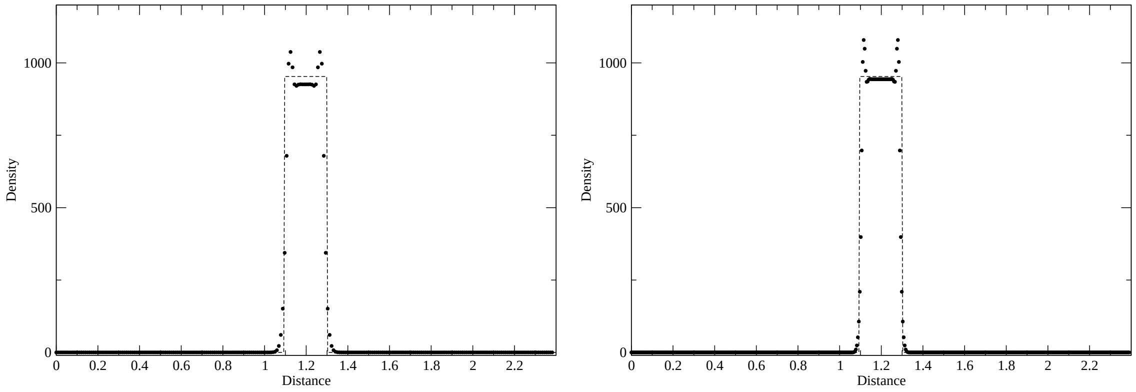

(c) 256^3 cells (d) 512^3 cells

Figure 3.12: Pressureless collapse with various mesh resolutions. The dotted line shows the density as computed by the analytical formula. The filled circles show the numerical result as evolved by GenASiS with slope limiter parameter $\theta = 2.0$.

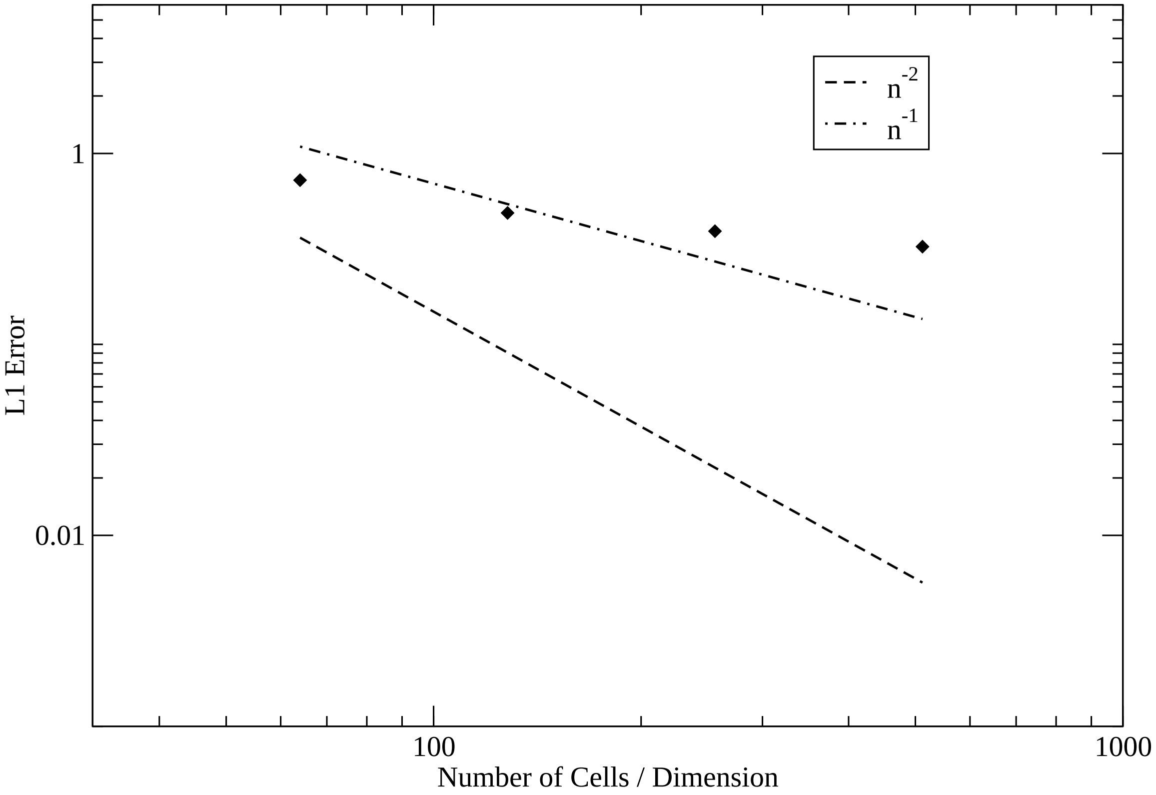

Figure 3.13: Convergence test of the pressureless collapse test problem. The L_1 norm relative error is plotted as function of mesh resolutions

The equations for hydrostatic equilibrium describe the balance of gravitational force that compresses the gas and pressure gradient force that expands the gas. This can be expressed as the differential equations:

$$\begin{aligned}\frac{dm}{dr} &= 4\pi r^2 \rho(r) dr, && (3.28)\\ \frac{dp}{dr} &= -G\frac{m(r)}{r^2}\rho(r), && (3.29)\end{aligned}$$

with may be combined to give the second-order differential equation:

$$\frac{1}{r^2}\frac{d}{dr}\left(\frac{r^2}{\rho(r)}\frac{dp}{dr}\right) = -4\pi G\rho(r), \tag{3.30}$$

where $m(r)$ is the enclosed mass as function of radius r. Equation 3.30 is closed when one specifies an equation of state to relate the pressure and density. In this case we consider a polytropic equation of state as described in §2.2.

By introducing the dimensionless variables ξ and θ, equation 3.30 can be rewritten to yield:

$$\frac{1}{\xi^2}\frac{d}{d\xi}\left(\xi^2\frac{d\theta}{d\xi}\right) + \theta^n = 0, \tag{3.31}$$

where

$$\rho = \rho_c\theta^n \qquad \xi = r\left(\frac{4\pi G}{(n+1)\kappa\rho_c^{(1-n)/n}}\right)^{\frac{1}{2}}, \tag{3.32}$$

where $\rho_c \equiv \rho(r=0)$ is the central density, and n is the polytropic index. This is known as the Lane–Emden equation.

In terms of the new variables, the differential equations may be solved using the boundary conditions:

$$\theta(0) = 1, \qquad \left.\frac{d\theta}{d\xi}\right|_{\xi=0} = 0, \tag{3.33}$$

which imply that the central density $\rho_c = \rho(0)$ be reproduced and that the pressure gradient $\frac{dP}{dr} = 0$ at the origin. Equation 3.32 can then be integrated from the origin, where $r = \xi = 0$, until the point where $\xi = \xi_1$ where θ vanishes. This variables can be transformed back to relate to physical quantities of interest. For example, the radius R and total mass M are given by:

$$R = a\xi_1 = \left(\frac{(n+1)\kappa}{4\pi G}\right)^{\frac{1}{2}} \rho_c^{\frac{1-n}{2n}} \xi_1, \tag{3.34}$$

$$M = 4\pi a^3 \rho_c \left[-\xi^2 \frac{d\theta}{d\xi}\right]_{\xi=\xi_1}. \tag{3.35}$$

For polytropic index $n = 1$, the Lane-Emden equation has analytical solution:

$$\theta = \frac{\sin\xi}{\xi}, \qquad \xi_1 = \pi. \tag{3.36}$$

Therefore by specifying the total mass M and radius R of a gaseous sphere one can get a profile of density and pressure as function of radius that stays in hydrostatic equilibrium.

Using the analytical solution previously described, we initialize the computational domain with a hydrostatic polytrope star of mass $1.5M_\odot$ with radius 10 kilometers and polytropic index $n = 1$—a configuration that mimics a typical neutron star—on a computational domain with its inner and outer boundary coordinate at ± 20 kilometers in all dimensions. The central density of this objects is about 2.34×10^{15} gm cm^{-3}, located at the origin of the coordinate. We evolve this initial condition numerically up to 10 milliseconds, a timescale that is relevant to physical simulation such as neutron star mergers. We compare the density and pressure profile at the end of the evolution with our initial condition and calculate the L_1 norm relative error.

The analytical solution produces continuous but steep density and pressure drops near the surface of the star that go to zero. This and the atmospheric density surrounding the star, which is set very close to zero, would often produce a high velocity matter but very small density, due to inaccuracy in the numerical computation. Left alone, this would cause the simulation to take very small time steps due to the CFL condition. To avoid this, for this test problem we put a condition in our code (which is adjustable through a runtime parameter) to set the velocity to zero for density below a certain threshold. This threshold is set to 3.0×10^5 gm cm^{-3}, ten orders of magnitude below the central density.

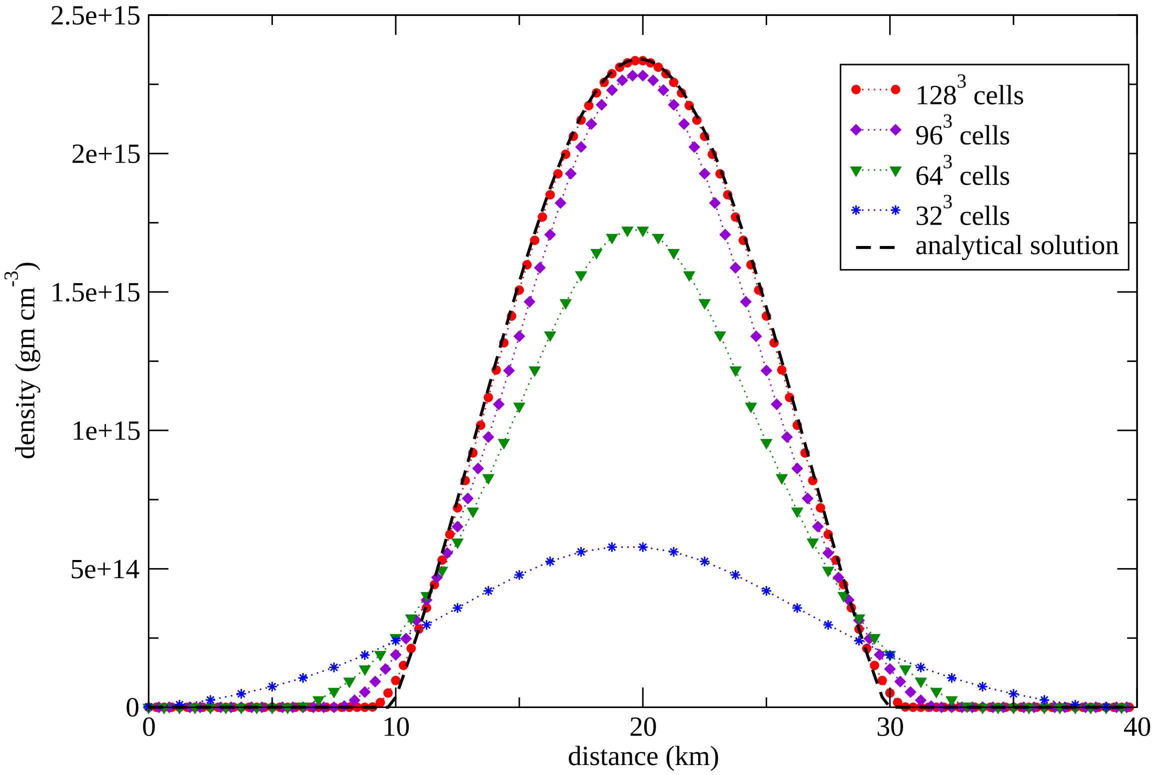

Figure 3.14: Density profile for hydrostatic polytrope. The symbols show the profiles for several different mesh resolution after being evolved for 10 milliseconds. The analytical solution is plotted as a dashed line for reference.

Figure 3.14 shows the density profile of this test problem with several different resolutions after being evolved for 10 milliseconds. We see that at the end of the simulation the star is slightly diffused out, accompanied by a decrease of the central density, as compared to its initial condition. This is more pronounced in the lower resolution simulation because of the higher inaccuracy in the balance between gravity and pressure gradient. Another way to see this is by plotting the central density of the star as a function of time. Figure 3.15 shows such a plot for several different resolutions. We see that for simulation with total number of cells less than 96^3, which correspond to approximately 0.41 kilometer spatial resolution, the central density has a decreasing trend. This gives a minimum of spatial resolution needed for any simulation that uses hydrostatic model.

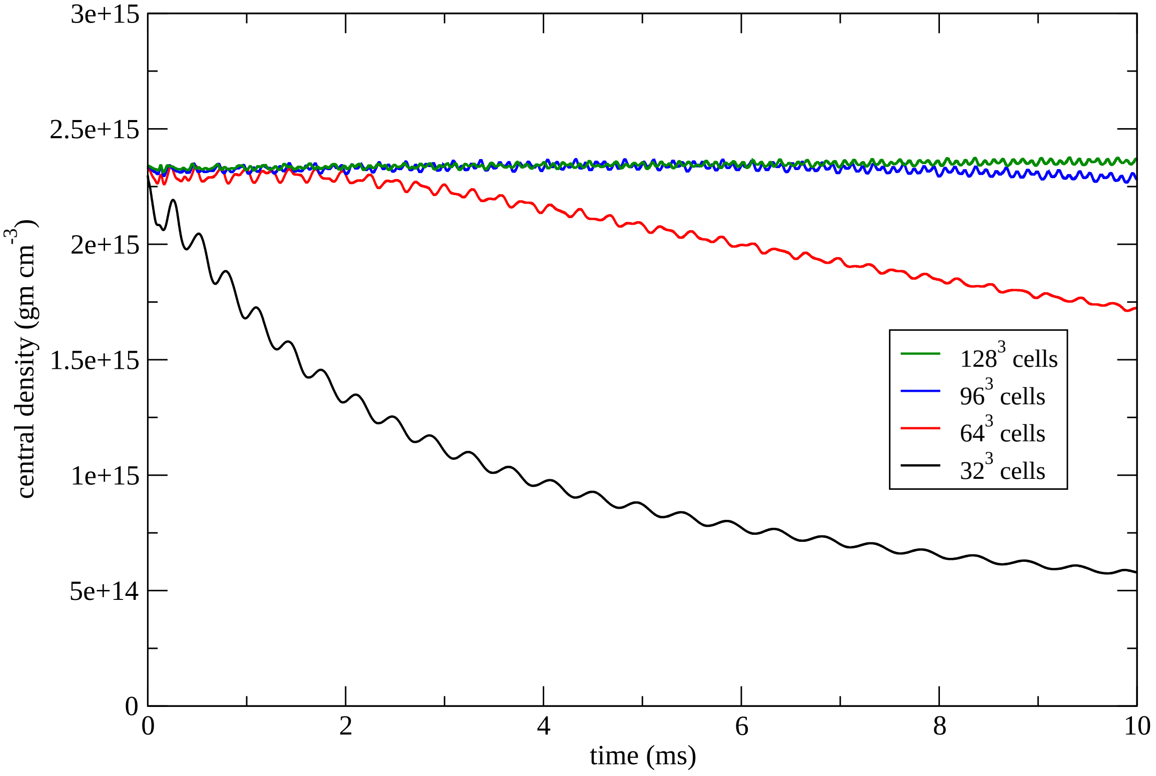

Figure 3.15: Central density of a hydrostatic polytrope plotted as function of time for several different resolutions.

The oscillation of the central density is expected, as a direct result of numerical solution with a finite resolution. Slight discrepancies in the balance between pressure gradient and gravity may cause the star to diffuse if the pressure gradient wins initially. As the star diffuses out, the pressure gradient becomes smaller. This tips the balance toward gravity, and let it counteract by compressing the star. As the gravity keeps compressing the star and increases the density, the stiff equation of state results in a large enough pressure gradient to counteract this, and so in the next time step it is able to counteract gravity and expand the star. This cycle continues, which results in the oscillation that we see in figure 3.15. For sufficient resolution, the net result should be a small and damped oscillation (albeit very slowly) around a constant central density. Figure 3.15 shows this property, and therefore shows that GenASiS is able to maintain hydrostatic equilibrium for some dynamical timescale provided that a sufficient mesh resolution is used.

We measure the error convergence of our code for this test problem. Figure 3.16 shows the L_1 norm relative error of the density profile at $t = 10$ milliseconds as compared to the initial density profile as function of resolutions. As expected for test problems with smooth flows, we get close to second-order error convergence.

3.4.6 Scaling

We measure the weak scaling of our solver by increasing the number of parallel processes as we increase the mesh resolution, thereby maintaining a constant amount of work assigned to each process. We use the homogeneous spheroid and binary spheroid test problems as representatives for the weak-scaling test. For each measurement we run the solver 2500 times and take the average to get the time per solve. The scaling tests were carried out on a Cray XT-4 with quad-core 2.1 GHz AMD Opteron 1354 (Budapest) processors and 8 GB of DDR2-800 memory per node. Figure 3.17 shows the weak scaling plot for the solver. The black square symbols with a solid line show the scaling for the homogeneous spheroid problem; the red circle symbols with a solid line indicate the scaling for the binary spheroid problem. For comparison, we also plot the theoretical weak-scaling curve for the algorithm

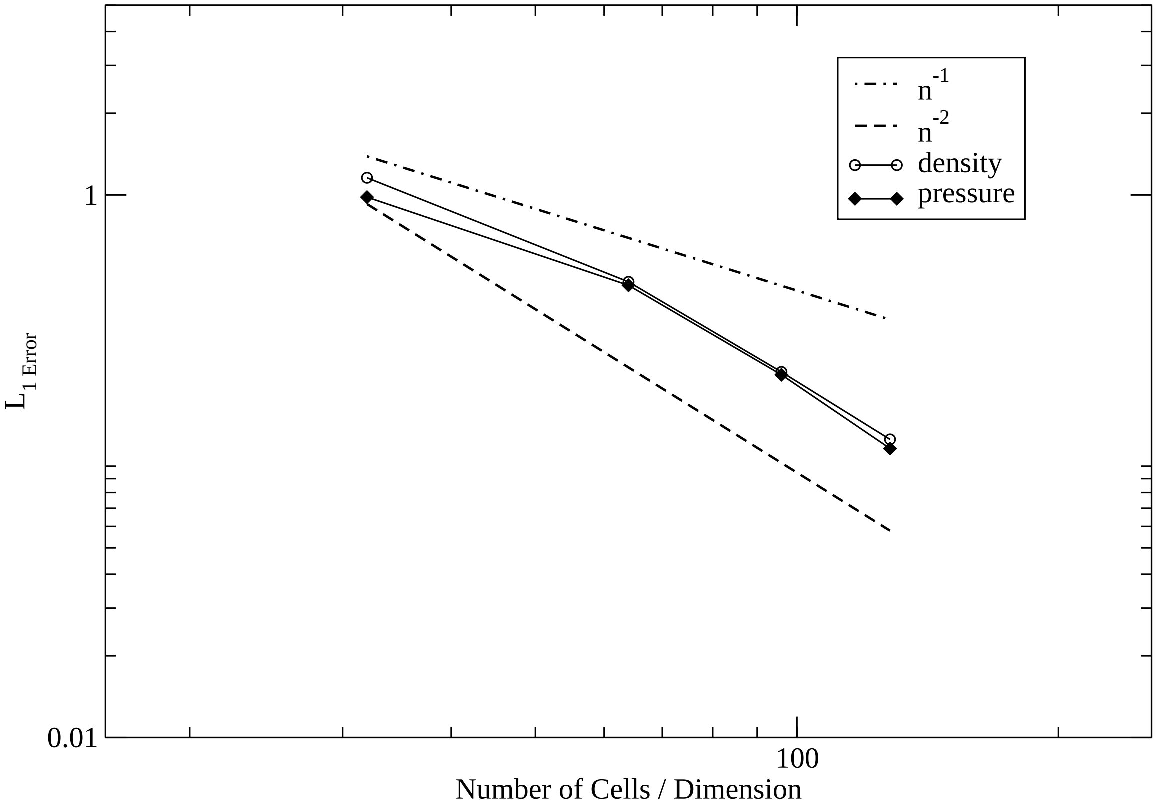

Figure 3.16: Convergence rate of hydrostatic polytrope test problem. The L_1 norm relative error of density, and pressure are plotted as function of mesh resolutions. Reference lines (dashed lines) for second-order and first-order convergence are also shown.

with $O(n\log(n))$ property as dotted black line for each test problem. Specifically, we plot $a(\log(n)/3 + \log(b))$ where n is the number of processes and a and b are constant. The rationale for this is as follow. To investigate weak scaling we maintain the amount of work per process. Therefore the algorithmic dependent on the number of cells to be transformed is a proportional to the number of processes. The second term, $\log(b)$ is this proportionality, where we have used the identity $\log(xy) = \log(x) + \log(y)$. We set $b = 2 \times 48$ for the homogeneous spheroid test problem since 48 is the number of cells per process in one dimension, and the factor 2 comes from the mesh doubling procedure. (For similar reason, $b = 2 \times 96$ for the binary spheroid test problem, since the binary spheroid test problem uses twice as many cells as the homogeneous spheroid problem to maintain the same effective resolution, as described in §3.4.3.) The factor $1/3$ on the first term comes from the fact that each process the transform on the pillars whose lengths are proportional to the cube root of the total number of processes. In other word, this is the number of cells to be transformed by each process in one dimension spanning the entire computational domain. Here we have used the identity $\log(n^{1/3}) = \log(n)/3$. The factor a is an arbitrary constant that could be hardware or algorithmic dependent, but independent of the number of operations. This affects the rate of increase or the slope of the algorithmic growth. We determine a in our theoretical curve simply by matching the first point of the theoretical curve to the experimental value. The experimental curves includes the computation and communication required to do a single solve. Therefore, if we assume that a remains constant in the hardware and algorithm as number of processes grows, we can attribute the increase in solution time for the experimental curves to the communication required to do the bricks to pillars construction (and vice-versa), and the transpose during multidimensional Fourier's transform. However, the slope of increase is relatively very small and we find this to be acceptable.

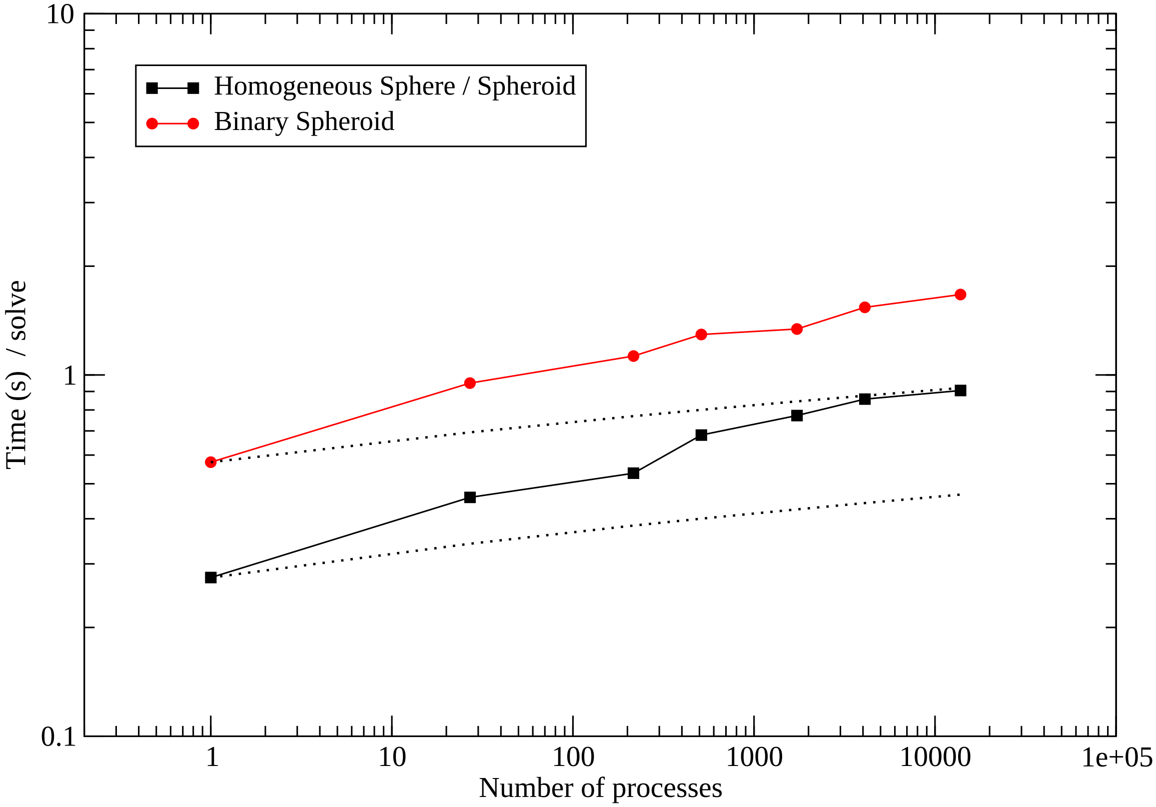

Figure 3.17: Weak scaling of the FFT Poisson solver. The homogeneous sphere and spheroid test problem assigns 48 × 48 × 48 cells per process, while the binary spheroid test problem assigns 96 × 48 × 48 cells per process due to the doubling of computational domain in one dimension.

3.5 Conclusion

We have described our implementation of a parallel solver for Poisson's equation of an isolated system on unigrid mesh. We use Fourier's transform as the solution method for Poisson's equation. We utilize a common protocol Message Passing Interface for communication across processes on distributed memory system to do a global solve on domain decomposed to multiple processes. We have also shown test problems, numerical convergence, and scaling behavior of our program. The implementation of this solver, coupled with magnetohydrodynamics, allows us to do simulations with self-gravity in GenASiS.

Chapter 4

Towards Adaptive Mesh Refinement

4.1 Introduction

Adaptive Mesh Refinement (AMR) is a technique to dynamically increase the spatial resolution of numerical simulations only where needed. This allows conservation of memory and computational cost of the simulations. This saving has become a necessity since advances in scientific modeling have reached a point where the required resolution in uniform meshes is taxing and even exceeding the resources of the largest supercomputers. AMR provides a way to ameliorate this computational demand by only increasing resolution where needed to capture the physical processes of interest.

There are two different approaches to AMR that have seen widespread usage in computational astrophysics. In block-structured AMR (Berger and Oliger, 1984; MacNeice, 2000), cells are organized as blocks of grid cells. The coarsest grid, consisting of several blocks, covers the entire computational domain. Each block may be refined, creating finer nested grids where higher resolution is required. Since each grid is composed of structured blocks, any single grid fluid flow solver may be re-used without modifications. However, block-structured AMR is often inflexible in capturing complex flow, leading to either possible overlapping of blocks at the same level of refinement and duplication of cells, or wasting of computational cells on smooth flow covered by the same refined blocks. A cell-by-cell AMR (Khokhlov, 1998), on the other hand, refines only individual cells. This flexibility leads to more saving of memory since each cell can be refined or coarsened as needed, independently of others. The cells are managed using ‘fully threaded tree’ data structures that provide a way to access and traverse the tree in parallel. This approach, however, requires maintenance of the tree and does not necessarily lend itself to straightforward re-use of solvers as in block-structured AMR. Care must also be taken in its implementation to avoid irregular memory referencing that can produce non-scalable code.

The flexibility and reduced memory footprint of cell-by-cell AMR have led us to adopt it in GenASiS. The current version of GenASiS does not have a mature implementation of AMR yet. However, we have developed some techniques and the necessary groundwork that will be necessary to have a scalable and fully functioning cell-by-cell AMR. In the next two sections, we discuss our approaches to the evolution of the fluid dynamics with AMR, and a scheme for the solution of the Poisson’s equation for mesh with refinements, necessary for simulations involving self-gravity.

4.2 Fluid Dynamics for Adaptive Mesh Refinement

4.2.1 Mesh and Tree Construction

Our implementation of cell-by-cell AMR uses the fully threaded tree (FTT) structure of Khokhlov (1998). In three-dimensional space, cells are organized in an oct-tree—a tree that has exactly eight children for each internal node that is not a leaf node—with each node of the tree representing a cell. A cell is either a refined or unrefined cell, which is represented by a parent node that has children, or a leaf node. The root of the tree covers the entire computational domain. Adaptivity is provided by splitting a cell into eight cells (bisection in each dimension), giving cells of higher resolution in that region of space. This is analogous to creating eight children for that particular node that represent the refined cell. This construction can be extended arbitrarily deep, providing arbitrary levels of refinement.

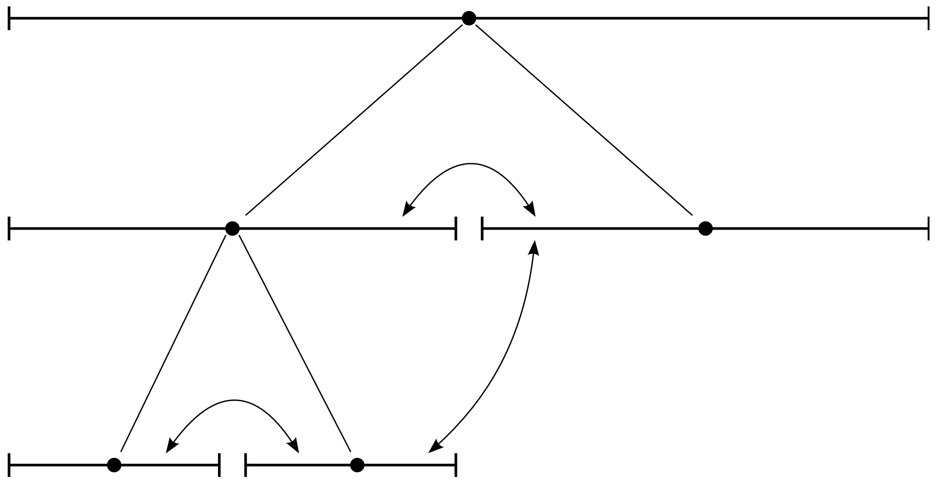

Figure 4.1: Logical relationship between cells in fully threaded binary tree (redrawn after Khokhlov (1998)). The width of the bar lines represent the size of the cells. Arrows represent pointers to neighbors, and lines represent pointers to children.

FTT structure adds more information to the tree to provide easy access to a cell's children, neighbors, and parents by threading the tree in all possible directions. This is necessary for an efficient algorithm of the refinement and integration of the hydrodynamics equations. For every cell in the tree, the following information is available: the level of the cell, a Boolean flag (true / false) indicating whether or not the cell has children, a pointer to the parent cell, pointers to its eight children, and pointers to its six neighbor. Figure 4.1 illustrates this logical relationship between cells in a binary tree (for simplification, but may be generalized to an oct-tree). The level of a cell corresponds to the depth of the node representing the cell in the tree. FTT provides neighbor information even for cells of different levels. However, the neighbor relationship is not reciprocal for cells of different levels, as seen in figure 4.1. Neighboring cells are not allowed to differ by more than one level.

Physical state information (e.g. the primitive and conserved variables of the magnetohydrodynamics) need to be associated with each cell in the oct-tree. In GenASiS, rather than keeping this information with the cells in the tree structure, the physical state resides in a flattened 'data' array as a large contiguous blocks of memory. Each cell in the tree is given a 'cell number' that refers to the array index of the data. This separation of ‘metadata’—the oct-tree representing the mesh structure—and the ‘data’ is intended to provide computational efficiency by concentrating the floating point operations on contiguous memory block. The mapping of the metadata to the data is afforded by constructing a large one-dimensional array of cells, neighbor cells (in each direction), cell faces, and cell edges from the connectivity information implicit in the tree structure. Floating point intensive operations such as integration of the hydrodynamics equations are done utilizing these arrays only, avoiding much of the pointer referencing necessary with tree traversal. These arrays are updated with the tree structure during refinement and coarsening.

In implementing the cell-by-cell AMR in GenASiS, we follow object-oriented design principle using the Fortran 95 standard. Cell is represented as an object `zoneType` that contains information such as its parent (if it exists), pointer to its children (if not a leaf cell), and its geometrical position in the computational domain. A collection of cells (four or eight of them, in two or three dimension respectively) are contained in an object `zoneArrayType`, representing a region of a spacelike slice. The children of a cell, which represent a refinement of that cell, can then be contained in a `zoneArrayType`. This structure can be extended arbitrarily deep. Figure 4.2 outlines this construction. The `zoneType` also hold a number as an index of its position in the flattened data block. The objects `zoneType` and `zoneArrayType` constitute the oct-tree. This object-oriented design allows operation on the tree to be done via well-controlled interfaces.

A composite mesh is built from the union of all leaf cells. The composite mesh therefore has multiple levels of refinement and varying resolutions covering the whole computational domain. As in §2.3, parallelization is done by decomposing the domain into multiple sub-domains. However, in the case of AMR, simple spatial decomposition may result in unbalanced workload on the processes since the same spatial region size may be covered by different numberss of cells due to refinement or coarsening. Therefore, domain decomposition is achieved instead by walking through the composite mesh similar to a Morton space-filling curve (Morton, 1966), which results in a mapping of the three-dimensional mesh to a one-dimensional curve representing a string of cells. The string

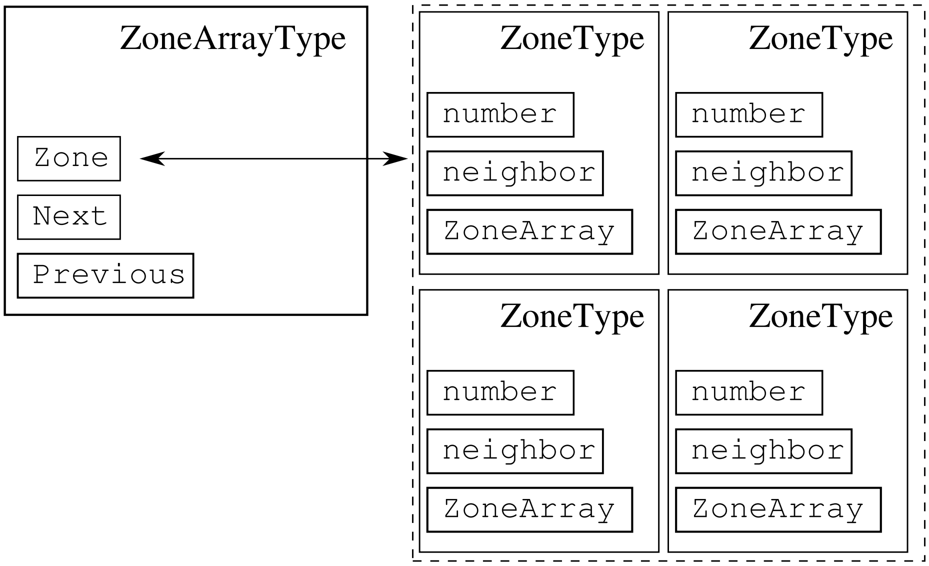

Figure 4.2: Illustration of data structure for the implementation of fully threaded tree.

is cut into pieces of uniform length. The pieces (representing cells) are distributed to the processes, yielding roughly equal workload for each process (since the workload depends only on the number of cells, regardless of the resolution of the cells). This load-balancing has to be maintained as simulation proceeds.

4.2.2 Fluid Evolution

The list of cell faces (cell interfaces) that we built from the connectivity of the tree allow fluxes to be evaluated at arbitrary order and even in parallel. For each cell face, fluxes are evaluated, changes to cell values on the left and right of the face are stored on an array to be applied, then fluxes are discarded. As in §2.3, fluxes on the interface for the ‘sent cells’ are calculated first, followed by by the fluxes across the faces of the ‘unsent cells’.

Since updates to the cell values are organized by cell faces, the evaluation of fluxes across cells of different sizes can be treated the same way as fluxes across cells with the same size. The illustration in figure 4.3 shows two possible situations for fluxes across cells of different sizes. Cells c_2 and c_4 share the same right neighbor, cell c_5, to form a fine-coarse interface. The fluxes F_{45} are evaluated as usual, and the change is applied to

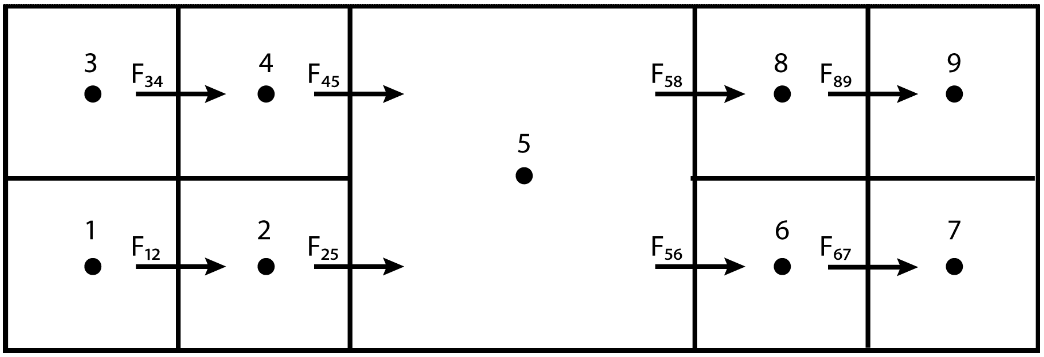

Figure 4.3: Illustration of flux evaluation across fine-coarse and coarse-fine interface on mesh with refinements.

the left cell (i.e. cell c_4) and right cell (i.e. cell c_5). Similarly, change due to flux F_{25} are applied to cell c_2 and cell c_5. Both fluxes, F_{25} and F_{45}, contribute to the change of cell c_5.

Coarse-fine interfaces are treated the same way. Fluxes F_{56} and F_{58} are evaluated, and changes for cell c_5, c_6, and c_8 are stored. The change for c_5 are the contributions due to F_{56} and F_{58}. One thing to note is the usage of cell face area to calculate the changes to the cell. On the coarse-fine or fine-coarse interfaces, the cells have different face area. For example, the right face area of c_4 is half that of left face area of c_5. Therefore the correct way to compute the changes on the cells with these interfaces is to use the minimum of the shared face area. For example, the change to c_5 due to F_{45} is F_{45} multiplied by the minimum of right face area of c_4 and left face area of c_5. Combined with the fluxes contribute to the change to c_5, this sums up to the correct change to c_5.

On meshes with refinements, the CFL condition also varies due to the variation of cell sizes on the mesh. We take the simple approach of using the synchronized timestep on the mesh just by taking the maximum timestep allowed by the CFL condition. This is more straightforward than the sub-cycling approach used by Khokhlov (1998). At the end of each timestep, we use the refinement criteria (see §4.2.3) to check whether any refinement or coarsening is needed on the mesh. This is iterated as many times as necessary until the criteria tell us that no further refinement is necessary, at which time we start the next timestep. As previously mentioned, for every refinement (or coarsening), the lists of neighbor cells, cell faces, and cell edges are rebuilt from the connectivity information in the oct-tree.

4.2.3 Refinement Criteria

Approaches to refinement criteria are mostly problem dependent. One approach is to measure the convergence of a solution, which allows control of the solution accuracy on-the-fly (Berger and Collela, 1989). Another approach is to use the gradient of some quantities to show where to expect large errors in the solution, and therefore tell us where refinements are needed (Aftosmis et al., 1995; Melton et al., 1995). We have used this latter approach for the refinement criteria in GenASiS.

The gradient of a variable χ is used to compute a refinement indicator ξ as

$$\xi_i^{\chi} = \max_{\forall j}\left(\frac{|N(\chi_i, j)| - |\chi_i|}{\max\left(|N(\chi_i, j)|, |\chi_i|\right)}\right), \tag{4.1}$$

where the index i indicates the cell index and the operator $N(\chi_i, j)$ gives the value of the variable χ at the neighbor cell j of cell i. The variable χ may represent mass density, pressure, velocity, internal energy, etc. For every cell, a single refinement indicator ξ is obtained as

$$\xi_i = \max_{\forall \chi}(\xi_i^{\chi}), \tag{4.2}$$

where $0 \leq \xi_i \leq 1$. Two predefined constant, ξ_{refine} and ξ_{coarsen}, are needed to indicate whether a leaf cell must be refined or coarsened. Refinements happen where $\xi_i > \xi_{\text{refine}}$, while coarsening is done for all $\xi_i < \xi_{\text{coarsen}}$. The latter can only be done if the cells were not just refined.

The refinement indicator ξ may fluctuate around critical values ξ_{refine} and ξ_{coarsen} that causes cell to be falsely refined or coarsened. To avoid this, ξ is smoothed before it is used for refinement. An analogy with the propagation of a reaction-diffusion front is appropriate for smoothing ξ (Khokhlov, 1998), where ξ is considered as a concentration of some reactant obeying

$$\frac{\partial \xi}{\partial \tilde{t}} = K \nabla^2 \xi + Q \tag{4.3}$$

for some fiducial time $\tilde{t}$. K is a constant diffusion coefficient, given the value $K = 2^{-2l} L^2$ where l is the level of refinement of the cell and L is the size of the computational domain. The reaction rate Q is

$$Q = \begin{cases} 1, & \text{if } 1 > \xi > \xi_{\text{refine}}, \\ 0, & \text{otherwise.} \end{cases} \tag{4.4}$$

Equation 4.3 has a steady-state form

$$-S_\xi \frac{\partial \xi}{\partial x} = \frac{\partial^2 \xi}{\partial x^2} + Q, \tag{4.5}$$

with a reaction front that moves with constant speed

$$S_\xi = 2^{-l} L \sqrt{\xi_{\text{refine}}} \tag{4.6}$$

and has a thickness

$$\delta_\xi \simeq \frac{2^{-l} L}{\sqrt{\xi_{\text{refine}}}}. \tag{4.7}$$

The refinement indicator ξ obtained from equations 4.2 for every cells are used as the initial value for equation 4.3, which describes a reaction front propagating outward from $\xi > \xi_{\text{refine}}$ with thickness 2-3 computational cells. As the front curvature decrease with time according to the Huygen-Fresnel principle, the boundaries where $\xi > \xi_{\text{refine}}$ becomes smoother. Isolated regions marked for refinement do not trigger reaction fronts themselves due to the diffusion in the equation, provided that the area is less than δ_ξ. This smoothing then avoids 'mesh thrashing', a fluctuation of refinement and coarsening.

4.2.4 Numerical Example

Here we show a numerical example of hydrodynamics evolution with AMR as currently implemented in GenASiS. We set up a two-dimensional (2D) counterpart of the Newtonian

Riemann shock tube problem previously described in §2.4.1. The discontinuity separating the 'left' and 'right' state is described by a line connecting the point $(0.5, 0.0)$ and $(0.0, 0.5)$. The rest of the initial conditions are similar to the one-dimensional shock tube problem.

To illustrate the AMR capability, we start with a very low uniform resolution of eight-by-eight cells covering the whole computational domain, as shown in figure 4.4a. In that figure, the left panel shows the domain decomposition into sixteen sub-domains owned by sixteen different processes, indicated by colors. The right panel shows the mass density. The discontinuity looks like a staircase due to the low resolution. The bar at the bottom indicates the simulation time, initially at $t = 0$. There is no refinement yet in this figure.

The refinement criteria of equation 4.1 is triggered due to the steep gradient. The mesh is refined before evolution even begins, as shown in figure 4.4b. Again on the left panel, the composite mesh is shown with the domain decomposition. The regularity of the decomposition is non-existent here due to refinement. The decomposition is such that every process (color) has roughly the same number of cells, regardless of resolution. Redistribution of cells is done by all the processes to form this decomposition from the initial decomposition in figure 4.4b. For example, after process 1 (red) refines its top right cell, it needs to send the newly created level-two cells, four of them, to process 2 (green), as shown, to maintain load-balancing. Similarly, process 2 (green), knowing that it will receive four newly created cells, sends 9 newly created level-two cells to process 3 (dark blue). Thus point-to-point communications happens to all processes for the redistribution of cells.

Note also the width of refined cells, which is roughly two-cells wide in each direction from the discontinuity line that triggered the refinement. This is due to the smoothing of the refinement indicator via reaction-diffusion of equation 4.3 to reduce mesh thrashing.

The refinement criteria is still satisfied due to the discontinuity, thus another refinement is triggered. In fact, this discontinuity triggers refinements until the maximum levels of refinement allowed (five levels) before the simulation proceeds. After every refinement,

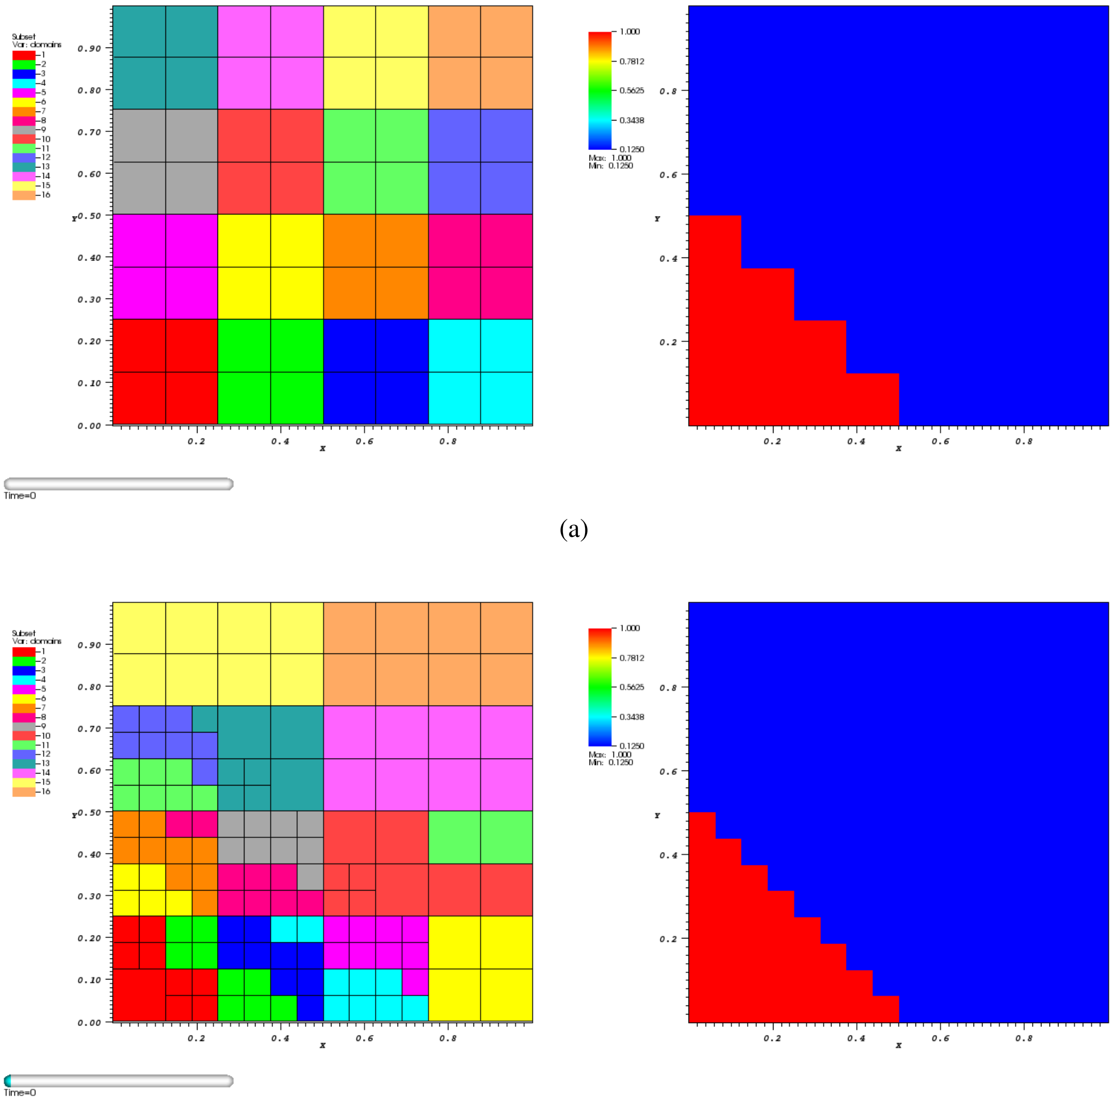

(a)

(b)

Figure 4.4: Initial conditions of the 2D shock tube problem with AMR shown in 4.4a. In this and subsequent figures, the left panels show the domain decomposition with colors to indicate processes that own the subdomains. The right panels shows the mass density. 4.4b shows the mesh after one level of refinement and redistribution of cells. The bar and label on the bottom left on each panel indicate the simulation time.

the cells are redistributed to maintain the load-balancing. This is shown in figures 4.5 and 4.6.

After reaching the maximum level of refinement (five level of refinements, giving six-level mesh), the simulation proceeds. Figure 4.7 shows the evolution of the problem at $t \simeq 0.001$ and $t \simeq 0.052$. We can see that the refinements tracks the shock front because of the discontinuity. In this current version, coarsening is not yet implemented, therefore the previously refined cells are not coarsened. Cells redistribution are being done throughout the simulation, maintaining load-balancing of the processes.

Figure 4.8a shows the simulation at late time when $t \simeq 0.327$. The simulation is ended at $t \simeq 0.6$, shown in figure 4.8b. Here the whole computational domain is fully refined to the maximum level of refinement, and thus the domain decomposition becomes regular again. This example shows the current AMR implementation for the evolution of fluid dynamic in GenASiS.

4.3 Poisson Solver for a Mesh with Refinements

The fast direct method to solve the Poisson's equation described in chapter 3 cannot be employed directly on the composite mesh with multiple level of refinements. In our implementation of AMR, the composite mesh is a union of leaf cells of the AMR tree and therefore has varying mesh resolutions over the computational domain (see figure 4.9). This violates the constraint of the Poisson's solver with the FFT method. The solution of the Poisson's equation however, needs to be available on every level of refinement. Furthermore, since the cells with higher level of refinement have higher resolution, a more accurate source for the Poisson's equation is also available on these cells and therefore a better solution (i.e potential) is desirable. To this end, we have developed a variant of multigrid method to obtain the solution of the Poisson's equation on a mesh with multiple levels of refinements in three dimensions. In this section we present the algorithm and numerical results for a static test problem.

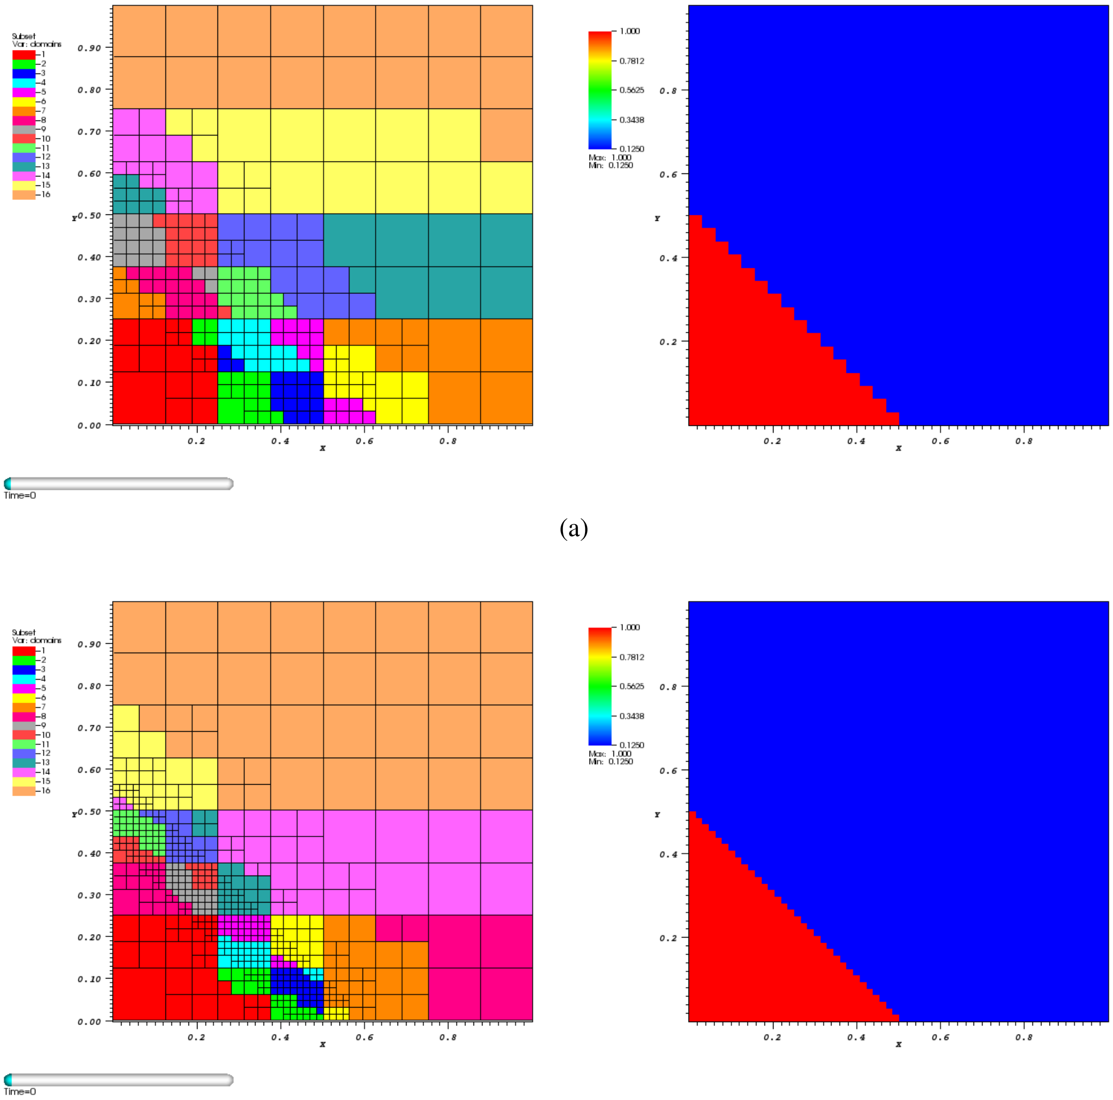

(a)

(b)

Figure 4.5: 2D shock tube problem with two and three levels of refinements.

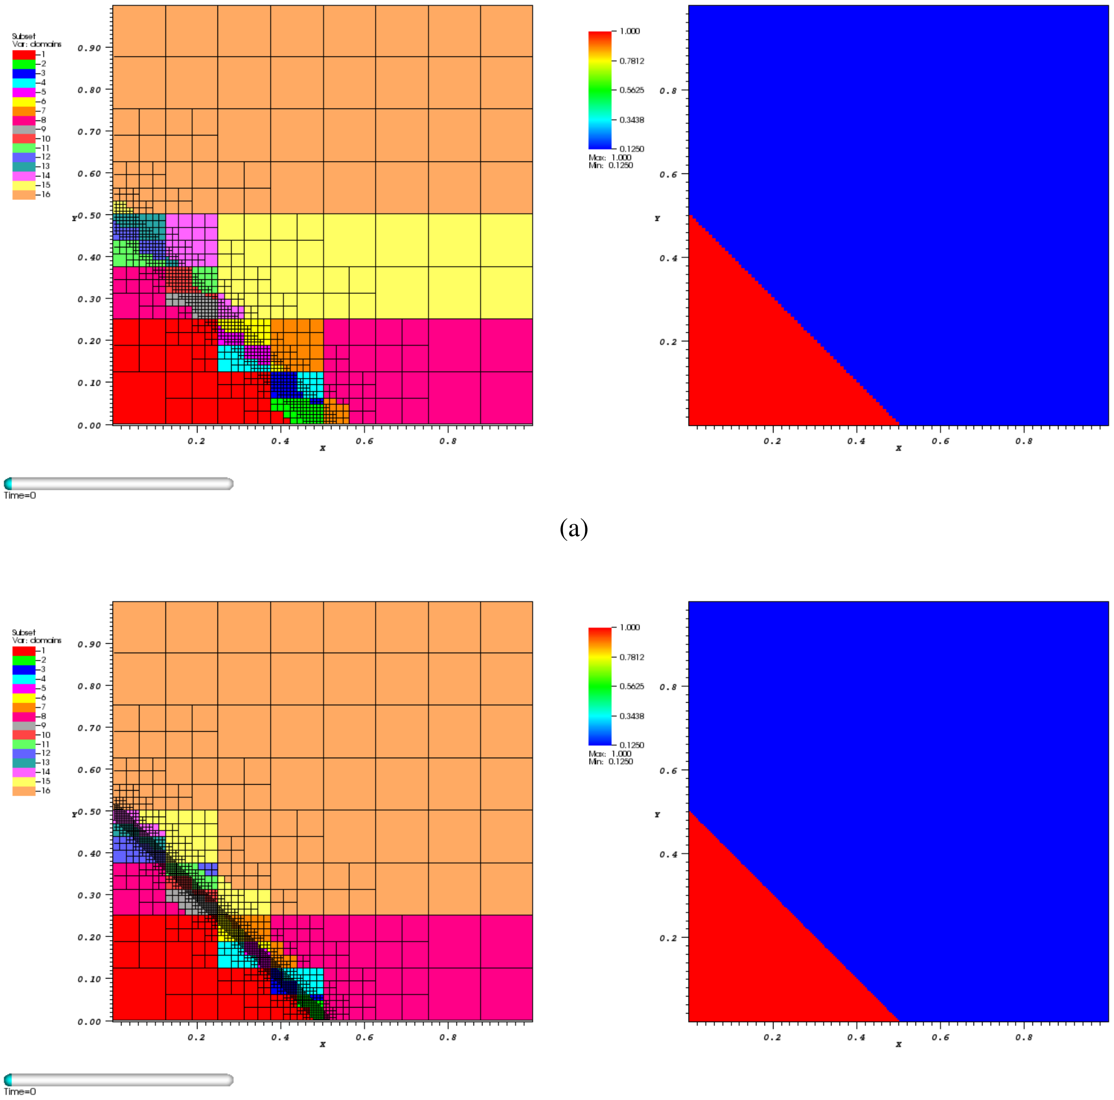

(a)

(b)

Figure 4.6: 2D shock tube problem with three and four levels of refinements.

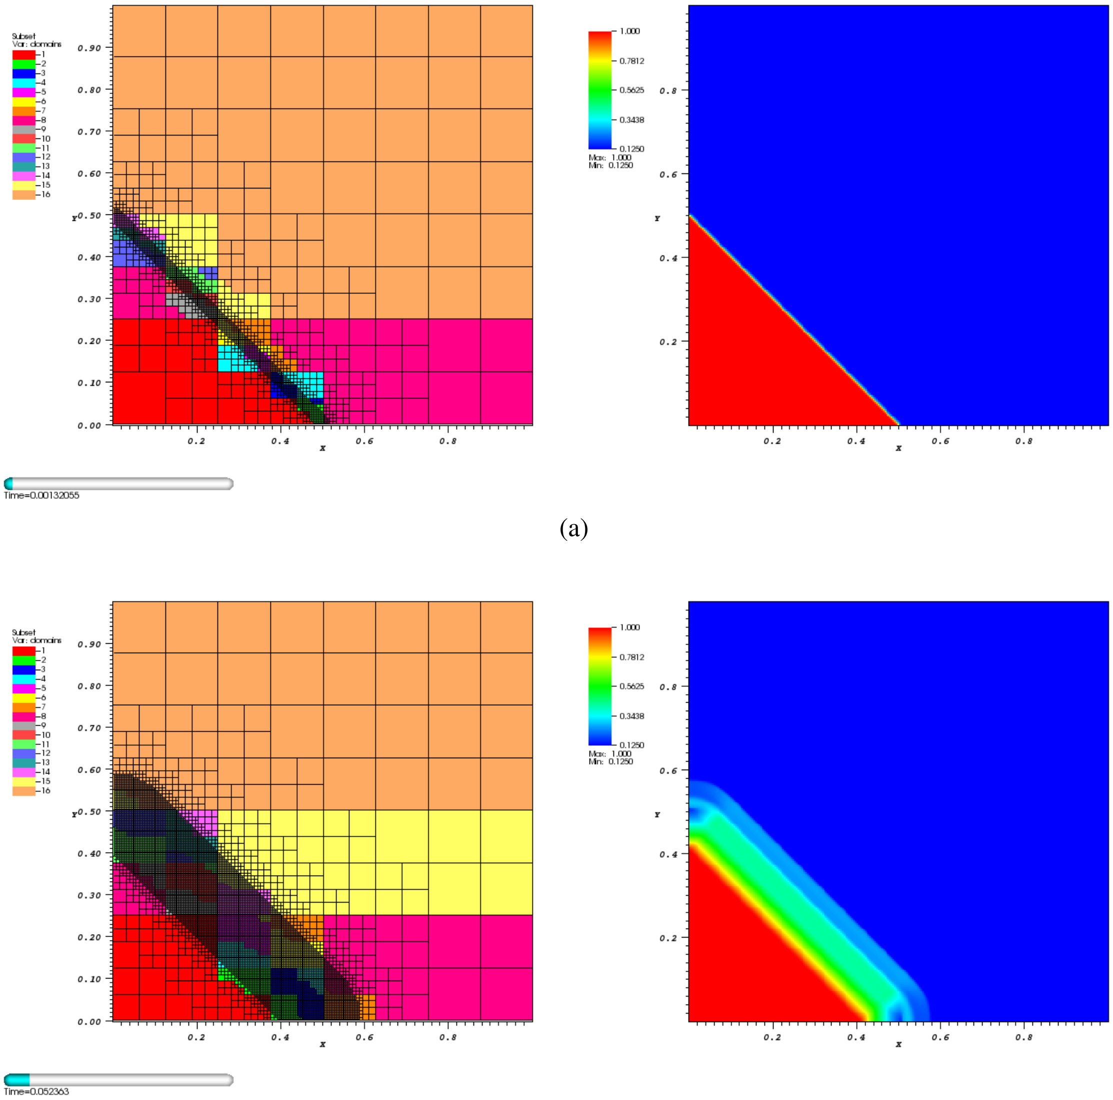

(a)

(b)

Figure 4.7: Evolution of 2D shock tube problem with AMR.

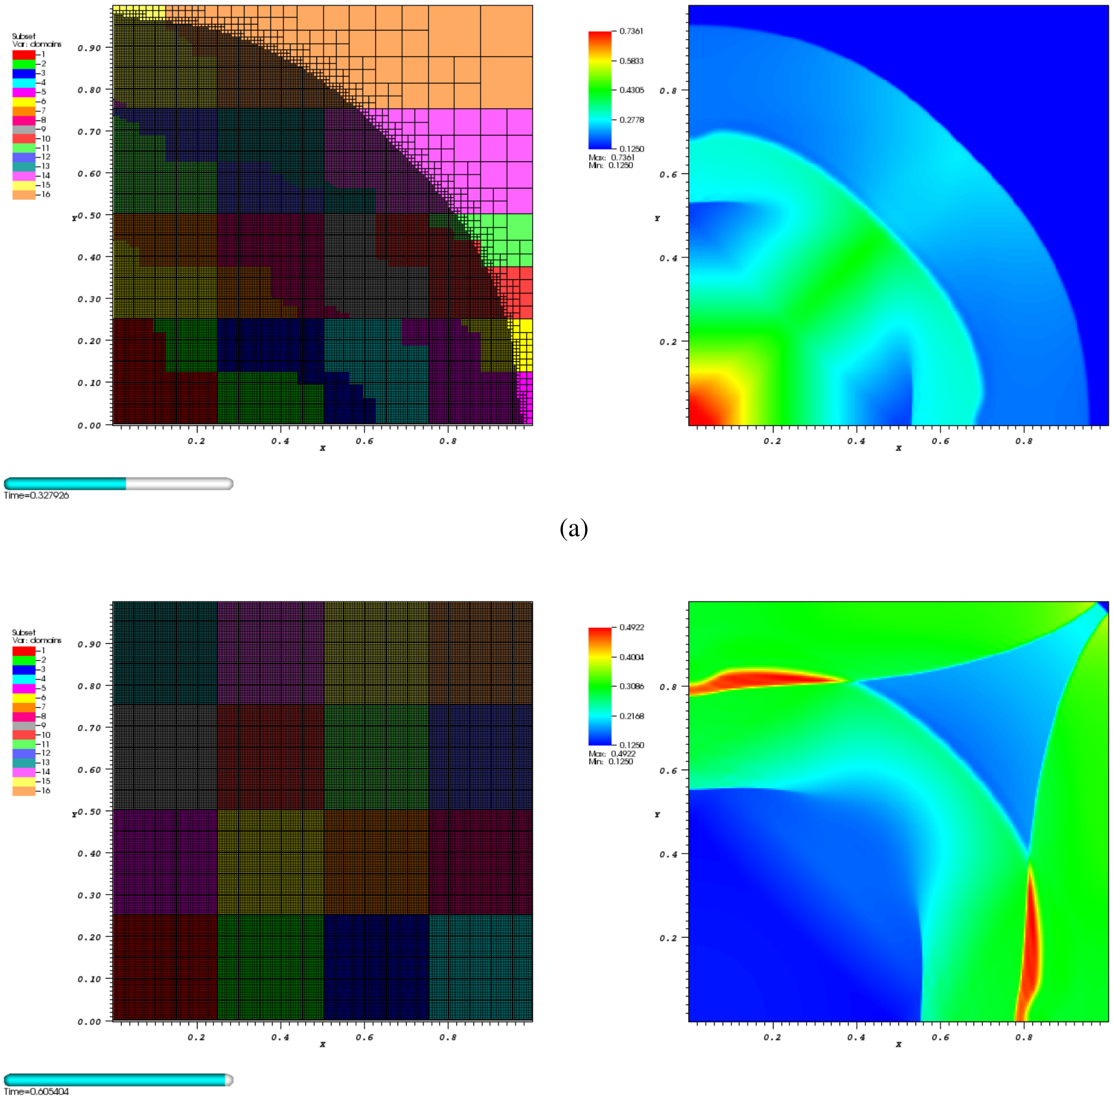

(a)

(b)

Figure 4.8: Evolution of 2D shock tube problem with AMR.

4.3.1 Construction and Definitions

Our method starts by creating a non-uniform 'level mesh' for every level of refinement. Non-uniform here means that for each mesh corresponding to a level of refinement, the cells that exist on the level mesh are either parent cells on that level (cells that have been refined) or leaf cells on that level. The cells themselves are still collinear (i.e. rectangular in Cartesian coordinate). Therefore, each level mesh has only 'islands' of cells. Let us call these cell the *proper* cells. For example, on the coarsest level mesh the cells cover the entire computational domain since each cell on this level has to be either a parent cell or a leaf cell. On each level mesh with higher level of refinement than the coarsest level, a layer of *exterior* cells is also created. These are guard cells surrounding the proper cells, which act as a boundary to the proper cells, and are created solely for this purpose. Figures 4.10 and 4.11 illustrate this construction.

The Poisson's equation 3.1 needs to be solved on every level mesh. On the coarsest mesh, we use the FFT method described in chapter 3, since the mesh is uniform. On other level meshes, we discretize the Laplacian operator in similar manner to equation 2.14. This yields a linear system for the values of Φ at the center of every cell on the level mesh. Boundary values are needed to solve this system, and they are provided by the exterior cells. The values of Φ at the center of the exterior cells are the prolongated values from the coarser level mesh. The linear system is cast into its matrix representation, distributed over several processes as each process fills in the portion of the matrix corresponding to its share of cells. The system then may be solved by inverting the matrix representation. We rely on PETSc library (Balay et al., 1997, 2008, 2009) to give us a programming interface to create a distributed matrix for the linear system and perform the matrix inversion in parallel over a large number of processes. PETSc uses the iterative method of a Krylov subspace with preconditioner (Balay et al., 2008; Freund et al., 2008) to solve the system. A good initial guess as the preconditioner is therefore at the heart of convergence of the iteration within the specified error tolerance. In our multigrid algorithm, we aim to exploit this fact by providing a good initial guess obtained from the values of the coarser level mesh via ‘prolongation’ (defined below). This is key to the efficiency of our multigrid algorithm for a mesh with refinements.

Let us define some terms before we continue with the algorithm. We use the notation M^l to indicate mesh of level l, with l_{min} and l_{max} indicate the minimum (coarsest) and maximum (finest) level of refinements. Similarly, the notation c^l_{ijk} indicates a cell at position i, j, k of refinement level l. The cell may be a leaf cell, or a parent cell, in which case it has eight children which we denote $c^{l+1}_{i'j'k'}$. This notation avoids the need to enumerate the children of c^l_{ijk}.

In our finite volume scheme, quantities are defined as cell-centered values. A value on refined cell c^l_{ijk} is ‘prolongated’ to its children $c^{l+1}_{i'j'k'}$ by reconstructing the value so that it exists as cell center values on all $c^{l+1}_{i'j'k'}$. The reconstruction is done by linear interpolation on each dimension. The reverse operator, a ‘restriction’, takes the an average of the cell-centered values of $c^{l+1}_{i'j'k'}$ (children cells) to create a cell-centered value in c^l_{ijk} (parent cell). Applying prolongation and restriction operator on a level mesh only means that the operation is done on every cells on that mesh that have children (for prolongation) or parent (for restriction).

We define a residual $R(\mathbf{x})$ as a measure of error of the approximate solution $\widetilde{\Phi}(\mathbf{x})$ relative to the ‘exact’ solution $\Phi(\mathbf{x})$ of the Poisson’s equation 3.1

$$R(\mathbf{x}) = S(\mathbf{x}) - \nabla^2 \widetilde{\Phi}(\mathbf{x}). \tag{4.8}$$

The residual satisfies the equation (Ricker, 2008)

$$\nabla^2 C(\mathbf{x}) = R(\mathbf{x}) \tag{4.9}$$

due to the linearity of the Poisson’s equation. Adding the correction $C(\mathbf{x})$ to the approximate solution $\widetilde{\Phi}(\mathbf{x})$ yields the correct solution $\Phi(\mathbf{x})$ of the original Poisson’s equation for the source $S(\mathbf{x})$.

4.3.2 Multigrid Algorithm

Now we are ready to describe the algorithm to obtain the solution of Poisson's equation on every level mesh. The listing 4.1 shows the pseudocode of such an algorithm. The algorithm begins by restricting the source from the highest level of refinement to the coarsest level on line 1-3 (see figures 4.10 - 4.12 for illustrations). Since the source is initially only defined on the composite mesh, when the level meshes were created the source exists only on leaf cells. Restriction makes sure that the source exists on all cells on every level mesh. On line 5, Poisson's equation is solved, either with the FFT method (on the coarsest level) or with the matrix inversion method (on other levels). The prolongation of the potential on line 9 makes sure that when line 5 is executed on the next iteration, there is already an initial guess (preconditioner) for the matrix inversion method. The loop block on line 11-14 computes the residual on all the level meshes that the main loop block starting on line 4 has gone through so far. Line 12 computes the residual on all the leaf cells on M^j, while line 13 guarantees that the residual on cells that are parents on M^j is also defined. For $j = i$ at the first iteration of the loop, this comes from the residual previously computed on line 10. On the leaf cells of the coarsest level mesh, we set the residual to zero by definition of our solution method. This restriction of the residual from the higher level of refinement to the coarsest level in a sense propagates the knowledge of better solution due to better resolution at the higher refinement level back to the coarsest level. This information is used later to compute the correction so that the solution on the coarser level, initially computed with the resolution available only on that level mesh, also benefits from the higher resolution that is available.

The iteration block on lines 15-23 solves for the corrections on every level mesh up to the level currently being worked on by the main loop. The corrections are added to the already solved potentials on every level mesh. Either the correction $C(\mathbf{x})$ or the potential $\Phi(\mathbf{x})$ is then prolongated depending whether or not we are at the last iteration of the loop. The prolongation of the correction gives the matrix inversion scheme to solve for the correction on line 16 a good initial guess. The prolongation of the potential gives the

Algorithm 4.1 Multigrid Poisson Solver

```
for $i = (l_{\max} - 1)$ to $l_{\min}$ do
    Restrict $S^{i+1}(\mathbf{x})$ to $S^{i}(\mathbf{x})$
end for
for $i = l_{\min}$ to $l_{\max}$ do
    Solve $\nabla^2 \Phi^{i}(\mathbf{x}) = S^{i}(\mathbf{x})$ obtaining $\widetilde{\Phi}^{i}$
    if $i == l_{\max}$ then
        return
    end if
    Prolong $\Phi^{i}(\mathbf{x})$ to $\Phi^{i+1}(\mathbf{x})$
    Compute the residual $R^{i+1}(\mathbf{x})$
    for $j = i$ to $l_{min}$ do
        Compute the residual $R^{j}(\mathbf{x})$
        Restrict $R^{j+1}(\mathbf{x})$ to $R^{j}(\mathbf{x})$
    end for
    for $j = l_{min}$ to $i$ do
        Solve $\nabla^2 C^{j}(\mathbf{x}) = R^{j}(\mathbf{x})$ obtaining $C^{j}(\mathbf{x})$
        Add correction: $\widetilde{\Phi}^{j}(\mathbf{x}) = \widetilde{\Phi}^{j}(\mathbf{x}) + C^{j}(\mathbf{x})$
        if $j < i$ then
            Prolong $C^{j}(\mathbf{x})$ to $C^{j+1}(\mathbf{x})$
        else
            Prolong $\Phi^{j}(\mathbf{x})$ to $\Phi^{j+1}(\mathbf{x})$
        end if
    end for
end for
```

matrix inversion scheme at the next loop iteration of the main loop on line 5 a good initial guess. By the time the last iteration of the main loop block (line 4-24) is reached, the potentials of every level mesh have been corrected multiple times utilizing all the available information of the sources at the higher level of refinement. The execution of line 21 before the last iteration of the main loop makes sure that there is already a very good initial guess for the potential when line 5 is executed at the very last iteration, that is, on the finest level mesh. We expect the residual to be the error tolerance that we specify for the inversion scheme in PETSc, and that the iteration for the matrix inversion converge very quickly. Line 7 ensures that we exit this routine.

By the time the algorithm in listing 4.1 exits, every level mesh has the potential as the solution to Poisson's equation. To map this back to the composite mesh, we simply replicate the potential found on every leaf cells on every level mesh to the composite mesh. This completes the algorithm to solve for Poisson's equation on mesh with refinements.

4.3.3 Numerical Example

To test our multigrid algorithm, we calculate the gravitational potential of a spherical uniform mass on a mesh with four levels of refinements distributed over eight processes. This problem was previously discussed on §3.4.1 for aa unigrid. Here we set the radius of the sphere to $R = 0.5$ and mass density to $\rho = 1$. The potential and the relative error of the potential as compared to the analytical solution on the composite mesh are shown in figure 4.13. Figure 4.14 shows the gravitational potential on each level mesh. The coarsest level mesh is covered by 16^3 over the entire computational domain. We see that the maximum refinement happens on the surface of the sphere to try to resolve the mass density cliff. The relative error distribution resembles the one previously seen on §3.4.1, with the finest level of refinement here at half the resolution.

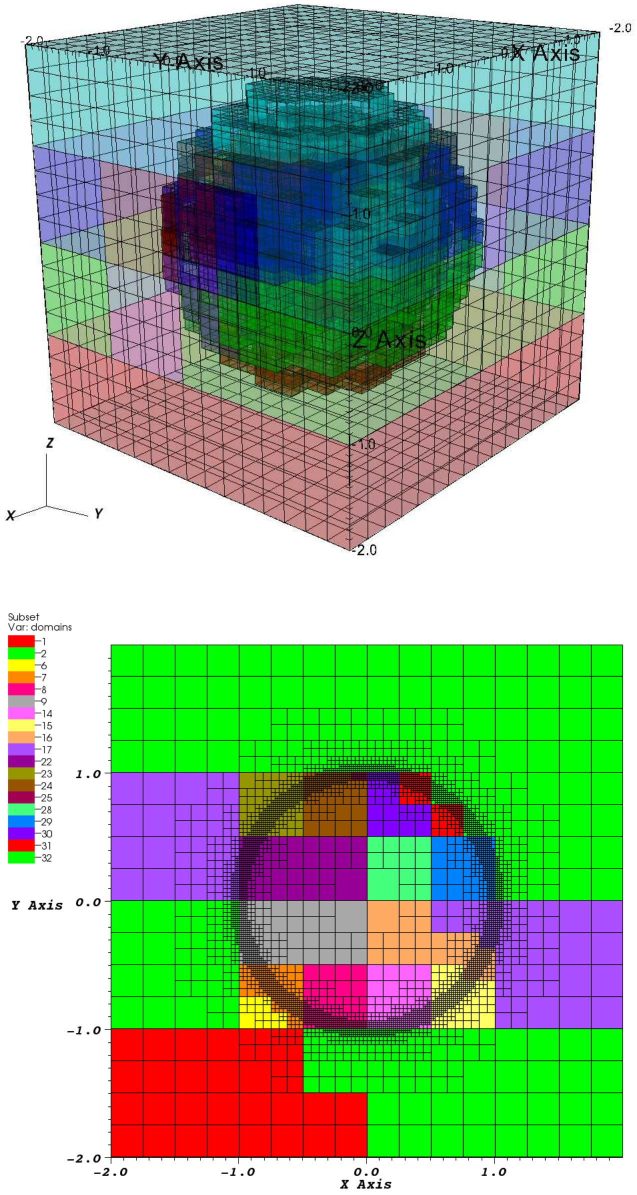

Figure 4.9: The composite mesh of a spherical uniform mass problem with five levels of refinement distributed over 64 processes. The colors indicate the different processes responsible for the cells. The top panel shows the whole mesh in three-dimensional perpective; the bottom panel shows a two-dimensional slice of of the mesh across an *xy*-plane.

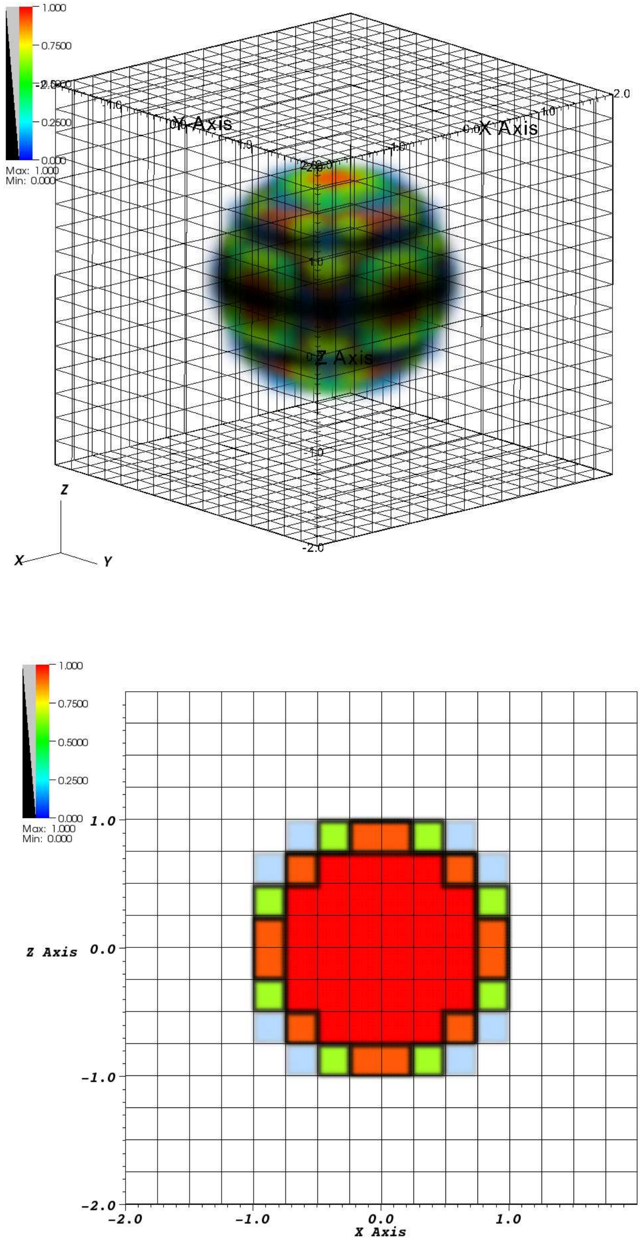

Figure 4.10: The coarsest level (level 1) of the level mesh. The top panel shows the mesh in three-dimensional perspective; the bottom panel shows a two-dimensional slice of of the mesh across an *xy*-plane. On this level mesh, the cells covers the entire computation and forms the uniform mesh. Solution of Poisson's equation on this level mesh can therefore be obtained using the FFT method.

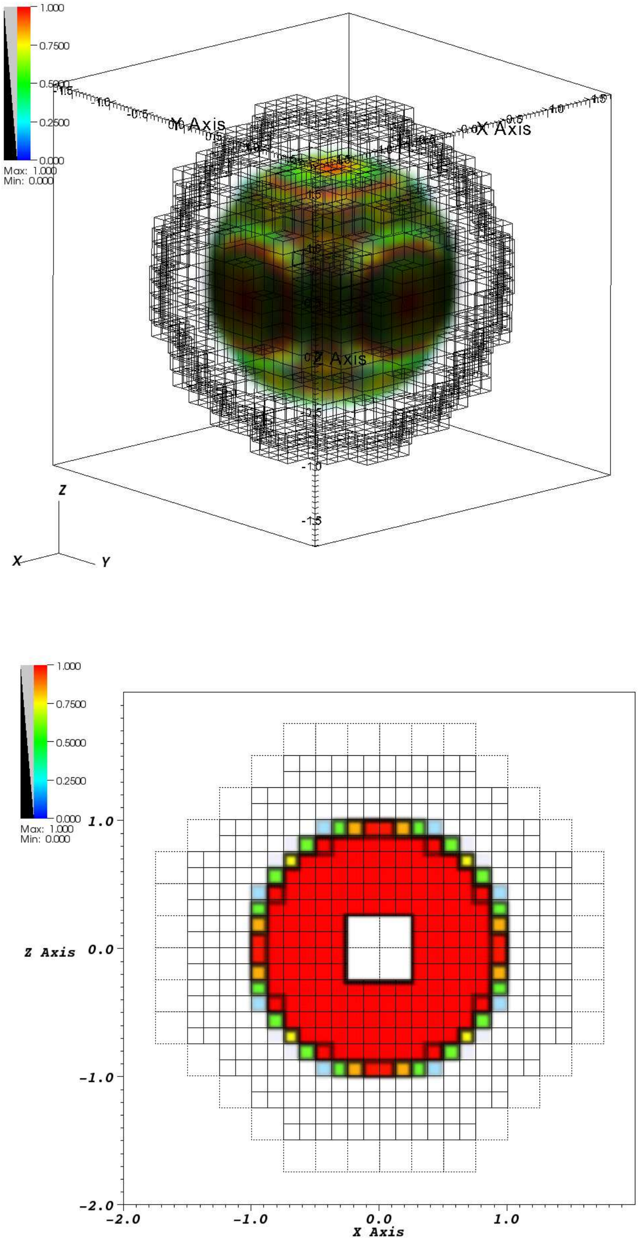

Figure 4.11: Level 2 of the level mesh. The top panel shows the mesh in three-dimensional perspective; the bottom panel shows a two-dimensional slice of of the mesh across an *xy*-plane. The *proper* cells are shown with solid line. The cells with dotted lines are the *exterior* cells that form the boundary around the proper cells. Mass density is plotted as a volume plot. The variations in density are caused by restricted values from higher-resolution level mesh.

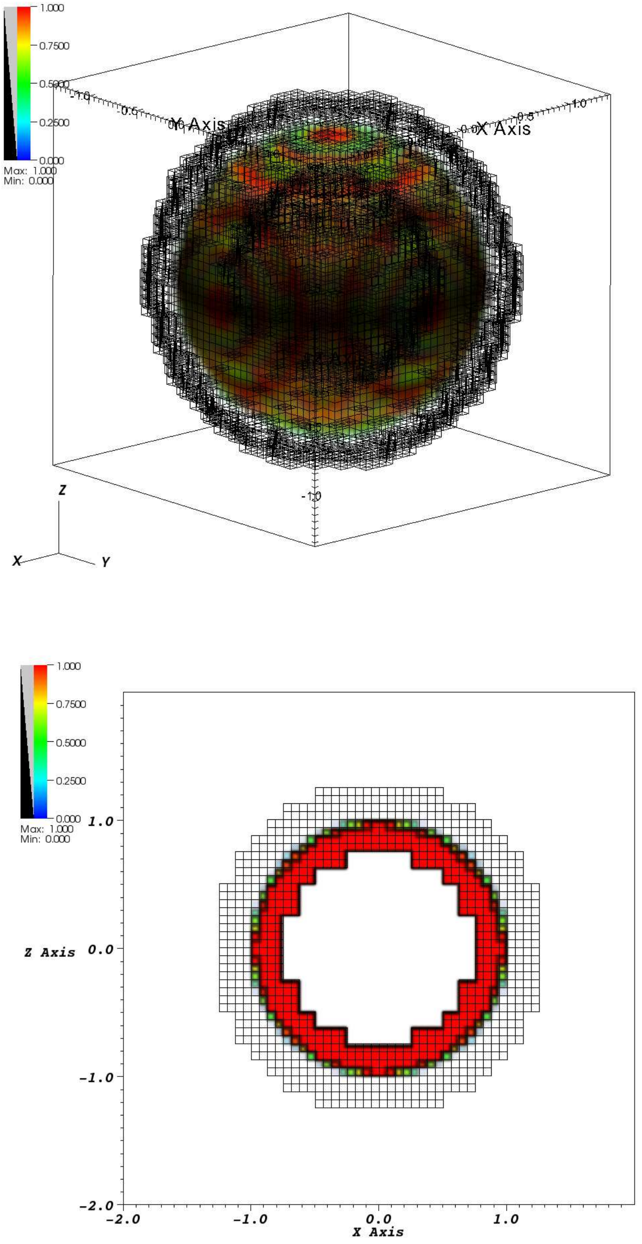

Figure 4.12: Level 3 of the level mesh. The top panel shows the mesh in three-dimensional perspective; the bottom panel shows a two-dimensional slice of of the mesh across an *xy*-plane. As shown in previous figure, cells form an island that does not necessarily cover the entire computational domain. Only the proper cells are shown here. Mass density is plotted as a volume plot. The variations in density are caused by restricted values from higher resolution level mesh.

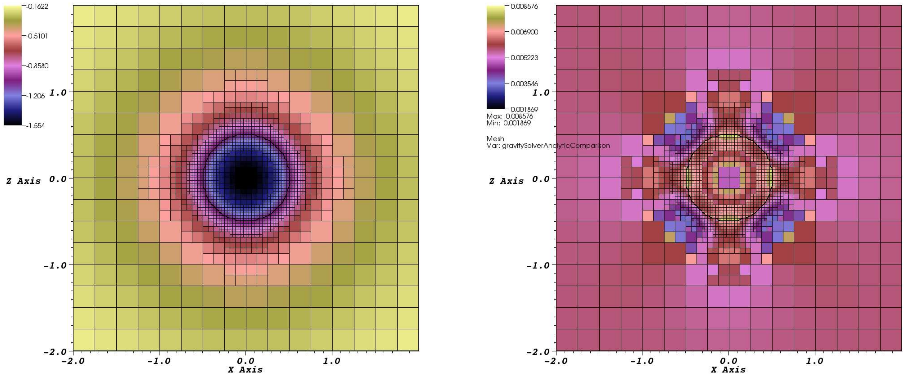

(a) (b)

Figure 4.13: Gravitational potential of spherical uniform mass mesh with refinements (left), and the relative error as compared to the analytical solution (right). The plots are a slice through the three-dimensional mesh. On each plot the composite mesh is shown. The black contour line represents the surface of the sphere at radius $R = 0.5$.

4.4 Conclusion and Outlook

In this chapter we have described the principle techniques for AMR in GenASiS. Further code developments are still necessary to make AMR a mature feature of the code. Missing features of the AMR includes mesh coarsening, a scheme to include the evolution of magnetic fields with AMR, and the coupling of the hydrodynamics and Poisson's solver for simulations involving self-gravity. All of these are still currently under development.

The scalability of the current AMR scheme is also an area under investigation. Because redistribution of cells is costly, it may be more advantageous to continue the simulation with some load-imbalance until certain threshold, at which point redistribution is triggered. The technique to do this is still being developed. The thresholds to know how much load-imbalance can be tolerated may be hardware dependent, and a way to quantitatively and systematically determine that needs to be investigated.

In the current version, the evolution of the fluid is done on the composite mesh with the timestep being limited by the CFL condition due to the highest resolution cells. Meanwhile, our multigrid scheme necessitates the creation of level meshes, which are then discarded

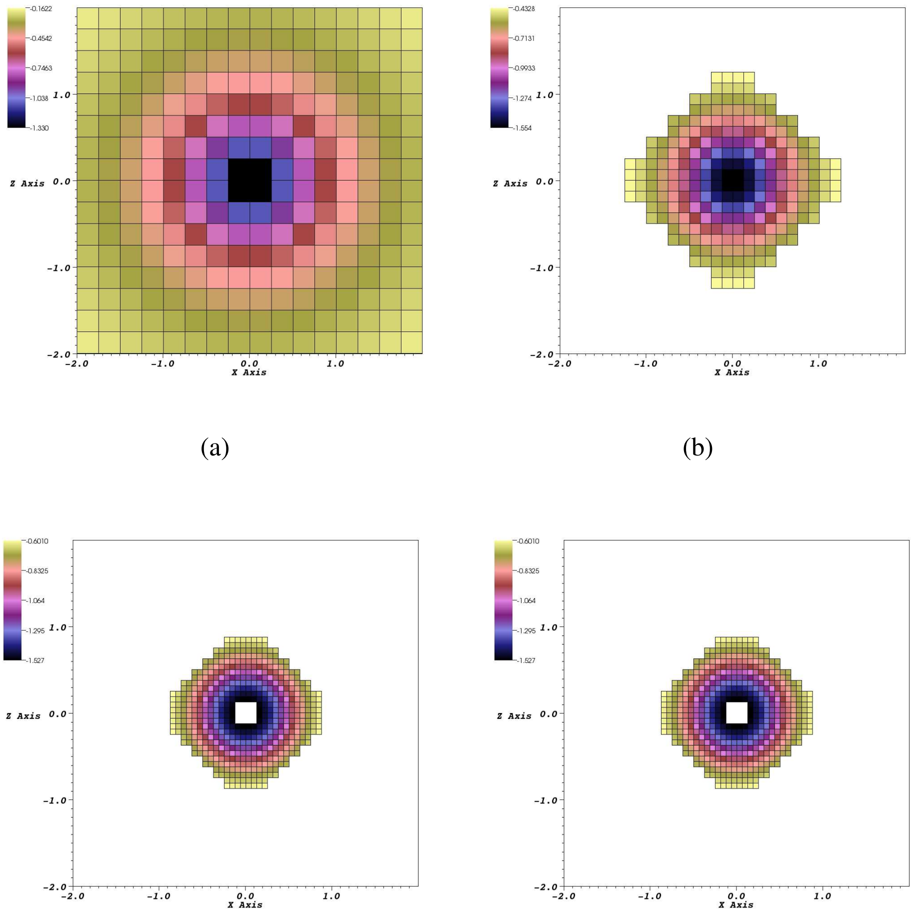

(a)

(b)

(c)

Figure 4.14: Gravitational potential of a spherical uniform mass distribution on level mesh computed with the multigrid algorithm. Levels one to four of the refinements are shown on the panel from left to right, top to bottom. The plots are a slice through the three-dimensional mesh.

when the solution of the Poisson's equation is obtained and mapped to the composite mesh. This mismatch is costly due to the extra transient memory needed to create the level meshes. Furthermore, during the development of the multigrid scheme, we serendipitously observed that the use of level meshes has properties that may simplify a lot of our constructions for the evolution of the fluid dynamics. In the language of object-oriented programming, we may consider a level mesh as an object on which a level by level evolution of the fluid may be done. Methods and interfaces for each level may then be re-used on many level meshes. Evolution across coarse-fine boundaries may also be treated using prolongations and restrictions. Each level may also be evolved with different timestepping, sub-cycling so that the coarser levels evolve with the largest timesteps as multiples of the timesteps of the finer levels. This may yield an overall increase of efficiency. These reasonings motivate us to investigate a modified scheme for the fluid dynamics to utilize level meshes. This line of thinking therefore is currently being pursued to see if it yields not only a more scalable AMR scheme, but also a reduction of code complexity.

Chapter 5

Merger of Binary Neutron Stars

5.1 Introduction

A neutron star is one possible outcome of stellar evolution. After a massive star runs out of fuel, it undergoes gravitational collapse and ejects its outer layers in a cataclysmic explosion, a phenomenon known as a core-collapse supernova. Depending on the mass of the progenitor, the remnant of this core-collapse supernova is either a neutron star or a black hole. A neutron star is therefore a very compact and dense object.

Typical neutron stars have roughly 1 to 1.5 solar masses $(M_\odot)$ compressed into an object of only about 20 to 10 km in radius. Most neutron stars were initially discovered as solitary objects known as pulsars (rotating neutron stars emitting periodic radio pulses). It was not until 1974 that a neutron star binary was discovered (two neutron stars in mutual orbit) by Hulse and Taylor (1975), a discovery which led to a Nobel Prize. Since then, additional discoveries of binary neutron stars in our galaxy have followed (Dewey et al., 1985; Anderson et al., 1990; Wolszczan, 1991).

Binary neutron stars have a finite lifetime due to their decaying orbit. As the stars orbit each other, their separation decreases and the two stars inspiral towards each other. This orbital decay is believed to be caused by the emission of gravitational radiation, as predicted by the theory of general relativity. Precise measurement of this orbital decay agrees with the prediction of general relativity (Taylor, 1994), indirectly confirming the existence of gravitational waves. Efforts to detect gravitational waves directly have begun, with large-scale gravitational detectors such as LIGO, VIRGO, and GEO-600 functioning and coming online. Neutron star mergers are one of the prime targets for these detectors (Abramovici et al., 1992; Bradaschia, 1990; Luck, 1997). However, theoretical waveforms and templates are needed for these detectors to be able to extract gravitational wave signals from background noise. Neutron star merger simulations and models provide such templates.

In this chapter we show current GenASiS capabilities to do such simulation yielding observables such as gravitational waveforms and spectra signatures. Here we present results of a full three-dimensional simulation of neutron star merger. We also show the tools we have developed to extract and analyze the physical observables: gravitational waveforms and spectra. As in other previous work, we use a simple polytropic equation of states to model the neutron stars. This simulation serves as a milestone in GenASiS development towards a more physically realistic neutron star merger simulations.

5.2 Gravitational Wave Radiation

Widely separated neutron stars inspiral, driven primarily by the loss of energy and momentum due to the gravitational wave emission. For nearly Newtonian sources, gravitational radiation can be calculated using the quadrupole approximation (Misner et al., 1973). Blanchet et al. (1990) introduced a formalism to include these effects in non-relativistic hydrodynamics flow without magnetic fields by adding source terms (in addition to the self-gravitation source term) to the hydrodynamics equations. We follow the specific implementation by Shibata et al. (1992) and Ruffert et al. (1996), and modify equations 2.1

- 2.3 to become:

$$\begin{aligned}
\frac{\partial\rho}{\partial t}+\frac{\partial}{\partial x^i}\left(\rho v^i\right) &= 0, && (5.1)\\
\frac{\partial S^j}{\partial t}+\frac{\partial}{\partial x^i}\left(\rho w^j v^i+p\delta^{ij}\right) &= -\rho\frac{\partial\Phi}{\partial x^j}-\rho\frac{\partial\Psi}{\partial x^j}, && (5.2)\\
\frac{\partial E}{\partial t}+\frac{\partial}{\partial x^i}\left(\left[e+p+\frac{1}{2}\rho w^j w^j\right]v^i\right) &= -\rho v^i\frac{\partial\Phi}{\partial x^i}+W. && (5.3)
\end{aligned}$$

The symbols have the same meaning as before in §2.1.1 , except for several new variables introduced here. We also redefine the conserved variables S and E in term of the new dynamic velocity w^j as:

$$\begin{aligned}
S^j &= \rho w^j, && (5.4)\\
E &= e+\frac{1}{2}\rho w^j w^j, && (5.5)
\end{aligned}$$

where w^j is given by the following relation:

$$v^j=w^j+\frac{4}{5}\frac{G}{c^5}\dddot{D}_{ij}w^i. \tag{5.6}$$

Ψ and W are the back-reaction potential and the energy source term due to the gravitational waves, respectively. They are given by

$$\begin{aligned}
\Psi &= \frac{2}{5}\frac{G}{c^5}\left(R-\dddot{D}_{ij}x^i\frac{\partial\Phi}{\partial x^j}\right), && (5.7)\\
W &= -\rho v^i\frac{\partial\Psi}{\partial x^i}+\frac{4}{5}\frac{G}{c^5}\dddot{D}_{ij}v^j\left(\rho\frac{\partial\Phi}{\partial x^i}+\frac{\partial p}{\partial x^i}\right), && (5.8)
\end{aligned}$$

where $\dddot{D}_{ij}$ is the third time derivative of the quadrupole moment:

$$\dddot{D}_{ij}=\mathrm{STF}\left[2\int dV\left(2p\frac{\partial v_i}{\partial x^j}+\frac{\partial\Phi}{\partial x^j}\left(x_i\frac{\partial\rho v^k}{\partial x^k}-2\rho v_i\right)-\rho x_i\frac{\partial\dot{\Phi}}{\partial x^j}\right)\right]. \tag{5.9}$$

The notation STF means *symmetric* and *trace free*, which is defined as

$$\mathrm{STF}\left[\chi_{ij}\right] \equiv \frac{1}{2}\chi_{ij} + \frac{1}{2}\chi_{ji} - \frac{1}{3}\delta_{ij}\chi_{kk}. \tag{5.10}$$

The gravitational potential, its time derivative, and R are described by Poisson equations:

$$\nabla^2\Phi = 4\pi G\rho, \tag{5.11}$$

$$\nabla^2\dot{\Phi} = -4\pi G\frac{\partial \rho v^i}{\partial x^i}, \tag{5.12}$$

$$\nabla^2 R = 4\pi G\dddot{D}_{ij}x^j\frac{\partial \rho}{\partial x^i}. \tag{5.13}$$

The conserved MHD variables are evolved as usual with the scheme described in chapter 2. After the primitive variables are recovered from the conserved ones at the end of each timestep, the kinematic velocities v^i are computed using equation 5.6. The source terms are updated in operator-splitting manner, and the Poisson's equation for the potentials (equations. 5.11-5.13) are solved using the scheme described in chapter 3. All spatial derivatives appearing in the above equations are computed using standard centered differences on the mesh.

The physical observables of interest are the waveforms and amplitudes of the gravitational wave radiation. By quadrupole approximations, the amplitudes for the two polarizations for an observer located at distance r perpendicular to the orbital plane are given by:

$$h_+ = \frac{G}{c^4}\frac{1}{r}\left(\ddot{D}_{xx} - \ddot{D}_{yy}\right), \tag{5.14}$$

$$h_\times = \frac{G}{c^4}\frac{2}{r}\ddot{D}_{xy}, \tag{5.15}$$

where the second time derivative of the quadrupole moment is computed from

$$\ddot{D}_{ij} = \mathrm{STF}\left[2\int dV\rho\left(v_i v_j - x_i\frac{\partial\Phi}{\partial x_j}\right)\right]. \tag{5.16}$$

The gravitational wave luminosity may be computed from the standard definition in the quadrupole formula

$$L = \frac{1}{5}\frac{G}{c^2}\dddot{D}_{ij}\dddot{D}_{ij}. \tag{5.17}$$

This is a non-average luminosity since averaging is not well-defined during the final phase of the merger when the the orbit decays and gets further from circular.

A useful comparison is to consider the analytical expressions for inspiral of a point-mass on an xy-plane due to the gravitational wave emission given by Misner et al. (1973). The separation of two point-mass objects with masses m_1 and m_2, total mass $\mathscr{M} = m_1 + m_2$, and reduced mass $\mu = m_1 m_2 / \mathscr{M}$ as function of time is:

$$a(t) = a_0 \left(1 - \frac{t}{\tau_0}\right)^{1/4}, \tag{5.18}$$

where a_0 is the initial separation at $t = 0$ and τ_0 is the inspiral time (i.e. the time to reach $a = 0$) given by

$$\tau_0 = \frac{5}{256}\frac{c^5}{G^3}\frac{{a_0}^4}{\mu \mathscr{M}^2}. \tag{5.19}$$

Lai (1994) gives the emitted gravitational waves as:

$$\begin{aligned} h_+ &= \frac{1}{r}\left(\ddot{D}^p_{xx} - \ddot{D}^p_{yy}\right), && (5.20) \\ h_\times &= \frac{2}{r}\ddot{D}^p_{xy}\cos\theta(t), && (5.21) \end{aligned}$$

where $\ddot{D}^p_{ij}$ is the second time derivative of the quadrupole moment for point mass, given by:

$$\begin{aligned} \ddot{D}^p_{xx} &= -2G\frac{\mu\mathscr{M}}{a}\cos(\theta(t) - \theta_0), \\ \ddot{D}^p_{yy} &= +2G\frac{\mu\mathscr{M}}{a}\cos(\theta(t) - \theta_0), \\ \ddot{D}^p_{xy} &= -2G\frac{\mu\mathscr{M}}{a}\sin(\theta(t) - \theta_0), \end{aligned} \tag{5.22}$$

for observer perpendicular to the orbital plane. The two waveform polarizations then may be reduced to

$$h_+ = -\frac{4}{r}\frac{\mathscr{M}\mu}{a(t)}\cos(\Phi(t)-\Phi_0), \tag{5.23}$$

$$h_\times = -\frac{4}{r}\frac{\mathscr{M}\mu}{a(t)}\sin(\Phi(t)-\Phi_0), \tag{5.24}$$

where $\Phi(t)$, the angle of the axis connecting the two point masses relative to the x-axis as function of time, is given by:

$$\Phi(t) = \frac{16}{5}c\tau_0\left(\frac{GM}{c^2a_0^3}\right)^{1/2}\left[1-\left(\frac{a(t)}{a_0}\right)^{5/2}\right]+\Phi_0. \tag{5.25}$$

For point masses initially on the x-axis, we set $\Phi_0 = \Phi(t=0) = 0$.

5.3 Simulation Setup

In this simulation we consider the merger of two $1.4M_\odot$ neutron stars with initial separation of $a_0 = 40$ km in the xy-plane. The stars are centered on $x = \pm 20$km. The neutron stars are modeled by “cold“ $\gamma = 2$ polytropes with initial profile produced by the method in §3.4.5. This gives us a profile with central density $\rho_c \simeq 2.18\times10^{15}$ g cm^{-3} and radius $r_s = 10$ km.

The initial separation was chosen to be wide enough so that tidal interaction between stars is negligible. At this separation $(a \gg r_s)$ the inspiral is mainly driven by the loss of energy and angular momentum due to the gravitational wave emissions. Initially the stars should follow the trajectory of point-mass inspiral. If we neglect the energy and angular momentum loss, the initial orbit is then the circular Keplerian orbit. Therefore we give the stars initial velocities of point-mass circular orbits:

$$v_y = \pm\sqrt{\frac{Gm}{2a_0}}, \tag{5.26}$$

where m is the mass of the star.

We set a low mass and internal energy density atmosphere in the computational domain of about 5 orders of magnitude lower than the star central density initially. During the simulation we also set the mass density floor to 1×10^5 g cm^{-3}. For cells with mass density lower than three times the mass density floor, we set the velocities to zero artificially to avoid problems with simulation timestep due to spurious high-velocity, low-density material. We use "ratcheting" boundary conditions which allow material to freely flow out of the computational domain but prevent matter to come onto the grid from outside. In the simulation run, the Courant number is set to $C = 0.3$ and the slope limiter parameter is $\theta = 1.9$. Our computational domain boundary is at ±40 km in all dimensions.

During the simulation, we only compute the quantity $\ddot{D}_{ij}$ rather than the waveforms, and write it to disk. The gravitational waves are then reconstructed in post-processing manner. Since disk activities are computationally costly, we only do enough sampling to capture a maximum wave frequency of 1×10^4 Hz. (For comparison, the frequency band of ground-based detectors is around $f \sim 10 - 10000$ Hz (Abramovici et al., 1992)). This means that rather than writing $\ddot{D}_{ij}$ to disk every simulation timestep, we only write it every $1/(2 \times 10^4)$ seconds (the Nyquist rate) in simulation time. During the evolution, if the next timestep is greater than required for the sampling, we reduce the timestep appropriately to hit the desired time.

We first show that our assumptions of the initial conditions are valid (i.e. negligible tidal interaction and point-mass approximation) by running a simulation with the back-reaction terms, energy and momentum loss terms due to the gravitational waves emission, turned off. This is done by removing the second term of the right hand side of equations 5.2 and 5.3. Without these terms, the stars should maintain circular orbits on a time scale greater than the dynamical time t_D for each star, where

$$t_D = \left(\frac{r_s^3}{Gm} \right)^{1/2} = 8.68 \times 10^{-5} \text{ s}. \tag{5.27}$$

This also shows that, with sufficient resolution, numerical viscosity in our code is low enough to be negligible.

Figure 5.1 shows the gravitational waves of this setup in simulation with 128^3 cells resolution. In this simulation, the stars finally coalesce around $t_D = 160$. The is caused by the numerical viscosity due to insufficient resolution. Increasing the resolution to 256^3 total cells removes this problem, as shown in figure 5.2. The circular orbit remains stable up to $t_D = 200$. The conservation of energy and angular momentum are also better conserved with this resolution (figure 5.3). These tests set the minimum resolution to use in the merger simulation.

5.4 Merger and Observables

The evolution of the inspiraling and merging of the neutron stars are shown as snapshots in figures 5.4, 5.5, and 5.6. Slices of mass density in *xy*-plane are plotted in all of the panels, together with a semi-transparent contour plot at $\rho = 1 \times 10^{12}$ gm cm^{-3}. A threshold is applied in the visualization to exclude the plotting of lower density. Minimum and maximum value of the density are shown below the legend bar on each panel.

The inspiral phase of the evolution is depicted in figure 5.4. The first panel of the figure shows the initial condition, when the stars are still in quasi-circular orbit, in counter-clockwise direction. Here we see the final two orbits that last about 4.6 ms before the stars start to touch each other. In this inspiral phase, initially the back-reaction terms due to the gravitational wave emissions dominate in driving the inspiraling as the system loses angular momentum and energy (see figure 5.8). As the stars get closer to each other, their tidal bulges grow and the tidal effects start to take over. The stars follow closely to the point-mass approximation initially, but diverge from point-mass approximation as tidal effect starts to take over (e.g. see figure 5.7).

Figure 5.5 shows the merger phase of the evolution. The merger phase happens fairly rapidly, lasting only about 1.2 ms. The first panel in the figure shows that the stars have already touched each other and formed a bar-like structure. Some mass and angular momentum is lost through the spiral arms that are formed due to the gravitational torque. As the central objects coalesce to a single massive object, the arms finally grow to form a

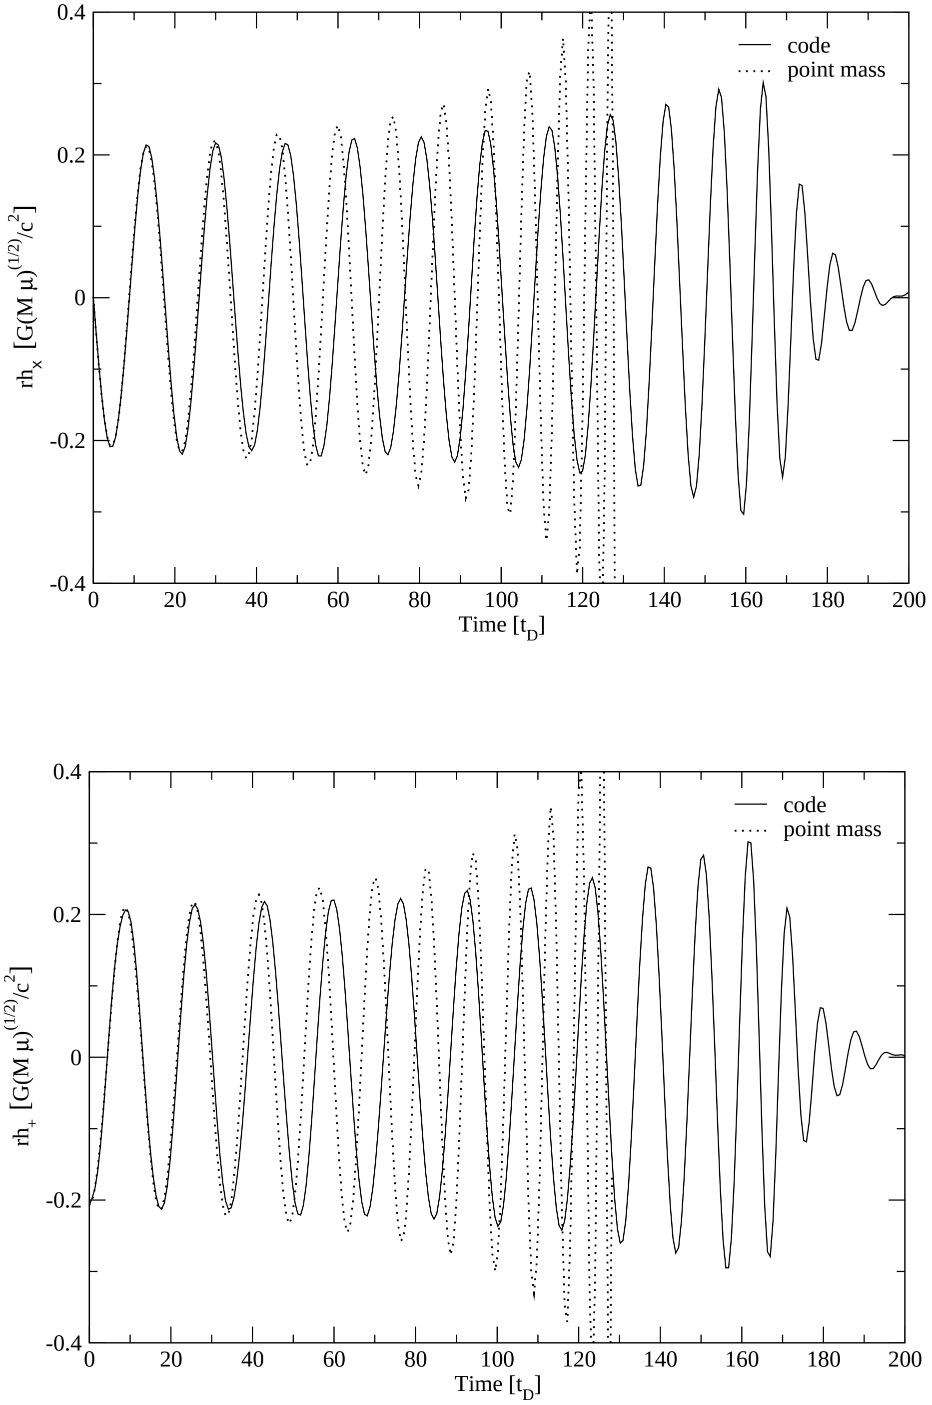

Figure 5.1: The two polarizations of gravitational waves of a binary neutron stars simulation without back-reaction terms, with mesh resolution of 128^3 cells. The dotted lines show the gravitational waves for inspiraling point-masses.

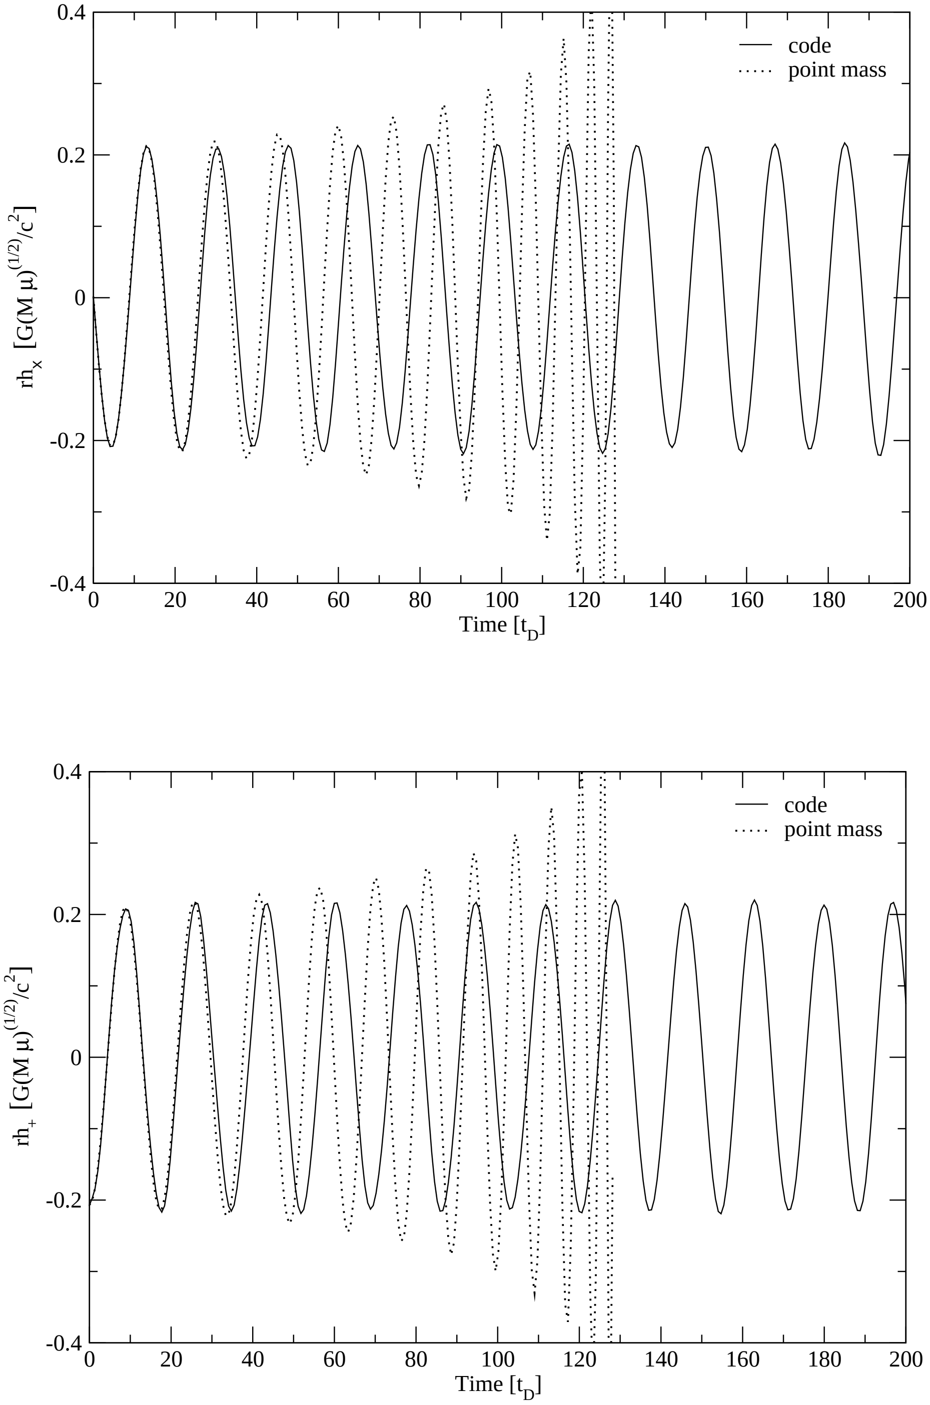

Figure 5.2: The two polarizations of gravitational waves of binary neutron stars simulation without back-reaction terms with a mesh resolution of 256^3 cells. The dotted lines show the gravitational waves for inspiraling point-masses.

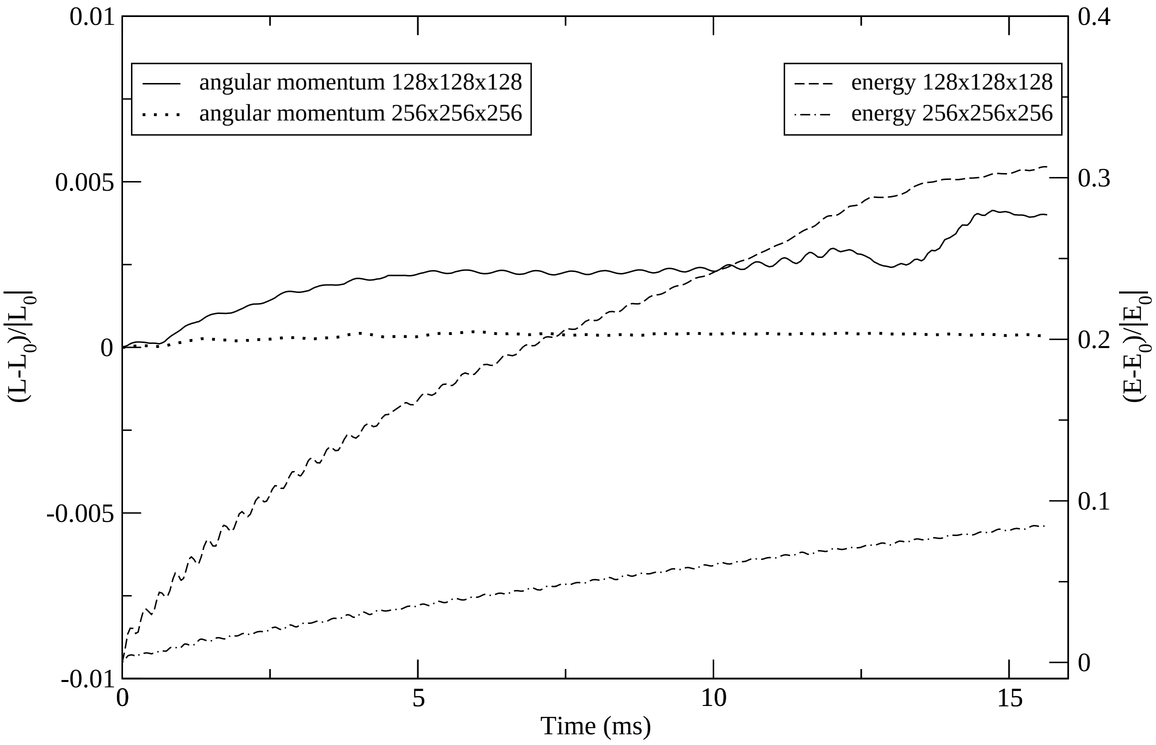

Figure 5.3: Conservation of angular momentum (left axis) and energy (right axis) plotted as relative difference to their initial value of simulation without back-reaction terms for 128^3 and 256^3 cells resolution.

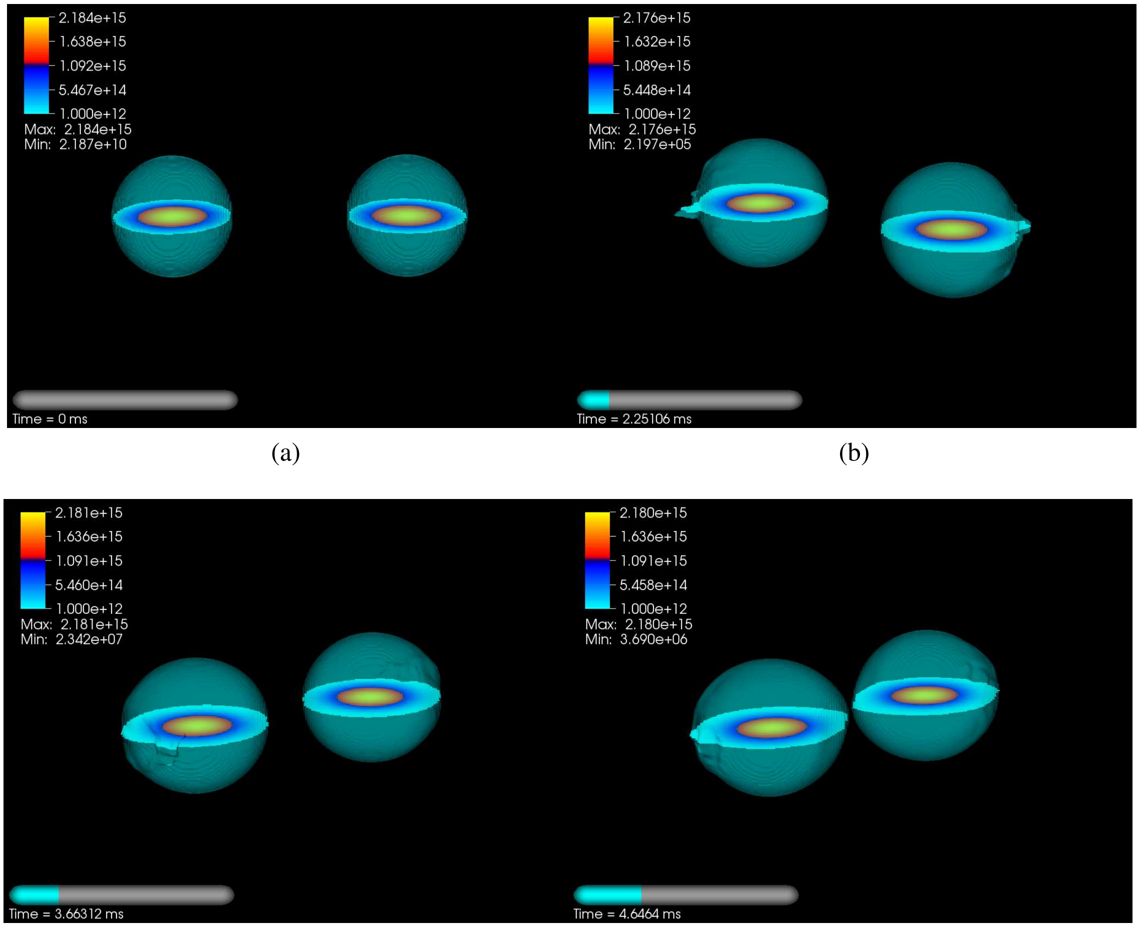

(a) (b)

(c) (d)

Figure 5.4: The last several orbits of the inspiraling phase of the binary neutron stars starting from the initial condition of the simulation as shown in panel (a). A mass density slice in *xy*-plane through the center of each star is plotted here, together with a semi-transparent contour plot of mass density at $\rho = 1 \times 10^{12}$ gm cm^{-3}. A visualization threshold operator is applied to exclude the plotting of lower value, although the minimum value on the grid is shown on the legend. The stars are orbiting in the counter-clockwise direction as viewed looking down from positive *z*-direction.

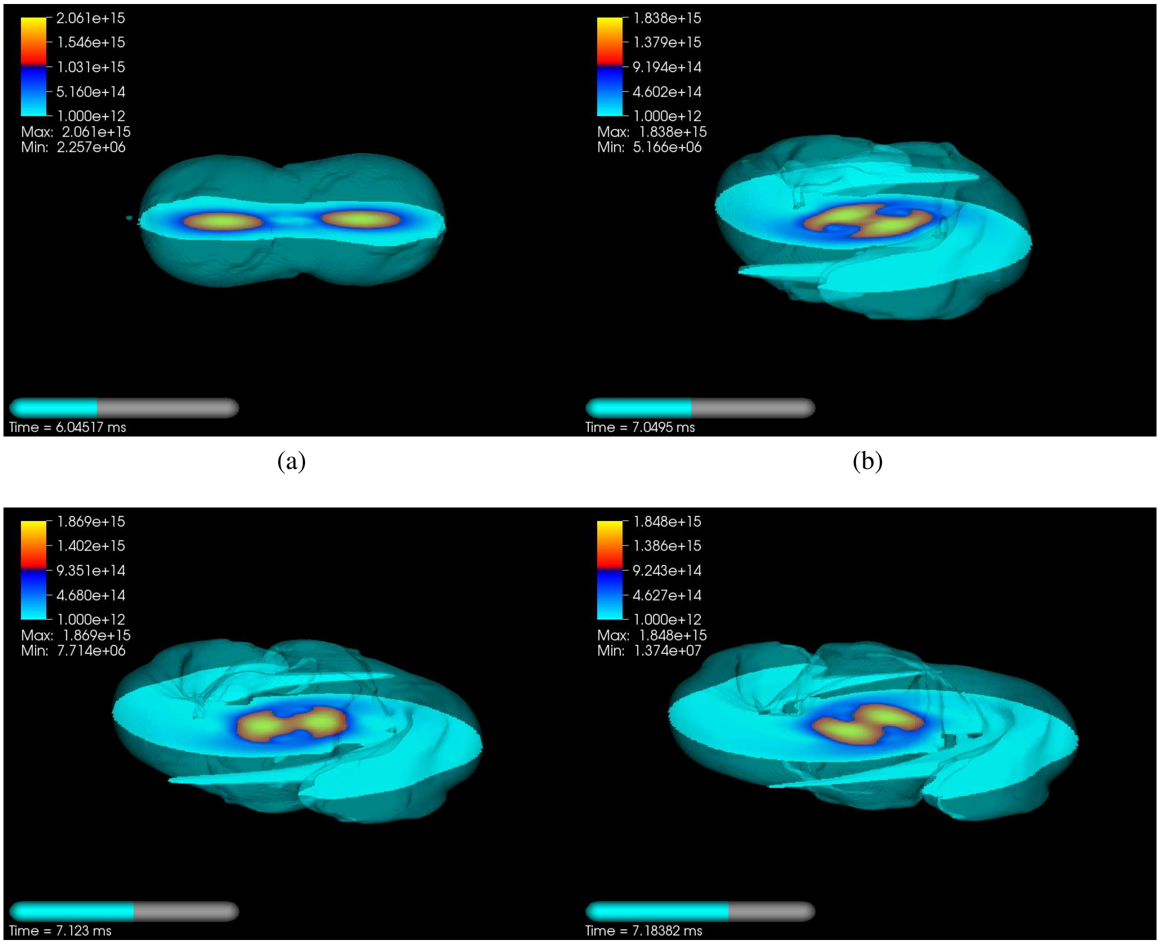

(a) (b)

(c) (d)

Figure 5.5: The merger phase of the neutron stars with the same plots as in figure 5.4. Panel (a) shows the plot after the initial contact of the two stars, followed by the coalescence of the stars within 1.2 ms, after which the two initial masses almost fully merge, as shown in panel (d).

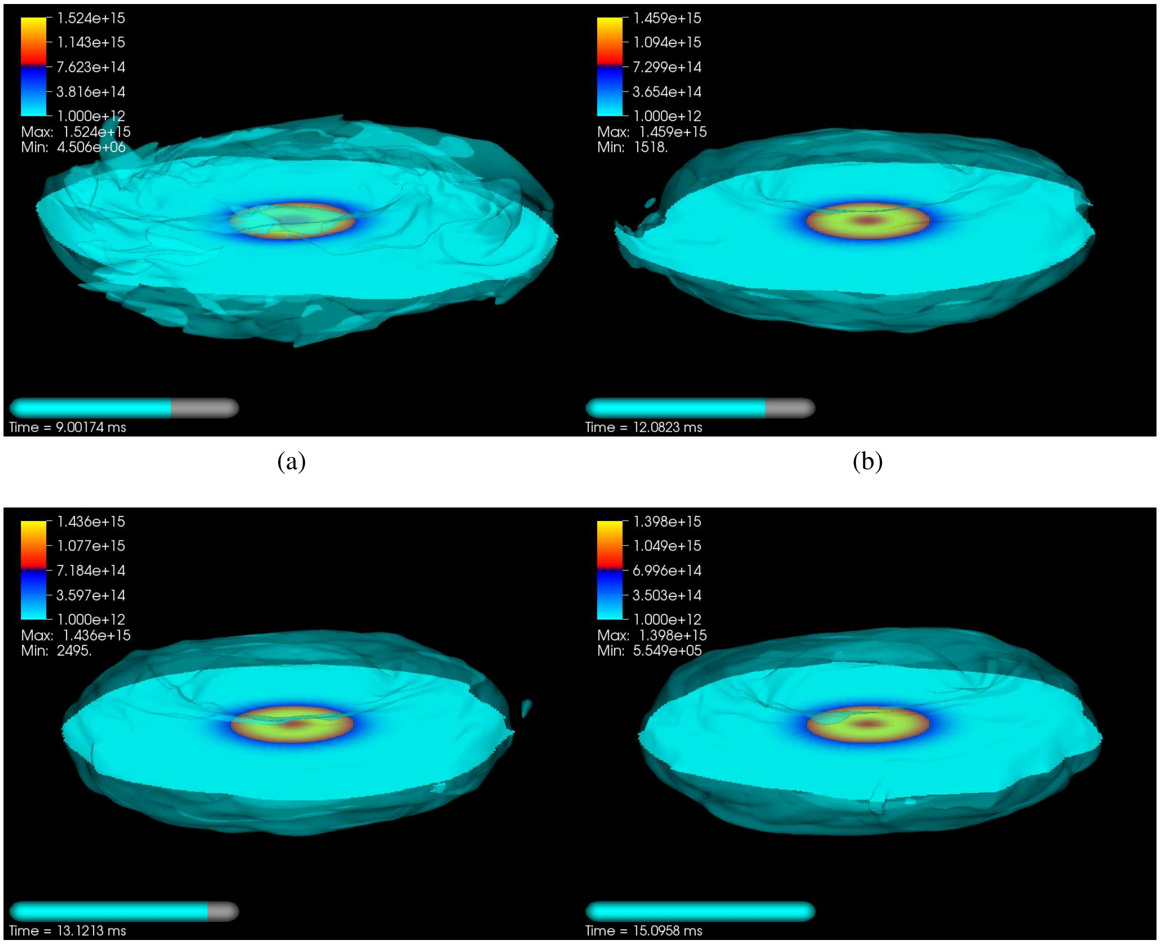

(a) (b)

(c) (d)

Figure 5.6: The final / ring-down phase of the neutron stars merger with the same plots as in figure 5.6. Here we see that a large rotating disk around a central object has formed from the merger.

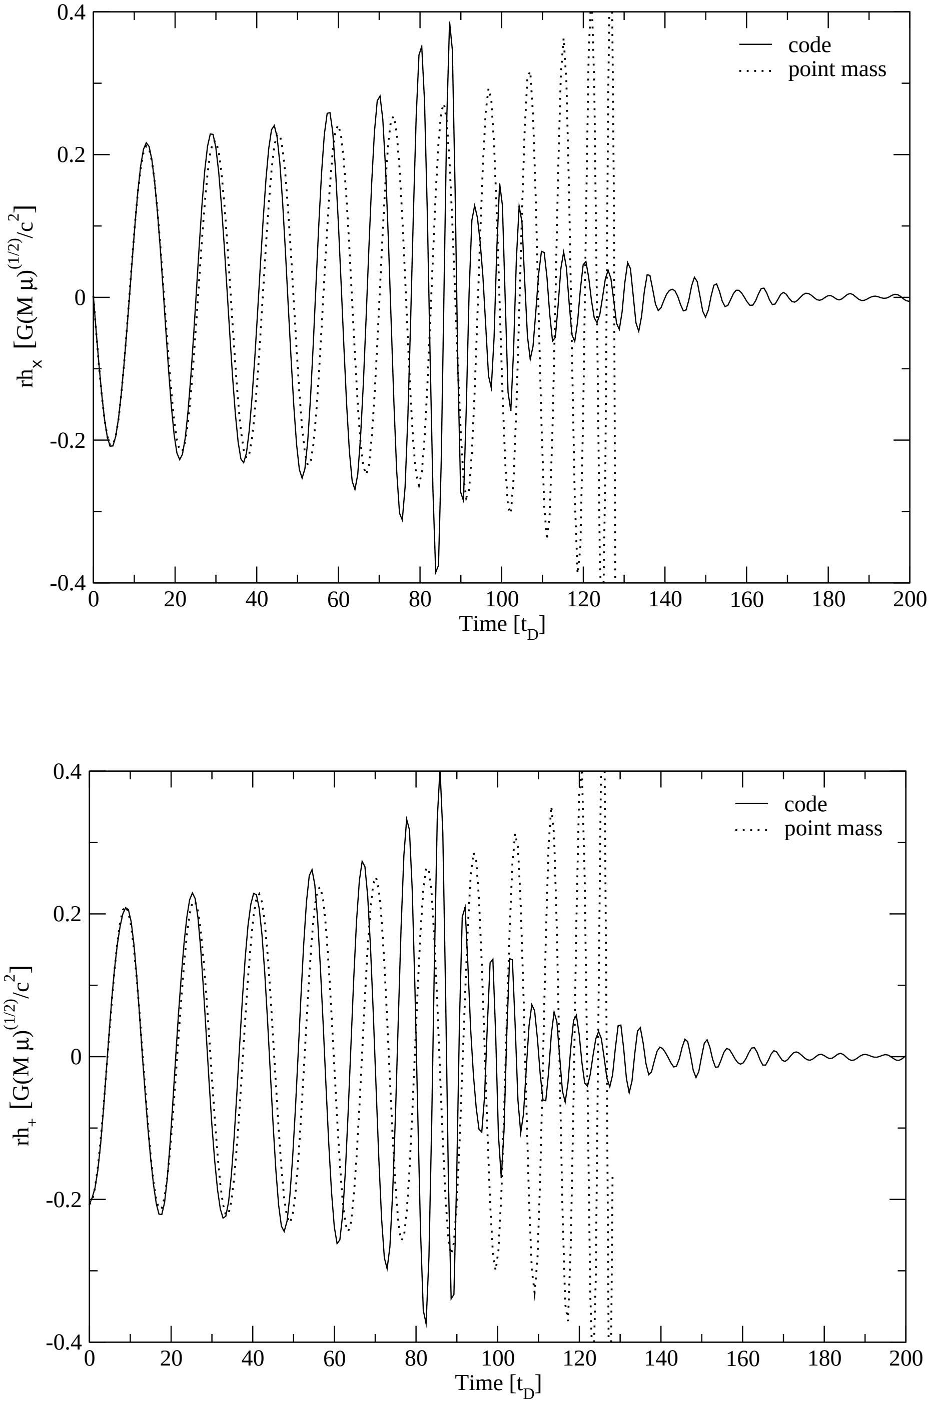

Figure 5.7: The two polarizations of gravitational waves for the merger simulation. The dotted lines show the gravitational waves of inspiralling point-mass.

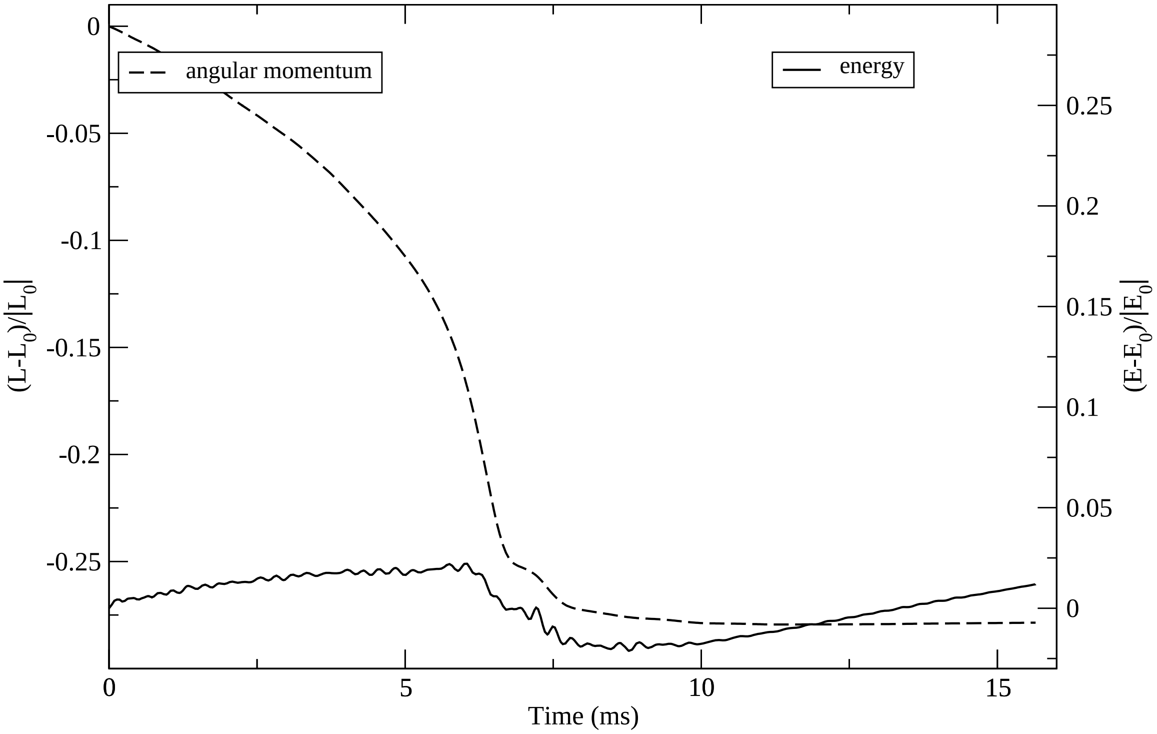

Figure 5.8: Angular momentum (left axis) and energy (right axis) plotted as relative difference to their initial value in the merger simulation.

disk around it. At the end of this phase (last panel in figure 5.5), the two stars are barely distinguishable.

The merger phase is followed by a "ring down" phase, shown in figure 5.6, as the spiral arms are fully grown and have formed a massive disk around the central object. The disk becomes increasingly axisymmetric, halting the production of the gravitational waves. No more loss of angular momentum and energy happens in this phase. We see in figure 5.8 that the system conserves its energy and momentum for $t > 10$ ms. Some differential rotation may still occur in the disk. It is likely that in this phase the central object collapses to form a black hole (Lattimer, 2000), although our simulation cannot show that due to it being Newtonian. How much mass can be supported and how long it takes before the central object collapses to form a black hole depends on the equation of state of the neutron stars, and is still open for investigation (Cook et al., 1994; Lattimer, 2000; Morrison et al., 2004).

Figure 5.7 shows the two gravitational waves polarizations produced by the system starting from the initial conditions of the simulation to the ring down phase. As previously mentioned, we see that the waves track the waves produced by a point-mass analytical calculation initially, but quickly goes out of phase as tidal effects start to dominate in the late inspiral phase. This causes the stars to spiral faster than the point-mass counterpart, manifested in the increase of frequency and amplitude of the waveforms starting around $t_D = 60 = 5.2$ ms. The waves reach the maximum at around $t_D = 85$, and quickly decreases as the object formed by the coalescence grows axisymmetric. By $t_D = 160$, the gravitational waves production practically ceases.

It is also instructive to look at the energy emitted per unit frequency interval dE/df for the gravitational waves. The spectrum data contains rich information and signatures that are not immediately obvious from the waveforms. Thorne (1989) gives the expression for dE/df in the form:

$$\frac{dE}{df} = \frac{c^3}{G}\frac{\pi}{2}4\pi r^2 f^2 \langle |\widetilde{h}_+(f)|^2 + |\widetilde{h}_\times(f)|^2 \rangle, \tag{5.28}$$

where the angle brackets denotes average over all sources, and $\widetilde{h}(f)$ is the Fourier transform of the function $h(t)$. Zhuge et al. (1994) shows that

$$\langle|\widetilde{h}_+(f)|^2+|\widetilde{h}_\times(f)|^2\rangle=\langle|\widetilde{h}_+(f)|^2\rangle+\langle|\widetilde{h}_\times(f)|^2\rangle, \tag{5.29}$$

where

$$\begin{aligned}\frac{c^8}{G^2}r^2\langle|\widetilde{h}_+(f)|^2\rangle &= \frac{4}{15}|\widetilde{\ddot{D}}_{xx}-\widetilde{\ddot{D}}_{zz}|^2+\frac{4}{15}|\widetilde{\ddot{D}}_{yy}-\widetilde{\ddot{D}}_{zz}|^2+\frac{1}{10}|\widetilde{\ddot{D}}_{xx}-\widetilde{\ddot{D}}_{yy}|^2\\ &\quad+\frac{14}{15}|\widetilde{\ddot{D}}_{xy}|^2+\frac{4}{15}|\widetilde{\ddot{D}}_{xz}|^2+\frac{4}{15}|\widetilde{\ddot{D}}_{yz}|^2, &(5.30)\\ \frac{c^8}{G^2}r^2\langle|\widetilde{h}_\times(f)|^2\rangle &= \frac{1}{6}|\widetilde{\ddot{D}}_{xx}-\widetilde{\ddot{D}}_{yy}|^2+\frac{2}{3}|\widetilde{\ddot{D}}_{xy}|^2+\frac{4}{3}|\widetilde{\ddot{D}}_{xz}|^2+\frac{4}{3}|\widetilde{\ddot{D}}_{yz}|^2. &(5.31)\end{aligned}$$

The energy spectrum for point-mass inspiral is given by Cutler and Flanagan (1994):

$$\frac{dE^p}{df}=\frac{\pi}{3}\frac{G\mu\mathscr{M}}{c}\left(\frac{c^3}{\pi G\mathscr{M}f}\right)^{1/3}. \tag{5.32}$$

We have computed the spectrum dE/df for our simulation using equation 5.28. We prepend the simulation data with $\ddot{D}^p_{ij}$ from point-mass formula (equation 5.22) before calculating the Fourier transform. By doing so, we have a long region of the inspiral phase in the frequency domain to ensure that we can see any cutoff frequency in our data. This is valid since initially the evolution of the simulation tracks the inspiral of the point mass (e.g. see figure 5.7).

Figure 5.9 shows the gravitational wave spectra dE/df with the solid line representing the spectrum from the simulation preprended with point-mass data in the low-frequency inspiral regime. The dotted line shows the spectrum of the point-mass inspiral calculated from the analytical formula. We can identify some features of this plot. As f increases, we see that dE/df first drops below the point-mass inspiral value at $f\sim 800$ Hz, and then again reaches a local minimum at $f\sim 1500$ Hz. We identify this as the onset of the regime where the tidal effect stars to dominate, which causes the stars to coalesce faster than the point-mass inspiral. This is followed by a peak at $f\sim 2200$ Hz and secondary peaks at

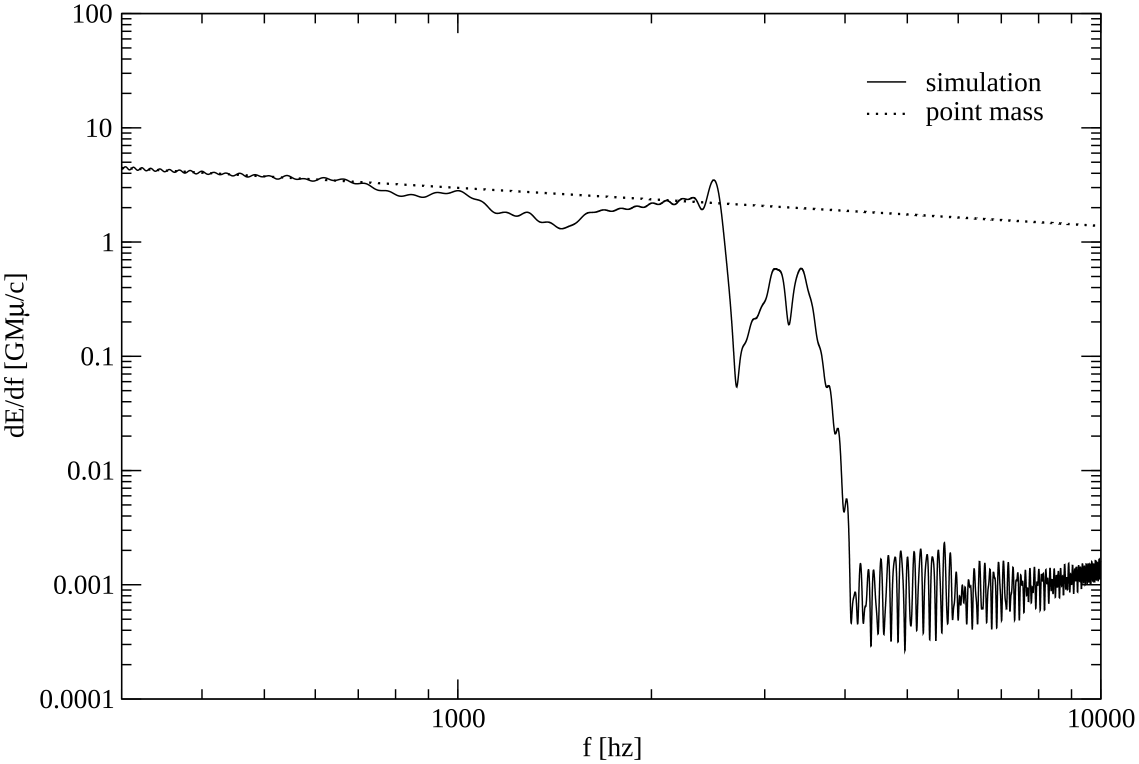

Figure 5.9: Gravitational wave spectrum dE/df. Solid line represent the spectrum from the simulation pre-prended with point-mass data in the low-frequency inspiral regime. The dotted line shows the spectrum of the point-mass inspiral calculated from the analytical formula.

$f \sim 3200$ Hz which we associate with the fairly rapid merger phase. The cutoff in dE/df happens rapidly as the gravitational waves production ceases.

5.5 Conclusion

In this chapter we have carried out a fully three-dimensional simulation of neutron star merger with GenASiS. Although our model is still fairly simple, this simulation serves as a milestone in GenASiS development toward a more realistic merger simulation. We have shown that we can extract physical observables such as gravitational waveforms and spectra. These kinds of templates is necessary for ground-based gravitational wave detectors data analysis. We have also shown that the spectra from such an event lies on the frequency range of f $10 - 10000$ Hz, the range covered by broadband detectors such as LIGO and VIRGO.

The relationship between mass and radius of neutron stars is determined by the nuclear equation of state. As such, knowledge of these relationships provide constraints on the form of the nuclear equation of state. In this simulation, we have used a simple model of polytropes for the equation of state of the neutron star. A parameterized studies of the equation of state covering the parameter spaces of mass and radius and their effects on the gravitational waveforms signature is necessary to provide a catalog of observables. In the event of a physical observation then, a matching in the catalog may gives us knowledge of the neutron star quantities such radius or mass, and thereby a better constraint on the equation of state (Lindblom, 1992; Cutler et al., 1993; Cutler and Flanagan, 1994). Such simulations using the Langragian smooth particle hydrodynamics (SPH) have been carried out by Zhuge et al. (1996). Our simulations have shown that we have the necessary tools and techniques to do a similar and more thorough study with Eulerian formalism even just using the current version of GenASiS (see chapter 1 for the motivation of Eulerian formalism). Such parameter studies will be the subject of future work.

Chapter 6

Generation of Magnetic Fields by the Stationary Accretion Shock Instability

In this chapter we explore the amplification of magnetic fields in the post-bounce environment of core-collapse supernova. These simulations demonstrate a suitable scientific application for GenASiS.

6.1 Introduction

Despite over four decades of intensive investigations, the details of the mechanism behind core-collapse supernovae remains elusive. In the modern paradigm for core-collapse supernovae, there is a phase of the stalled shock after the core bounce, followed by the revival of the shock that finally results in an explosion that disrupts the outer layers of the star, leaving behind a new neutron star. The details of the stalled shock revival are at the heart of the core-collapse supernova mechanism, and therefore subject to considerable research efforts. Recent studies have shown that the stalled shock itself is unstable to non-radial perturbation (Blondin et al., 2003; Blondin and Mezzacappa, 2006), and this instability, dubbed the "stationary accretion shock instability" (SASI) may play critical role in reviving the stalled shock that leads to explosions. Recent multiphysics supernova simulations have confirmed the existence of the SASI (Bruenn et al., 2006; Buras et al., 2006; Burrows et al., 2006; Scheck et al., 2008; Marek and Janka, 2009). Furthermore, the importance of SASI has also been noted for its role in improving the conditions for neutrino energy deposition in the post-shock gas (Scheck et al., 2008; Marek and Janka, 2009), and in explaining the velocity distribution of young pulsars (Scheck et al., 2006).

Although the role of magnetic fields in supernova explosion mechanisms remains unclear, recent interest in its relevance in the context of core-collapse supernova has significantly increased. The discovery of magnetars (Duncan and Thompson, 1992)—neutron stars with very strong magnetic field (in the order of 10^{14}–10^{15} G)—, the theoretical discovery of magnetorotational instability (MRI) (Balbus and Hawley, 1991), and the observation of collimated jets in supernovae associated with gamma-ray bursts (Woosley and Bloom, 2006) help sparks this interest. The generally observed asphericity of core-collapse supernova explosions (Wang et al., 2001b) was also long thought to be caused by a magnetic field with rotation, although with the recent discovery of SASI this assumption may be partially challenged. Recent works therefore have included magnetic fields in the modeling of core-collapse supernovae (Kotake et al., 2004; Ardeljan et al., 2005; Obergaulinger et al., 2006; Thompson et al., 2005; Shibata et al., 2006b; Burrows et al., 2007; Cerdá-Durán et al., 2007; Endeve et al., 2007; Suzuki et al., 2008; Mikami et al., 2008; Takiwaki et al., 2009).

The magnetohydrodynamics studies of core-collapse supernovae mentioned above rely on the progenitor having rapid rotation as the necessary ingredient for magnetic field amplification during collapse and post-bounce. The rotation acts as both energy reservoir and mechanical agent for the amplification via compression during collapse and winding of the field during the post-bounce phase. Strong magnetic fields, orders of magnitude stronger than the initial field of the progenitor star (Heger et al., 2005), are necessary if the magnetic field is to have any significant effect on the ensuing dynamics, hence the amplification. Most progenitors however are expected to have only modest initial rotation and initial field strength. This is one of the reasons that magnetohydrodynamics had been largely ignored for the modeling of core-collapse supernovae in the past.

The recent discovery of SASI and its effect in generating phenomena previously attributed to the rotation of progenitor, particularly the asphericity of the explosion and pulsar spin (Blondin and Mezzacappa, 2007), raise the question whether it can also act as an agent of magnetic field amplification in the absence of initial rotation. Many of the past supernova simulations that include magnetic fields have either been done in only a single quadrant for the computational domain or only followed the post-bounce evolution up to a few tens of milliseconds. Both of these constraints necessarily exclude the development of SASI. Most studies have also been carried out with axial symmetry imposed, which excludes the development of the spiral mode of SASI (Blondin and Mezzacappa, 2007). Therefore in our study, we extend the models of Blondin et al. (2003) and Blondin and Mezzacappa (2007) to include initially weak magnetic field, comparable in strength to the progenitor star. We perturb the initial condition and follow the evolution of magnetic fields as SASI develops. Our results show that in the three-dimensional (3D) models, the flows are eventually dominated by the $m = 1$ (spiral) mode that amplify the magnetic field via flux tube stretching. Although the amplification of magnetic fields in our simulation is ultimately limited by numerical resistivity due to finite spatial resolution, our simulations suggest that SASI-generated magnetic fields are unlikely to be dynamically important for the dynamics of core-collapse supernovae. Nevertheless, our simulations suggest a mechanism for the magnetization of proto-neutron star (PNS) in absence of initial progenitor rotation. What follows is a summary of the study we previously published in Endeve et al. (2010).

6.2 Model and Numerical Setup

6.2.1 Initial Conditions

We set up our models in an idealized post-bounce supernova environment. The magnetized fluid is described by the ideal magnetohydrodynamics of equations 2.33 – 2.37, and modified to include gravity as source term in the conservation of momentum and energy equations as:

$$
\begin{aligned}
\frac{\partial D}{\partial t}+\frac{\partial}{\partial x^i}\left(\rho v^i\right) &= 0, && (6.1)\\
\frac{\partial S^j}{\partial t}+\frac{\partial}{\partial x^i}\left(\rho v^j v^i+p\delta^{ij}-B^jB^i\right) &= -\rho\frac{\partial \Phi}{\partial x^j}, && (6.2)\\
\frac{\partial E}{\partial t}+\frac{\partial}{\partial x^i}\left(\left[e+p+\frac{1}{2}\rho v^j v^j+\frac{1}{2}B^jB^j\right]v^i-B^i(\mathbf{B}\cdot\mathbf{v})\right) &= -\rho v^i\frac{\partial \Phi}{\partial x^i}, && (6.3)\\
\frac{\partial \mathbf{B}}{\partial t} &= \nabla\times(\mathbf{v}\times\mathbf{B}), && (6.4)\\
\nabla\cdot\mathbf{B} &= 0. && (6.5)
\end{aligned}
$$

The variables have the same meaning as before (see § 2.1.3), with the gravitational potential Φ given by the point-mass formula $\Phi = -GM/r$, where G is Newton's gravitational constant, M is the mass of the central object, and r is the radial distance from the center of the star. We use a polytropic equation of state with adiabatic index $\Gamma = 4/3$. We also decompose the total energy E as $E = e_{\text{int}} + e_{\text{kin}} + e_{\text{mag}}$ where $e_{\text{int}} = p/(\Gamma-1)$ is the internal energy density, $e_{\text{kin}} = (\rho\mathbf{v}\cdot\mathbf{v})/2$ is the kinetic energy density, and $e_{\text{mag}} = (\mathbf{B}\cdot\mathbf{B})/2$ is the magnetic energy density (also often referred as the magnetic pressure).

We follow a similar setup of a post-bounce stalled supernova shock as in Blondin et al. (2003) and Blondin and Mezzacappa (2007). The mass of the central object is set to $M = 1.2\,M_{\odot}$, and is set to be constant throughout the simulation. A steady-state, spherically symmetric accretion shock is placed at $r = R_{\text{sh}} = 200$ km with accretion rate set to $0.36\,M_{\odot}\,\text{s}^{-1}$. This accretion rate is maintained throughout the simulation. Although this accretion rate is large enough to increase the mass of the central object over the timescales we consider, we artificially set the rate of mass change of the central object to $\partial M/\partial t = 0$ so that a steady state solution of unperturbed initial condition may be constructed and compare against for this study. Ahead of the shock, matter falls into the shock with free-fall speed. We use highly supersonic Mach number of 300 to set the pressure in this pre-shock gas. The hydrodynamic state just inside the shock is then determined from the Rankine-Hugoniot jump condition (Courant and Friedrichs, 1977). The structure of the fluid just inside the shock down to the inner boundary of $r = R_{\text{PNS}} = 40$ km is obtained by solving the Bernoulli equation. A "cutout" boundary if used as inner boundary of our grid and may loosely be interpreted as the surface of the proto-neutron star (PNS). Fluid is allowed to flow through the cutout inner boundary in a manner that in our experience maintains the steady state solution in simulations without non-radial perturbation: the fluid velocity is held fixed to its initial value just inside the inner boundary; the power laws for mass density ($\rho \propto r^{-3}$) and pressure ($P \propto r^{-4}$) obtained from the Bernoulli equation are used to dynamically interpolate values from cells just outside the inner boundary to ghost cells just inside the boundary.

The strength and topology of the magnetic field in a supernova progenitor is not known with confidence. For our models, we simply consider a purely radial initial field. This choice is consistent with steady-state initial condition. Also collapse tends to drag any initial higher-order multipole moments of the field into a more radial configuration. Consider an initial poloidal field of about 10^6 G for the progenitor (Heger et al., 2005). During the collapse the mass density increase about 5 orders of magnitude, and the magnetic field strength increases roughly as $B \propto \rho^{2/3}$, since the field is "frozen-in" to the fluid. This gives a field strength of roughly a few times 10^9 G in the collapse core. In our base model we set $B_0 = 10^{12}$ G, where B_0 is the field strength at $r = R_{\mathrm{PNS}}$. This initial field is not expected to have any effect on the early development of the development of SASI since both the ratio of the magnetic pressure to the fluid pressure and the ratio of the magnetic energy density to the kinetic energy density are small everywhere inside the shock. The value of the former is $e_{\mathrm{mag}}/P < 2 \times 10^{-11}$, and the latter is $e_{\mathrm{mag}}/e_{\mathrm{kin}} < 2 \times 10^{-8}$. We have also varied B_0 in some of our models.

The magnetic field is added to the fluid in a "split monopole" fashion with $\mathbf{B} = B_r \mathbf{e_r}$ with $B = \mathrm{sign}(\cos\theta) \times B_0 (R_{\mathrm{PNS}}/r)^2$, where θ is the polar angle in spherical coordinate system. The magnetic field has positive and negative polarity in the northern and southern hemisphere respectively, implying existence of a thin current sheet in the equatorial plane. At the outer boundary the magnetic field is held constant throughout the simulations.

In one of our models, we add rotation about the z-axis to the initial condition by setting the azimuthal velocity to $v_\phi = l \sin\theta / r$ where l is the (constant) specific angular momentum. This method of initializing the rotation is similar to the one used by Iwakami et al. (2009), except that we include rotation from the onset of the simulation rather than introduce it in the nonlinear evolution. Since our inner boundary is spherical, this study is limited to include only models with rotation rates that do not result in significant deviation from spherical symmetry. Specifically, in the rotating model we set the specific angular momentum to $l = 1.5 \times 10^{15}$ cm^2 s^{-1}. This is a relatively slow rotation, but consistent with the stellar evolution calculation by Heger et al. (2005).

6.2.2 Steady-state Standing Accretion Shock

Before we begin the study of SASI with magnetic field, we needed to demonstrate that our code is able to maintain the steady-state solution of the standing accretion shock when there is no non-radial perturbation of the initial conditions. We choose a resolution that is high enough to give acceptable results for maintaining the steady-state solution. From experiments, we found this to be a cell width of $\Delta l \approx 1.56$ km. Thus, this is the resolution we use as our base resolution in this study. We will present also results where we vary the resolution. At the late times in the calculations however, we do not find convergence because highly nonlinear flows have structures on all grid scales, although global values (e.g. total kinetic energy inside the shock) are qualitatively the same for models with various resolutions.

With our base resolution of $\Delta l \approx 1.56$ km, we evolve the steady-state solution to the standing accretion shock up to the timescale comparable to the time between bounce and shock revival in a core-collapse supernova. This duration is about 1 second. This will show that our code is able to integrate the MHD equation correctly on this timescale. Due to numerical errors and slight mismatch of the analytical solution and the discretization, we expect transients from the analytical initial condition before the system finally settles to a new configuration that matches the steady-state solution for the discretized MHD equation.

The code's ability to maintain the steady-state solution is illustrated in figure 6.1. In this figure, total values on the grid, obtained by volume integrals over the computational domain, of internal energy density e_{int}, kinetic energy density e_{kin}, magnetic energy density e_{mag}, and gravitational energy density $e_{\text{grav}} = \rho\Phi$ are plotted as E_{int}, E_{kin}, E_{mag}, and E_{grav}, respectively. The accumulated total magneto-fluid energy flux density $F_{\text{fluid}} = E + p + \mathbf{B} \cdot \mathbf{B}/2$, and the gravitational energy flux density $F_{\text{grav}} = \rho\Phi\mathbf{v}$ that have been lost from the grid through the inner boundary and added to the grid from the inflow of the outer boundary, are also tracked. These are plotted as $\mathscr{F}^{-}_{\text{fluid}}$ and $\mathscr{F}^{-}_{\text{grav}}$ for the total magneto-fluid energy and gravitational energy lost through the inner boundary, and $\mathscr{F}^{+}_{\text{fluid}}$ and $\mathscr{F}^{+}_{\text{grav}}$ added to the grid due to the inflow outer boundary.

The initial transients period can be seen in figure 6.1 from time $t = 0$ to $t = 200$ ms. During this period the shock radius increases from $R_{\text{sh}} = 200$ km to $R_{\text{sh}} \approx 207$ km, where it settles to the configuration. The expansion of the shock radius increases the shocked volume, which correspond to a small net increase in E_{kin}, as seen in the figure. Beyond this initial transients period, the energies in the computational domain remain constant. The sum of all the curves in figure 6.1, plotted as thick solid black line, is constant within numerical precision. This shows the code's ability to maintain the steady-state solution and correctly integrate the MHD equations on the timescale of interest.

6.3 Magnetic Field Amplification

6.3.1 Magnetic Field Evolution in SASI

In this section we describe and explain the amplification of magnetic field due to SASI-induced flows in our three-dimensional models.

The development of SASI needs to be initialized by a non-radial perturbation. In the axisymmetric case, Blondin et al. (2003) showed that the qualitative features of SASI evolution do not depends on the details of the perturbation. We confirm this result. In the three-dimensional case, eventually SASI is dominated by the $m = 1$ (spiral) mode, regardless of the perturbation, although the details of the early evolution may vary.

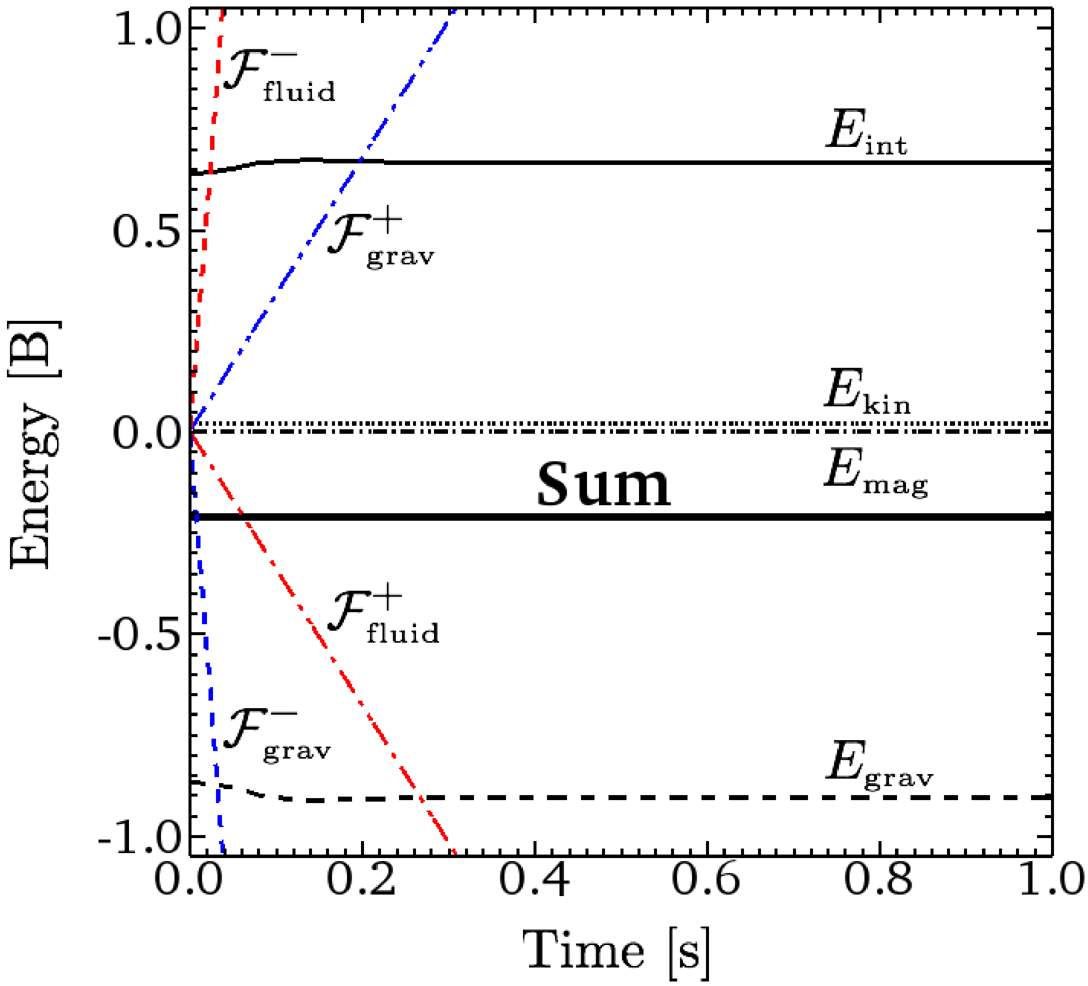

Figure 6.1: Conservation of energy for steady-state standing accretion shock. Total values on the grid of internal energy (black solid line), kinetic energy (black dotted line), magnetic energy (black dash-dot line) and gravitational energy are plotted over time. The energy unit is in Bethe, with 1 B = 10^{51} erg. The magneto-fluid energy and gravitational energy lost from the grid through the inner boundary are also plotted as red and blue dashed lines, respectively, while the magneto-fluid and gravitational energy coming into the grid due to inflow from the outer boundary are plotted as red and blue dot-dashed lines, respectively. The sum of all these are plotted as thick solid black line, showing conservation of energy within numerical precision.

Table 6.1: Tabular overview of three-dimensional SASI models.

Model Name	B_0 [G]	l [cm^2 s^{-1}]	Perturbation	Spatial Resolution [km]
3DB12Al	1×10^{12}	0	axisymmetric	$300/128$
3DB12Am	1×10^{12}	0	axisymmetric	$300/192$
3DB12Ah	1×10^{12}	0	axisymmetric	$300/256$
3DB12Rm	1×10^{12}	0	random	$300/192$
3DB12ΩRm	1×10^{12}	1.5×10^{15}	random	$300/192$
3DB10Rm	1×10^{10}	0	random	$300/192$

In our simulations, we use two different perturbations to initial the SASI. The first one, dubbed "axisymmetric" perturbation, is done by introducing two tori in the pre-shock gas, one in the northern and one in the southern hemisphere, whose density is increased by 20% compared to the non-perturbed steady-state flow. Axisymmetric perturbation may favor the $l = 1$ sloshing mode of SASI. For comparison therefore, we also use "random pressure" perturbation in our simulations, where a small amplitude (1%) random pressure is introduced inside the shock to initiate SASI.

For our 3D models, we have chosen to use the initial magnetic field $B_0 = 10^{12}$ G for most of our models. We also have one model with $B_0 = 10^{10}$ G for comparison. We also vary the spatial resolution of our simulations to investigate the effect of resolution on the magnetic field amplification. One of the models has an initial rotation with specific angular momentum $l = 1.5\times10^{15}$ cm^2 s^{-1} about the z-axis. Table 6.1 gives an overview of all our models.

6.3.1.1 Reference Model with Axisymmetric Perturbation

We begin by exploring the evolution of magnetic field in model `3DB12Am` as a reference model. The axisymmetric perturbation in this model initiate the sloshing mode of the SASI in the early state of the its development. Figure 6.2 provides an overview of the magnetic field evolution of this model.

Plotted in figure 6.2 are the magnitude of the magnetic fields overlaid with density contours at selected times. The axisymmetric perturbation initiate the sloshing mode in

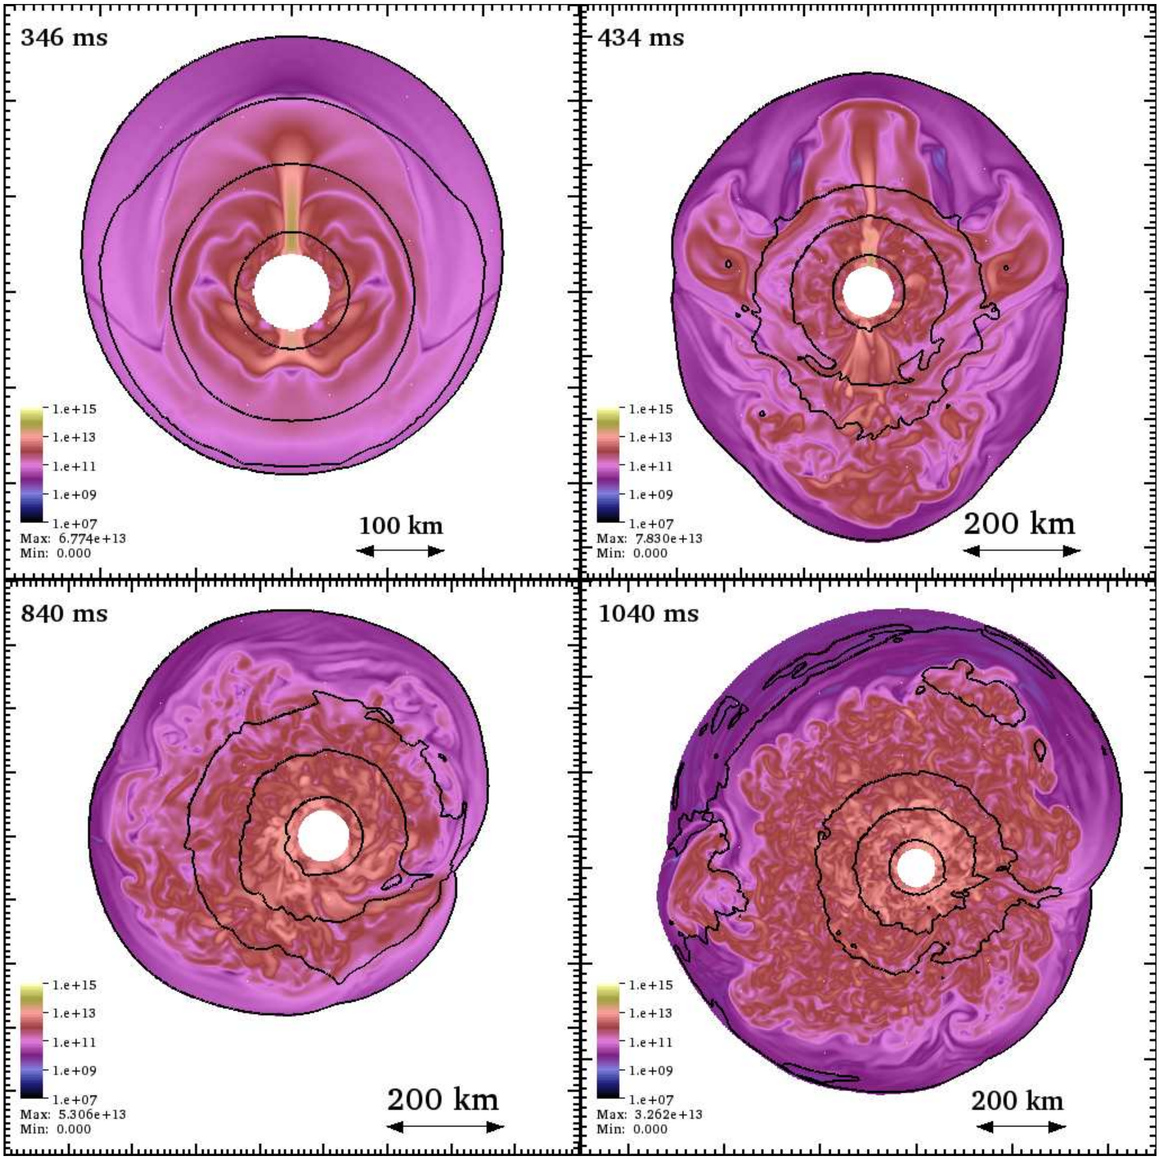

Figure 6.2: Snapshots showing the evolution and distribution of the magnitude of the magnetic field in a slice through SASI model `3DB12Am`. The times of the snapshots are indicated in the upper left corner of each panel. The color scale gives the magnetic field magnitude (in G). Note also that the sides of the top left panel are 100 km, while the other panels are 200 km. The top two panels show a slice of the xz-plane through the origin, while the two bottom panels show the evolution through a slicing plane whose normal vector is parallel to the total angular momentum vector of the flow between the PNS and the shock surface. Contour of constant density are plotted as black lines, starting with the innermost, they denote $\rho = 10^{10}$, 10^{9}, 3×10^{8}, and 6×10^{7} g cm^{-3}. The last contour is visible only on the lower right panel.

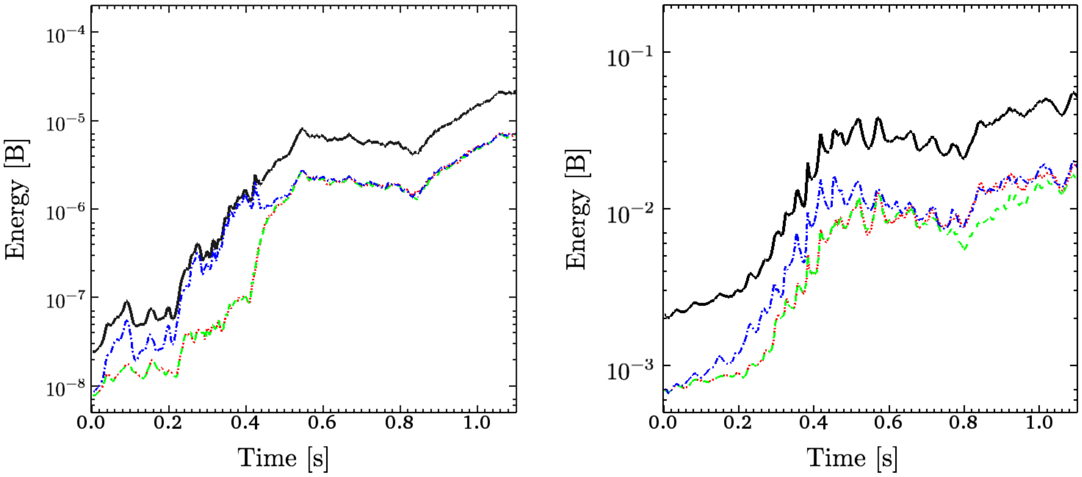

Figure 6.3: Evolution of magnetic and kinetic energies in a 3D model perturbed with the axisymmetric perturbation (model 3DB12Am). In the left panel we plot the total magnetic energy inside the shock E_{mag} (solid, thick black line), the individual components $E_{\mathrm{mag},x}$, $E_{\mathrm{mag},y}$, and $E_{\mathrm{mag},z}$ (dotted red, dashed green, and dash-dot blue lines, respectively). We have also plotted the total magnetic energy inside the shock for a 2D axisymmetric model (model 2DB12Am; solid grey line). In the right panel we plot the evolution of the kinetic energy inside the accretion shock: Total (solid black line), and the individual components $E_{\mathrm{kin},x}$ (dotted red line), $E_{\mathrm{kin},y}$ (dashed greed line), and $E_{\mathrm{kin},z}$ (dash-dot blue line).

the early time of the evolution, as shown in the upper left panel of this figure. The total magnetic energy and kinetic energy inside the shock are plotted over time in figure 6.3 on left and right panel, respectively. Several distinct stages can be identified in this simulation.

The evolution begins right after the perturbation has been applied to the initial condition. Until about $t = 200$ ms, there is an early oscillatory period with relic transients from the initial perturbation. The axisymmetric perturbation favors the $l = 1$ sloshing mode. In this early stage, the shock remains quasi-spherical, but its overall position shifts up and down relative to the PNS. This sloshing can be seen clearly in the upper left panel of figure 6.2. There is also a symmetry along the z-axis at this stage. On figure 6.3 this symmetry is shown by curves of the x- and y- components of both the magnetic and kinetic energy that fall on top of each other.

The SASI-induced flows result in magnetic field amplification around the axial symmetry. Between 200 to about 420 ms, there is a build up of the magnetic field around the z-axis, as can be seen from the structure of the magnetic field on the left upper panel of 6.2. The left panel of 6.3 shows that most of the magnetic energy is concentrated more and more in the z-component during this period. At 400 ms we find that $E_{\text{mag}} \approx E_{\text{mag},z}$ where $E_{\text{mag},z} = B_z^2/2$.

As the accretion shock sloshes up and down, the infalling materials hit the shock at an oblique angle. This introduces lateral velocities to the matter inside the shock. The build up of this leads to more vigorous non-axisymmetric flows. At $t = 420$ ms non-axisymmetric modes takes over and disrupt the magnetic field structure along the z-axis. On the top right panel of 6.2 a disrupted relic of this structure can be seen. The magnetic energy is distributed equally among the components at $t \geq 480$ ms with $E_{\text{mag},x} \approx E_{\text{mag},y} \approx E_{\text{mag},z} \approx E_{\text{mag}}/3$ (see also left panel of figure 6.3). The magnetic energy continues to grow as SASI develops non-linearly until about $t = 540$ ms.

This is followed by a period extending for about 300 ms where the magnetic energy between the PNS and shock surface shows a declining trend. This decline is caused by multiple factors: the disruption of the concentration of magnetic energy through the surface of the PNS, the shrinking of the shock volume V_{sh}, and a slowed magnetic energy generation rate while the post-shock flow rearranges. The magnetic energy density, $E_{\text{mag}}/V_{\text{sh}}$, continue to grow slowly during this phase. The total kinetic energy shows the same declining trend in this period. The kinetic energy exhibits a growing trend due to the sloshing mode up to about $t = 615$ ms, at which time it becomes less organized. At this point the components of the kinetic energy becomes roughly equal: $E_{\text{kin},x} \approx E_{\text{kin},y} \approx E_{\text{kin},z}$ where $E_{\text{kin},\chi} = \rho v_\chi^2/2$

After this “pausing” phase, a gradual increase of the the magnetic energy generation can be seen for $t > 830$ ms. The magnetic energy continues to grow until near the end of the simulation, where it appears to be leveling off at a value of about 2.5×10^{-5} B, almost 3 orders of magnitude higher than the initial value. Prior to this increase of magnetic energy, at around $t = 800$ ms a clear spiral mode pattern emerges in the fluid flow. Figure 6.4 illustrates this, where we plot polytropic constant (a proxy for the fluid entropy) at two instances separated by two full revolution about the PNS. The two panels in figure 6.4 correspond to the two lower panels of figure 6.2. The spiral mode of the SASI generates a

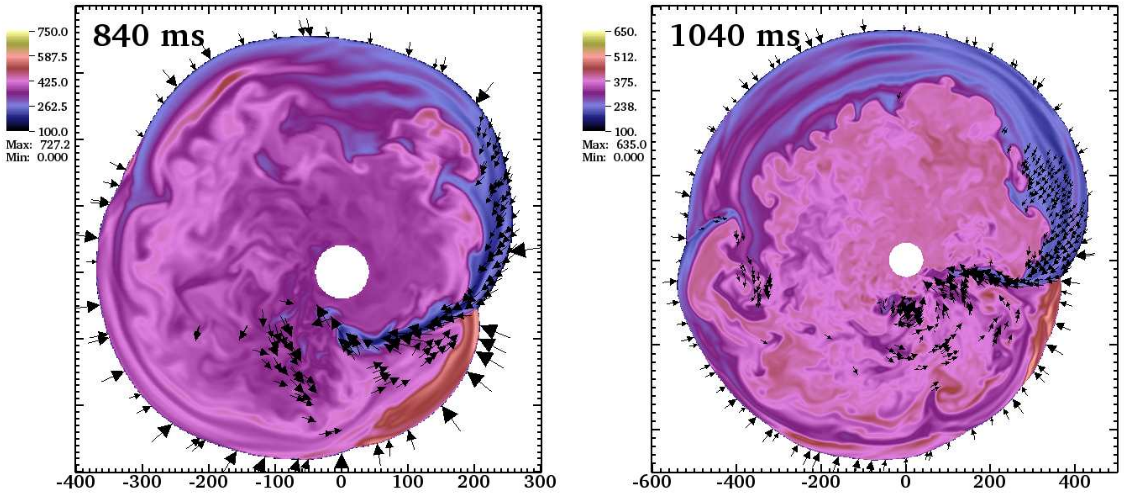

Figure 6.4: Snapshots showing the polytropic constant $(\kappa = p/\rho^{\Gamma})$ at selected times during the spiral mode of SASI. Velocity vectors where $|\mathbf{v}| \geq c_s = \sqrt{\Gamma p/\rho}$ are over overlaid on both plots. The two selected times of the plots correspond to the two lower panels of figure 6.2

.

significant amount of angular momentum about the PNS, shown in figure 6.5. At the end of the simulation, the angular momentum is 3.4×10^{47} g cm^{-2} s^{-1}, which is consistent with what reported by Blondin and Mezzacappa (2007).

Figure 6.4 shows the triple point where an internal shock inside the accretion is connected to the accretion shock (Blondin and Mezzacappa, 2007). Ahead of the triple point, a plunging supersonic stream penetrates down toward the PNS, as indicated by the black velocity vectors. This flow introduces a shear that may be susceptible to the Kevin-Helmholtz instability and other fluid instabilities associated with velocity shear. This often results in turbulent flows. The shearing region connected to the triple point generates fluid vorticity—a local measure of the rate of rotation of the fluid—which has been pointed out to be helpful for magnetic field generation (Mee and Brandenburg, 2006). In figure 6.6 we plot the fluid vorticity at selected times correspond to the two bottom panels of figure 6.2. A clear similarity can be seen for the distribution of the magnetic field and the fluid velocity at these times.

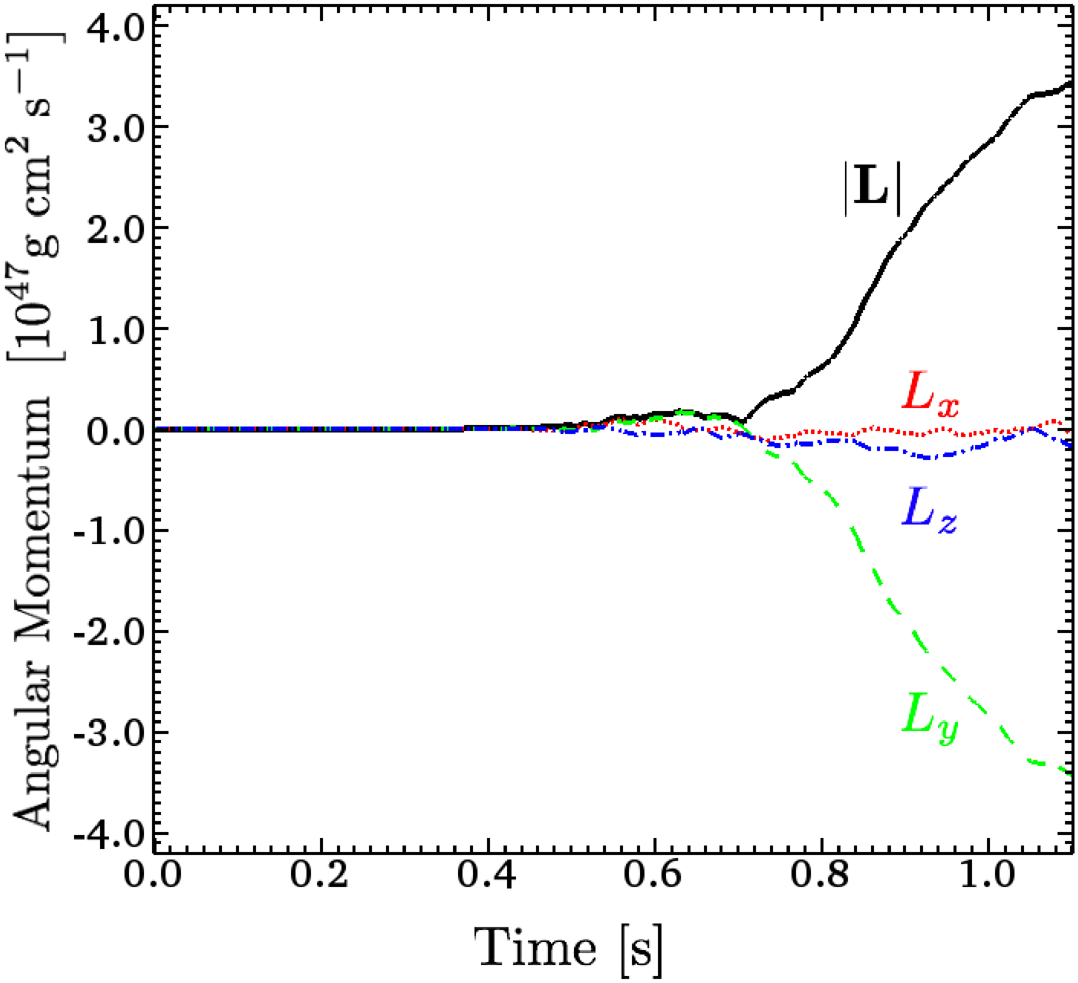

Figure 6.5: Angular momentum of the matter between the shock surface and the PNS for model `3DB12Am`. The total angular momentum $|\mathbf{L}| = \sqrt{L_x^2 + L_y^2 + L_z^2}$ is plotted as solid black line; the individual components L_x, L_y, L_z are plotted as dotted red line, dashed green line and dash-dot blue line, respectively.

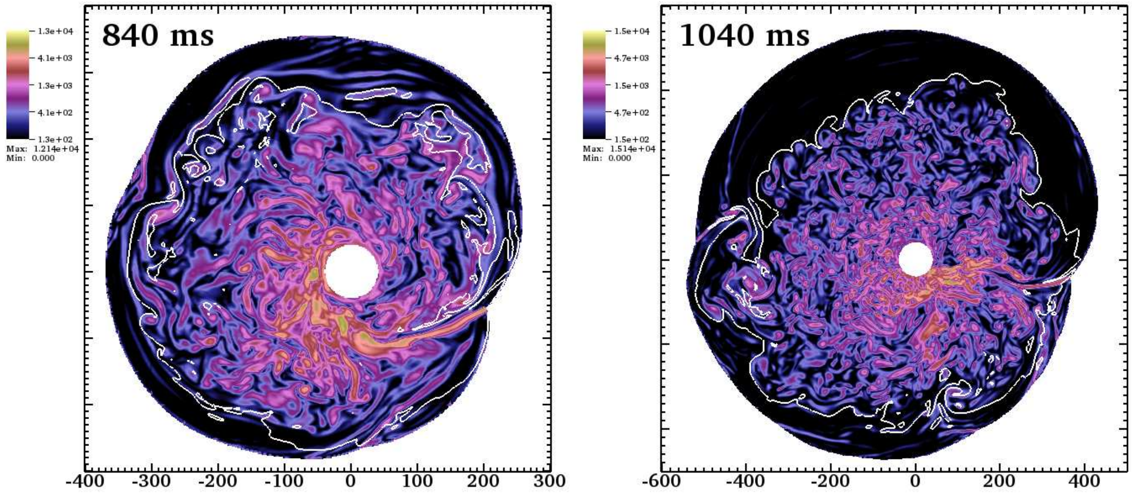

Figure 6.6: The distribution of fluid vorticity $|\omega| = |\nabla \times \mathbf{v}|$ at late times of SASI. The selected times for these panels correspond to the two bottom panels of figure 6.2. Vorticity is plotted in units of s^{-1}. The white contours are plotted where the magnitude of the magnetic field is 6×10^{10} G and 4×10^{10} G

At the end of the simulation, the magnetic field evolves into an intermittent "flux rope" structure. The bulk of the magnetic energy is stored in fields with strength around 10^{12} G, although there are extended region with magnetic field strength that exceeds 10^{13} G. Figure 6.7 shows the distribution of the magnetic field magnitude as function of radius. The magnetic energy, however, is still far below both the internal and kinetic energies of the post-shock flows. In some regions, the magnetic energy density reaches up to about 10% of the kinetic energy density, and 10% of the fluid pressure intermittently, but the magnetic energy does not seem to impact the dynamics of the SASI evolution in any significant way.

6.3.1.2 Model with Random Pressure Perturbation

The previous model `3DB12Am` with axisymmetric perturbation quickly developed the $l = 1$ sloshing mode due the nature of the perturbation. To complement this model and investigate any dependence of the perturbation method on the magnetic field generation, we have computed three different models in addition to the model `3DB12Am` at the same spatial resolution. The SASI in these models are initiated with random pressure perturbation.

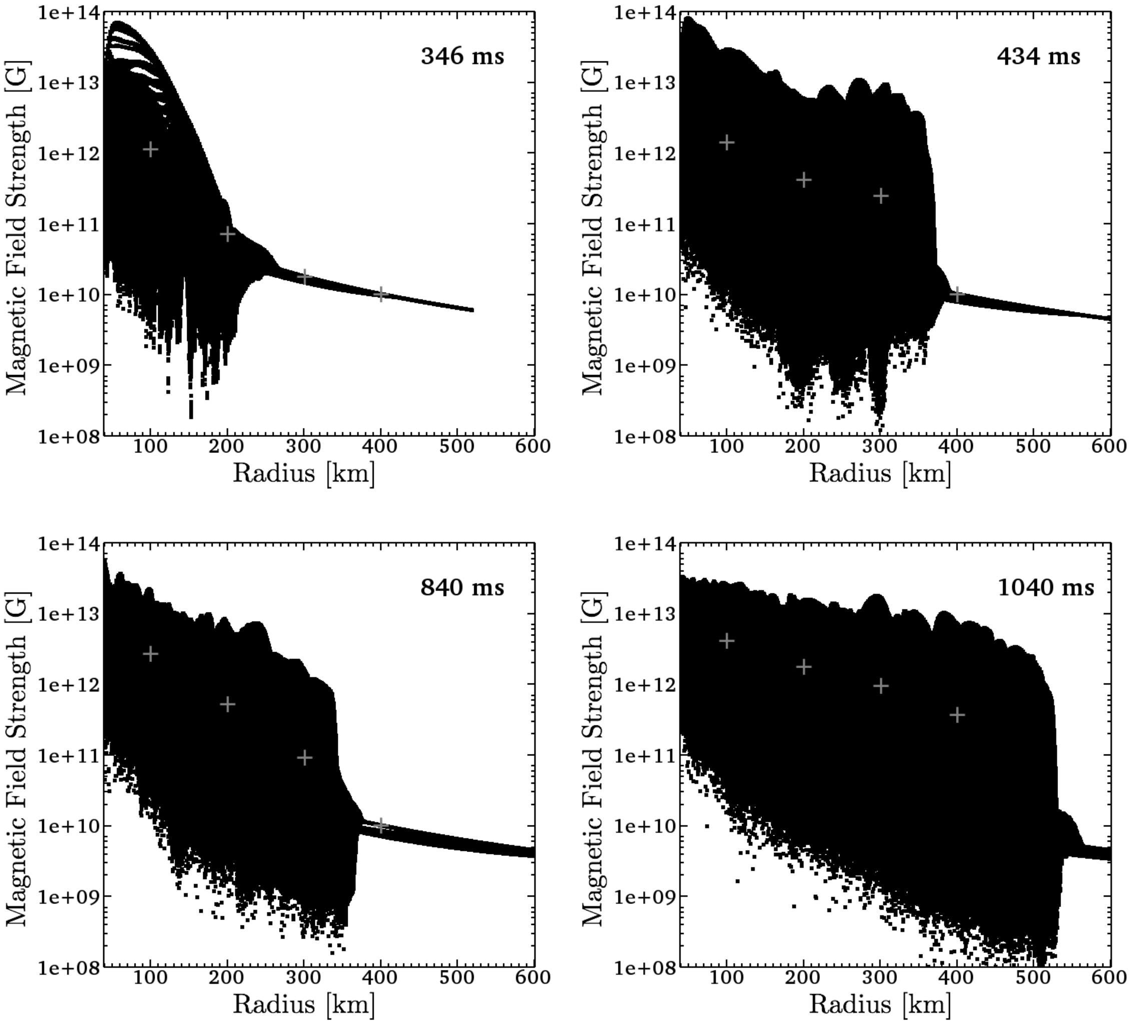

Figure 6.7: Scatter plot of the magnetic field magnitude as a function of the radius for model `3DB12Am`. The selected times correspond to the snapshots of figure 6.2. Plus signs denote the RMS value of the magnetic field in spherical shells bounded by $r^{\pm} = r \pm 25$ km, with $r = 100$, 200, 200, and 400 km.

They are one non-rotating model with initial magnetic field $B_0 = 10^{12}$ G (`3DB12Rm`), another non-rotating model with weaker initial magnetic field $B_0 = 10^{10}$ G (`3DB10Rm`), and a rotating model with initial magnetic field $B_0 = 10^{12}$ G and initial specific angular momentum $l = 1.5 \times 10^{15}$ cm^2 s^{-1} about the z-axis (`3DB12ΩRm`).

Figure 6.8 shows a brief overview of the results of these numerical experiments. In this figure, we plot the relative change in total magnetic energy between the PNS and the shock over time (upper left panel), total angular momentum between the PNS and the shock (upper right panel), kinetic energy of the flow between the PNS and the shock (lower left panel), and the average shock radius $\bar{R}_{\rm sh}$. The first quantities are integrated over the volume bounded by the inner boundary (the cutout) and the accretion shock surface. The average shock radius is defined as the radius of a sphere whose volume is equal to that encompassed by the shock. Models `3DB12Am`, `3DB12Rm`, `3DB12ΩRm`, and model `3DB10Rm` are represented by black, red, blue, and green lines, respectively.

Several features common to all models here can be identified. The magnetic energy shows a period of exponential increase followed by a period of less vigorous growth in all models. At the end of the simulations, the magnetic energy gets amplified by a factor of a few thousands. In all models, the spiral mode eventually dominates the SASI. This is consistent with the results of Blondin and Mezzacappa (2007). The magnetic energy growth appears to be responsive to the spiral mode of SASI. In all models the angular momentum reaches similar values on the order of 10^{47} g cm^2 s^{-1}, sufficient to impact the rotation rate of the PNS. The kinetic energy also reaches the same levels in all models.

A correlation can be seen of the initiation magnetic energy growth (upper left panel) with the onset of the nonlinear phase of SASI in all models (lower left panel). In both models `3DB12Rm` (red) and model `3DB12ΩRm` (blue), the trends of the magnetic energy growth continue for a while even after the kinetic energy starts to level off around 500 ms and 800 ms for those models. All the models with random perturbation shows both the sloshing and spiral mode early in the evolution, while the model with axisymmetric perturbation does not show the spiral mode until later. The model with initial rotation enters (`3DB12ΩRm`) enters the nonlinear state sooner, and the magnetic energy grows

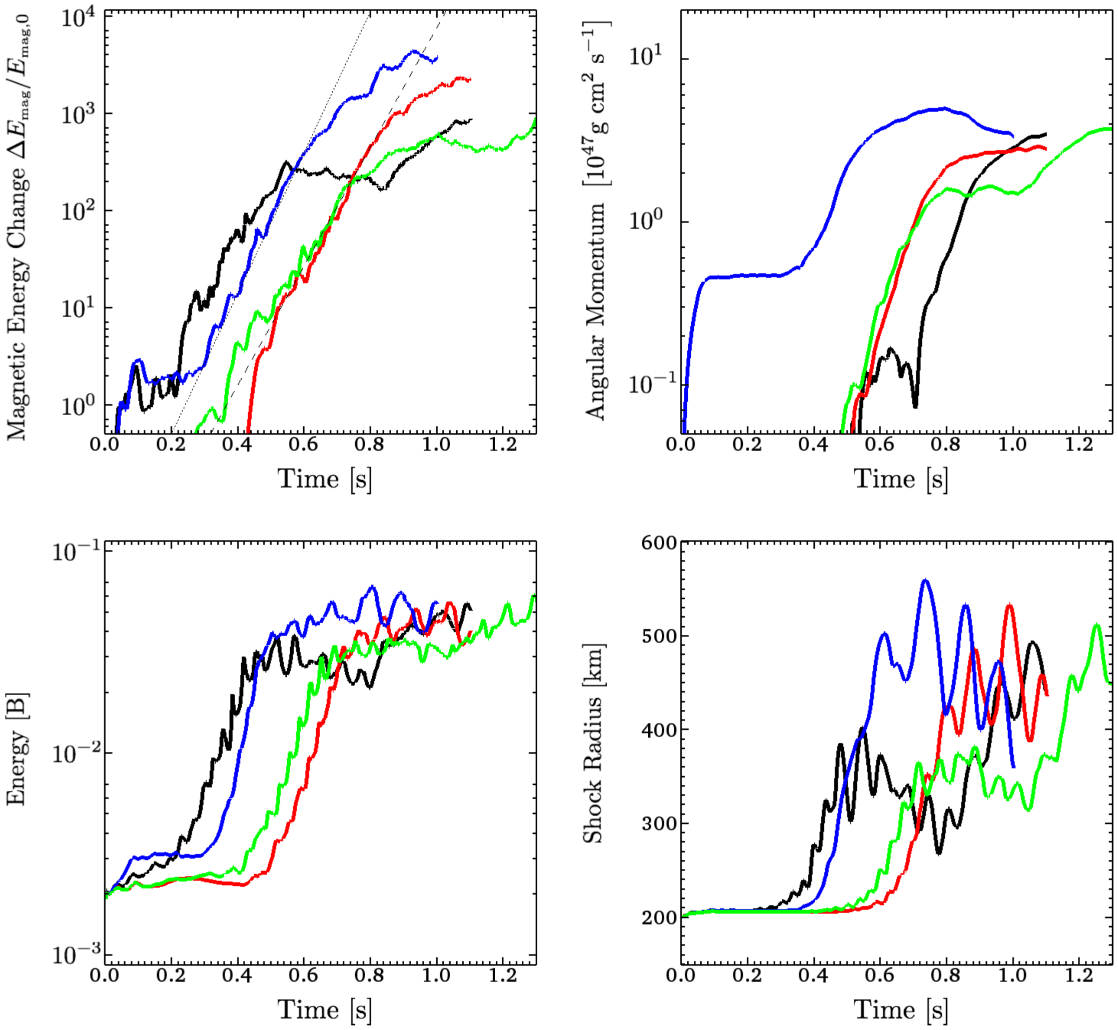

Figure 6.8: Overview of all SASI 3D models at 300/192 km spatial resolution. The plots shows shows the relative change in total magnetic energy between the PNS and the shock over time (upper left panel), total angular momentum between the PNS and the shock (upper right panel), kinetic energy of the flow between the PNS and the shock (lower left panel), and the average shock radius $\bar{R}_{\rm sh} = (3V_{\rm sh}/4\Pi)^{1/3}$ (lower right panel). For all panels, the results are plotted for the model `3DB12Am` (black), model `3DB12Rm` (red), model `3DB12ΩRm` (blue), and model `3DB10Rm` (green). The initial magnetic energy is 2.3×10^{-12} B for model with $B_0 = 10^{10}$ G and 2.3×10^{-8} B for model with $B_0 = 10^{12}$ G. On the upper left panel, reference lines for exponential growth with e-folding times of 71 ms and 60 ms are drawn as thin dashed and dotted blue lines, respectively.

faster than the other non-rotating models. It also develop more directly into the spiral mode pattern. During the simulation, the model with initial rotation seems to somewhat maintain its angular momentum close to the original axis of rotation (to within an angle of $\arctan(\sqrt{L_x^2 + L_y^2}/L_z \lesssim 0.2$ rad), while the non-rotating models have their angular momentum vectors change direction in a seemingly random fashion when they enter the SASI spiral mode.

From these comparisons, we can draw two initial conclusions. The fact that model with the weaker field strength (`3DB10Rm`) develops similarly to the other models seems to indicate that the magnetic field has little effect on the nonlinear evolution of the SASI. Secondly, the growth rate of the magnetic fields at the end of the simulations seem to indicate that the magnetic energy will not attain significantly higher levels on a timescale that is relevant for core-collapse supernovae.

6.3.2 Mechanisms for Magnetic Field Amplification

In previous subsection we have seen that SASI-induced flows are capable of amplifying the magnetic field inside the shock. Preliminary observation (see figure 6.8) suggests that the increase of the magnetic energy is at the expense of the kinetic energy. When the SASI is in the nonlinear stage, there is a significant kinetic energy of the flow available as the immediate source of the magnetic energy. In this subsection, we explain the mechanisms for magnetic field amplification.

Let us consider the evolution of the scalar magnetic energy density to help us pinpoint the source of the magnetic field amplification. From the magnetic induction equation (equation 6.4), it can be shown that the magnetic energy density evolves according to

$$\frac{\partial e_{\mathrm{mag}}}{\partial t} = \mathbf{B} \cdot \left[(\mathbf{B} \cdot \nabla)\mathbf{v} - (\mathbf{v} \cdot \nabla)\mathbf{B} + \mathbf{v}\nabla \cdot \mathbf{B} - \mathbf{B}\nabla \cdot \mathbf{v} - \nabla \times (\eta \mathbf{J}) \right], \tag{6.6}$$

where the first, second, and fourth terms on the right hand side are, conventionally, said to represent magnetic field evolution due to stretching, advection, and compression, respectively. The third term is the magnetic monopoles term which vanishes analytically but is retained here. A dissipative term is also added as the last term in equation 6.6, containing the scalar resistivity η. The dissipative term should only appear when the non-ideal electric field $-\mathbf{v} \times \mathbf{B} + \eta \mathbf{J}$ is used in the magnetic induction equation, where the current density is obtained from Ampère's law, $\mathbf{J} = \nabla \times \mathbf{B}$.

Although we are only concerned with ideal MHD here, the numerical method for solving the magnetic induction equation contains dissipative terms that manifest as numerical resistivity in the regions where magnetic field varies significantly over a few computational grid cells. Physically, magnetic energy dissipation is extremely small in the supernova environment. However, to achieve such realism is computationally prohibitive and not feasible in numerical simulations of the type presented here. Therefore it is important to consider the effect of numerical resistivity in our analysis.

The MHD Poynting theorem can also be used to rewrite equation 6.6 as

$$\frac{\partial e_{\mathrm{mag}}}{\partial t} + \nabla \cdot \left[\mathbf{P} + \eta \mathbf{J} \times \mathbf{B}\right] = -\mathbf{v} \cdot (\mathbf{J} \times \mathbf{B}) - \eta \mathbf{J} \cdot \mathbf{J}, \tag{6.7}$$

where the Poynting vector is $\mathbf{P} = (\mathbf{E} \times \mathbf{B}) = [\mathbf{v}(\mathbf{B} \cdot \mathbf{B}) - \mathbf{B}(\mathbf{B} \cdot \mathbf{v})]$. This equation shows that the magnetic energy density increases through work done against the Lorentz force (first term on the right-hand side), provided that it overcomes any losses due to Joule dissipation, i.e. resistivity (second term on the right-hand side). The Lorentz work term $W_{\mathrm{L}} = -\mathbf{v} \cdot (\mathbf{J} \times \mathbf{B})$ can be either positive or negative; the dissipative term $Q_{\mathrm{J}} = \eta \mathbf{J} \cdot \mathbf{J}$ can only decrease the magnetic energy.

To identify the mechanisms responsible for magnetic energy growth in our simulations, we compare the individual terms on the right-hand side of equation 6.6 through magnetic energy growth rates due to stretching, advection, and compression, defined respectively as

$$\sigma_{\nabla \mathbf{v}} = \frac{2\langle \mathbf{B} \cdot [(\mathbf{B} \cdot \nabla)\mathbf{v}]\rangle}{\langle \mathbf{B} \cdot \mathbf{B}\rangle}, \tag{6.8}$$

$$\sigma_{\mathbf{v} \cdot \nabla} = -\frac{2\langle \mathbf{B} \cdot [(\mathbf{v} \cdot \nabla)\mathbf{B}]\rangle}{\langle \mathbf{B} \cdot \mathbf{B}\rangle}, \tag{6.9}$$

$$\sigma_{\nabla \cdot \mathbf{v}} = -\frac{2\langle \mathbf{B} \cdot [\mathbf{B}\nabla \cdot \mathbf{v}]\rangle}{\langle \mathbf{B} \cdot \mathbf{B}\rangle}. \tag{6.10}$$

Angle brackets denote an average over the volume encompassed by the shock. The growth rate due to magnetic monopoles is defined as

$$\sigma_{\nabla \cdot \mathbf{B}} = -\frac{2\langle \mathbf{B} \cdot [\mathbf{v}\nabla \cdot \mathbf{B}]\rangle}{\langle \mathbf{B} \cdot \mathbf{B}\rangle}, \tag{6.11}$$

which we included in our analysis as a consistency check and show that it remains small in our simulations. For ideal MHD, the magnetic growth rate $\sigma_{e_{\mathrm{mag}}} = \langle e_{\mathrm{mag}}^{-1} \rangle \langle \partial e_{\mathrm{mag}} / \partial t \rangle$ equals the sum of the rates in equations 6.8 – 6.11. We can also define the rate due to the work done against the Lorentz force W_{L} in SASI-driven flows. This is effectively a conversion of kinetic energy into magnetic energy. This rate is given by

$$\sigma_{\mathbf{J} \times \mathbf{B}} = -\frac{2\langle \mathbf{v} \cdot (\mathbf{J} \times \mathbf{B})}{\langle \mathbf{B} \cdot \mathbf{B}\rangle}. \tag{6.12}$$

We plot the quantities from equations 6.8 – 6.12 for model `3DB12Am`, `3DB12Rm`, and `3DB12ΩRm` in figures 6.9 – 6.11.

In model `3DB12Am` the magnetic energy receives significant contributions from compression up to about $t = 430$ ms due to to the SASI sloshing mode that tends to converge the fluid flows at the temporary symmetry axis. For $t \geq 460$ ms, the stretching term dominates, with $\sigma_{\mathbf{J} \times \mathbf{B}} \approx \sigma_{\nabla \mathbf{v}}$. Contributions from compression are small in this phase. In all models the contribution from advection is mostly negative, and near zero at late times. This is expected since the role of advection is to drag the magnetic field into the PNS (inner boundary) with the fluid flows. The time-averaged rates for model `3DB12Am` over the interval 460 ms to 1.1 s (the nonlinear stage of SASI) are $\langle \sigma_{\nabla \mathbf{v}} \rangle \approx 435$ s^{-1}, $\langle \sigma_{\mathbf{v} \cdot \nabla} \rangle \approx -19$ s^{-1}, $\langle \sigma_{\nabla \cdot \mathbf{v}} \rangle \approx 112$ s^{-1},

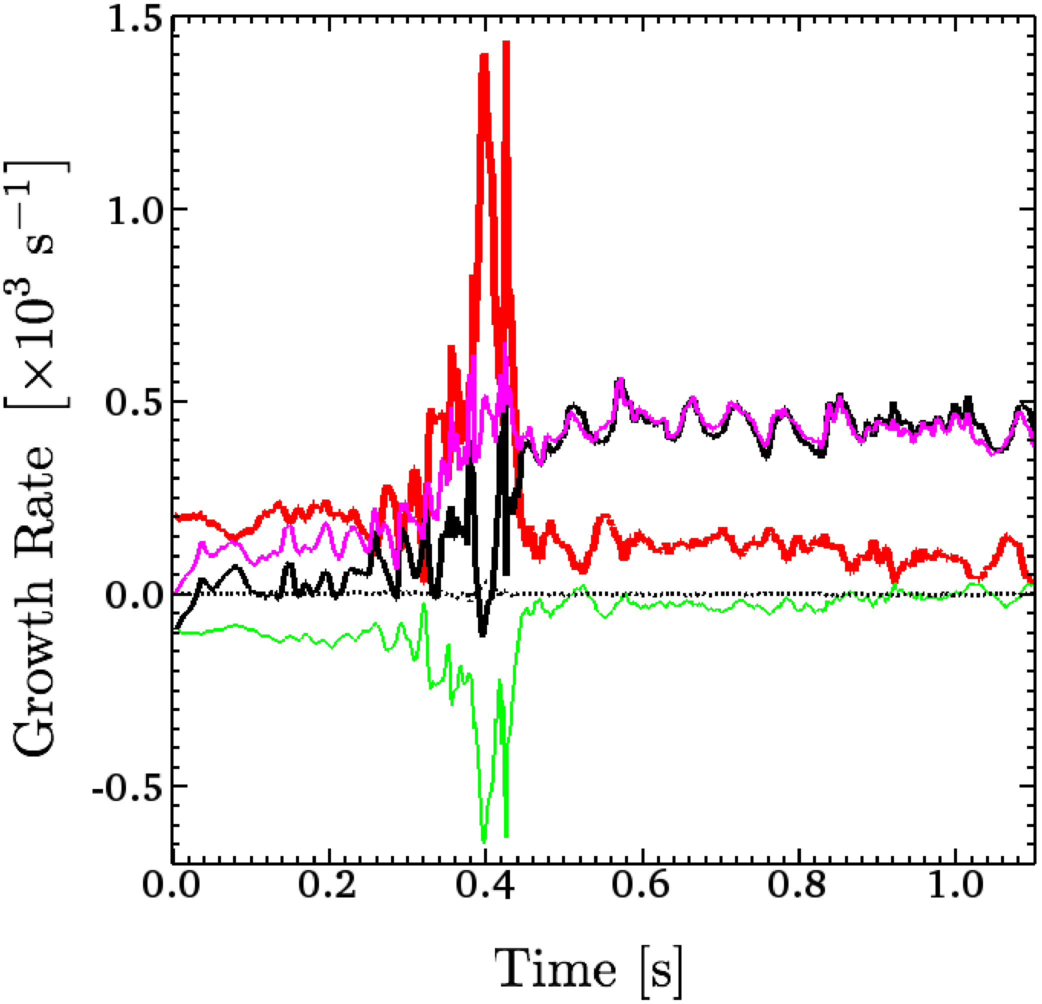

Figure 6.9: Magnetic energy growth rates for model `3DB12Am`. These quantities are plotted: stretching $\sigma_{\nabla \mathbf{v}}$ (black curve), advection $\sigma_{\mathbf{v}\cdot\nabla}$ (green curve), compression $\sigma_{\nabla\cdot\mathbf{v}}$ (red curve), and Lorentz work $\sigma_{\mathbf{J}\times\mathbf{B}}$ (magenta curve). Rate growth due to magnetic monopoles $\sigma_{\nabla\cdot\mathbf{B}}$ is also plotted as thin black dotted curve; it remains small throughout the simulation.

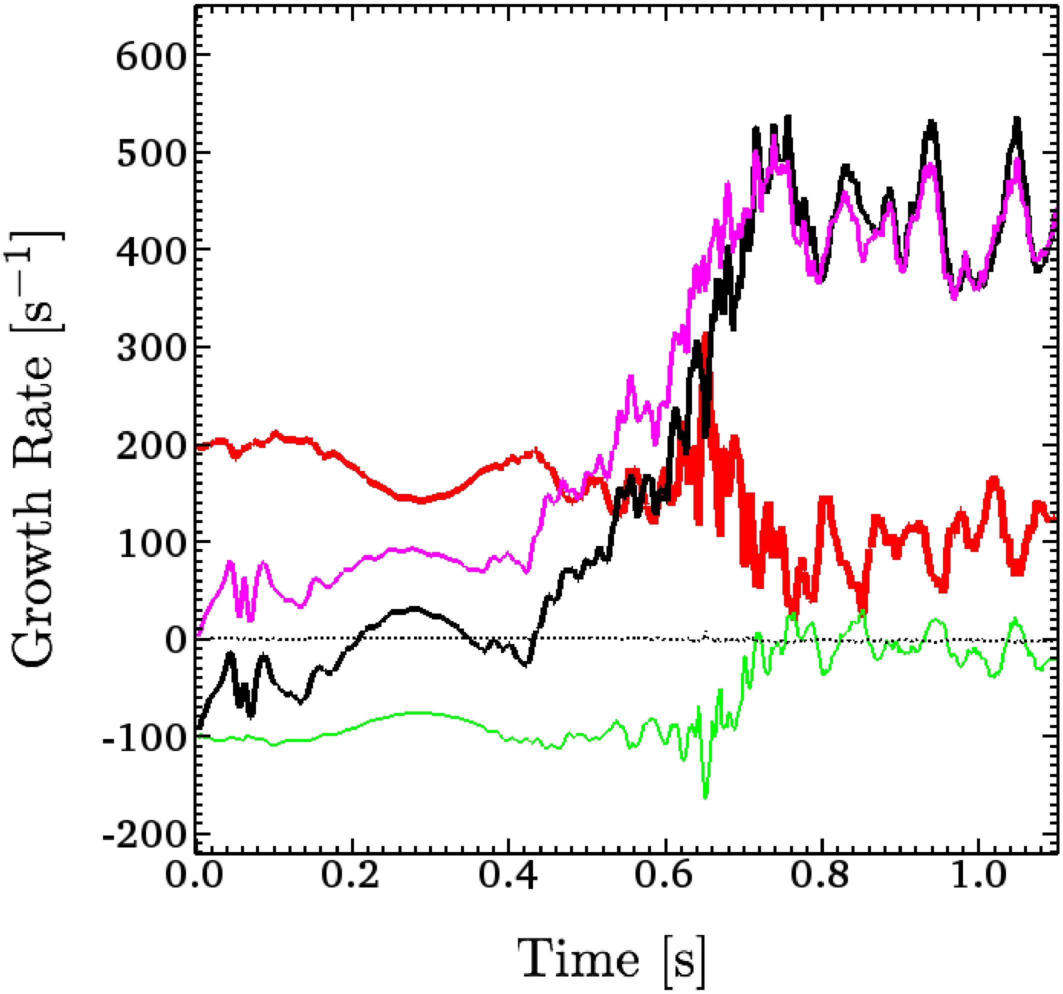

Figure 6.10: Magnetic energy growth rates for model `3DB12Rm`. The same quantities are plotted as in figure 6.9.

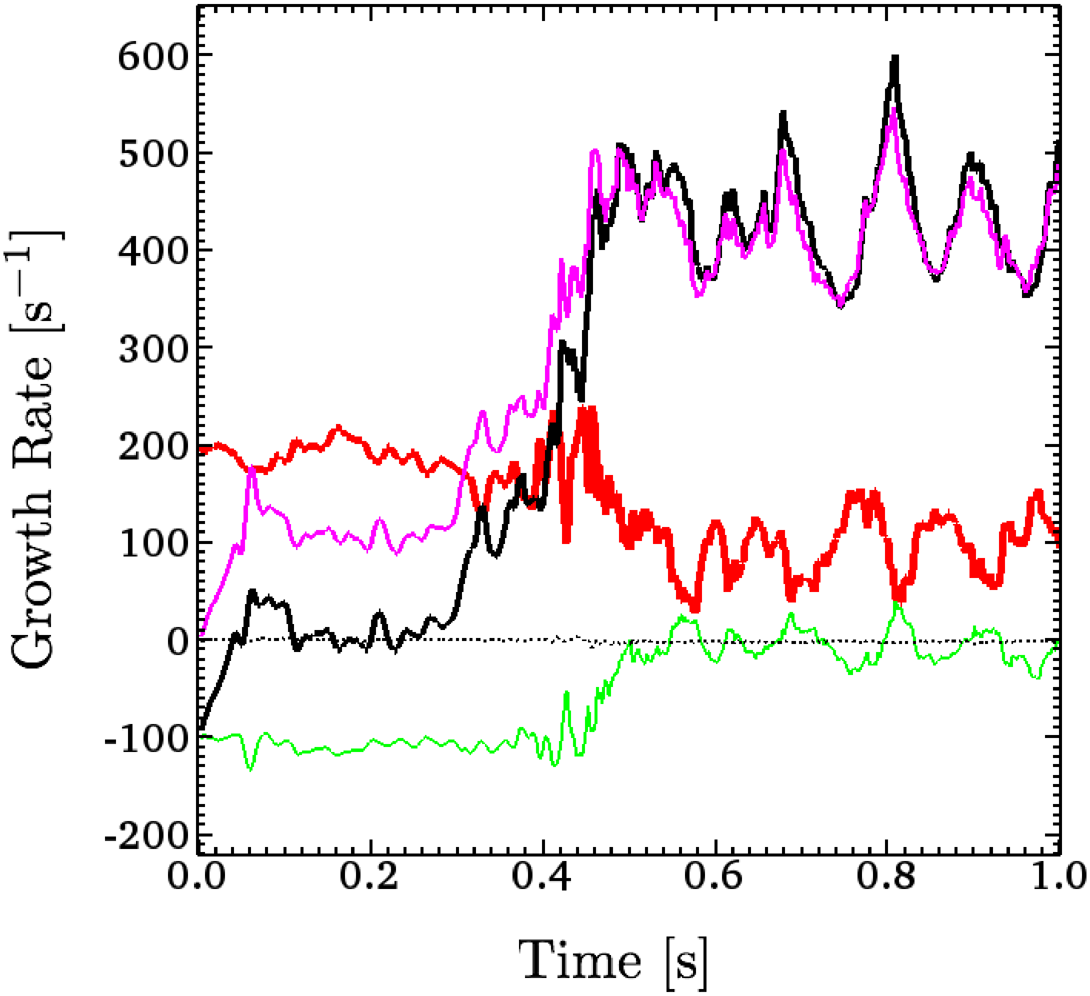

Figure 6.11: Magnetic energy growth rates for model `3DB1ΩRm`. The same quantities are plotted as in figure 6.9.

and $\langle \sigma_{\mathbf{J}\times\mathbf{B}} \rangle \approx 434$ s^{-1}. In all 3D models, the bulk of the magnetic energy is generated during the nonlinear stage covered by these intervals.

The two other models with random perturbation, model `3DB12Rm` and `3DB12ΩRm`, exhibit very similar behavior in the nonlinear regime. The rate at this stage are dominated by the stretching rate with $\sigma_{\nabla\mathbf{v}} \approx \sigma_{\mathbf{J}\times\mathbf{B}}$. They however do not show large spikes in the compression rate at earlier time that was seen for the compression rate of model `3DB12Am`. We postulate that these spikes are the result from the sloshing mode induced by the axisymmetric perturbation. For comparisons, we also list the computed time-average rates for the random perturbation models. In the interval from 700 ms to 1.1 s for model `3DB12Rm` we found $\langle \sigma_{\nabla\mathbf{v}} \rangle \approx 435$ s^{-1}, $\langle \sigma_{\mathbf{v}\cdot\nabla} \rangle \approx -9$ s^{-1}, $\langle \sigma_{\nabla\cdot\mathbf{v}} \rangle \approx 97$ s^{-1}, and $\langle \sigma_{\mathbf{J}\times\mathbf{B}} \rangle \approx 422$ s^{-1}. For model `3DB12ΩRm` the time averaged rates are $\langle \sigma_{\nabla\mathbf{v}} \rangle \approx 434$ s^{-1}, $\langle \sigma_{\mathbf{v}\cdot\nabla} \rangle \approx -6$ s^{-1}, $\langle \sigma_{\nabla\cdot\mathbf{v}} \rangle \approx 97$ s^{-1}, and $\langle \sigma_{\mathbf{J}\times\mathbf{B}} \rangle \approx 420$ s^{-1} over the interval from 500 ms to 1 s.

In all of our models, we see that in the nonlinear stage the magnetic energy generation is a result of the net work done against the Lorentz force by the fluid motion form the SASI-induced flows. This is essentially a conversion of the kinetic energy to the magnetic energy. The most dominant mechanism of the magnetic energy growth rate is via the stretching of the magnetic field, as we find that $\sigma_{\mathbf{J}\times\mathbf{B}} \approx \sigma_{\nabla\mathbf{v}}$ in all models. The SASI spiral mode gives a persistent shear flow inside the accretion shock, which generates fluid vorticity and triggers secondary fluid instabilities in a turbulent flow. In a turbulent flow, separation of two fluid elements grows exponentially with time. If these fluid elements were initially connected by a weak magnetic field, the growing separation results in the stretching of the field since in the case of ideal MHD the magnetic field are frozen in the fluid. The stretching decreases the cross-sectional area and thus, due to flux conservation, strengthens the magnetic field and increases the magnetic energy (Ott, 1998).

In §6.3.1 we have seen that the magnetic field evolves into a complicated flux rope structure in the nonlinear stage of the SASI. The total magnetic energy levels off at a level that is still well below the kinetic energy of the flows beneath the shock. However, here we see that at the end of the simulations the magnetic energy growth rate $\sigma_{\mathbf{J}\times\mathbf{B}}$ remains constant. These therefore suggest that the magnetic energy growth does not stop because of the dynamical equipartition of the magnetic field with the fluid on any spatial scale. On the other hand, the bulk of the magnetic energy is concentrated on spatial scales where numerical resistivity inevitably plays a role. To investigate this, we proceed with varying the spatial resolution in our simulations.

6.3.3 Effects of Spatial Resolution Variation

Our analysis from §6.3.2 suggests that the spatial resolution eventually plays a role in stopping the growth of the magnetic field in SASI. To further investigate this, we present results where we vary the spatial resolution of our model with initial field $B_0 = 10^{12}$ G and axisymmetric perturbation. These are models `3DB12Al`, `3DB12Am`, and `3DB12Ah` (see also table 6.1).

The three models evolve in qualitatively similar manner from the hydrodynamics point of view. The initial perturbation drives the $l = 1$ sloshing mode, leading to a temporary build up of magnetic field along the symmetry axis and its partial dissruption. In all three models, non-axisymmetric modes appear around $t = 400$ ms, followed by a rearrangement of the flow that leads to a prominent spiral mode and a post-shock flow with significant angular momentum about the PNS. The integrated kinetic energy and total angular momentum between the PNS and the shock evolve similarly with time in all three models. At the end of the simulation, they reached values of about $E_{\rm kin} \approx (3-4) \times 10^{-2}$ B and $|\mathbf{L}| \approx (3-4) \times 10^{47}$ cm^2 s^{-1}, respectively. However, the same thing cannot be said for the magnetic field evolution.

The evolution of the magnetic field of the three models is very sensitive to the spatial resolution. Figure 6.12 illustrates the sensitivity of the magnetic field to the spatial resolution. On the two upper panels the distribution of the magnetic field magnitude at the end of the simulation are shown for models with the lowest (`3DB12Al`) and highest (`3DB12Ah`) resolutions at the end of the simulations. The upper right panel, showing model `3DB12Ah`, reveals a highly intermittent magnetic field in larger region that are much stronger than model `3DB12Al` showed on the upper left panel.

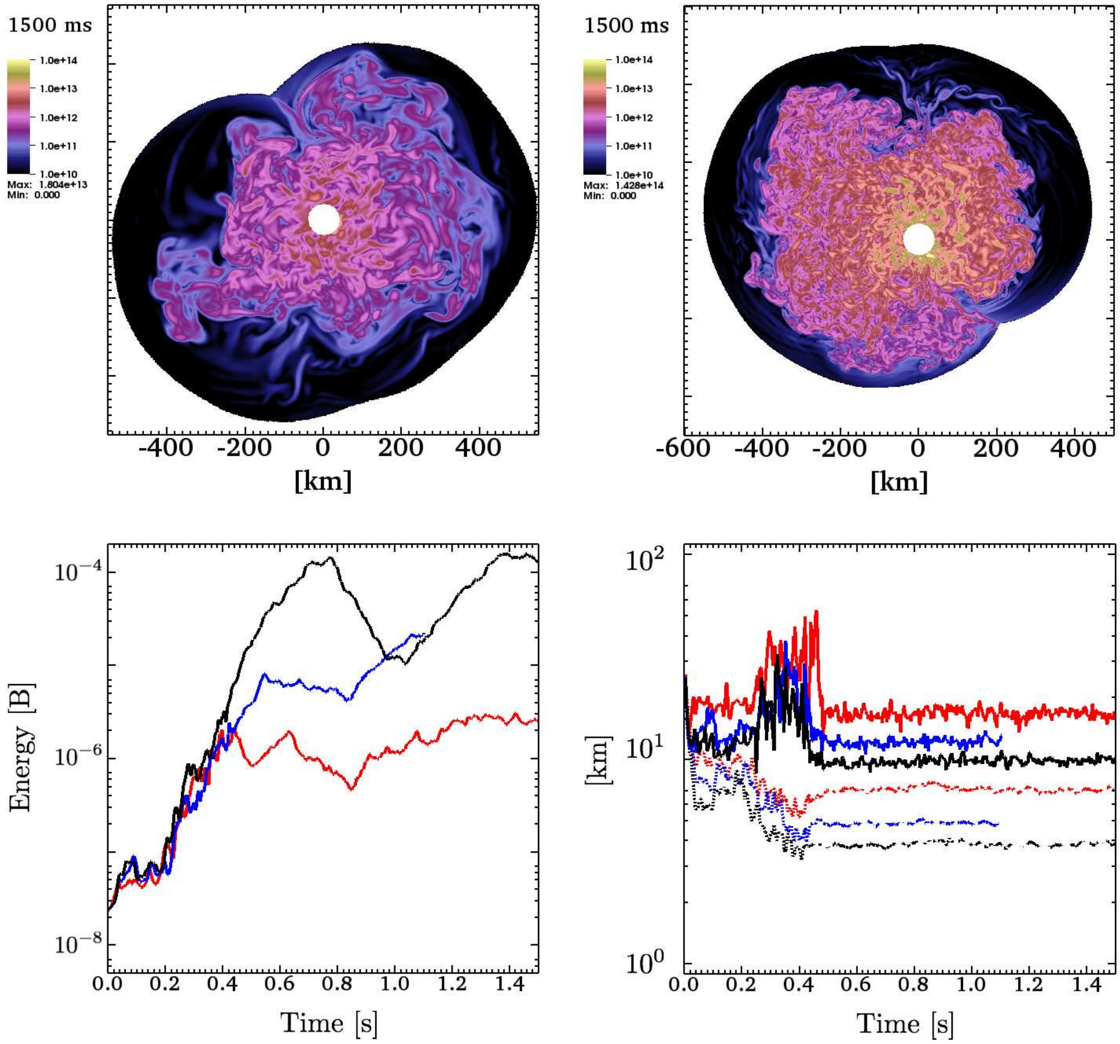

Figure 6.12: Results from SASI model with axisymmetric perturbation where spatial resolution has been varied. The two upper panels show the magnitude of the magnetic field at $t = 1500$ ms for model with $\Delta l \approx 2.34$ km (`3DB12Al`, upper left panel) and $\Delta l \approx 1.17$ km (`3DB12Ah`, upper right panel). The orientation of the plots is such that the normal of the slicing plane is parallel to the total angular momentum of the flow between the PNS and the shock surface. The two lower panels show results from models with the cell width $\Delta l = 2.34$ km, 1.56 km, and 1.17 km as red, blue, and black lines, respectively. The lower left panel shows the total magnetic energy between the shock and the PNS. The lower right panel shows the magnetic curvature radius $\lambda_{\rm c}$ (upper three solid lines) and the magnetic rms scale $\lambda_{\rm rms}$ (lower three dotted lines).

The lower left panel of figure 6.12 shows the total magnetic energy integrated over the volume between the surface of the PNS and shock. The plot shows that the magnetic energy is initially insensitive to resolution for $t \leq 400$ ms, when the axisymmetric phase dominates. When the flow becomes more turbulent ($t > 400$ ms, the evolution of the magnetic energy begin to diverge. At the end of the simulations, the magnetic energy of the highest resolution is up to 2 orders of magnitude higher than in the lowest resolution model, with only a factor of 2 difference in spatial resolution. In all models, stretching dominates the magnetic field amplification mechanism, with $\sigma_{\nabla \mathbf{v}} \approx \sigma_{\mathbf{J}\times\mathbf{B}}$ for $t \geq 500$ ms. The stretching term remains at a nearly constant level at late times in all three models. It increases somewhat with increasing spatial resolution (about 65% from `3DB12Al` to `3DB12Ah`).

To help us quantify the structure of the magnetic field, we plot also two useful characteristic scales in figure 6.12, the magnetic curvature radius,

$$\lambda_{\mathrm{c}} = \sqrt{\frac{\langle B^4 \rangle}{\langle |(\mathbf{B}\cdot\nabla)\mathbf{B}|^2 \rangle}}, \tag{6.13}$$

and the so-called magnetic rms scale,

$$\lambda_{\mathrm{rms}} = \sqrt{\frac{\langle B^2 \rangle}{\langle |\nabla B|^2 \rangle}}. \tag{6.14}$$

The magnetic curvature radius measure how sharply the magnetic field is bent (Ryu et al., 2000), while the magnetic rms scale provides a measure of the thickness of magnetic flux tubes when the magnetic field has evolved into a highly intermittent flux tube structure (Brandenburg and Subramanian, 2005).

The magnetic curvature radius and the rms scale reach their optimum values during the axisymmetric sloshing mode, peaking at around $t \approx 420$ ms. In this state, they respectively measure the length and the thickness of the flux tube. At later time ($t \geq 500$ ms), when the non-axisymmetric spiral mode dominates, they settle to a nearly constant level. Both scales decrease with increasing resolution, with the ratio $\lambda_{\mathrm{rms}}/\Delta l$ nearly unchanged in all three models, and the ratio $\lambda_{\rm c}/\Delta l$ increases slightly with increasing resolution (about 20% from `3DB12Al` to `3DB12Ah`).

From these simulations, it is apparent that the magnetic field amplification is sensitive to the spatial resolution. As spatial resolution increases, smaller scales of the turbulent flow due to the SASI become available for the magnetic field to develop via the stretching mechanism that serves to amplify the magnetic field.

In the highest resolution of our simulation, we find that the growth rate due to the stretching and the Lorentz work remain constant and larger than zero up to the end of the simulation. The magnetic energy grows at the expense of the kinetic energy in the flow, and it can be expected to grow until a dynamical equilibrium is established. Our simulations suggest that this has not been reached. Figure 6.13 illustrates this, where we plot the distribution of cells where the ratio of magnetic-to-kinetic energy $\beta_{\rm kin}^{-1} = e_{\rm mag}/e_{\rm kin} = v_{\rm A}/|\mathbf{v}|^2$ is greater than or equal to 10^{-3} (red curve), 10^{-2} (green curve), 10^{-1} (blue curve), and 1 (magenta curve) for model `3DB12Ah`, with the Alvèn speed $v_{\rm A} = B/\sqrt{\rho}$. We find that at the end of the simulation the bulk of the magnetic energy (97%) is stored in cells where $\beta_{\rm kin}^{-1} \geq 10^{-3}$, with progressive smaller fraction is in cells where $\beta_{\rm kin}^{-1} \geq 10^{-2}$ and $\beta_{\rm kin}^{-1} \geq 10^{-1}$. Only about 2% of the total magnetic energy is in cells where $\beta_{\rm kin}^{-1} \geq 10^{-1}$. This shows that dynamical equilibrium is not the limiting factor for the growth of magnetic energy. The bulk of the magnetic energy becomes concentrated on the smallest available spatial scale (as determined by spatial resolution), and amplification ceases because thinner flux tubes due to stretching cannot be resolved. Numerical diffusion is therefore, in effect, limiting the further growth of the magnetic energy. Therefore the growth of magnetic energy has not converged in the spatial range covered by our simulations.

6.4 Conclusions and Future Work

We have presented three-dimensional simulations of idealized MHD models of a stalled supernova shock. Our simulations demonstrate that SASI-driven flows are able to significantly amplify the magnetic field beneath the the shock surface. In our highest

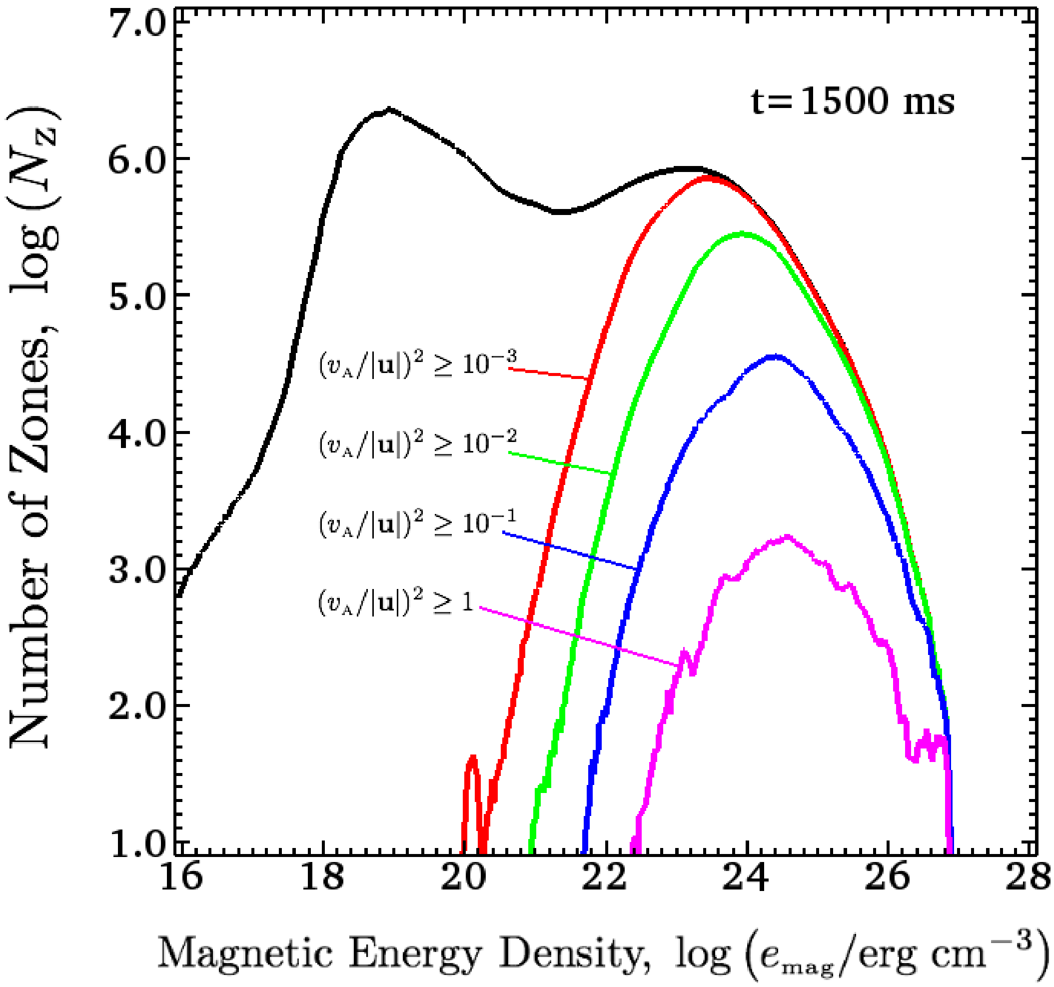

Figure 6.13: Distribution of magnetic energy in zones beneath the shock for model `3DB12Ah`. Plotted are curves representing subset of zones where the ratio of magnetic-to-kinetic energy $e_{\mathrm{mag}}/e_{\mathrm{kin}}$ is greater than or equal to 10^{-3} (red), 10^{-2} (green), 10^{-1} (blue), and 1 (magenta).

resolution model, this amplification results to almost 4 orders of magnitude increase in magnetic energy. The dominant mechanism of this amplification is the stretching of the field that develops into highly intermittent "flux rope" structure in the final stage of the simulation. The stretching of the magnetic field is facilitated by turbulent flows driven by the SASI spiral mode. The increase of the magnetic energy occurs at the expense of the kinetic energy of the flow beneath the shock.

The final magnetic energy and field strength remain sensitive to the spatial resolution in our numerical simulations. As the magnetic field evolves into the flux rope structure, the average thickness of the flux rope is limited by the spatial resolution—being resolved by only a few computational cells—and numerical resistivity eventually limits the growth of the magnetic field strength. This effect is, of course, non-physical in our simulations since the fluid between the PNS and the shock is expected to behave as a nearly perfect electrical conductor. Although we are unable to make exact predictions about the magnetic field in a presumed saturated state, we can expect that it may eventually become dynamically significant on relatively small spatial scales where the drag of the fluid on the flux ropes is balanced by the tension of the flux ropes (Thompson and Duncan, 1993).

In our models the kinetic energy is the energy reservoir for the amplification of the magnetic energy. The SASI-driven flows increase the magnetic energy from the work done against the Lorentz force. Therefore given infinite spatial resolution, the magnetic energy in our models is still limited by the kinetic energy of SASI-driven post-shock flows, which is a few of times 10^{-2} B. This leads us to believe that a SASI-generated magnetic field from non-rotating or slowly rotating progenitor is not sufficient to become important for global dynamics of core-collapse supernovae. Larger energy reservoir such as relic angular momentum from progenitor rotation appears to be needed for MHD effects to have significant role in driving the supernova explosion.

Past studies have concluded that rapid rotation is a necessary ingredient for magnetic field amplification during collapse and post-bounce. Our models show that SASI-driven flow alone can amplify the magnetic field in the post-bounce environment, in contrast to conclusions reached in the past core-collapse supernova simulations that included MHD (for example by Leblanc and Wilson (1970) and Symbalisty (1984)). Although this by itself does not impact the global dynamics of core-collapse supernovae significantly, an interesting question for further work is whether, combined with moderate rotation of the progenitor stars, the magnetic field may be sufficiently amplified to impact the pre-explosion dynamics as rotational energy becomes available to be harnessed. In other words, SASI may be able to extend the range of progenitors (rotation rates) for which magnetic fields play a role in their explosion dynamics.

In the simulations we carried out here we excised the PNS from the computational domain. This allows us to study the important magnetohydrodynamics aspects of the stalled supernova shock without the penalty of small time steps due to the high density (hence, high characteristic speed) in the PNS. However, excision of the PNS prevents us from fully exploring an important result from Blondin and Mezzacappa (2007), who concluded that the angular momentum accreted through the inner boundary would spin up the PNS to observed typical pulsar spin rates, by direct observation in simulations. Therefore in future work we plan to include the PNS in similar simulations.

The inclusion of the PNS requires that gravitational potential to be computed numerically and integrated into the solution of the MHD equations. This will be accomplished using the solver for Poisson's equation we have developed and described in chapter 3.

By including the PNS in future simulations, we expect to also address questions related to magnetic field dynamics. The simulations presented here demonstrated that SASI-induced flows can result in significant magnetic field amplification. An implication for this is that the PNS may attain strong magnetic fields due to the settling of magnetized material onto PNS. By including the PNS into future simulations, we will be able to investigate the degree of PNS magnetization that we can expect from SASI-induced magnetic field amplification. Since neutron star magnetic field is an observable, knowledge of its creation and structure may provide important insight into the dynamics of core-collapse supernovae. Additionally, the predicted spin up of the PNS due to the spiral mode of the SASI may result in favorable condition for further magnetic field amplification in the PNS. These questions will be explored in future simulations that include PNS.

Chapter 7

Conclusions and Outlook

GenASiS is a next-generation astrophysics simulation system that is currently under development. Its purpose is to explore and investigate astrophysical phenomena via large-scale computer simulations, with the primary emphasis on the simulations of core-collapse supernovae and neutron star mergers. In this manuscript we have provided a detailed description of the current version of GenASiS with emphasis on those aspects of GenASiS that the author has had a significant role in developing. We have shown the equations solved by the code and their numerical implementation. We have also shown some early scientific results that were accomplished by simulations with GenASiS.

In chapter 2 we have shown the implementation of Newtonian and special relativistic MHD in GenASiS. Test problems were performed to validate our implementation of the numerical scheme. We also described the parallel implementation, and showed its performance and weak scaling. Weak-scalability is critical for practical execution of large simulations such as the ones described in chapter 5 and 6.

In developing the MHD module for GenASiS, Dr. Eirik Endeve led the work on the Newtonian implementation, while I was responsible for the special relativistic version. I was also responsible for implementing and validating the code with the test problems and their documentation as described in §2.4. To help in guarding the code against unintended introduction of software bugs and regression of functionality as code development proceeds, I also developed an automated regression testing suite, utilizing test problems and validations that have been done previously.

So far we have only implemented special relativistic MHD, but general relativistic MHD needs to be included in GenASiS. Several HLL-type central schemes for general relativistic MHD (with constrained transport for the evolution of magnetic fields) exist and have been demonstrated to yield accurate results for relativistic test problems (for examples, Gammie et al. (2003); Del Zanna et al. (2007); Noble et al. (2006); see also Font (2008) for a review of relativistic MHD schemes). As discussed in §1.2, the code modularity in GenASiS architecture will allow this new extension to be added without replacing or invalidating the currently implementated Newtonian and special relativistic MHD. The addition of general relativistic MHD to GenASiS will be the subject of future work.

We described an FFT-based parallel solver for Poisson's equation in chapter 3. I was primarily responsible for its development. Coupled with the magnetohydrodyanmics equations, the solution to Poisson's equation describes a Newtonian self-gravitating fluid. This solver is general enough that it can be used to solve any Poisson equation with isolated boundary condition. We demonstrated the weak-scaling of the solver. As with MHD, a scalable solver is essential for any large-scale simulations with a large number of processes. We have used this solver in the simulations described in chapter 5. We will use this solver for any future work involving Newtonian self-gravity. In particular, we will use it for future SASI simulations that include the proto-neutron star, as described in §6.4. As a contribution to the community, this solver will be released and made freely available as a software library (Budiardja et al., 2010). The ease with which we are able to decouple this module from the rest of GenASiS is also a testament to the modularity of GenASiS architecture.

So far we have only considered Newtonian gravity in the code (with back-reaction potential and energy source terms added to include the effects of gravitational wave emissions, for example as in §5.2, where needed). However, considering the compaction of neutron stars, $GM/Rc^2 \approx 0.1 - 0.2$, we expect that general relativistic effects are important to their evolution. General relativity has also been a standard consideration since the early days of supernova simulations with spherical symmetry (for example, May and White (1966); Colgate and White (1966)). Therefore the inclusion of general relativity in GenASiS is essential. Toward that goal, we are currently implementing the BSSNOK formulation for numerical relativity (Alcubierre, 2008; Nakamura et al., 1987; Shibata, 1995; Baumgarte, 1998) in GenASiS. Coupled with general relativistic MHD, this will make GenASiS a fully consistent relativistic numerical code.

Although the AMR feature in GenASiS is still under development, we have shown its current functionality and described the necessary scheme for a fully-working AMR in chapter 4. Dr. Christian Cardall and I have spent considerable efforts in developing the necessary infrastructure for AMR in GenASiS. I have also developed the solver for Poisson's equation in mesh with refinements, as described in §4.3. Further work for AMR is still necessary to make this a mature feature of GenASiS, as explained in §4.4.

In chapter 5, we have shown results of neutron star merger simulations with GenASiS. We have also developed the necessary tools in GenASiS for the extraction of gravitational wave signatures for such simulations. This particular project was primarily my work. Theoretical templates for simulations such as these are necessary for matched filtering procedure to extract signals from gravitational wave detectors data sets. Beyond these simulations, future work would include more physical realism by using nuclear equations of state already implemented in GenASiS. Magnetic fields should also be included in future simulations, and their amplification as a viable central engine for GRB should be investigated. Further down the road, more accurate microphysics for the modeling of neutron star mergers needs to include better approximations, if not the full solutions, to the neutrino transport problem. This would shed light on the mechanisms of short-hard GRB.

We investigated the amplification of magnetic fields due to SASI in chapter 6. These results, which we have subsequently published in Endeve et al. (2010), showed that GenASiS is already capable of performing large-scale scientific simulations, although the code is still under development. I contributed to this accomplishment through the building of the code, its test problems and verifications, and its input/output module that enables such simulations. I also contributed to the analysis of the simulation data. Future work, building on our experience from these simulations, will include the proto-neutron star with the SASI, as was explained in §6.4.

Core-collapse supernovae are largely neutrino events, since the bulk of gravitational binding energy is released as neutrino bursts on the energy scale of 10^{53} ergs. As such, any supernova simulation code should have a good treatment of the radiation (neutrino) transport problem. Since neutrinos interact weakly with matter, neutrino detection provides an excellent opportunity to probe the physics in the collapsed core of a supernova explosion. This requires a theoretical understanding of the neutrinos evolution during the event such as their trajectories, spectra, luminosity, and oscillations. The investigation of neutrino-heating mechanism to revive the stalled shock in the post-bounce evolution also requires some treatment of the radiation transport problem.

Neutrinos may also play crucial role in the mergers of neutron stars, and presumably in the generations of gamma-ray bursts. Although in this manuscript we have not discussed the treatment of radiation transport in GenASiS, from its conception GenASiS was designed to eventually solve six-dimensional (three space and three momentum space dimensions) radiation hydrodynamics (Cardall et al., 2006). The current version of GenASiS does not have the radiation transport integrated into the code yet, although some work has been done in multidimensional formulation of conservative relativistic radiative transfer (Cardall and Mezzacappa, 2003; Cardall et al., 2005). Early studies with GenASiS have used a simplified treatment of the neutrino transport (Endeve et al., 2007). The inclusion of multidimensional radiation transport in GenASiS presents major challenges (see Cardall et al. (2006) for a review) and therefore will be a major accomplishment in the future.

GenASiS has made a promising start as a next-generation astrophysics simulation system. We have built the foundation for the inclusion of all physics relevant to simulations of core-collapse supernovae and neutron star mergers. The similarities of the physics involved in both of these problems, combined with the versatility of our code, invite us to investigate both. Although much work remains to be done, we are confident that this new tool will enable us to advance the state-of-the-art simulations of core-collapse supernovae and neutron star mergers.

Bibliography

Bibliography

Abbott, B., Abbott, R., Adhikari, R., Ageev, A., Agresti, J., Ajith, P., Allen, B., Allen, J., Amin, R., Anderson, S., Anderson, W., Araya, M., Armandula, H., Ashley, M., Asiri, F., Aufmuth, P., Aulbert, C., Babak, S., Balasubramanian, R., Ballmer, S., Barish, B., Barker, C., Barker, D., Barnes, M., Barr, B., Barton, M., Bayer, K., Beausoleil, R., Belczynski, K., Bennett, R., Berukoff, S., Betzwieser, J., Bhawal, B., Bilenko, I., Billingsley, G., Black, E., Blackburn, K., Blackburn, L., Bland, B., Bochner, B., Bogue, L., Bork, R., Bose, S., Brady, P., Braginsky, V., Brau, J., Brown, D., Bullington, A., Bunkowski, A., Buonanno, A., Burgess, R., Busby, D., Butler, W., Byer, R., Cadonati, L., Cagnoli, G., Camp, J., Cannizzo, J., Cannon, K., Cantley, C., Cao, J., Cardenas, L., Carter, K., Casey, M., Castiglione, J., Chandler, A., Chapsky, J., Charlton, P., Chatterji, S., Chelkowski, S., Chen, Y., Chickarmane, V., Chin, D., Christensen, N., Churches, D., Cokelaer, T., Colacino, C., Coldwell, R., Coles, M., Cook, D., Corbitt, T., Coyne, D., Creighton, J., Creighton, T., Crooks, D., Csatorday, P., Cusack, B., Cutler, C., Dalrymple, J., DAmbrosio, E., Danzmann, K., Davies, G., Daw, E., DeBra, D., Delker, T., Dergachev, V., Desai, S., DeSalvo, R., Dhurandhar, S., Di Credico, A., Díaz, M., Ding, H., Drever, R., Dupuis, R., Edlund, J., Ehrens, P., Elliffe, E., Etzel, T., Evans, M., Evans, T., Fairhurst, S., Fallnich, C., Farnham, D., Fejer, M., Findley, T., Fine, M., Finn, L., Franzen, K., Freise, A., Frey, R., Fritschel, P., Frolov, V., Fyffe, M., Ganezer, K., Garofoli, J., Giaime, J., Gillespie, A., Goda, K., Goggin, L., González, G., Goß ler, S., Grandclément, P., Grant, A., Gray, C., Gretarsson, A., Grimmett, D., Grote, H., Grunewald, S., Guenther, M., Gustafson, E., Gustafson, R., Hamilton, W., Hammond, M., Hanna, C., Hanson, J., Hardham, C., Harms, J., Harry, G., Hartunian, A., Heefner, J., Hefetz, Y., Heinzel, G., Heng, I., Hennessy, M., Hepler, N., Heptonstall, A., Heurs, M., Hewitson, M., Hild, S., Hindman, N., Hoang, P., Hough, J., Hrynevych, M., Hua, W., Ito, M., Itoh, Y., Ivanov, A., Jennrich, O., Johnson, B., Johnson, W., Johnston, W., Jones, D., Jones, G., Jones, L., Jungwirth, D., Kalogera, V., Katsavounidis, E., Kawabe, K., Kells, W., Kern, J., Khan, A., Killbourn, S., Killow, C., Kim, C., King, C., King, P., Klimenko, S., Koranda, S., Kötter, K., Kovalik, J., Kozak, D., Krishnan, B., Landry, M., Langdale, J., Lantz, B., Lawrence, R., Lazzarini, A., Lei, M., Leonor, I., Libbrecht, K., Libson, A., Lindquist, P., Liu, S., Logan, J., Lormand, M., Lubinski, M., Lück, H., Luna, M., Lyons, T., Machenschalk, B., MacInnis, M., Mageswaran, M., Mailand, K., Majid, W., Malec, M., Mandic, V., Mann, F., Marin, A., Márka, S., Maros, E., Mason, J., Mason, K., Matherny, O., Matone, L., Mavalvala, N., McCarthy, R., McClelland, D., McHugh, M., McNabb, J., Melissinos, A., Mendell, G., Mercer, R., Meshkov, S., Messaritaki, E., Messenger, C., Mikhailov, E., Mitra, S., Mitrofanov, V., Mitselmakher, G., Mittleman, R., Miyakawa, O., Mohanty, S., Moreno, G., Mossavi, K., Mueller, G., Mukherjee, S., Murray, P., Myers, E., Myers, J., Nagano, S., Nash, T., Nayak, R., Newton, G., Nocera, F., Noel, J., Nutzman, P., Olson, T., OReilly, B., Ottaway, D., Ottewill, A., Ouimette, D., Overmier, H., Owen, B., Pan, Y., Papa, M., Parameshwaraiah, V., Parameswariah, C., Pedraza, M., Penn, S., Pitkin, M., Plissi, M., Prix, R., Quetschke, V., Raab, F., Radkins, H., Rahkola, R., Rakhmanov, M., Rao, S., Rawlins, K., Ray-Majumder, S., Re, V., Redding, D., Regehr, M., Regimbau, T., Reid, S., Reilly, K., Reithmaier, K., Reitze, D., Richman, S., Riesen, R., Riles, K., Rivera, B., Rizzi, A., Robertson, D., Robertson, N., Robinson, C., Robison, L., Roddy, S., Rodriguez, A., Rollins, J., Romano, J., Romie, J., Rong, H., Rose, D., Rotthoff, E., Rowan, S., Rüdiger, A., Ruet, L., Russell, P., Ryan, K., Salzman, I., Sandberg, V., Sanders, G., Sannibale, V., Sarin, P., Sathyaprakash, B., Saulson, P., Savage, R., Sazonov, A., Schilling, R., Schlaufman, K., Schmidt, V., Schnabel, R., Schofield, R., Schutz, B., Schwinberg, P., Scott, S., Seader, S., Searle, A., Sears, B., Seel, S., Seifert, F., Sellers, D., Sengupta, A., Shapiro, C., Shawhan, P., Shoemaker, D., Shu, Q., Sibley, A., Siemens, X., Sievers, L., Sigg, D., Sintes, A., Smith, J., Smith, M., Sneddon, P., Spero, R., Spjeld, O., Stapfer, G., Steussy, D., Strain, K., Strom, D., Stuver, A., Summerscales, T., Sumner, M., Sung, M., Sutton, P., Sylvestre, J., Tanner, D., Tariq, H., Tarallo, M., Taylor, I., Taylor, R., Thorne, K., Tibbits, M., Tilav, S., Tinto, M., Tokmakov, K., Torres, C., Torrie, C., Traylor, G., Tyler, W., Ugolini, D., Ungarelli, C., Vallisneri, M., van Putten, M., Vass, S., Vecchio, A., Veitch, J., Vorvick, C., Vyachanin, S., Wallace, L., Walther, H., Ward, H., Ward, R., Ware, B., Watts, K., Webber, D., Weidner, A., Weiland, U., Weinstein, A., Weiss, (2006). Joint LIGO and TAMA300 search for gravitational waves from inspiralling neutron star binaries. *Physical Review D*, 73(10). 5

Abbott, B., Abbott, R., Adhikari, R., Ageev, A., Allen, B., Amin, R., Anderson, S., Anderson, W., Araya, M., Armandula, H., Ashley, M., Asiri, F., Aufmuth, P., Aulbert, C., Babak, S., Balasubramanian, R., Ballmer, S., Barish, B., Barker, C., Barker, D., Barnes, M., Barr, B., Barton, M., Bayer, K., Beausoleil, R., Belczynski, K., Bennett, R., Berukoff, S., Betzwieser, J., Bhawal, B., Bilenko, I., Billingsley, G., Black, E., Blackburn, K., Blackburn, L., Bland, B., Bochner, B., Bogue, L., Bork, R., Bose, S., Brady, P., Braginsky, V., Brau, J., Brown, D., Bullington, A., Bunkowski, A., Buonanno, A., Burgess, R., Busby, D., Butler, W., Byer, R., Cadonati, L., Cagnoli, G., Camp, J., Cantley, C., Cardenas, L., Carter, K., Casey, M., Castiglione, J., Chandler, A., Chapsky, J., Charlton, P., Chatterji, S., Chelkowski, S., Chen, Y., Chickarmane, V., Chin, D., Christensen, N., Churches, D., Cokelaer, T., Colacino, C., Coldwell, R., Coles, M., Cook, D., Corbitt, T., Coyne, D., Creighton, J., Creighton, T., Crooks, D., Csatorday, P., Cusack, B., Cutler, C., DAmbrosio, E., Danzmann, K., Daw, E., DeBra, D., Delker, T., Dergachev, V., DeSalvo, R., Dhurandhar, S., Credico, A., Díaz, M., Ding, H., Drever, R., Dupuis, R., Edlund, J., Ehrens, P., Elliffe, E., Etzel, T., Evans, M., Evans, T., Fairhurst, S., Fallnich, C., Farnham, D., Fejer, M., Findley, T., Fine, M., Finn, L., Franzen, K., Freise, A., Frey, R., Fritschel, P., Frolov, V., Fyffe, M., Ganezer, K., Garofoli, J., Giaime, J., Gillespie, A., Goda, K., González, G., Goß ler, S., Grandclément, P., Grant, A., Gray, C., Gretarsson, A., Grimmett, D., Grote, H., Grunewald, S., Guenther, M., Gustafson, E., Gustafson, R., Hamilton, W., Hammond, M., Hanson, J., Hardham, C., Harms, J., Harry, G., Hartunian, A., Heefner, J., Hefetz, Y., Heinzel, G., Heng, I., Hennessy, M., Hepler, N., Heptonstall, A., Heurs, M., Hewitson, M., Hild, S., Hindman, N., Hoang, P., Hough, J., Hrynevych, M., Hua, W., Ito, M., Itoh, Y., Ivanov, A., Jennrich, O., Johnson, B., Johnson, W., Johnston, W., Jones, D., Jones, L., Jungwirth, D., Kalogera, V., Katsavounidis, E., Kawabe, K., Kawamura, S., Kells, W., Kern, J., Khan, A., Killbourn, S., Killow, C., Kim, C., King, C., King, P., Klimenko, S., Koranda, S., Kötter, K., Kovalik, J., Kozak, D., Krishnan, B., Landry, M., Langdale, J., Lantz, B., Lawrence, R., Lazzarini, A., Lei, M., Leonor, I., Libbrecht, K., Libson, A., Lindquist, P., Liu, S., Logan, J., Lormand, M., Lubinski, M., Lück, H., Lyons, T., Machenschalk, B., MacInnis, M., Mageswaran, M., Mailand, K., Majid, W., Malec, M., Mann, F., Marin, A., Márka, S., Maros, E., Mason, J., Mason, K., Matherny, O., Matone, L., Mavalvala, N., McCarthy, R., McClelland, D., McHugh, M., McNabb, J., Mendell, G., Mercer, R., Meshkov, S., Messaritaki, E., Messenger, C., Mitrofanov, V., Mitselmakher, G., Mittleman, R., Miyakawa, O., Miyoki, S., Mohanty, S., Moreno, G., Mossavi, K., Mueller, G., Mukherjee, S., Murray, P., Myers, J., Nagano, S., Nash, T., Nayak, R., Newton, G., Nocera, F., Noel, J., Nutzman, P., Olson, T., OReilly, B., Ottaway, D., Ottewill, A., Ouimette, D., Overmier, H., Owen, B., Pan, Y., Papa, M., Parameshwaraiah, V., Parameswaran, A., Parameswariah, C., Pedraza, M., Penn, S., Pitkin, M., Plissi, M., Prix, R., Quetschke, V., Raab, F., Radkins, H., Rahkola, R., Rakhmanov, M., Rao, S., Rawlins, K., Ray-Majumder, S., Re, V., Redding, D., Regehr, M., Regimbau, T., Reid, S., Reilly, K., Reithmaier, K., Reitze, D., Richman, S., Riesen, R., Riles, K., Rivera, B., Rizzi, A., Robertson, D., Robertson, N., Robison, L., Roddy, S., Rollins, J., Romano, J., Romie, J., Rong, H., Rose, D., Rotthoff, E., Rowan, S., Rüdiger, A., Russell, P., Ryan, K., Salzman, I., Sandberg, V., Sanders, G., Sannibale, V., Sathyaprakash, B., Saulson, P., Savage, R., Sazonov, A., Schilling, R., Schlaufman, K., Schmidt, V., Schnabel, R., Schofield, R., Schutz, B., Schwinberg, P., Scott, S., Seader, S., Searle, A., Sears, B., Seel, S., Seifert, F., Sengupta, A., Shapiro, C., Shawhan, P., Shoemaker, D., Shu, Q., Sibley, A., Siemens, X., Sievers, L., Sigg, D., Sintes, A., Smith, J., Smith, M., Sneddon, P., Spero, R., Stapfer, G., Steussy, D., Strain, K., Strom, D., Stuver, A., Summerscales, T., Sumner, M., Sutton, P., Sylvestre, J., Takamori, A., Tanner, D., Tariq, H., Taylor, I., Taylor, R., Thorne, K., Tibbits, M., Tilav, S., Tinto, M., Tokmakov, K., Torres, C., Torrie, C., Traylor, G., Tyler, W., Ugolini, D., Ungarelli, C., Vallisneri, M., van Putten, M., Vass, S., Vecchio, A., Veitch, J., Vorvick, C., Vyachanin, S., Wallace, L., Walther, H., Ward, H., Ware, B., Watts, K., Webber, D., Weidner, A., Weiland, U., Weinstein, A., Weiss, R., Welling, H., Wen, L., Wen, S., Whelan, J., Whitcomb, S., Whiting, B., Wiley, S., Wilkinson, C., Willems, P., Williams, P., Williams, R., Willke, B., Wilson, A., Winjum, B., Winkler, W., Wise, S., Wiseman, A., Woan, G., Woods, D., Wooley, R., Worden, J., Wu, W., Yakus (2005). Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Physical Review D*, 72(8). 5

Abbott, B., Abbott, R., Adhikari, R., Ageev, A., Allen, B., Amin, R., Anderson, S., Anderson, W., Araya, M., Armandula, H., Asiri, F., Aufmuth, P., Aulbert, C., Babak, S., Balasubramanian, R., Ballmer, S., Barish, B., Barker, D., Barker-Patton, C., Barnes, M., Barr, B., Barton, M., Bayer, K., Beausoleil, R., Belczynski, K., Bennett, R., Berukoff, S., Betzwieser, J., Bhawal, B., Bilenko, I., Billingsley, G., Black, E., Blackburn, K., Bland-Weaver, B., Bochner, B., Bogue, L., Bork, R., Bose, S., Brady, P., Braginsky, V., Brau, J., Brown, D., Brozek, S., Bullington, A., Buonanno, A., Burgess, R., Busby, D., Butler, W., Byer, R., Cadonati, L., Cagnoli, G., Camp, J., Cantley, C., Cardenas, L., Carter, K., Casey, M., Castiglione, J., Chandler, A., Chapsky, J., Charlton, P., Chatterji, S., Chen, Y., Chickarmane, V., Chin, D., Christensen, N., Churches, D., Colacino, C., Coldwell, R., Coles, M., Cook, D., Corbitt, T., Coyne, D., Creighton, J., Creighton, T., Crooks, D., Csatorday, P., Cusack, B., Cutler, C., DAmbrosio, E., Danzmann, K., Davies, R., Daw, E., DeBra, D., Delker, T., DeSalvo, R., Dhurandhar, S., Díaz, M., Ding, H., Drever, R., Dupuis, R., Ebeling, C., Edlund, J., Ehrens, P., Elliffe, E., Etzel, T., Evans, M., Evans, T., Fallnich, C., Farnham, D., Fejer, M., Fine, M., Finn, L., Flanagan, E., Freise, A., Frey, R., Fritschel, P., Frolov, V., Fyffe, M., Ganezer, K., Giaime, J., Gillespie, A., Goda, K., González, G., Goß ler, S., Grandclément, P., Grant, A., Gray, C., Gretarsson, A., Grimmett, D., Grote, H., Grunewald, S., Guenther, M., Gustafson, E., Gustafson, R., Hamilton, W., Hammond, M., Hanson, J., Hardham, C., Harry, G., Hartunian, A., Heefner, J., Hefetz, Y., Heinzel, G., Heng, I., Hennessy, M., Hepler, N., Heptonstall, A., Heurs, M., Hewitson, M., Hindman, N., Hoang, P., Hough, J., Hrynevych, M., Hua, W., Ingley, R., Ito, M., Itoh, Y., Ivanov, A., Jennrich, O., Johnson, W., Johnston, W., Jones, L., Jungwirth, D., Kalogera, V., Katsavounidis, E., Kawabe, K., Kawamura, S., Kells, W., Kern, J., Khan, A., Killbourn, S., Killow, C., Kim, C., King, C., King, P., Klimenko, S., Kloevekorn, P., Koranda, S., Kötter, K., Kovalik, J., Kozak, D., Krishnan, B., Landry, M., Langdale, J., Lantz, B., Lawrence, R., Lazzarini, A., Lei, M., Leonhardt, V., Leonor, I., Libbrecht, K., Lindquist, P., Liu, S., Logan, J., Lormand, M., Lubinski, M., Lück, H., Lyons, T., Machenschalk, B., MacInnis, M., Mageswaran, M., Mailand, K., Majid, W., Malec, M., Mann, F., Marin, A., Márka, S., Maros, E., Mason, J., Mason, K., Matherny, O., Matone, L., Mavalvala, N., McCarthy, R., McClelland, D., McHugh, M., McNamara, P., Mendell, G., Meshkov, S., Messenger, C., Mitrofanov, V., Mitselmakher, G., Mittleman, R., Miyakawa, O., Miyoki, S., Mohanty, S., Moreno, G., Mossavi, K., Mours, B., Mueller, G., Mukherjee, S., Myers, J., Nagano, S., Nash, T., Naundorf, H., Nayak, R., Newton, G., Nocera, F., Nutzman, P., Olson, T., OReilly, B., Ottaway, D., Ottewill, A., Ouimette, D., Overmier, H., Owen, B., Papa, M., Parameswariah, C., Parameswariah, V., Pedraza, M., Penn, S., Pitkin, M., Plissi, M., Pratt, M., Quetschke, V., Raab, F., Radkins, H., Rahkola, R., Rakhmanov, M., Rao, S., Redding, D., Regehr, M., Regimbau, T., Reilly, K., Reithmaier, K., Reitze, D., Richman, S., Riesen, R., Riles, K., Rizzi, A., Robertson, D., Robertson, N., Robison, L., Roddy, S., Rollins, J., Romano, J., Romie, J., Rong, H., Rose, D., Rotthoff, E., Rowan, S., Rüdiger, A., Russell, P., Ryan, K., Salzman, I., Sanders, G., Sannibale, V., Sathyaprakash, B., Saulson, P., Savage, R., Sazonov, A., Schilling, R., Schlaufman, K., Schmidt, V., Schofield, R., Schrempel, M., Schutz, B., Schwinberg, P., Scott, S., Searle, A., Sears, B., Seel, S., Sengupta, A., Shapiro, C., Shawhan, P., Shoemaker, D., Shu, Q., Sibley, A., Siemens, X., Sievers, L., Sigg, D., Sintes, A., Skeldon, K., Smith, J., Smith, M., Sneddon, P., Spero, R., Stapfer, G., Strain, K., Strom, D., Stuver, A., Summerscales, T., Sumner, M., Sutton, P., Sylvestre, J., Takamori, A., Tanner, D., Tariq, H., Taylor, I., Taylor, R., Thorne, K., Tibbits, M., Tilav, S., Tinto, M., Tokmakov, K., Torres, C., Torrie, C., Traeger, S., Traylor, G., Tyler, W., Ugolini, D., Vallisneri, M., van Putten, M., Vass, S., Vecchio, A., Vorvick, C., Vyachanin, S., Wallace, L., Walther, H., Ward, H., Ware, B., Watts, K., Webber, D., Weidner, A., Weiland, U., Weinstein, A., Weiss, R., Welling, H., Wen, L., Wen, S., Whelan, J., Whitcomb, S., Whiting, B., Willems, P., Williams, P., Williams, R., Willke, B., Wilson, A., Winjum, B., Winkler, W., Wise, S., Wiseman, A., Woan, G., Wooley, R., Worden, J., Yakushin, I., Yamamoto, H., Yoshida, S., Zawischa, I., Zhang, L., Zotov, N., Zucker, M., and Zweizig, J. (2004a). Analysis of LIGO data for gravitational waves from binary neutron stars. *Physical Review D*, 69(12). 5

Abbott, B., Abbott, R., Adhikari, R., Ageev, A., Allen, B., Amin, R., Anderson, S. B., Anderson, W. G., Araya, M., and Armandula, H. (2004b). Detector description and performance for the first coincidence observations between LIGO and GEO. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 517(1-3):154–179. 4

Abbott, B. P., Abbott, R., Acernese, F., Adhikari, R., Ajith, P., Allen, B., Allen, G., Alshourbagy, M., Amin, R. S., Anderson, S. B., Anderson, W. G., Antonucci, F., Aoudia, S., Arain, M. A., Araya, M., Armandula, H., Armor, P., Arun, K. G., Aso, Y., Aston, S., Astone, P., Aufmuth, P., Aulbert, C., Babak, S., Baker, P., Ballardin, G., Ballmer, S., Barker, C., Barker, D., Barone, F., Barr, B., Barriga, P., Barsotti, L., Barsuglia, M., Barton, M. A., Bartos, I., Bassiri, R., Bastarrika, M., Bauer, T. S., Behnke, B., Beker, M., Benacquista, M., Betzwieser, J., Beyersdorf, P. T., Bigotta, S., Bilenko, I. A., Billingsley, G., Birindelli, S., Biswas, R., Bizouard, M. A., Black, E., Blackburn, J. K., Blackburn, L., Blair, D., Bland, B., Boccara, C., Bodiya, T. P., Bogue, L., Bondu, F., Bonelli, L., Bork, R., Boschi, V., Bose, S., Bosi, L., Braccini, S., Bradaschia, C., Brady, P. R., Braginsky, V. B., Brau, J. E., Bridges, D. O., Brillet, A., Brinkmann, M., Brisson, V., Van Den Broeck, C., Brooks, A. F., Brown, D. A., Brummit, A., Brunet, G., Budzyski, R., Bulik, T., Bullington, A., Bulten, H. J., Buonanno, A., Burmeister, O., Buskulic, D., Byer, R. L., Cadonati, L., Cagnoli, G., Calloni, E., Camp, J. B., Campagna, E., Cannizzo, J., Cannon, K. C., Canuel, B., Cao, J., Carbognani, F., Cardenas, L., Caride, S., Castaldi, G., Caudill, S., Cavaglià, M., Cavalier, F., Cavalieri, R., Cella, G., Cepeda, C., Cesarini, E., Chalermsongsak, T., Chalkley, E., Charlton, P., Chassande-Mottin, E., Chatterji, S., Chelkowski, S., Chen, Y., Chincarini, A., Christensen, N., Chung, C. T. Y., Clark, D., Clark, J., Clayton, J. H., Cleva, F., Coccia, E., Cokelaer, T., Colacino, C. N., Colas, J., Colla, A., Colombini, M., Conte, R., Cook, D., Corbitt, T. R. C., Corda, C., Cornish, N., Corsi, A., Coulon, J.-P., Coward, D., Coyne, D. C., Creighton, J. D. E., Creighton, T. D., Cruise, A. M., Culter, R. M., Cumming, A., Cunningham, L., Cuoco, E., Danilishin, S. L., D'Antonio, S., Danzmann, K., Dari, A., Dattilo, V., Daudert, B., Davier, M., Davies, G., Daw, E. J., Day, R., De Rosa, R., DeBra, D., Degallaix, J., del Prete, M., Dergachev, V., Desai, S., DeSalvo, R., Dhurandhar, S., Fiore, L. D., Lieto, A. D., Emilio, M. D. P., Virgilio, A. D., Díaz, M., Dietz, A., Donovan, F., Dooley, K. L., Doomes, E. E., Drago, M., Drever, R. W. P., Dueck, J., Duke, I., Dumas, J.-C., Dwyer, J. G., Echols, C., Edgar, M., Effler, A., Ehrens, P., Espinoza, E., Etzel, T., Evans, M., Evans, T., Fafone, V., Fairhurst, S., Faltas, Y., Fan, Y., Fazi, D., Fehrmann, H., Ferrante, I., Fidecaro, F., Finn, L. S., Fiori, I., Flaminio, R., Flasch, K., Foley, S., Forrest, C., Fotopoulos, N., Fournier, J.-D., Franc, J., Franzen, A., Frasca, S., Frasconi, F., Frede, M., Frei, M., Frei, Z., Freise, A., Frey, R., Fricke, T., Fritschel, P., Frolov, V. V., Fyffe, M., Galdi, V., Gammaitoni, L., Garofoli, J. A., Garufi, F., Gemme, G., Genin, E., Gennai, A., Gholami, I., Giaime, J. A., Giampanis, S., Giardina, K. D., Giazotto, A., Goda, K., Goetz, E., Goggin, L. M., González, G., Gorodetsky, M. L., Goß ler, S., Gouaty, R., Granata, M., Granata, V., Grant, A., Gras, S., Gray, C., Gray, M., Greenhalgh, R. J. S., Gretarsson, A. M., Greverie, C., Grimaldi, F., Grosso, R., Grote, H., Grunewald, S., Guenther, M., Guidi, G., Gustafson, E. K., Gustafson, R., Hage, B., Hallam, J. M., Hammer, D., Hammond, G. D., Hanna, C., Hanson, J., Harms, J., Harry, G. M., Harry, I. W., Harstad, E. D., Haughian, K., Hayama, K., Heefner, J., Heitmann, H., Hello, P., Heng, I. S., Heptonstall, A., Hewitson, M., Hild, S., Hirose, E., Hoak, D., Hodge, K. A., Holt, K., Hosken, D. J., Hough, J., Hoyland, D., Huet, D., Hughey, B., Huttner, S. H., Ingram, D. R., Isogai, T., Ito, M., Ivanov, A., Jaranowski, P., Johnson, B., Johnson, W. W., Jones, D. I., Jones, G., Jones, R., de La Jordana, L. S., Ju, L., Kalmus, P., Kalogera, V., Kandhasamy, S., Kanner, J., Kasprzyk, D., Katsavounidis, E., Kawabe, K., Kawamura, S., Kawazoe, F., Kells, W., Keppel, D. G., Khalaidovski, A., Khalili, F. Y., Khan, R., Khazanov, E., King, P., Kissel, J. S., Klimenko, S., Kokeyama, K., Kondrashov, V., Kopparapu, R., Koranda, S., Kowalska, I., Kozak, D., Krishnan, B., Królak, A., Kumar, R., Kwee, P., Penna, P. L., Lam, P. K., Landry, M., Lantz, B., Lazzarini, A., Lei, H., Lei, M., Leindecker, N., Leonor, I., Leroy, N., Letendre, N., Li, C., Lin, H., Lindquist, P. E., Littenberg, T. B., Lockerbie, N. A., Lodhia, D., Longo, M., Lorenzini, M., Loriette, V., Lormand, M., Losurdo, G., Lu, P., Lubinski, M., Lucianetti, A., Lück, H., Machenschalk, B., MacInnis, M., Mackowski, J.-M., Mageswaran, M., Mailand, K., Majorana, E., Man, N., Mandel, I., Mandic, V., Mantovani, M., Marchesoni, F., Marion, F., Márka, S., Márka, Z., Markosyan, A., Markowitz, J., Maros, E., Marque, J., Martelli, F., Martin, I. W., Martin, R. M., Marx, J. N., Mason, K., Masserot, A., Matichard, F., Matone, L., Matzner, R. A., Mavalvala, N. (2010). Searches for Gravitational Waves from Known Pulsars with Science Run 5 LIGO Data. *The Astrophysical Journal*, 713(1):671–685. 5

Abramovici, A., Althouse, W. E., Drever, R. W. P., Gürsel, Y., Kawamura, S., Raab, F. J., Shoemaker, D., Sievers, L., Spero, R. E., Thorne, K. S., Vogt, R. E., Weiss, R., Whitcomb, S. E., and Zucker, M. E. (1992). LIGO: The Laser Interferometer Gravitational-Wave Observatory. *Science (New York, N.Y.)*, 256(5055):325–333. 5, 129, 134

Acernese, F., Amico, P., Al-Shourbagy, M., Aoudia, S., Avino, S., Babusci, D., Ballardin, G., Barillé, R., Barone, F., Barsotti, L., Barsuglia, M., Beauville, F., Bizouard, M. A., Boccara, C., Bondu, F., Bosi, L., Bradaschia, C., Braccini, S., Brillet, A., Brisson, V., Brocco, L., Buskulic, D., Calloni, E., Campagna, E., Cavalier, F., Cavalieri, R., Cella, G., Chassande-Mottin, E., Corda, C., Clapson, A.-C., Cleva, F., Coulon, J.-P., Cuoco, E., Dattilo, V., Davier, M., Rosa, R. D., Fiore, L. D., Virgilio, A. D., Dujardin, B., Eleuteri, A., Enard, D., Ferrante, I., Fidecaro, F., Fiori, I., Flaminio, R., Fournier, J.-D., Frasca, S., Frasconi, F., Freise, A., Gammaitoni, L., Gennai, A., Giazotto, A., Giordano, G., Giordano, L., Gouaty, R., Grosjean, D., Guidi, G., Hebri, S., Heitmann, H., Hello, P., Holloway, L., Kreckelbergh, S., Penna, P. L., Loriette, V., Loupias, M., Losurdo, G., Mackowski, J.-M., Majorana, E., Man, C. N., Mantovani, M., Marchesoni, F., Marion, F., Marque, J., Martelli, F., Masserot, A., Mazzoni, M., Milano, L., Moins, C., Moreau, J., Morgado, N., Mours, B., Pai, A., Palomba, C., Paoletti, F., Pardi, S., Pasqualetti, A., Passaquieti, R., Passuello, D., Perniola, B., Piergiovanni, F., Pinard, L., Poggiani, R., Punturo, M., Puppo, P., Qipiani, K., Rapagnani, P., Reita, V., Remillieux, A., Ricci, F., Ricciardi, I., Ruggi, P., Russo, G., Solimeno, S., Spallicci, A., Stanga, R., Taddei, R., Tombolato, D., Tonelli, M., Toncelli, A., Tournefier, E., Travasso, F., Vajente, G., Verkindt, D., Vetrano, F., Viceré, A., Vinet, J.-Y., Vocca, H., Yvert, M., and Zhang, Z. (2005). Status of Virgo. *Classical and Quantum Gravity*, 22(18):S869–S880. 4

Aftosmis, M. J., Melton, J. E., and Berger, M. J. (1995). Adaptation and Surface Modeling for Cartesian Mesh Methods. *AIAA Paper*, 95-1725-CP. 107

Agertz, O., Moore, B., Stadel, J., Potter, D., Miniati, F., Read, J., Mayer, L., Gawryszczak, A., Kravtsov, A., Nordlund, A. k., Pearce, F., Quilis, V., Rudd, D., Springel, V., Stone, J., Tasker, E., Teyssier, R., Wadsley, J., and Walder, R. (2007). Fundamental differences between SPH and grid methods. *Monthly Notices of the Royal Astronomical Society*, 380(3):963–978. 14

Alcubierre, M. (2008). *Introduction to 3+1 numerical relativity*. Oxford University Press. 182

Anderson, M., Hirschmann, E., Lehner, L., Liebling, S., Motl, P., Neilsen, D., Palenzuela, C., and Tohline, J. (2008). Magnetized Neutron-Star Mergers and Gravitational-Wave Signals. *Physical Review Letters*, 100(19):1–4. 12

Anderson, S. B., Gorham, P. W., Kulkarni, S. R., Prince, T. A., and Wolszczan, A. (1990). Discovery of two radio pulsars in the globular cluster M15. *Nature*, 346(6279):42–44. 10, 128

Anile, A. M. (1989). *Relativistic fluids and magneto-fluids: with applications in astrophysics ...* Cambridge University Press. 30

Anile, A. M. and Pennisi, S. (1987). On the Mathematical Structure of Test Relativistic Magneto-fluid Dynamics. *Annales del I'lnstitut Henri Poincare*, 46:127. 30

Ardeljan, N. V., Bisnovatyi-Kogan, G. S., and Moiseenko, S. G. (2005). Magnetorotational supernovae. *Monthly Notices of the Royal Astronomical Society*, 359(1):333–344. 149

Arnett, W. D., Bahcall, J. N., Kirshner, R. P., and Woosley, S. E. (1989). Supernova 1987A. *Annual Review of Astronomy and Astrophysics*, 27(1):629–700. 8

Balay, S., Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H. (2008). PETSc Users Manual. 116

Balay, S., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H. (2009). PETSc Web page: http://www.mcs.anl.gov/petsc. 116

Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F. (1997). *Efficient Management of Parallelism in Object Oriented Numerical Software Libraries*, pages 163–202. Birkhäuser Press, Boston. 116

Balbus, S. A. and Hawley, J. F. (1991). A powerful local shear instability in weakly magnetized disks. I - Linear analysis. II - Nonlinear evolution. *The Astrophysical Journal*, 376:214. 149

Balsara, D. (2001). Total Variation Diminishing Scheme for Relativistic Magnetohydrodynamics. *The Astrophysical Journal Supplement Series*, 132(1):83–101. 53, 57

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. 2

Barmin, A., Kulikovskiy, A. G., and Pogorelov, N. V. (1996). Shock-Capturing Approach and Nonevolutionary Solutions in Magnetohydrodynamics. *Journal of Computational Physics*, 126(1):77–90. 53

Baron, E., Cooperstein, J., and Kahana, S. (1985). Type II supernovae in 12M_{cirdot} and 15M_{cirdot} stars: The equation of state and general relativity. *Physical Review Letters*, 55(1):126–129. 7

Baumgarte, T., Brady, P., Creighton, J., Lehner, L., Pretorius, F., and DeVoe, R. (2008). Learning about compact binary merger: The interplay between numerical relativity and gravitational-wave astronomy. *Physical Review D*, 77(8):1–19. 5

Baumgarte, T. W. (1998). Numerical integration of Einsteins field equations. *Physical Review D*, 59(2). 182

Belczynski, K., Kalogera, V., and Bulik, T. (2002). A Comprehensive Study of Binary Compact Objects as Gravitational Wave Sources: Evolutionary Channels, Rates, and Physical Properties. *The Astrophysical Journal*, 572(1):407–431. 5

Belczynski, K., Kalogera, V., Rasio, F. A., Taam, R. E., Zezas, A., Bulik, T., Maccarone, T. J., and Ivanova, N. (2008). Compact Object Modeling with the StarTrack Population Synthesis Code. *The Astrophysical Journal Supplement Series*, 174(1):223–260. 10

Belczynski, K., Perna, R., Bulik, T., Kalogera, V., Ivanova, N., and Lamb, D. Q. (2006). A Study of Compact Object Mergers as Short GammaRay Burst Progenitors. *The Astrophysical Journal*, 648(2):1110–1116. 4

Berger, M. and Collela, P. (1989). Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82(1):64–84. 107

Berger, M. and Oliger, J. (1984). Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53(3):484–512. 101

Blanchet, L., Damour, T., and Schaefer, G. (1990). Post-Newtonian Hydrodynamics and Post-Newtonian Gravitational Wave Generation for Numerical Relativity. *Royal Astronomical Society, Monthly Notices*, 242:289–305. 129

Blondin, J. M. and Mezzacappa, A. (2006). The Spherical Accretion Shock Instability in the Linear Regime. *The Astrophysical Journal*, 642(1):401–409. 8, 148

Blondin, J. M. and Mezzacappa, A. (2007). Pulsar spins from an instability in the accretion shock of supernovae. *Nature*, 445(7123):58–60. 150, 151, 160, 164, 179

Blondin, J. M., Mezzacappa, A., and DeMarino, C. (2003). Stability of Standing Accretion Shocks, with an Eye toward CoreCollapse Supernovae. *The Astrophysical Journal*, 584(2):971–980. 8, 148, 150, 151, 154

Bradaschia, C. (1990). The VIRGO Project: A wide band antenna for gravitational wave detection. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 289(3):518–525. 5, 129

Brandenburg, A. and Subramanian, K. (2005). Astrophysical magnetic fields and nonlinear dynamo theory. *Physics Reports*, 417(1-4):1–209. 175

Brio, M. and Wu, C. (1988). An upwind differencing scheme for the equations of ideal magnetohydrodynamics. *Journal of Computational Physics*, 75(2):400–422. 43, 44

Bruenn, S. W., Dirk, C. J., Mezzacappa, A., Hayes, J. C., Blondin, J. M., Hix, W. R., and Messer, O. E. B. (2006). Modeling core collapse supernovae in 2 and 3 dimensions with spectral neutrino transport. *Journal of Physics: Conference Series*, 46:393–402. 149

Budiardja, R. D., Cardall, C. Y., and Mezzacappa, A. (2010). Parallel FFT-based Poisson Solver for Isolated Three-dimensional Systems. *In preparation to be submitted to Computer Physics Communications*. 14, 181

Buras, R., Janka, H.-T., Rampp, M., and Kifonidis, K. (2006). Two-dimensional hydrodynamic core-collapse supernova simulations with spectral neutrino transport. *Astronomy and Astrophysics*, 457(1):281–308. 8, 149

Burrows, A., Dessart, L., Livne, E., Ott, C. D., and Murphy, J. (2007). Simulations of Magnetically Driven Supernova and Hypernova Explosions in the Context of Rapid Rotation. *The Astrophysical Journal*, 664(1):416–434. 149

Burrows, A., Hayes, J., and Fryxell, B. A. (1995). On the Nature of Core-Collapse Supernova Explosions. *The Astrophysical Journal*, 450:830. 8

Burrows, A., Livne, E., Dessart, L., Ott, C. D., and Murphy, J. (2006). A New Mechanism for CoreCollapse Supernova Explosions. *The Astrophysical Journal*, 640(2):878–890. 149

Burrows, A., Walder, R., Ott, C. D., and Livne, E. (2004). Rotating Core Collapse and Bipolar Supernova Explosions. page 11. 9

Cardall, C., Lentz, E., and Mezzacappa, A. (2005). Conservative special relativistic radiative transfer for multidimensional astrophysical simulations: Motivation and elaboration. *Physical Review D*, 72(4). 183

Cardall, C. and Mezzacappa, A. (2003). Conservative formulations of general relativistic kinetic theory. *Physical Review D*, 68(2). 183

Cardall, C. Y., Razoumov, A. O., Endeve, E., and Mezzacappa, A. (2006). *The Long Term: Six-dimensional Core-collapse Supernova Models*. World Scientific Publishing Company, London, England. 183

Cerdá-Durán, P., Font, J. A., and Dimmelmeier, H. (2007). General relativistic simulations of passive-magneto-rotational core collapse with microphysics. *Astronomy and Astrophysics*, 474(1):169–191. 149

Chandrasekhar, S. (1987). *Ellipsoidal Figures of Equilibrium*. Dover. 82

Chatterjee, S., Vlemmings, W. H. T., Brisken, W. F., Lazio, T. J. W., Cordes, J. M., Goss, W. M., Thorsett, S. E., Fomalont, E. B., Lyne, A. G., and Kramer, M. (2005). Getting Its Kicks: A VLBA Parallax for the Hyperfast Pulsar B1508+55. *The Astrophysical Journal*, 630(1):L61–L64. 8

Colgate, S. and Johnson, M. (1960). Hydrodynamic Origin of Cosmic Rays. *Physical Review Letters*, 5(6):235–238. 7

Colgate, S. A. and White, R. H. (1966). The Hydrodynamic Behavior of Supernovae Explosions. *The Astrophysical Journal*, 143:626. 182

Cook, G. B., Shapiro, S. L., and Teukolsky, S. A. (1994). Rapidly rotating neutron stars in general relativity: Realistic equations of state. *The Astrophysical Journal*, 424:823. 144

Costa, E., Frontera, F., Heise, J., Feroci, M., Zand, J. I., and F (1997). Discovery of an X-ray afterglow associated with the γ-ray burst of 28 February 1997. *Nature*, (February 1997):783–785. 3

Courant, R. and Friedrichs, K. O. (1977). *Supersonic flow and shock waves*. Springer. 40, 151

Cutler, C., Apostolatos, T., Bildsten, L., Finn, L., Flanagan, E., Kennefick, D., Markovic, D., Ori, A., Poisson, E., Sussman, G., and Thorne, K. (1993). The last three minutes: Issues in gravitational-wave measurements of coalescing compact binaries. *Physical Review Letters*, 70(20):2984–2987. 147

Cutler, C. and Flanagan, E. (1994). Gravitational waves from merging compact binaries: How accurately can one extract the binarys parameters from the inspiral waveform? *Physical Review D*, 49(6):2658–2697. 145, 147

Del Zanna, L. and Bucciantini, N. (2002). An efficient shock-capturing central-type scheme for multidimensional relativistic flows I. Hydrodynamics. *Astronomy and Astrophysics*, 390(3):1177–1186. 16, 21, 24, 44, 46, 53

Del Zanna, L., Bucciantini, N., and Londrillo, P. (2003). An efficient shock-capturing central-type scheme for multidimensional relativistic flows II. Magnetohydrodynamics. *Astronomy and Astrophysics*, 400(2):397–413. 16, 17, 26, 28, 30, 53, 57, 64

Del Zanna, L., Zanotti, O., Bucciantini, N., and Londrillo, P. (2007). ECHO: a Eulerian conservative high-order scheme for general relativistic magnetohydrodynamics and magnetodynamics. *Astronomy and Astrophysics*, 473(1):11–30. 181

Dewey, R. J., Taylor, J. H., Weisberg, J. M., and Stokes, G. H. (1985). A search for low-luminosity pulsars. *The Astrophysical Journal*, 294:L25. 10, 128

Donat, R. (1998). A Flux-Split Algorithm Applied to Relativistic Flows. *Journal of Computational Physics*, 146(1):58–81. 44

Dorr, F. W. (1970). The Direct Solution of the Discrete Poisson Equation on a Rectangle. *SIAM Review*, 12(2):248–263. 67

Duez, M., Liu, Y., Shapiro, S., Shibata, M., and Stephens, B. (2006). Collapse of Magnetized Hypermassive Neutron Stars in General Relativity. *Physical Review Letters*, 96(3). 12

Duncan, R. C. and Thompson, C. (1992). Formation of very strongly magnetized neutron stars - Implications for gamma-ray bursts. *The Astrophysical Journal*, 392:L9. 149

Eastwood, J. and Brownrigg, D. (1979). Remarks on the solution of poisson's equation for isolated systems. *Journal of Computational Physics*, 32(1):24–38. 67, 70

Einfeldt, B. (1988). On Godunov-Type Methods for Gas Dynamics. *SIAM Journal on Numerical Analysis*, 25(2):294. 16

Endeve, E., Cardall, C. Y., Budiardja, R. D., and Mezzacappa, A. (2007). MHD models of stellar core collapse with GenASiS. *Journal of Physics: Conference Series*, 78:012016. 149, 183

Endeve, E., Cardall, C. Y., Budiardja, R. D., and Mezzacappa, A. (2010). Generation of Magnetic Fields by The Stationary Accretion Shock Instability. *The Astrophysical Journal*, 713(2):1219–1243. 15, 150, 182

Evans, C. R. and Hawley, J. F. (1988). Simulation of magnetohydrodynamic flows - A constrained transport method. *The Astrophysical Journal*, 332:659. 17

Faber, J., Grandclément, P., and Rasio, F. (2004). Mergers of irrotational neutron star binaries in conformally flat gravity. *Physical Review D*, 69(12). 11

Figer, D. F., Najarro, F., Geballe, T. R., Blum, R. D., and Kudritzki, R. P. (2005). Massive Stars in the SGR 1806-20 Cluster. *The Astrophysical Journal*, 622(1):L49–L52. 9

Fishman, G. J. and Meegan, C. A. (1995). Gamma-Ray Bursts. *Annual Review of Astronomy and Astrophysics*, 33(1):415–458. 2, 3

Font, J. A. (2008). Numerical Hydrodynamics and Magnetohydrodynamics in General Relativity. *Living Reviews in Relativity*, 22. 181

Fox, D. B., Frail, D. A., Price, P. A., Kulkarni, S. R., Berger, E., Piran, T., Soderberg, A. M., Cenko, S. B., Cameron, P. B., Gal-Yam, A., Kasliwal, M. M., Moon, D.-S., Harrison, F. A., Nakar, E., Schmidt, B. P., Penprase, B., Chevalier, R. A., Kumar, P., Roth, K., Watson, D., Lee, B. L., Shectman, S., Phillips, M. M., Roth, M., McCarthy, P. J., Rauch, M., Cowie, L., Peterson, B. A., Rich, J., Kawai, N., Aoki, K., Kosugi, G., Totani, T., Park, H.-S., MacFadyen, A., and Hurley, K. C. (2005). The afterglow of GRB 050709 and the nature of the short-hard gamma-ray bursts. *Nature*, 437(7060):845–50. 4

Frail, D. A., Kulkarni, S. R., Bloom, J. S., Djorgovski, S. G., Gorjian, V., Gal, R. R., Meltzer, J., Sari, R., Chaffee, F. H., Goodrich, R., Frontera, F., and Costa, E. (1999). The Radio Afterglow and the Host Galaxy of the X-Rayrich GRB 981226. *The Astrophysical Journal*, 525(2):L81–L84. 3

Frail, D. A. and Taylor, G. B. (1997). The radio afterglow from the -ray burst of 8 May 1997. *Nature*, (May):261–263. 3

Freund, R. W., Golub, G. H., and Nachtigal, N. M. (2008). Iterative solution of linear systems. *Acta Numerica*, 1:57. 116

Frey, R. E. (2007). *LIGO: Status and Recent Results*. AIP. 5

Frigo, M. and Johnson, S. (2005). The Design and Implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231. 68, 71

Fryer, C. L. (2004). *Stellar collapse*. Springer. 9

Fryer, C. L. and Heger, A. (2000). CoreCollapse Simulations of Rotating Stars. *The Astrophysical Journal*, 541(2):1033–1050. 9

Gaensler, B. M., McClure-Griffiths, N. M., Oey, M. S., Haverkorn, M., Dickey, J. M., and Green, A. J. (2005). A Stellar Wind Bubble Coincident with the Anomalous X-Ray Pulsar 1E 1048.1-5937: Are Magnetars Formed from Massive Progenitors? *The Astrophysical Journal*, 620(2):L95–L98. 9

Galama, T., Vreeswijk, P., Paradijs, J. V., and C (1998). An unusual supernova in the error box of the γ-ray burst of 25 April 1998. *Nature*, 395(April):670–672. 3

Gammie, C. F., McKinney, J. C., and Tóth, G. (2003). HARM: A Numerical Scheme for General Relativistic Magnetohydrodynamics. page 38. 181

Gehrels, N., Sarazin, C. L., O'Brien, P. T., Zhang, B., Barbier, L., Barthelmy, S. D., Blustin, A., Burrows, D. N., Cannizzo, J., Cummings, J. R., Goad, M., Holland, S. T., Hurkett, C. P., Kennea, J. A., Levan, A., Markwardt, C. B., Mason, K. O., Meszaros, P., Page, M., Palmer, D. M., Rol, E., Sakamoto, T., Willingale, R., Angelini, L., Beardmore, A., Boyd, P. T., Breeveld, A., Campana, S., Chester, M. M., Chincarini, G., Cominsky, L. R., Cusumano, G., de Pasquale, M., Fenimore, E. E., Giommi, P., Gronwall, C., Grupe, D., Hill, J. E., Hinshaw, D., Hjorth, J., Hullinger, D., Hurley, K. C., Klose, S., Kobayashi, S., Kouveliotou, C., Krimm, H. A., Mangano, V., Marshall, F. E., McGowan, K., Moretti, A., Mushotzky, R. F., Nakazawa, K., Norris, J. P., Nousek, J. A., Osborne, J. P., Page, K., Parsons, A. M., Patel, S., Perri, M., Poole, T., Romano, P., Roming, P. W. A., Rosen, S., Sato, G., Schady, P., Smale, A. P., Sollerman, J., Starling, R., Still, M., Suzuki, M., Tagliaferri, G., Takahashi, T., Tashiro, M., Tueller, J., Wells, A. A., White, N. E., and Wijers, R. A. M. J. (2005). A short gamma-ray burst apparently associated with an elliptical galaxy at redshift z = 0.225. *Nature*, 437(7060):851–4. 3, 4

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: Letters*, 399(1):L164–L168. 12

Gropp, W. and Lusk, E. (2010). The Message Passing Interface (MPI) standard: http://www.mcs.anl.gov/research/projects/mpi/. 32, 72

Gropp, W., Lusk, E., and Skjellum, A. (1999). *Using MPI - 2nd Edition: Portable Parallel Programming with the Message Passing Interface (Scientific and Engineering Computation)*. The MIT Press, 2 edition. 32, 72

Hansen, B. and Phinney, E. (1997). The pulsar kick velocity distribution. *Monthly Notices of Royal Astronomical Society*, 291:569. 8

Harry, G., Houser, J., and Strain, K. (2002). Comparison of advanced gravitational-wave detectors. *Physical Review D*, 65(8). 5

Harten, A., Lax, P. D., and Leer, B. V. (1983). On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Review*, 25(1):35. 16

Heger, A., Woosley, S. E., and Spruit, H. C. (2005). Presupernova Evolution of Differentially Rotating Massive Stars Including Magnetic Fields. *The Astrophysical Journal*, 626(1):350–363. 149, 152, 153

Herant, M., Benz, W., Hix, W. R., Fryer, C. L., and Colgate, S. A. (1994). Inside the supernova: A powerful convective engine. *The Astrophysical Journal*, 435:339. 8

Hjorth, J., Watson, D., Fynbo, J. P. U., Price, P. A., Jensen, B. L., Jø rgensen, U. G., Kubas, D., Gorosabel, J., Jakobsson, P., Sollerman, J., Pedersen, K., and Kouveliotou, C. (2005). The optical afterglow of the short gamma-ray burst GRB 050709. *Nature*, 437(7060):859–61. 4

Hobbs, G., Lorimer, D. R., Lyne, A. G., and Kramer, M. (2005). A statistical study of 233 pulsar proper motions. *Monthly Notices of the Royal Astronomical Society*, 360(3):974–992. 10

Hockney, R. W. (1970). The Potential Calculation and Some Applications. *Methods in Computational Physics*, 9:135–211. xiii, 67, 70, 71

Hockney, R. W. and Eastwood, J. W. (1989). *Computer Simulation Using Particles*. CRC Press. 67, 70

Hulse, R. A. and Taylor, J. H. (1975). Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195:L51. 9, 128

Hunter, C. (1962). The Instability of the Collapse of a Self-Gravitating Gas Cloud. *The Astrophysical Journal*, 136:594. 88

Hurley, K. (1992). Receding from our grasp. *Nature*, 357(6374):112–113. 2

Hurley, K., Dingus, B. L., Mukherjee, R., Sreekumar, P., Kouveliotou, C., Meegan, C., Fishman, G. J., Band, D., Ford, L., Bertsch, D., Cline, T., Fichtel, C., Hartman, R., Hunter, S., Thompson, D. J., Kanbach, G., Mayer-Hasselwander, H., von Montigny, C., Sommer, M., Lin, Y., Nolan, P., Michelson, P., Kniffen, D., Mattox, J., Schneid, E., Boer, M., and Niel, M. (1994). Detection of a γ-ray burst of very long duration and very high energy. *Nature*, 372(6507):652–654. 2

Iwakami, W., Kotake, K., Ohnishi, N., Yamada, S., and Sawada, K. (2009). Effects of Rotation on Standing Accretion Shock Instability in Nonlinear Phase for Core-Collapse Supernovae. *The Astrophysical Journal*, 700(1):232–242. 153

James, R. (1977). The solution of poisson's equation for isolated source distributions. *Journal of Computational Physics*, 25(2):71–93. 67

Janka, H., Langanke, K., Marek, A., Martinezpinedo, G., and Muller, B. (2007). Theory of core-collapse supernovae. *Physics Reports*, 442(1-6):38–74. 8

Janka, H. and Müller, E. (1996). Neutrino heating, convection, and the mechanism of Type-II supernova explosions. *Astronomy and Astrophysics*, 306:167. 8

Khokhlov, A. (1998). Fully Threaded Tree Algorithms for Adaptive Refinement Fluid Dynamics Simulations. *Journal of Computational Physics*, 143(2):519–543. xvi, 102, 103, 106, 107

Klebesadel, R. W., Strong, I. B., and Olson, R. A. (1973). Observations of Gamma-Ray Bursts of Cosmic Origin. *The Astrophysical Journal*, 182:L85. 2

Kotake, K., Sawai, H., Yamada, S., and Sato, K. (2004). Magnetorotational Effects on Anisotropic Neutrino Emission and Convection in CoreCollapse Supernovae. *The Astrophysical Journal*, 608(1):391–404. 149

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*, 413:L101. 2, 3

Kulkarni, S. R., Djorgovski, S. G., Odewahn, S. C., Bloom, J. S., Gal, R. R., Koresko, C. D., and Harrison, F. A. (1999). The afterglow , redshift and extreme energetics of the g-ray burst of. *Nature*, 398(APRIL):389–394. 3

Kurganov, A., Noelle, S., and Petrova, G. (2001). Semidiscrete Central-Upwind Schemes for Hyperbolic Conservation Laws and Hamilton–Jacobi Equations. *SIAM Journal on Scientific Computing*, 23(3):707. 16

Kurganov, A. and Tadmor, E. (2000). New High-Resolution Central Schemes for Nonlinear Conservation Laws and ConvectionDiffusion Equations. *Journal of Computational Physics*, 160(1):241–282. 16, 20

Lai, D. (1994). Resonant oscillations and tidal heating in coalescing binary neutron stars. *Royal Astronomical Society, Monthly Notices*, 270:611–629. 132

Lamb, D. Q. (1995). The Distance Scale to Gamma-Ray Bursts. *Publications of the Astronomical Society of the Pacific*, 107:1152. 2

Landau, L. D. and Lifshitz, E. M. (1959). *Fluid mechanics*, volume 6. Pergamon Press. 18, 22

Landau, L. D., Pitaevskii, L. P., and Lifshitz, E. (1984). *Electrodynamics of Continuous Media*. Butterworth-Heinemann, 2 edition. 24

Laney, C. B. (1998). *Computational gasdynamics*. Cambridge University Press. 14

Lattimer, J. (2000). Nuclear matter and its role in supernovae, neutron stars and compact object binary mergers. *Physics Reports*, 333-334(1):121–146. 144

Lattimer, J. and Swesty, D. F. (1991). A generalized equation of state for hot, dense matter. *Nuclear Physics, Section A*, 535(2):331–376. 31

Lattimer, J. M. and Prakash, M. (2004). The physics of neutron stars. *Science (New York, N.Y.)*, 304(5670):536–42. 9

Leblanc, J. M. and Wilson, J. R. (1970). A Numerical Example of the Collapse of a Rotating Magnetized Star. *The Astrophysical Journal*, 161:541. 179

Lee, W. H. and Ramirez-Ruiz, E. (2007). The progenitors of short gamma-ray bursts. *New Journal of Physics*, 9(1):17–17. 4

Lee, W. H., RamirezRuiz, E., and Page, D. (2005). Dynamical Evolution of Neutrinocooled Accretion Disks: Detailed Microphysics, Leptondriven Convection, and Global Energetics. *The Astrophysical Journal*, 632(1):421–437. 12

Leonard, D. C., Filippenko, A. V., Ganeshalingam, M., Serduke, F. J. D., Li, W., Swift, B. J., Gal-Yam, A., Foley, R. J., Fox, D. B., Park, S., Hoffman, J. L., and Wong, D. S. (2006). A non-spherical core in the explosion of supernova SN 2004dj. *Nature*, 440(7083):505–7. 8

LeVeque, R. J., Steiner, O., und Astronomie, S. G. f. A., and Gautschy, A. (1998). *Computational methods for astrophysical fluid flow, Volume 199*. Springer. 14

Lindblom, L. (1992). Determining the nuclear equation of state from neutron-star masses and radii. *The Astrophysical Journal*, 398:569. 147

Liu, Y. T., Shapiro, S. L., Etienne, Z. B., and Taniguchi, K. (2008). General relativistic simulations of magnetized binary neutron star mergers. *Physical Review D*, 78(2):024012. 12

Londrillo, P. and Del Zanna, L. (2000). HighOrder Upwind Schemes for Multidimensional Magnetohydrodynamics. *The Astrophysical Journal*, 530(1):508–524. 16

Londrillo, P. and Del Zanna, L. (2004). On the divergence-free condition in Godunov-type schemes for ideal magnetohydrodynamics: the upwind constrained transport method. *Journal of Computational Physics*, 195(1):17–48. 17, 57

Lucas-Serrano, A., Font, J. A., Ibez, J. M., and Mart, J. M. (2004). Assessment of a high-resolution central scheme for the solution of the relativistic hydrodynamics equations. *Astronomy and Astrophysics*, 428(2):703–715. 16

Luck, H. (1997). The GEO600 Project. *Classical and Quantum Gravity*, 14:246–246. 5, 129

MacFadyen, A. I. and Woosley, S. E. (1999). Collapsars: GammaRay Bursts and Explosions in Failed Supernovae. *The Astrophysical Journal*, 524(1):262–289. 3

MacNeice, P. (2000). PARAMESH: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126(3):330–354. 101

Mandel, I., Brown, D. A., Gair, J. R., and Miller, M. C. (2008). Rates and Characteristics of Intermediate Mass Ratio Inspirals Detectable by Advanced LIGO. *The Astrophysical Journal*, 681(2):1431–1447. 5

Marek, A. and Janka, H.-T. (2009). Delayed Neutrino-Driven Supernova Explosions Aided by the Standing Accretion-Shock Instability. *The Astrophysical Journal*, 694(1):664–696. 149

Martí, J. M. and Müller, E. (1994). The analytical solution of the Riemann problem in relativistic hydrodynamics. *Journal of Fluid Mechanics*, 258:317. 46

Martí, J. M. and Müller, E. (2003). Numerical Hydrodynamics in Special Relativity. *Living Rev. Relativity*, 6(7). xi, 46, 47, 49, 51

Matheson, T., Garnavich, P. M., Stanek, K. Z., Bersier, D., Holland, S. T., Krisciunas, K., Caldwell, N., Berlind, P., Bloom, J. S., Bolte, M., Bonanos, A. Z., Brown, M. J. I., Brown, W. R., Calkins, M. L., Challis, P., Chornock, R., Echevarria, L., Eisenstein, D. J., Everett, M. E., Filippenko, A. V., Flint, K., Foley, R. J., Freedman, D. L., Hamuy, M., Harding, P., Hathi, N. P., Hicken, M., Hoopes, C., Impey, C., Jannuzi, B. T., Jansen, R. A., Jha, S., Kaluzny, J., Kannappan, S., Kirshner, R. P., Latham, D. W., Lee, J. C., Leonard, D. C., Li, W., Luhman, K. L., Martini, P., Mathis, H., Maza, J., Megeath, S. T., Miller, L. R., Minniti, D., Olszewski, E. W., Papenkova, M., Phillips, M. M., Pindor, B., Sasselov, D. D., Schild, R., Schweiker, H., Spahr, T., ThomasOsip, J., Thompson, I., Weisz, D., Windhorst, R., and Zaritsky, D. (2003). Photometry and Spectroscopy of GRB 030329 and Its Associated Supernova 2003dh: The First Two Months. *The Astrophysical Journal*, 599(1):394–407. 3

May, M. and White, R. (1966). Hydrodynamic Calculations of General-Relativistic Collapse. *Physical Review*, 141(4):1232–1241. 182

Mazets, E. P., Golenetskii, S. V., Ilyinskii, V. N., Guryan, Y. A., Aptekar, R. L., Panov, V. N., Sokolov, I. A., Sokolova, Z. Y., and Kharitonova, T. V. (1982). Cosmic gamma-ray burst spectroscopy. *Astrophysics and Space Science*, 82(2):261–282. 2

McCray, R. (1993). Supernova 1987A Revisited. *Annual Review of Astronomy and Astrophysics*, 31(1):175–216. 8

Mee, A. J. and Brandenburg, A. (2006). Turbulence from localized random expansion waves. *Monthly Notices of the Royal Astronomical Society*, 370:415–419. 160

Meegan, C. A., Fishman, G. J., Wilson, R. B., Paciesas, W. S., Pendleton, G. N., Horack, J. M., Brock, M. N., and Kouveliotou, C. (1992). Spatial distribution of γ-ray bursts observed by BATSE. *Nature*, 355(6356):143–145. 2

Melton, J. E., Berger, M. J., Aftosmis, M. J., and Wong, M. D. (1995). 3D Applications of a Cartesian Grid Euler Method. *AIAA Paper*, 95-0853. 107

Meszaros, P. (2002). Theories of Gamma-Ray Bursts. *Annual Review of Astronomy and Astrophysics*, 40(1):137–169. 4

Meszaros, P. and Rees, M. J. (1992). Tidal heating and mass loss in neutron star binaries - Implications for gamma-ray burst models. *The Astrophysical Journal*, 397:570. 4

Mészáros, P. and Rees, M. J. (1997). Poynting Jets from Black Holes and Cosmological Gamma-Ray Bursts. *The Astrophysical Journal*, 482(1):L29–L32. 4

Metzger, M. R., Djorgovski, S. G., Kulkarni, S. R., Steidel, C. C., Adelberger, K. L., Frail, D. A., Costa, E., and Frontera, F. (1997). Spectral constraints on the redshift of the optical counterpart to the big gamma-ray burst of 8 May 1997. *Nature*, 387(JUNE):878–880. 3

Mezzacappa, A. (2005). ASCERTAINING THE CORE COLLAPSE SUPERNOVA MECHANISM: The State of the Art and the Road Ahead. *Annual Review of Nuclear and Particle Science*, 55(1):467–515. 8

Mezzacappa, A., Calder, A. C., Bruenn, S. W., Blondin, J. M., Guidry, M. W., Strayer, M. R., and Umar, A. S. (1998). An Investigation of Neutrinodriven Convection and the Core Collapse Supernova Mechanism Using Multigroup Neutrino Transport. *The Astrophysical Journal*, 495(2):911–926. 8

Mikami, H., Sato, Y., Matsumoto, T., and Hanawa, T. (2008). Threedimensional Magnetohydrodynamical Simulations of a CoreCollapse Supernova. *The Astrophysical Journal*, 683(1):357–374. 149

Miller, M., Gressman, P., and Suen, W.-M. (2004). Towards a realistic neutron star binary inspiral: Initial data and multiple orbit evolution in full general relativity. *Physical Review D*, 69(6):1–21. 11

Misner, C. W., Thorne, K. S., and Wheeler, J. A. (1973). *Gravitation*. W. H. Freeman. 129, 132

Monaghan, J. J. (1992). Smoothed Particle Hydrodynamics. *Annual Review of Astronomy and Astrophysics*, 30(1):543–574. 14

Morrison, I. A., Baumgarte, T. W., and Shapiro, S. L. (2004). Effect of Differential Rotation on the Maximum Mass of Neutron Stars: Realistic Nuclear Equations of State. *The Astrophysical Journal*, 610(2):941–947. 144

Morton, G. M. (1966). A Computer Oriented Geodetic Data Base and a New Technique in File Sequencing. 104

MPI-Forum (2010). Message Passing Interface (MPI) Forum Home Page: http://www.mpi-forum.org/. 32, 72

Myong, R. and Roe, P. L. (1998). On Godunov-Type Schemes for Magnetohydrodynamics 1. A Model System. *Journal of Computational Physics*, 147(2):545–567. 53

Nakamura, T., Oohara, K., and Kojima, Y. (1987). General Relativistic Collapse to Black Holes and Gravitational Waves from Black Holes. *Progress of Theoretical Physics Supplement*, (90):1–218. 182

Nakar, E., GalYam, A., and Fox, D. B. (2006). The Local Rate and the Progenitor Lifetimes of ShortHard GammaRay Bursts: Synthesis and Predictions for the Laser Interferometer GravitationalWave Observatory. *The Astrophysical Journal*, 650(1):281–290. 5

Noble, S. C., Gammie, C. F., McKinney, J. C., and Del Zanna, L. (2006). Primitive Variable Solvers for Conservative General Relativistic Magnetohydrodynamics. *The Astrophysical Journal*, 641(1):626–637. 29, 181

Obergaulinger, M., Aloy, M. A., and Müller, E. (2006). Axisymmetric simulations of magneto-rotational core collapse: dynamics and gravitational wave signal. *Astronomy and Astrophysics*, 450(3):1107–1134. 149

Oechslin, R. and Janka, H.-T. (2006). Torus formation in neutron star mergers and well-localized short gamma-ray bursts. *Monthly Notices of the Royal Astronomical Society*, 368(4):1489–1499. 11

Oechslin, R., Janka, H.-T., and Marek, A. (2007). Relativistic neutron star merger simulations with non-zero temperature equations of state. *Astronomy and Astrophysics*, 467(2):395–409. 11

Oechslin, R., Rosswog, S., and Thielemann, F.-K. (2002). Conformally flat smoothed particle hydrodynamics application to neutron star mergers. *Physical Review D*, 65(10). 11

Ott, C. D. (2009). Probing the core-collapse supernova mechanism with gravitational waves. *Classical and Quantum Gravity*, 26(20):204015. 6

Ott, E. (1998). Chaotic flows and kinematic magnetic dynamos: A tutorial review. *Physics of Plasmas*, 5(5):1636. 172

Paczynski, B. (1995). How Far Away Are Gamma-Ray Bursters? *Publications of the Astronomical Society of the Pacific*, 107:1167. 2

Paczynski, B. and Rhoads, J. E. (1993). Radio Transients from Gamma-Ray Bursters. *The Astrophysical Journal*, 418:L5. 3

Piran, T. (2005). The physics of gamma-ray bursts. *Reviews of Modern Physics*, 76(4):1143–1210. 4

Pons, J., Martí, J. M., and Müller, E. (2000). The exact solution of the Riemann problem with non-zero tangential velocities in relativistic hydrodynamics. *Journal of Fluid Mechanics*, 422:125–139. 49

Prakash, M., Lattimer, J., Pons, J., Steiner, A., and Reddy, S. (2001). *Evolution of a Neutron Star from Its Birth to Old Age*, pages 364–423. Springer, Berlin / Heidelberg. 9

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. (1986). *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press. 24, 30

Price, D. J. and Rosswog, S. (2006). Producing ultrastrong magnetic fields in neutron star mergers. *Science (New York, N.Y.)*, 312(5774):719–22. 11

Prochaska, J. X., Bloom, J. S., Chen, H., Foley, R. J., Perley, D. A., RamirezRuiz, E., Granot, J., Lee, W. H., Pooley, D., Alatalo, K., Hurley, K., Cooper, M. C., Dupree, A. K., Gerke, B. F., Hansen, B. M. S., Kalirai, J. S., Newman, J. A., Rich, R. M., Richer, H., Stanford, S. A., Stern, D., and van Breugel, W. J. M. (2006). The Galaxy Hosts and LargeScale Environments of ShortHard GammaRay Bursts. *The Astrophysical Journal*, 642(2):989–994. 4

Quilis, V. (2000). Gone with the Wind: The Origin of S0 Galaxies in Clusters. *Science*, 288(5471):1617–1620. 14

Ricker, P. M. (2008). A Direct Multigrid Poisson Solver for OctTree Adaptive Meshes. *The Astrophysical Journal Supplement Series*, 176(1):293–300. 82, 117

Roe, P. L. (1986). Characteristic-Based Schemes for the Euler Equations. *Annual Review of Fluid Mechanics*, 18(1):337–365. 20

Rosswog, S. (2003). Neutron Star Binaries as Central Engines of GRBs. 662(14 April 2003):220–222. 4

Rosswog, S. and Davies, M. B. (2002). High-resolution calculations of merging neutron stars - I. Model description and hydrodynamic evolution. *Monthly Notices of the Royal Astronomical Society*, 334(3):481–497. 11

Rosswog, S. and Liebendörfer, M. (2003). High-resolution calculations of merging neutron stars - II. Neutrino emission. *Monthly Notices of the Royal Astronomical Society*, 342(3):673–689. 11

Ruffert, M., Janka, H., and Schafer, G. (1996). Coalescing neutron stars A step towards physical models – I. Hydrodynamic evolution and gravitational-wave emission. *Astronomy & Astrophysics*, 311:532–566. 11, 129

Ruffert, M. and Janka, H.-T. (2001a). Coalescing neutron stars - A step towards physical models – II. Neutrino emission, neutron tori, and gamma-ray bursts. *Astronomy and Astrophysics*, 380(02):544–577. 11

Ruffert, M. and Janka, H.-T. (2001b). Coalescing neutron stars - A step towards physical models – III. Improved numerics and different neutron star masses and spins. *Astronomy and Astrophysics*, 380(02):544–577. 11

Ryu, D., Jones, T. W., and Frank, A. (2000). The Magnetohydrodynamic KelvinHelmholtz Instability: A Threedimensional Study of Nonlinear Evolution. *The Astrophysical Journal*, 545(1):475–493. 175

Sadowski, A., Belczynski, K., Bulik, T., Ivanova, N., Rasio, F. A., and OShaughnessy, R. (2008). The Total Merger Rate of Compact Object Binaries in the Local Universe. *The Astrophysical Journal*, 676(2):1162–1169. 5

Scheck, L., Janka, H.-T., Foglizzo, T., and Kifonidis, K. (2008). Multidimensional supernova simulations with approximative neutrino transport. *Astronomy and Astrophysics*, 477(3):931–952. 149

Scheck, L., Kifonidis, K., Janka, H.-T., and Müller, E. (2006). Multidimensional supernova simulations with approximative neutrino transport. *Astronomy and Astrophysics*, 457(3):963–986. 149

Schneid, E. J., Bertsch, D. L., Dingus, B. L., Fichtel, C. E., Hartman, R. C., Hunter, S. D., Kanbach, G., Kniffen, D. A., Lin, Y. C., Mattox, J. R., Mayer-Hasselwander, H. A., Michelson, P. F., von Montigny, C., Nolan, P. L., Sreekumar, P., and Thompson, D. J. (1995). EGRET Measurements of Energetic Gamma Rays from the Gamma-Ray Bursts of 1992 June 22 and 1994 March 1. *The Astrophysical Journal*, 453:95. 2

Setiawan, S., Ruffert, M., and Janka, H.-T. (2006). Three-dimensional simulations of non-stationary accretion by remnant black holes of compact object mergers. *Astronomy and Astrophysics*, 458(2):553–567. 12

Shen, H., Toki, H., Oyamatsu, K., and Sumiyoshi, K. (1998). Relativistic equation of state of nuclear matter for supernova and neutron star. *Nuclear Physics A*, 637(3):435–450. 31, 32

Shibata, M. (1995). Evolution of three-dimensional gravitational waves: Harmonic slicing case. *Physical Review D*, 52(10):5428–5444. 182

Shibata, M., Duez, M., Liu, Y., Shapiro, S., and Stephens, B. (2006a). Magnetized Hypermassive Neutron-Star Collapse: A Central Engine for Short Gamma-Ray Bursts. *Physical Review Letters*, 96(3). 12

Shibata, M., Liu, Y., Shapiro, S., and Stephens, B. (2006b). Magnetorotational collapse of massive stellar cores to neutron stars: Simulations in full general relativity. *Physical Review D*, 74(10). 149

Shibata, M., Nakamura, T., and OOHARA, K. (1992). Coalescence of Spinning Binary Neutron Stars of Equal Mass. *Progress of Theoretical Physics*, 88(6):1079–1095. 129

Shibata, M. and Taniguchi, K. (2006). Merger of binary neutron stars to a black hole: Disk mass, short gamma-ray bursts, and quasinormal mode ringing. *Physical Review D*, 73(6). 12

Shibata, M., Taniguchi, K., and Uryu, K. (2003). Merger of binary neutron stars of unequal mass in full general relativity. *Physical Review D*, 68(8). 11

Shibata, M., Taniguchi, K., and Uryu, K. (2005). Merger of binary neutron stars with realistic equations of state in full general relativity. *Physical Review D*, 71(8). 12

Shibata, M. and Uryu, K. (2002). Gravitational Waves from the Merger of Binary Neutron Stars in a Fully General Relativistic Simulation. *Progress of Theoretical Physics*, 107(2):265–303. 11

Shu, C. and Osher, S. (1988). Efficient implementation of essentially non-oscillatory shock-capturing schemes. *Journal of Computational Physics*, 77(2):439–471. 16, 22

Sod, G. (1978). A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *Journal of Computational Physics*, 27(1):1–31. 38

Soderberg, A. M., Berger, E., Kasliwal, M., Frail, D. A., Price, P. A., Schmidt, B. P., Kulkarni, S. R., Fox, D. B., Cenko, S. B., GalYam, A., Nakar, E., and Roth, K. C. (2006). The Afterglow, Energetics, and Host Galaxy of the ShortHard GammaRay Burst 051221a. *The Astrophysical Journal*, 650(1):261–271. 4

Stone, J. M. and Norman, M. L. (1992). ZEUS-2D: A Radiation Magnetohydrodynamics Code for Astrophysical Flows in Two Space Dimensions: I. The Hydrodynamic Algorithms and Tests. *The Astrophysical Journal Supplement Series*, 80:753–790. 88

Suzuki, T. K., Sumiyoshi, K., and Yamada, S. (2008). Alfvén WaveDriven Supernova Explosion. *The Astrophysical Journal*, 678(2):1200–1206. 149

Swarztrauber, P. N. (490). The Methods of Cyclic Reduction, Fourier Analysis and the FACR Algorithm for the Discrete Solution of Poisson's Equation on a Rectangle. *SIAM Review*, 19(3):1977. 67

Symbalisty, E. M. D. (1984). Magnetorotational iron core collapse. *The Astrophysical Journal*, 285:729. 179

Taam, R. E. and Sandquist, E. L. (2000). Common Envelope Evolution of Massive Binary Stars. *Annual Review of Astronomy and Astrophysics*, 38(1):113–141. 10

Tagoshi, H., Kanda, N., Tanaka, T., Tatsumi, D., Telada, S., Ando, M., Arai, K., Araya, A., Asada, H., Barton, M., Fujimoto, M.-K., Fukushima, M., Futamase, T., Heinzel, G., Horikoshi, G., Ishizuka, H., Kamikubota, N., Kawabe, K., Kawamura, S., Kawashima, N., Kojima, Y., Kozai, Y., Kuroda, K., Matsuda, N., Matsumura, S., Miki, S., Mio, N., Miyakawa, O., Miyama, S., Miyoki, S., Mizuno, E., Moriwaki, S., Musha, M., Nagano, S., Nakagawa, K., Nakamura, T., Nakao, K.-i., Numata, K., Ogawa, Y., Ohashi, M., Ohishi, N., Okutomi, A., Oohara, K.-i., Otsuka, S., Saito, Y., Sasaki, M., Sato, S., Sekiya, A., Shibata, M., Shirakata, K., Somiya, K., Suzuki, T., Takahashi, R., Takamori, A., Taniguchi, S., Tochikubo, K., Tomaru, T., Tsubono, K., Tsuda, N., Uchiyama, T., Ueda, A., Ueda, K.-i., Ueda, K., Waseda, K., Watanabe, Y., Yakura, H., Yamamoto, K., and Yamazaki, T. (2001). First search for gravitational waves from inspiraling compact binaries using TAMA300 data. *Physical Review D*, 63(6). 5

Takahashi, R. and The TAMA Collaboration (2004). Status of TAMA300. *Classical and Quantum Gravity*, 21(5):S403–S408. 4

Takiwaki, T., Kotake, K., and Sato, K. (2009). Special Relativistic Simulations of Magnetically Dominated Jets in Collapsing Massive Stars. *The Astrophysical Journal*, 691(2):1360–1379. 149

Taylor, J. H. (1994). Binary pulsars and relativistic gravity. *Reviews of Modern Physics*, 66(3):711–719. 4, 129

Taylor, J. H. and Weisberg, J. M. (1989). Further experimental tests of relativistic gravity using the binary pulsar PSR 1913 + 16. *The Astrophysical Journal*, 345:434. 4

Thompson, C. and Duncan, R. C. (1993). Neutron star dynamos and the origins of pulsar magnetism. *The Astrophysical Journal*, 408:194. 178

Thompson, T. A., Quataert, E., and Burrows, A. (2005). Viscosity and Rotation in CoreCollapse Supernovae. *The Astrophysical Journal*, 620(2):861–877. 8, 149

Thorne, K. (1989). *Three hundred years of gravitation*. Cambridge University Press. 144

Tóth, G. (2000). The B=0 Constraint in Shock-Capturing Magnetohydrodynamics Codes. *Journal of Computational Physics*, 161(2):605–652. 57

van Den Heuvel, E. (2006). Evolution of X-ray binaries: Achievements and advances. *Advances in Space Research*, 38(12):2667–2672. 10

van Leer, B. (1977). Towards the ultimate conservative difference scheme III. Upstream-centered finite-difference schemes for ideal compressible flow. *Journal of Computational Physics*, 23(3):263–275. 20

van Paradijs, J., Groot, P. J., Galama, T., Kouveliotou, C., Strom, R. G., Telting, J., Rutten, R. G. M., Fishman, G. J., Meegan, C. A., Pettini, M., Tanvir, N., Bloom, J., Pedersen, H., Nø rdgaard Nielsen, H. U., Linden-Vø rnle, M., Melnick, J., van Der Steene, G., Bremer, M., Naber, R., Heise, J., in't Zand, J., Costa, E., Feroci, M., Piro, L., Frontera, F., Zavattini, G., Nicastro, L., Palazzi, E., Bennet, K., Hanlon, L., and Parmar, A. (1997). Transient optical emission from the error box of the γ-ray burst of 28 February 1997. *Nature*, 386(6626):686–689. 3

van Paradijs, J., Kouveliotou, C., and Wijers, R. A. M. J. (2000). Gamma-Ray Burst Afterglows. *Annual Review of Astronomy and Astrophysics*, 38(1):379–425. 3

Voss, R. and Tauris, T. M. (2003). Galactic distribution of merging neutron stars and black holes - prospects for short gamma-ray burst progenitors and LIGO/VIRGO. *Monthly Notices of the Royal Astronomical Society*, 342(4):1169–1184. 10

Wang, L., Howell, D. A., Hoflich, P., and Wheeler, J. C. (2001a). Bipolar Supernova Explosions. *The Astrophysical Journal*, 550(2):1030–1035. 8

Wang, L., Howell, D. A., Hoflich, P., and Wheeler, J. C. (2001b). Bipolar Supernova Explosions. *The Astrophysical Journal*, 550(2):1030–1035. 149

Weiler, K. W., Panagia, N., Montes, M. J., and Sramek, R. A. (2002). Radio Emission from Supernovae and Gamma-Ray Bursters. *Annual Review of Astronomy and Astrophysics*, 40(1):387–438. 3

Wheeler, J. C. and Akiyama, S. (2006). *Magnetic Fields in Core Collapse Supernovae: Possibilities and Gaps*, page 156. World Scientific Publishing Company. 9

Willke, B. (2007). GEO600: status and plans. *Classical and Quantum Gravity*, 24(19):S389–S397. 5

Willke, B., Aufmuth, P., Aulbert, C., Babak, S., Balasubramanian, R., Barr, B. W., Berukoff, S., Cagnoli, G., Cantley, C. A., Casey, M. M., Chelkowski, S., Churches, D., Colacino, C. N., Crooks, D. R. M., Cutler, C., Danzmann, K., Davies, R., Dupuis, R. J., Elliffe, E., Fallnich, C., Freise, A., Goß ler, S., Grant, A., Grote, H., Grunewald, S., Harms, J., Heinzel, G., Heng, I. S., Hepstonstall, A., Heurs, M., Hewitson, M., Hild, S., Hough, J., Ingley, R., Itoh, Y., Jennrich, O., Jones, R., Hutter, S. H., Kawabe, K., Killow, C., Kötter, K., Krishnan, B., Leonhardt, V., Lück, H., Machenschalk, B., Malec, M., Mercer, R. A., Messenger, C., Mohanty, S., Mossavi, K., Mukherjee, S., Nagano, S., Newton, G. P., Papa, M. A., Perreur-Lloyd, M., Pitkin, M., Plissi, M. V., Quetschke, V., Re, V., Reid, S., Ribichini, L., Robertson, D. I., Robertson, N. A., Rowan, S., Rüdiger, A., Sathyaprakash, B. S., Schilling, R., Schnabel, R., Schutz, B. F., Seifert, F., Sintes, A. M., Smith, J. R., Sneddon, P. H., Strain, K. A., Taylor, I., Torrie, C. I., Ungarelli, C., Vecchio, A., Ward, H., Weiland, U., Welling, H., Williams, P., Winkler, W., Woan, G., and Zawischa, I. (2004). Status of GEO 600. *Classical and Quantum Gravity*, 21(5):S417–S423. 4

Wilson, J., Mathews, G., and Marronetti, P. (1996). Relativistic numerical model for close neutron-star binaries. *Physical review D: Particles and fields*, 54(2):1317–1331. 11

Wolszczan, A. (1991). A nearby 37.9-ms radio pulsar in a relativistic binary system. *Nature*, 350(6320):688–690. 10, 128

Woosley, S. and Bloom, J. (2006). The Supernova-Gamma-Ray Burst Connection. *Annual Review of Astronomy and Astrophysics*, 44(1):507–556. 3, 149

Woosley, S. and Janka, H. T. (2006). The Physics of Core-Collapse Supernovae. page 17. 9

Woosley, S. E. (1993). Gamma-ray bursts from stellar mass accretion disks around black holes. *The Astrophysical Journal*, 405:273. 3, 4

Yakunin, K. N., Marronetti, P., Mezzacappa, A., Bruenn, S. W., Lee, C.-T., Chertkow, M. A., Hix, W. R., Blondin, J. M., Lentz, E. J., Messer, O. E. B., and Yoshida, S. (2010). Gravitational Waves from Core Collapse Supernovae. page 10. 6

Zhang, W. and MacFadyen, A. I. (2006). RAM: A Relativistic Adaptive Mesh Refinement Hydrodynamics Code. *The Astrophysical Journal Supplement Series*, 164(1):255–279. 46, 51

Zhuge, X., Centrella, J., and McMillan, S. (1994). Gravitational radiation from coalescing binary neutron stars. *Physical Review D*, 50(10):6247–6261. 145

Zhuge, X., Centrella, J. M., and McMillan, S. L. W. (1996). Gravitational Radiation from the Coalescence of Binary Neutron Stars: Effects Due to the Equation of State, Spin, and Mass Ratio. *Physics*, page 39. 147

Zou, W. Z., Hobbs, G., Wang, N., Manchester, R. N., Wu, X. J., and Wang, H. X. (2005). Timing measurements and proper motions of 74 pulsars using the Nanshan radio telescope. *Monthly Notices of the Royal Astronomical Society*, 362(4):1189–1198. 8

Vita

Vita goes here...